

IMPROVING PERIPHERAL IV CATHETERIZATION THROUGH ROBOTICS:
FROM SIMPLE ASSISTIVE DEVICES
TO A FULLY-AUTONOMOUS SYSTEM

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Reuben D. Brewer

June 2015

© 2015 by Reuben Daniel Brewer. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/gh215qd3853>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

J Salisbury, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mark Cutkosky, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Thomas Kenny

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Curt Salisbury

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Intravenous (IV) catheterization is a medical procedure wherein a metal needle is used to puncture a vein and introduce a plastic catheter that remains in the vein for the delivery of medicinal fluids. Nearly 250 million IV insertions take place in the United States annually, and 28% of those insertions fail on the first attempt in normal adults, with appreciably higher failure rates in children. Failed insertions commonly cause bruising and pain but can also lead to more serious complications, including long-term nerve damage, local tissue damage, and sclerosis of the veins. There are two main difficulties in IV insertion: locating the veins, which are often difficult to see or feel, and moving the needle precisely into a vein that is a small, movable target.

We present our robotic IV insertion system, HaemoBot, as a possible solution that offers enhanced vein-detection sensory abilities and precise movement of the needle. HaemoBot can be teleoperated by a human to increase his or her sensory and motor abilities without ceding intelligent control over the insertion, or it can insert autonomously. We envision our system being used to treat people in remote or hostile locations where a human practitioner could not be physically present or in hospitals to increase the success and through-put of practitioners.

The highest-level, most far-reaching contribution of our work was to answer affirmatively the question of whether a robot can insert an IV autonomously by demonstrating that our system can 1. autonomously locate and calculate a 6D needle pose to insert on a venous bifurcation and 2. autonomously insert an IV into a realistic phantom of a venous bifurcation by detecting the force penetration-event “pop”.

We performed a clinical experiment with instrumented IV needles to record the process and mechanics of actual IV insertions, and this allowed us to determine a complete list of human IV insertion steps that could be translated into robot steps, determine the normative force, speed, and angles at key moments in IV insertions, and prove that a human can pick out the needle and catheter penetration events into the skin and veins from the force data. We further developed three novel explanations for why bifurcations can be easier insertion targets than straight sections of vein.

We proved that a robot can locate a vein and compute a 6D needle pose autonomously. To achieve this, we developed novel Near Infrared (NIR) transillumination and stereo camera hardware capable of generating uniquely accurate, high-resolution, high-contrast, and three-dimensional hand/vein point-clouds and developed novel computer-vision algorithms for locating, characterizing, and tracking venous bifurcations in a 2D image and then generating a 6D needle pose in the 3D point-cloud.

We proved that a robot can autonomously perform all of the complicated 1-handed and 2-handed mechanical steps of IV insertion on realistic hand, arm, and vein geometries (not just rubber tubes) using similar sensory cues, like the force penetration-event “pop”, that we observed with human practitioners. During this process, we developed a patented method for shining a targeting-laser through an IV needle while maintaining sterility and a patent-pending design for a Remote Center of Motion that has many superior fabrication, installation, and performance properties over existing designs.

We distilled the best lessons from our work on HaemoBot into a simple assistive device that provides more-accurate, stereo vein-visualization than currently available, makes IV insertion easier by reducing it

from a 2-handed to 1-handed procedure, and theoretically allows for self-insertion. We designed our device to fill unmet clinical needs that we established during our testing of many different commercially-available vein-finding devices.

We addressed an unmet need for realistic but inexpensive hand/vein phantoms for testing HaemoBot. We developed novel methodologies and tools for fabricating an artificial hand/vein phantom that used the NIR vein images and point-clouds from our vision system to copy real patient anatomy almost exactly, including seamless, properly-sized, latex bifurcations that had a realistic force profile and a silicone hand that had accurate anatomy from the macrostructure down to the fingerprints.

Our work is unique from other research in robotic phlebotomy and IV insertion in that we are the first to:

1. Develop the necessary degrees-of-freedom for treating the needle and catheter as two, separate parts,
2. Design a system that allows for either fully-teleoperated or fully-autonomous control of all steps, including the simple ones (such as cleaning the insertion site),
3. Develop a complete step-by-step model of IV insertion and establish normative values for variables like force, insertion speed, and pitch angle at key skin and vein penetration-events,
4. Develop a 2-armed robot that can perform all of the motions that we observed human practitioners performing, including all of the preparatory, assistive, and actual needle-insertion steps,
5. Employ a mechanical Remote Center of Motion as a safer way of setting the needle's pose,
6. Develop a high-Degree-of-Freedom, NIR-transillumination system and control algorithm that are able to provide uniquely accurate, high-resolution, high-contrast, and 3D vein point-clouds,
7. Employ a higher-Degree-of-Freedom force/torque sensor with non-penetration-detection goals in mind,
8. Target venous bifurcations and to explain technically why this is a useful, safe strategy, and
9. Develop an insertion phantom that includes venous bifurcations.

Acknowledgements

- My wife, Kris Brewer. You made this happen; all of it. You deserve this degree as much as I do. I dedicate this dissertation, degree, and rest of my life to you.
- My East Coast and West Coast families. For encouraging me on the phone and in person at all hours of the day and night and picking up the slack with Ellie, Ronen, and Bruno.
- The technical geniuses who taught me 99% of what I know. You are the masters; I am your padawan. Words cannot express my gratitude for your friendship and tutelage.
 - Morgan Quigley, Ph.D.
 - Joan Savall, Ph.D.
 - Curt Salisbury, Ph.D.
 - Kevin Loewke, Ph.D.
 - Sonny Chan, Ph.D.
 - Francois Conti, Ph.D.
- Research partners
 - Tony Pratkanis (wrote all of the ROS Stereo Camera code)
 - Aubrey Shapero (re-wrote 2D vein-finding algorithm and extended it to 3D)
 - Fred Mihm, M.D., RJ Ramamurthi, M.D., and Sumit Singh, M.D. (clinical experiment)
 - Scott Delp, Ph.D. for providing the Motion Capture equipment and space.
- Advisor and Committee Members
 - Ken Salisbury, Ph.D.
 - Mark Cutkosky, Ph.D.
 - Curt Salisbury, Ph.D.
 - Thomas Kenny, Ph.D.
 - Allison Okamura, Ph.D.
 - Chris Contag, Ph.D.
- Funders
 - NSF Graduate Research Fellowship
 - NDSEG Graduate Fellowship
 - Stanford Graduate Fellowship
 - Timothy Francis Kennedy Memorial Scholarship, Stanford Dept. of Mechanical Engineering
 - Intuitive Surgical Grant
 - Department of Computer Science Teaching Grant
 - Anonymous Angel Funders
 - The Bank of Kris Brewer, Member FDIC

Contents

1	Introduction	1
1.1	Motivation	1
1.1.1	Explaining the Basic Steps of IV Insertion	2
1.2	IV Insertion Sites We Are Targeting	3
1.3	Related Work	5
1.4	Thesis Content and Contributions	9
2	Understanding IV Insertion	11
2.1	Motivation and Goals	11
2.2	Clinical Experiment Design	11
2.2.1	Overall Needle Instrumentation and Ergonomic Design	13
2.2.1.1	Comparison to Other Studies	15
2.2.2	IV Needle and Sensor Details	15
2.2.2.1	IV Needle Choice and Modifications	15
2.2.2.2	ATI Nano17 6-Axis Force/Torque Sensor	16
2.2.2.3	Flexiforce 1-Axis Force Sensor	16
2.2.2.4	Intravenous Blood Pressure Transducer	18
2.2.2.5	3-Axis Digital Accelerometer	18
2.2.2.6	Motion Capture System for Measuring the 6D Needle Motions	19
2.2.2.6.1	Motion Capture Camera and Marker Placement	20
2.2.2.6.2	Motion Capture System Calibration	20
2.2.2.6.3	Synchronization of Motion Capture Data With Other Sensors	21
2.2.2.6.4	Calculating Needle Pose From Marker Positions	21
2.2.2.6.5	Establishing the Hand Frame	22
2.2.2.7	Touch-Sensors for Manually-Coding Sections of the Insertion	24
2.2.2.8	Video Cameras	24
2.2.3	Experimental Procedure	24
2.3	Results	25
2.3.1	The 27 Steps of Human IV Insertion	26
2.3.2	Quantitative Results and “Key Discoveries”	27
2.3.2.1	How We Parsed the Data	28
2.3.2.2	Needle Forces, Torques	32
2.3.2.2.1	Comparison of Axial Force Results to Other Studies	32
2.3.2.3	Intravenous Blood Pressure	33
2.3.2.4	Needle Distance and Speed	35

2.3.2.5	Needle Angles	37
2.3.2.6	Accelerometer	37
2.3.2.7	Variable Dependence on Age, Sex	39
2.3.3	Observed Clinical Failure Modes and Recovery Techniques	40
2.3.3.1	Failure Mode: Catheter Stuck Outside Vein Wall	40
2.3.3.2	Failure Mode: Back-Wall Penetration	40
2.3.3.3	Failure Mode: Valves Causing Blockage, Back-Wall Penetration	41
2.3.3.4	Possible Failure Mode: Venospasm and Disappearing Veins	42
2.3.4	Technical Explanations for Conventional Wisdom on Bifurcations, Traction	42
2.4	Discussion and Thesis Contributions	44
3	Vein-Imaging Hardware and Software	45
3.1	Motivation and Overview	45
3.1.1	The Clinical Problem	45
3.1.2	Vein Sensing Modalities and Prior Work	47
3.1.3	Using Near Infrared (NIR) Light For Vein-Visualization	48
3.1.3.1	Explanation of Why Near Infrared Light Shows Veins	48
3.1.3.2	Transillumination vs Reflectance-Based Camera/Light Configurations	49
3.1.3.3	How Light Scattering in Transillumination Hides Arteries	50
3.1.3.4	NIR Light Shows Only The Vein's Interior	51
3.2	Computer Vision Hardware	52
3.2.1	NIR Stereo Camera	52
3.2.1.1	Commercially-Available 3D Depth Cameras and Why We Built Our Own	52
3.2.1.2	Required Depth-Resolution	53
3.2.1.2.1	Hand Vein Inner Diameter	53
3.2.1.2.2	Calculating the Depth-Resolution	55
3.2.1.3	Monocular Camera	55
3.2.1.4	Texture Projection	56
3.2.1.5	Stereo Camera Controller Electronics and Code	57
3.2.1.6	Verifying Correct 3D Geometry from Stereo Camera	58
3.2.2	NIR LED Grid	59
3.2.2.1	LED Grid Mechanical and Optical Design	60
3.2.2.2	Using Cross Polarization to Reduce Stray Light.	62
3.2.2.3	LED Grid Electrical Design	63
3.2.2.4	LED Grid Manual Brightness-Setting GUI	63
3.2.2.5	LED Grid Automated Brightness-Setting Algorithm	63
3.3	Computer Vision Software	66
3.3.1	2D Bifurcation-Detection Algorithm	66

3.3.1.1	Related Work	67
3.3.1.2	Image Enhancement	67
3.3.1.3	Identifying Wrist Veins	67
3.3.1.4	Tracking Along Veins and Detecting Bifurcations	69
3.3.1.5	Centering the Bifurcation Estimate	69
3.3.1.6	Error-Checking	71
3.3.1.7	Calculating the Insertion Point	72
3.3.1.8	Weighting Multiple Bifurcations	72
3.3.1.9	Tracking the Bifurcation During Hand/Vein Movement	72
3.3.1.10	Experimental Methods and Results	74
3.3.2	6D Needle Pose Generation Algorithm	75
3.4	Discussion and Thesis Contributions	77
4	Development of the HaemoBot Robotic IV System	78
4.1	Motivation and Overview	78
4.1.1	Unique Contribution 1: Performing an Insertion with Two Arms	78
4.1.2	Unique Contribution 2: Moving the Catheter Independently of the Needle	82
4.1.3	Unique Contribution 3: Implementing Inherent, Mechanical Safety	82
4.2	HaemoBot, the Final Design	82
4.2.1	Overall Design	82
4.2.2	NeedleBot	83
4.2.2.1	Translational XYZ	83
4.2.2.1.1	Z-Axis	84
4.2.2.1.1.1	Z-Axis, Spring Counterbalance System	84
4.2.2.2	Rotation, Yaw	86
4.2.2.3	Rotation, Pitch: Novel Remote Center of Motion (RCM) Design	88
4.2.2.3.1	HaemoBot RCM Operation Explained	89
4.2.2.3.1.1	Novel Wire Management System for RCM	91
4.2.2.3.2	HaemoBot RCM Calibration Explained	96
4.2.2.3.2.1	Measuring the IV Needles	96
4.2.2.3.2.2	Pitch Angle and Needle Length Calibration	96
4.2.2.3.2.3	Accuracy Analysis of the RCM	98
4.2.2.4	Needle Driver	99
4.2.2.4.1	Laser-Targeting Through the IV Needle	99
4.2.2.4.2	Needle-Holder	101
4.2.2.4.3	Roll DOF	101
4.2.2.4.4	Catheter Insertion DOF	103
4.2.2.4.4.1	Catheter Insertion DOF, Spline Details	106

4.2.2.4.4.2	Catheter Insertion DOF, Stringpot	106
4.2.2.4.4.3	Catheter Insertion DOF, Electrical Wiring Design	107
4.2.2.4.5	Catheter Grab-and-Release DOFs	108
4.2.2.4.5.1	Distal-Contact-Mechanism	108
4.2.2.4.5.2	Proximal-Drive-Assembly	110
4.2.2.4.6	Force/Torque Sensor	112
4.2.2.4.7	Accelerometer	114
4.2.2.4.8	Blood-Flash and First-Person-Shooter (FPS) Cameras	114
4.2.2.4.9	Needle Length Adjustment Mechanism	116
4.2.2.4.10	Electronics and Communication	118
4.2.3	AssistoBot	119
4.2.3.1	“AssistoBot-Arm-Box-Assembly”: Arm Restraint and Tourniquet	120
4.2.3.2	“AssistoBot-Hand-Plate-Assembly”: Hand Restraint	120
4.2.3.2.1	“Hand-Plate”	120
4.2.3.2.2	Finger Restraint Mechanism	121
4.2.3.3	“AssistoBot-Arm”: Applying Forces, Alcohol, and Air	122
4.2.3.3.1	Kinematics and Drive Details	122
4.2.3.3.1.1	Planar XY Mechanism	122
4.2.3.3.1.2	Final-Link Parallelism Mechanism	123
4.2.3.3.1.3	Z-Axis Mechanism	126
4.2.3.3.2	Foam-Tip	126
4.2.3.3.3	Force Sensor	126
4.2.3.3.4	Air and IPA Sprayers	127
4.2.3.3.4.1	IPA Sprayer	127
4.2.3.3.4.2	Air Sprayer	129
4.2.3.3.5	Overhead NIR Camera	129
4.2.3.3.6	Control Electronics and Communication	129
4.2.4	Interacting With the Robot	130
4.2.4.1	Robot Controller GUI	130
4.2.4.2	Teleoperation Inputs	130
4.2.5	Robot Base and Electronics	131
4.2.6	Controlling the Robot	134
4.2.6.1	Proof-of-Concept Autonomous Insertion	134
4.2.6.2	Proof-of-Concept Teleoperated Insertion	136
4.3	Discussion and Thesis Contributions	138

5	Development of a Simple Assistive Device for IV Insertion	139
5.1	Motivation and Overview	139
5.1.1	Goal 1: Reliable, Stereo NIR Vein-Visualization Without False-Vein-Detection	141
5.1.1.1	Comparison With Commercially-Available Vein-Finding Devices	141
5.1.1.1.1	AccuVein AV400	141
5.1.1.1.2	Christie Digital Flex, Luminetx VeinViewer	142
5.1.1.1.3	VueTek Veinsite	144
5.1.1.1.4	Evena Medical Eyes-On Glasses	144
5.1.1.1.5	InfraRed Imaging Systems (INRIS) VascularViewer	145
5.1.1.1.6	TransLite VeinLite	145
5.1.1.1.7	Philips Wee Sight	145
5.1.2	Goal 2: A “Helping Hand” to Replace the Practitioner’s Non-Dominant Hand	146
5.2	Final Design	147
5.2.1	NIR-Goggles	147
5.2.1.1	Lens Caps, Longpass-Filtering, and Compatibility with Room Lighting	148
5.2.2	LightWand	150
5.2.2.1	Light Source	150
5.2.2.1.1	Transillumination vs Reflectance-Based Imaging Configurations	150
5.2.2.1.2	Laser Choice	151
5.2.2.1.3	Eye and Skin Safety	151
5.2.2.1.3.1	Passive Safety Through Beam Divergence	151
5.2.2.1.3.2	Active Safety Through Mechanical Interlocks	152
5.2.2.1.4	Controlling the Laser’s Brightness Via Slider-Pot-Set PWM	153
5.2.2.2	Geometry and Ergonomic Design	153
5.2.2.3	3D-Printing the Main Structure	155
5.2.2.4	Silicone Sleeve	155
5.2.2.5	Silicone Cap	158
5.2.2.6	Device Mounting	158
5.2.2.7	Traction and Restraint Bar	158
5.2.2.8	Tourniquet	159
5.2.2.9	Tape Dispenser	159
5.2.2.10	Alcohol Dispenser	161
5.2.2.11	Battery	161
5.2.3	Cost Estimate	162
5.3	Discussion and Thesis Contributions	162

6	Development of Anatomically Correct, Hyper-Realistic Hand Vein Phantoms	163
6.1	Motivation and Overview	163
6.2	Potential Insertion Targets	163
6.2.1	Live Human Hand	163
6.2.2	Deceased Human Hand (Cadaver)	165
6.2.3	Live Animal Vein	165
6.2.4	Deceased Animal Vein	165
6.2.5	Artificial Human Hand/Vein (Phantom)	166
6.3	Fabrication of Our Custom Phantom	168
6.3.1	Hand Fabrication	168
6.3.2	Vein Fabrication	169
6.3.2.1	Vein Material	170
6.3.2.2	Core Material	171
6.3.2.3	Making the Vein-Core Geometry	171
6.3.2.3.1	Molding the Vein-Core From 3D Point-Cloud, NIR Vein-Data	171
6.3.2.3.2	Splicing/Tracing the Vein-Core Against a 2D, NIR Vein Image	172
6.3.2.4	Defeating the Drip: Obtaining Uniform Thickness Though Constant Rotation	174
6.3.2.5	Closing Vein Segment Ends and Filling With Artificial Blood	177
6.4	Discussion and Thesis Contributions	177
7	Conclusion, Contributions, and Future Work	178
	Bibliography	184

List of Tables

5.1 Comparison of Commercially-Available Vein-Finding Devices.	142
--	-----

List of Figures

1.1 Example Injuries from IV Insertion.	2
1.2 Details of IV Needle.	3
1.3 Most-Basic Steps of IV Insertion.	4
1.4 We Are Only Targeting Venous Bifurcations in The Hand.	5
2.1 Design Details of the “Nano-needle” and “Simple-needle” Instrumented IV Needles.	12
2.2 Our “Hook-And-Gap” Natural Ergonomics vs Eriksson-Study Unnatural Ergonomics.	14
2.3 Motion Capture and Experimental Setup.	19
2.4 Synchronization of Motion Capture Data With Other Sensors Via “Impulse Alignment”.	21
2.5 Calibrating the IV Needle’s Local Motion Capture Frame and Establishing the Hand Frame.	23
2.6 Improper Vein Closure During Needle Withdrawal Leads to Bleeding.	29
2.7 Using Video to Identify “Pops” and Prove Tissue Acceleration During “Pops”.	29
2.8 Results: Axial Insertion Force vs Time.	31
2.9 Results: Intravenous Blood Pressure vs Time.	34
2.10 Results: Needle Tip Trajectory.	34
2.11 Results: Needle Tip Velocity vs Time, Axial Insertion Force vs Needle Tip Velocity.	36
2.12 Results: Pitch Angle, Acceleration vs Time.	38
2.13 Technical Explanations for Choosing Bifurcations and Applying Skin Traction.	43
3.1 NIR Vein-Imaging System.	46
3.2 NIR Absorption of Blood, Transillumination vs Reflectance-Based Imaging Configurations.	48
3.3 Example of Light Scattering Hiding Deeper Structures/Arteries.	50
3.4 NIR Light Shows Venous Blood But Not Vein Walls.	51
3.5 Setting Stereo Depth-Resolution Based on Hand Vein Inner Diameter.	53
3.6 The Effects of Overhead, Fluorescent Room Lights on NIR Image/Point-Cloud Quality.	56
3.7 The Effects of Stereo Camera Laser Texture Projection.	57
3.8 Stereo Camera and LED Grid Controller Electronics.	58
3.9 Verifying the Accuracy of Point-Cloud Geometries.	59
3.10 Need For an LED Grid: Inadequate Lighting From a Single LED.	60
3.11 LED Subassembly Details.	61

3.12	Effects of Cross-Polarization on Stray Light.	62
3.13	LED Grid Software Control.	64
3.14	LED Grid, Example Point-Clouds Generated By The Automated Brightness-Setting Algorithm.	65
3.15	2D Bifurcation-Detection: Image Preprocessing and Vein Detection.	68
3.16	2D Bifurcation-Detection: First-Pass Bifurcation Detection, Clustering, and Centering Details.	70
3.17	2D Bifurcation-Detection: Error Checking and Calculating Insertion Point/Angle.	71
3.18	2D Bifurcation-Detection: Tracking with the Centering Algorithm.	72
3.19	2D Bifurcation-Detection: Examples of Successes and Failures.	73
3.20	6D Needle Pose Generation: Finding Neighbors in Point-Cloud and Fitting a Tangent Plane.	75
3.21	6D Needle Pose Generation: Results.	76
4.1	HaemoBot, Overall View, Picture Set 1.	80
4.2	HaemoBot, Overall View, Picture Set 2.	81
4.3	Z-Axis Details.	85
4.4	Yaw Details.	87
4.5	RCM, Earlier Designs/Prototypes.	87
4.6	RCM, Novel Design and Operation Explained.	90
4.7	RCM, Pitch Design Details.	92
4.8	RCM, Wire Management System, Picture Set 1.	94
4.9	RCM, Wire Management System, Picture Set 2.	95
4.10	RCM, Pitch Angle and Needle Length Calibration.	97
4.11	Laser-Targeting Through The IV Needle.	100
4.12	Needle Driver, Needle-Holder and Roll DOF Details.	102
4.13	Needle Driver, Catheter Insertion DOF.	104
4.14	Needle Driver, Catheter Insertion DOF, Design Details.	105
4.15	Needle Driver, Catheter Grab-and-Release DOFs, Demonstration.	109
4.16	Needle Driver, Catheter Grab-and-Release DOFs, Drive Details.	111
4.17	Needle Driver, Force/Torque Sensor.	113
4.18	Needle Driver, Blood-Flash and First-Person-Shooter (FPS) Cameras.	115
4.19	Needle Driver, Needle Length Adjustment Mechanism.	117
4.20	AssistoBot's Main Subassemblies	119
4.21	"AssistoBot-Hand-Plate-Assembly", Hand Restraint.	121
4.22	"AssistoBot-Arm": Design Details of the Z-Axis and Final-Link Parallelism Mechanism.	124
4.23	"AssistoBot-Arm": Design Details of the Planar XY Mechanism.	125
4.24	"AssistoBot-Arm", Head Design Details.	128
4.25	GUI and Teleoperated Control of the Robot.	131
4.26	Robot Base and Electronics.	132
4.27	Examples of "Pop" Detection Using Axial Insertion Force.	135

4.28	Proof-of-Concept Autonomous Insertion.	137
5.1	eye.V.ease, Two Types of Help: NIR Vision and “Helping Hand”.	140
5.2	Demonstrations of Commercially-Available Vein-Finding Devices.	143
5.3	eye.V.ease, Practitioner’s View Through the NIR-Goggles.	149
5.4	eye.V.ease, Design for Laser Safety.	152
5.5	eye.V.ease, Laser Brightness Control.	154
5.6	eye.V.ease in Use on Ankles and Feet.	156
5.7	eye.V.ease, Design Details, Picture Set 1.	157
5.8	eye.V.ease, Design Details, Picture Set 2.	160
6.1	Our Custom Phantom: Latex Bifurcation on Silicone Hand.	164
6.2	Commercial Off-The-Shelf IV Phantoms.	166
6.3	Silicone Venous Network Patterned On Author’s Anatomy.	169
6.4	Force vs Time Profile For Automatic Insertion Into Our Custom, Latex Vein Phantom.	170
6.5	Making the Soluble-Wax Core.	173
6.6	Rotocaster, Motion Sequence.	175
6.7	Rotocaster Design Details.	176
7.1	Chapter 2 (“Think”) Contribution Summary.	179
7.2	Chapter 3 (“See”) Contribution Summary.	180
7.3	Chapter 4 (“Do”) Contribution Summary.	181
7.4	Chapter 5 (“Do Simpler”) Contribution Summary.	182
7.5	Chapter 6 (“Do, But Practice First”) Contribution Summary.	183

Chapter 1: Introduction

1.1 Motivation

Intravenous (IV) catheterization is a medical procedure wherein a flexible plastic tube, or catheter, is inserted into a vein for the delivery of medicinal fluids. The catheter initially surrounds a needle that punctures the wall of a vein so that the catheter can be slid off of the needle and into the vein, whereupon the needle is removed. To prevent confusion, it is important to disambiguate IV catheterization from other, similar vascular access procedures. Vascular access consists of any procedure wherein a needle is used to puncture a vein or artery. The most common forms of vascular access are intravenous (IV) catheter insertion for the delivery of fluids, phlebotomy sticks for extracting venous blood samples, and arterial sticks for extracting arterial blood samples. IV catheterization is unique in using a 2-part insertion device, consisting of a sharp, metal needle and a softer, plastic catheter that remains in the vein, while both phlebotomy and arterial sticks use just a single, sharp metal needle and leave nothing in the vessel.

Nearly 250 million IV insertions take place in the United States annually, and 28% of those insertions fail on the first attempt in normal adults, with appreciably higher failure rates in children [56] [38] [55] [46]. Failed insertions commonly cause bruising and pain, but can also lead to more serious complications, including long-term nerve damage [5] [40], local tissue damage, and sclerosis of the veins. Figure 1.1A shows a typical bruise from a failed insertion, in this case produced by the needle inserting all the way through the vein's front and back walls. Figure 1.1B shows an extreme case of local tissue damage due to extravasation, or the leaking of a caustic drug out of the vein and into the surrounding tissue. Extravasation typically occurs either from a false presumption of successful IV insertion when the needle is actually only partially inserted or when a properly-inserted catheter pulls partially or entirely out of the vein due to patient movement or improper securement of the catheter to the patient's hand/arm. Extravasation occurs more commonly during chemotherapy/cancer treatment because some common chemotherapy drugs (vesicants) are caustic, and the drugs are often administered via a powerful pump that increases the likelihood of leakage due to increased fluid pressure. There are other failure modes for IV insertion with complications ranging in severity between those of bruising and extravasation, but their discussion is unnecessary. Our working-philosophy is that even temporary bruising is sufficiently traumatic to patients for us to work on improving the IV insertion process.

There are two main difficulties in IV insertion: locating the vein and precisely moving the needle into the vessel. Practitioners often have great difficulty seeing or feeling small veins in women or children's hands/arms or veins that lie beneath a layer of fat, as found in obese patients and many infants. It is somewhat easier to feel and see the veins in healthy men and many elderly patients, given their thinning skin, but it can still be difficult. Precisely moving a small needle into a vein that is not much bigger than the needle can be challenging, particularly in smaller patients such as infants and children. Further, the veins are often moving targets as they roll underneath the pressure of the IV needle, particularly in elderly patients.

Robotic IV insertion is a possible solution that could increase the success rate through precise movement of the needle and enhanced sensory abilities for vein detection and targeting. Detection of the veins through



Figure 1.1: Example Injuries from IV Insertion. A: Bruise from back-wall penetration. Used with permission of Brad Gerten. B: Extravasation. Used with permission of [1] and Lorelle Wuerz, BS, RN.

various sensors could allow a robot to target veins that human practitioners otherwise could not localize. We envision that such a system could be teleoperated by a human to increase his or her motor and sensory abilities without ceding intelligent control to the system, or it could be operated autonomously. An IV-robot could be used to treat people in remote or hostile locations where a human practitioner could not be physically present or in hospitals to increase the success and through-put of nurses and anesthesiologists.

1.1.1 Explaining the Basic Steps of IV Insertion

The IV catheter, as shown in Figure 1.2, is composed of three separate parts: a needle, catheter, and filter plug. The needle is a hollow, metal tube with a sharp, beveled tip for penetrating the skin and vein. The catheter is a thin-walled tube made of a biocompatible plastic and slides concentrically over the needle. During insertion, the catheter and needle move as one piece to penetrate the skin and vein wall, but only the catheter is left in the patient's vein. The catheter's tip ends before the needle bevel begins so that the needle bevel can penetrate tissue without the delicate catheter tip being damaged. The catheter has a Luer Lock screw thread, or "hub", at one end that allows for IV bags to be attached to the catheter for the introduction of fluids. The back of the needle, or flash-chamber, is a translucent space where blood can pool and be seen through the side of the needle. The filter plug presses into the back of the needle and contains an air-permeable, blood-impermeable filter to allow the air to escape but trap the blood. Allowing air to escape through the back of the needle is necessary for blood to enter the flash-chamber and confirm successful insertion, and trapping that blood in the needle is necessary for sterility.

In Chapter 2, we will detail the 27 steps of human IV insertion that we observed in our clinical experiment. However, we now present a simplified, 4-step explanation of insertion into just the vein (ignoring the skin) to serve as starting background for the reader. Figure 1.3 shows the following, most-basic steps of IV insertion:

1. The catheter and needle start the insertion as a single unit that moves as one piece, and the assembly is held roughly 15 degrees (or lower) above the hand surface.
2. The catheter and needle are inserted as a single unit until they penetrate the vein wall. As the needle presses into and deforms the vein, the axial insertion force rises steadily. At the moment of penetration, the practitioner feels a needle-"pop" as the insertion force drops suddenly because it has entered the




Figure 1.2: Details of IV Needle.

vein's hollow center, or lumen. Because the catheter's tip does not extend to the needle bevel, the assembly must be pushed a small distance beyond the point at which the needle penetrated the vein for the catheter to enter the vein as well. As the edge of the catheter tip is pushed beyond the edge of the hole created by the needle, the axial force on the needle decreases suddenly again, creating the secondary catheter-"pop". After needle penetration, blood enters the needle and collects in its flash-chamber, an occurrence referred to as "blood-flash". The combination of tactile "pop" and visual "blood-flash" communicates successful needle penetration of the vein to the practitioner.

3. Once these sensory cues confirm that the catheter tip is in the vein, the rest of the catheter is inserted while the needle is held stationary as a guidewire that prevents the catheter from buckling.
4. Once the catheter has been inserted until only its "hub" remains above the skin, the needle is withdrawn.

1.2 IV Insertion Sites We Are Targeting

The danger and difficulty of IV insertion, whether performed by a human or robot, depend in part on the location of the vein and its anatomical structure. The terms "Central IV", "Central Venous Catheter", and "Central Line" all refer to insertions in large veins in the neck, chest, and groin. A "Central Line" is typically used for long-term (multi-week) delivery of fluids. The terms "Peripheral IV", "Peripheral Venous Catheter", and "IV" (typically) refer to insertions in the hands, arms, legs, and feet for delivery of fluids that lasts for only a matter of hours or days. Legs and feet are only used as insertion sites when a suitable site cannot be found in the hands or arms, and this typically happens with IV drug users who have extensive venous scarring from frequent insertions. A last insertion site that defies either the label of "central" or "peripheral" is the

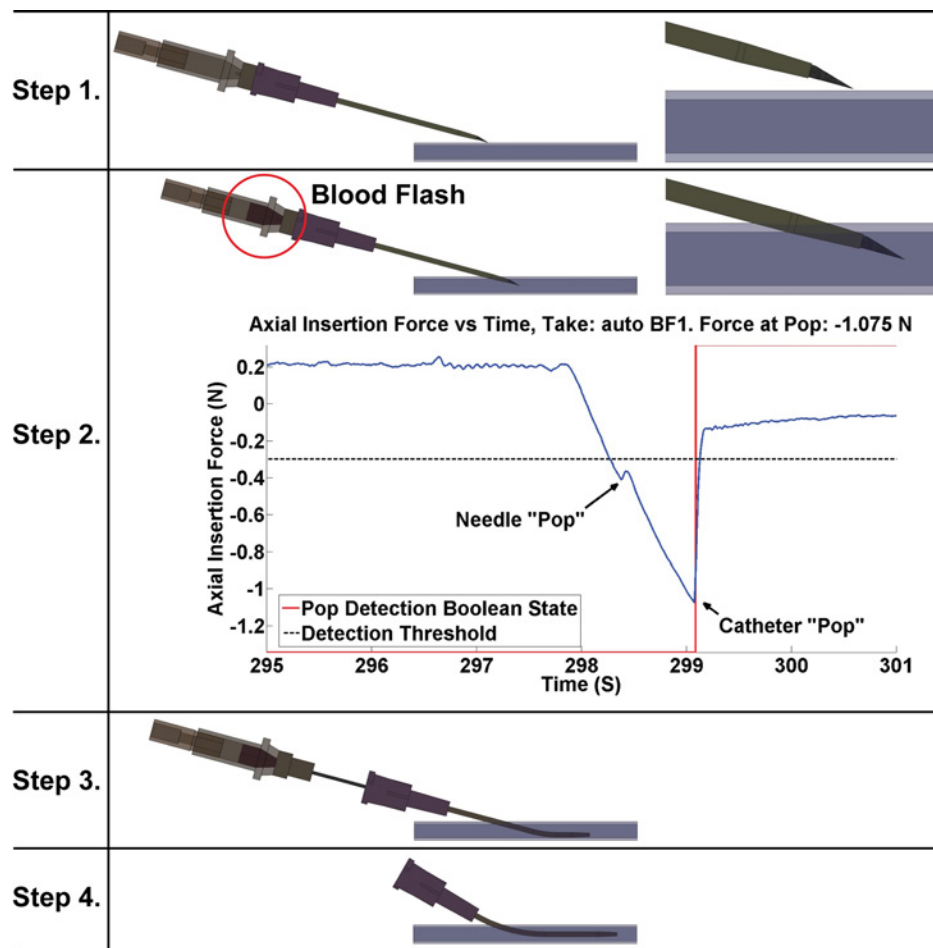


Figure 1.3: Most-Basic Steps of IV Insertion. Skin insertion not shown for simplicity. The axial-insertion-force-vs-time curve shown in Step 2 was recorded during our robot's autonomous insertion into an artificial vein phantom, as will be discussed further in Chapter 4.

case of inserting into a vein in a neonate or infant's scalp, which can sometimes be easier to find than in the typical peripheral access regions. We are restricting our focus to the insertion of Peripheral IV's because their risk of serious complications is much lower than for non-Peripheral sites.

In most adults, Peripheral IV's are inserted in the hand or antecubital fossa (inside of the elbow). However, inserting in the hand has many benefits over inserting in the antecubital fossa, including:

1. The conventional wisdom of many practitioners is that inserting at a venous bifurcation minimizes the vein rolling as opposed to inserting along the midsection of a single vein [52]. Figure 1.4A shows such a bifurcation. The dorsum of the hand contains far more bifurcations than does the antecubital fossa.
2. If the insertion fails in the hand, another insertion can be attempted proximally (towards the heart) without pumping fluid through the failed insertion site and possibly leaking into the surrounding tissue. If the insertion fails in the antecubital fossa, another insertion cannot be performed in the hand because then the fluids could leak in the failed antecubital fossa insertion site.
3. If the patient bends his or her elbow with an IV inserted in the antecubital fossa, the catheter can

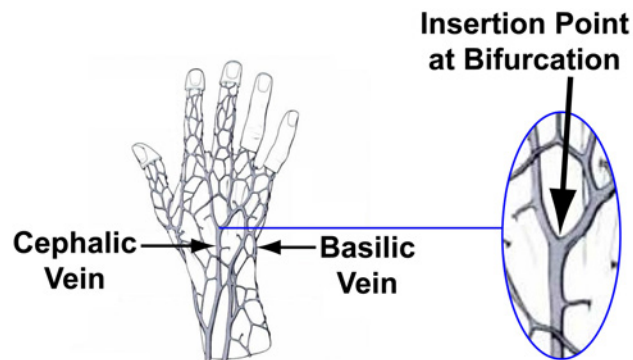


Figure 1.4: We Are Only Targeting Venous Bifurcations in The Hand.

dislodge or kink permanently, requiring another insertion.

4. Placing the IV in the hand keeps the IV line further from the Operating Room table and surgical site.
5. Less of the arm must be inserted into the robot to access the hand than for the antecubital fossa.
6. The thumb can be used as a placement jig for the hand. That is, we can keep the hand flat and the thumb against a fixture with a known location to ensure proper placement of the hand. Such fixturing is harder for the elbow.
7. The hand is thin enough to allow for Near Infrared (NIR) imaging via transillumination (shining all the way through the tissue), whereas the arm is too thick for transillumination and must use reflected light instead. This adds optical complexity and yields potentially less-accurate results.
8. The hand veins are more superficial than in the antecubital fossa, so traction is more effective.
9. The veins and arteries are located in different planes in the hand, so the chances of accidentally sticking an artery are negligible. In the antecubital fossa, the veins and arteries are close together, so it is easier to hit an artery by accident.

There are a few benefits to inserting in the antecubital fossa instead of the hand, including:

1. The main vein in the antecubital fossa is more consistently located across patients than in the hands.
2. The main vein in the antecubital fossa rolls less than in the hands.
3. The veins in the arm contain fewer “valves” than in the hand. Similar to the heart valves, vein valves are internal flaps of tissue that ensure venous blood only flows away from the extremities back towards the heart. Encountering a valve during an insertion often leads to a failed insertion because the catheter cannot proceed past the valve.

Ultimately, we chose venous bifurcations in the hand as the easiest target for insertion.

1.3 Related Work

Although this dissertation provides some novel approaches and insights into robotic IV insertion, much good work has already been, and continues to be, accomplished by other researchers, as detailed below:

In [47] and [58], the authors tested penetration-event-detection on veins in rabbit ears as part of their development of an automated blood-sampling device. They proved that the force “pop” could be detected

and used to halt the uniaxial insertion motion at the correct depth automatically. Their work on an automatic blood-sampling device continued in [57], when they developed a system for detecting blood-flash based on using the electrical conductivity of blood to complete a circuit for which they measured impedance. This work did not culminate in an actual robotic prototype.

In [68] and [69], the authors developed a robotic system for automated blood-sampling. The system was comprised of a powered XYZ translation system with a uniaxial force sensor along the Z-axis. Either a palpation probe or needle could be attached to the force sensor. A passive, rotational degree-of-freedom (DOF) at the robot's base allowed for manual rotation of the entire XYZ gantry about the arm for setting the insertion angle of the needle above surface of the skin. Their robot's automatic-vein-detection algorithm palpated the antecubital fossa at positions along its width and used stiffness measurements above a certain threshold to identify veins. Once the robot had located a vein, the robot paused for a human operator to swap the palpation probe for a needle and rotate the Z-axis from its initial vertical, probing pose to 30 degrees above the surface of the arm. The robot would then insert the needle along its Z-axis until it detected a peak in the force signal, at which point the robot would halt and request that the human operator attach a vacuum-tube to the needle. The robot's success rate for detecting vein location on a live human was 78%. The authors were able to perform proof-of-concept insertions on vein phantoms (synthetic silicone vein filled with fluid and placed underneath the skin in tissue samples of chicken/lamb).

In [48], [49], [50], and [51], the authors developed structured light and reflectance-based, Near-Infrared (NIR) imaging hardware and computer-vision algorithms for detecting and mapping veins in 3D towards the eventual goal of developing an automated IV insertion robot. The deoxygenated haemoglobin in veins absorbs infrared light more than surrounding tissues, so veins appear as dark on a light background under infrared light [18]. This work did not culminate in an actual robotic prototype.

In [16], the authors developed and tested a system for performing automatic venipuncture which was comprised of 6 active DOFs (planar XY, Z-needle-depth-limit, pitch, yaw, and prismatic-insertion along the needle axis) and an ultrasound probe for vein detection. The ultrasound probe was swept across the surface of the hand to detect the vein, and the needle was then aligned with the detected vein. A uniaxial force sensor was used to detect the force "pop"s and halt insertion automatically. The authors developed an insertion phantom that contained latex skin, an agarose fat layer, and a silicone vein, and their penetration-detection algorithm was able to detect "pop"s when penetrating the skin and vein layers. They performed speed-versus-force tests on their phantom, as well as deceased porcine skin and fat, and established that insertion force and tissue deformation increased with increasing speed. This discovery led them to conclude that slower speeds would lead to safer insertions. Their tests also showed that insertion force through their phantom skin and vein decreased as the insertion pitch angle increased.

In [60], the authors investigated tactile detection of veins towards automated blood-sampling through measuring the tissue elasticity in the forearm. They mounted a voice-coil-based, linear actuator as an end-effector on a KUKA KR-3 6-DOF robotic arm and used the voice-coil to measure elasticity at a variety of positions. Their approach was unique from [68] and [69] in that they measured the dynamic response of the

voice-coil carriage as it was dropped onto the tissue below, measuring how far into the tissue it settled and how quickly it settled in order to determine the tissue elasticity. With this method they achieved a 40% success rate for identifying vein location on a 50-person test group. This same research group has also investigated the use of NIR light and computer vision algorithms for visualizing veins [59] [32], but their methods and results were quite similar to other vision approaches described above. This work did not culminate in an actual robotic prototype.

In [39], the authors developed a voice-coil-based, tactile-sensing probe that measured tissue compliance to detect vein locations. They also instrumented a needle with a force sensor and proved that they could automatically detect vein penetration when the needle was inserted at constant velocity along a single DOF that was pre-aligned with their phantom vein. This work did not culminate in an actual robotic prototype that combined the vein-detection with needle insertion.

VeeBot, LLC [53] [27] is a startup company seeking to develop an automated blood-sampling robot. While written details of their system are sparse, their promotional video shows an Epson robotic arm that carries a custom manipulator-head in a macro-mini configuration to allow for 6-DOF (XYZ, yaw, pitch, and prismatic-insertion along the needle axis) positioning of the needle. They use NIR imaging as the primary vein-detection method and confirm the selected vein's adequate blood-flow with ultrasound, achieving a vein-identification success-rate of 83%. They rely on the ultrasound's depth-measurement of the vein to insert the needle using only position-control and no force sensing. The company has not announced a date for clinical testing of their device.

VascuLogic [12] is a startup company that has developed a proof-of-concept robot for automated blood-sampling. Their system appears to be the most successful and complete demonstration of automated blood-sampling to-date. They use a combination of stereo-NIR imaging and ultrasound to build a 3D, segmented map of the detected veins in the antecubital fossa and track the veins in real-time throughout the procedure. While it calculates and suggests the optimal insertion site, it allows the human operator to make the final selection. A 5-DOF (XY, yaw, pitch, and prismatic-insertion along the needle axis) robot guides the needle into the vein based on position control and real-time tracking of the needle tip using the stereo camera. To prevent accidental needle sticks to the human operator, the robot automatically picks up and attaches the needle via an electromagnet and discards it into a sharps-container at the end of the procedure. The device automatically applies an analgesic patch to the insertion site to reduce pain. Cleaning the insertion site and applying the tourniquet are left to the human operator. In 270 trial-insertions on a standard IV phantom, their device had a 100% success-rate and average completion-time of 30 seconds. In trials on 288 different customized-phantoms that the company fabricated to simulate a wide range of patient skin tones, BMI, age, and vein size, their device had a 98% success-rate and 35 second average completion-time. In a partial-test on 100 live, human patients, their device was able to detect at least one useable vein in the antecubital fossa, forearm, or wrist in 99 of the 101 patients. In proof-of-concept tests, their device was able to insert on the tail veins of adult Sprague Dawley rats, similar to human pediatric veins, with a 100% success-rate and a 1.5 minute average completion-time. The company has acquired IRB-approval from Rutgers University to test

their device on live, human patients beginning in May 2015, and they advertise that they will begin selling their system in 2017.

In looking at these efforts, common research themes emerge that are different from the approach we took:

1. Most of the efforts were to create a robot for blood-sampling, not IV catheterization. Those that examined IV catheterization did not address the problem of handling the needle and catheter as separate parts, but treated them as a single unit, as in blood-sampling. We are the first researchers to treat the needle and catheter as two, separate parts and to provide the additional degrees-of-freedom and more-complicated insertion strategy that are necessary to move the needle and catheter separately.
2. All of the efforts were to create a robot that sensed the vein and performed the insertion autonomously with the human operator performing only simple, manual tasks for convenience and simplicity, like attaching a vacuum tube or setting a rotation axis. We are the first researchers to design a system that allows for either fully-teleoperated control or fully-autonomous control of all steps, including the simple ones (e.g. cleaning the insertion site and applying the tourniquet) that other systems left to the human operator.
3. None of the efforts performed clinical experiments with actual IV insertions performed by human practitioners on live, human patients to establish normative insertion techniques/steps and values for key variables. We are the first researchers to perform such a controlled experiment and to develop a complete step-by-step model of IV insertion and establish normative values for variables like insertion speed, pitch angle, and force at key penetration events (“pop”s).
4. All of the efforts envisioned a single-armed robot with a maximum of 6 DOFs performing the procedure. We are the first researchers to develop a 2-armed robot (with 17 mechanical DOFs and 23 optical DOFs) that can perform all of the motions that we observed human practitioners performing during insertions, including all of the preparatory (cleaning, tourniquetting, and restraining), assistive (applying traction and clamping the vein/catheter at the end of the insertion), and actual needle-insertion steps.
5. None of the efforts employed kinematic designs for intrinsic mechanical safety in controlling the pose/rotation of the needle. We are the first to employ a mechanical Remote Center of Motion (RCM) as a safer way of setting the needle’s pose. Further, we developed a novel RCM design that has many benefits over existing designs.
6. Most of the efforts employed reflectance-based NIR imaging and computer vision for vein detection, and those that used tactile-sensing experienced low detection success rates and expressed practical concerns about its viability. We are the first to develop a high-Degree-of-Freedom, NIR-transillumination system and control algorithm that are able to provide optimized lighting that results in clearer, potentially more anatomically-accurate, vein images. Our ability to produce a better image from better hardware reduces the need for complicated computer vision filtering and vein-segmentation techniques used to analyze the image.
7. Most of the efforts successfully employed uniaxial force-sensing for detecting needle penetration of the vein. We are the first to employ a higher-DOF force/torque sensor with non-penetration-detection

goals in mind, including absolute-encoding of the pitch-angle and detection of collisions/vein-rolling.

8. All but one of the efforts targeted veins exclusively in the antecubital fossa, with one effort targeting straight sections of veins in the hand. We are the first researchers to target venous bifurcations and to explain technically why this is a useful, safer strategy. Accordingly, we are also the first to develop an insertion phantom that includes venous bifurcations.

In the remaining chapters we will detail additional, novel aspects of our approach that build on the current state-of-the-art for robotic IV insertion.

1.4 Thesis Content and Contributions

The fundamental, driving question behind our research was whether or not we could raise the IV insertion success rate of a robot beyond that of a human practitioner. However, this question assumed implicitly that a robot could, in fact, insert an IV autonomously, which until this thesis had remained simply speculation. To answer whether or not a robot could insert an IV autonomously, we first had to establish exactly how a human practitioner inserts an IV and translate that process into robotic actions. We used the paradigm, “Think. See. Do.” to help us categorize and explain the human’s actions which we must replicate in an IV-robot.

1. “Think.” The practitioner thinks about the insertion strategy and sensory cues for determining success.
2. “See.” The practitioner locates the vein, often by seeing it visually.
3. “Do.” The practitioner performs the mechanical insertion of the needle as a 2-handed effort.

The following list contains the questions that we asked (and answered) according to this “Think. See. Do.” paradigm, along with which chapter contains their description:

- “Think”. Chapter 2.
 - What is the complete list of steps for an IV insertion?
 - What are the mechanics of human IV insertion (e.g. sensory cues, forces, motions, and intra-venous blood pressure)?
 - Can a human parse data from an insertion to identify the needle and catheter penetration-events into the skin and vein? We need to show how/if humans can do this before teaching the robot.
 - What are the failure modes of IV insertion, and how could we prevent or recover from them?
 - Can we explain technically the conventional wisdom that targeting bifurcations and applying skin traction can make insertions easier?
- “See”. Chapter 3.
 - Can a robot locate a vein and compute a 6D needle pose by itself?
- “Do”. Chapter 4.
 - Can we make a robot perform all of the many, complicated 1-handed and 2-handed mechanical steps of IV insertion by itself?
 - Can we make the robot perform these steps based on similar sensory cues (e.g. the force “pop”) that we observed with human practitioners?
 - Can we do all of this on realistic hand, arm, and vein geometries (not just rubber tubes).

In the course of our research, our paradigm expanded to include two additional items: “Do Simpler” and “Do, But Practice First”. “Do Simpler” refers to our desire to simplify our robotic system while preserving its most useful elements. We wanted to see our robotic system go beyond serving as a research platform to being deployed clinically, and the most realistic option was to reduce its complexity and cost. “Do, But Practice First” refers to our need for a realistic training phantom (artificial vein) that we could use for developing and testing our device, as well as for teaching practitioners to use it.

- “Do Simpler”. Chapter 5.
 - Could lessons we learned from developing the robot be used in a simple assistive device to make IV insertion easier for humans?
 - Could we develop a tool to allow someone to insert an IV on themselves (Make a 2-handed procedure 1-handed)?
- “Do, But Practice First”. Chapter 6.
 - Can we develop a non-cadaver phantom that is more realistic and less expensive than commercially-available options?
 - Can we make hand and vein phantoms with realistic force “pop” characteristics and patient-specific, realistic anatomical size and geometry?

Chapter 2: Understanding the Procedure and Mechanics of IV Insertion

2.1 Motivation and Goals


Before we could design a robotic solution to IV insertion, we needed to gain a more thorough understanding of IV insertion itself. Towards this aim, we partnered with the Stanford University Hospital’s Department of Anesthesiology to perform a clinical experiment (Internal Review Board Protocol #17585) wherein we recorded video and mechanical data from instrumented needles on IV insertions performed by human practitioners on live human patients. Our broad goals were to answer the following questions:

1. What is the complete list of steps for an IV insertion, including all preparatory, insertion, and post-insertion components? We have to know the details of the procedure before we tell the robot how to do it. The complexity of these steps will determine what degrees of freedom are needed by the robot.
2. What are the mechanics of human IV insertion (forces, torques, accelerations, motions, and intravenous blood pressure)? This quantitative data will allow us to generate engineering design specifications for the robot’s required degrees of freedom, actuators, and sensors. The data will also give specifics on the needle-tissue penetration-events (force “pops”) that will allow the robot to insert autonomously.
3. Can a human parse data from a real IV insertion to identify the needle and catheter penetration-events into the skin and veins? Before we can program a robot to parse data automatically, we must first show a human’s ability and method for parsing such data.


2.2 Clinical Experiment Design

The clinical experiment was conducted in the Motion Capture Facility of the Stanford Neuromuscular Biomechanics Lab in the James H. Clark Center. Dr. Fred Mihm, Professor of Anesthesia at Stanford University Medical School and Co-director of the Intensive Care Units at Stanford University Hospital, was the lead practitioner and helped to design and conduct the study. Four other practitioners, including two Stanford anesthesiology-fellows, one Stanford pediatric anesthesiologist, and one non-Stanford pediatrician, were selected based on their experience and research interests. Volunteer patients were recruited through personal contacts and given informed consent per the IRB protocol.

The overall experimental setup, as shown in Figure 2.3, consisted of a central hand-table with two swivel-stools on either side, one for the practitioner and one for the patient. The hand-table consisted of a Near Infrared (NIR) transilluminator upon which the patient would place his or her hand and a heavy base to provide a stable platform. The NIR transilluminator could be used to show the patient’s veins during the insertion and also provided a small ledge over which the patient could curl his or her fingers in a standard hand-pose often used in IV insertions (if the practitioner wanted this pose). To prevent blood contamination, we covered the NIR transilluminator with a disposable transparency and the rest of the table with brown construction paper, both of which were replaced in between patients.



(A)



(B)

Figure 2.1: A: "Nano-needle". The "Hook-And-Gap" geometry allows the 6-axis force/torque sensor to be placed behind the finger-hold to maintain the natural ergonomics of a normal, unaltered IV needle while measuring forces as if it was located between the needle and fingers. B: "Simple-needle". This needle measures only 1-axis of force but is lighter-weight and somewhat more ergonomic than the "Nano-needle".

2.2.1 Overall Needle Instrumentation and Ergonomic Design

We used two different types of instrumented IV needles, referred to hereafter as the “Nano-needle” and “Simple-needle”. The “Nano-needle” provided more complete, accurate measurements with a heavier, more complicated design than the “Simple-needle”. Figure 2.1 shows the details of each instrumented needle. Each volunteer was asked to undergo two insertions, a “Nano-needle” insertion in one hand and a “Simple-needle” insertion in the other hand.

The “Nano-needle” insertions used the following sensors, with the “on-needle” label referring to sensors that were physically connected to the instrumented needle:

1. ATI Nano17 6-axis force/torque sensor (**on-needle**).
2. Tekscan Flexiforce 1-axis piezoresistive force sensor (**on-needle**).
3. Intravenous blood pressure transducer (off-needle).
4. 3-axis digital accelerometer (**on-needle**).
5. Motion capture system to measure the 6-DOF motions (markers **on-needle**, cameras off-needle).
6. Four touch-sensors for manually coding subsections of the insertion (off-needle).
7. Two video cameras, one visible-light and one NIR-light, recording at 30Hz (off-needle).

It is worth noting that there were two force sensors: the 6-axis ATI Nano17 and the 1-axis Flexiforce sensor. We included the Nano17 sensor because we were uncertain as to which force/torque axes would provide useful data, and the Nano17 also provided extremely accurate, precise measurements. However, we suspected, and hoped to prove, that all of the penetration-event force information could be deduced from the axial direction alone and that a simpler, less expensive, and more compact single-axis sensor like the Flexiforce could be used adequately instead of the larger, heavier, and more expensive Nano17. If so, the Flexiforce sensor could be treated as disposable and incorporated into simple assistive devices or the robot more easily than the Nano17. We included both sensors to cover all contingencies and test all theories.

We intentionally designed the “Nano-needle” to preserve the grasp of a normal, unaltered IV needle, as shown in Figures 2.2A-B, such that the practitioner could grasp it with three or four fingers, left or right-handed. Our “Hook-And-Gap” geometry is what enabled this achievement by allowing the 6-axis force/torque sensor to be placed behind the finger-hold to maintain natural ergonomics while measuring forces as if it was located between the needle and fingers. If we had we used the simplest force-sensing geometry, the practitioner would have held onto one side of the force sensor, and the IV needle would have been attached to the other side of the sensor, with nothing in between the practitioner’s fingers and the IV needle but the force sensor and simple attachment parts. However, such a configuration would have placed the practitioner’s fingers unusually far back on the IV needle’s base/finger-hold and added a larger diameter near the IV needle’s base, potentially changing the nature of the IV insertion (“observer error”). Instead, we shortened the finger-hold portion of a normal IV needle, attached the shortened needle to a “hook” structure that routed around the hand to connect to the back end of the force sensor, and replaced what had been the finger-hold portion of the IV needle with a “dummy-finger-hold” that connected to the front side of the force sensor and extended to within 1 millimeter of the “hook”, leaving a gap between the two so that the needle




Figure 2.2: A-B: The “Nano-needle” uses a “Hook-And-Gap” geometry that allows us to measure insertion forces on the needle without changing the normal way in which a practitioner holds the IV needle. A: The instrumented needle’s green and red-highlighted components are connected only through the ATI Nano17 6-axis force/torque sensor. C: In [22], the instrumented needle required a whole-handed grasp to the side of the needle instead of a finger-grasp along the needle’s axis, changing the ergonomics of insertion and potentially the mechanics. Reprinted with permission.

and “dummy-finger-hold” were connected only through the “Hook-And-Gap” support structure that measured forces via the inline Nano17 sensor. We built the “Hook-And-Gap” support structure out of laser-cut wood (5-ply Baltic birch plywood) and thin-walled, aluminum tubing to maximize its strength-to-weight ratio. Figure 2.2A highlights the two force-sensing sides of this “Hook-And-Gap” geometry in red and green, with the same red-green transition occurring at the gap as occurs in the middle of the force sensor (between its mounting and sensing plates), showing visually that the two sides of the gap are serving as two sides of our force sensor. The diameter of the “dummy-finger-hold” was only 12.6% larger than the normal IV needle base/finger-hold, and the position of the grasp relative to the rest of the IV needle was unchanged. Our practitioners’ only ergonomic concern was that they were not allowed to rest an index finger on the actual IV needle because it would alter the measurements. We observed that in normal, clinical practice, some practitioners rest an index finger on the needle behind the catheter so that their finger is prepositioned to push the catheter into the vein in a single-handed action after penetrating the vein. Such a restriction is unavoidable for force measurement, and our practitioners categorized this as only a “slight” concern on our questionnaire.

Although the “Nano-needle” closely matched the finger-grasp of a normal IV needle, it was considerably heavier. Our chosen IV needle (Jelco 20-gauge, 1.25inch) weighed only 0.8g, and our “Nano-needle” weighed 26.47g. We weighed a variety of heavier IV needles (some with metal Luer Locks, and some with additional safety-features that added mass) and found that the heaviest (3.8g) was still considerably lighter than our instrumented needle. As such, we were concerned that the added mass might affect the insertion, even though our practitioners felt that the mass was not problematic. To provide a lighter-weight alternative to the “Nano-needle”, we designed the “Simple-needle”, which included all of the same sensors with the exception

of the ATI Nano17 force/torque sensor. Although we had not yet proven that the Flexiforce would provide adequate force measurement on its own, we reasoned that if the “Nano-needle” insertions showed that the Flexiforce provided data comparable to the Nano17, then all of the data from the “Simple-needle” insertions would be validated and allow us to examine the forces for insertions performed with a lighter-weight needle. As shown in Figure 2.1B, the “Simple-needle” did not require the “Hook-And-Gap” geometry, instead replacing the finger-hold portion of a normal IV needle with an identically-sized, plastic “dummy-finger-hold” that was glued to one side of the Flexiforce sensor, and press-fitting the shortened IV needle into a Delrin puck that was glued to the other side of the Flexiforce sensor. Removing the ATI Nano17 and accompanying “Hook-And-Gap” support structure allowed us to lighten the “Simple-needle” to only 8.1g. The practitioners found the weight and ergonomics of the “Simple-needle” somewhat closer to a normal IV needle than the “Nano-needle” but expressed the same concern about being unable to place their index finger on the needle.

2.2.1.1 Comparison to Other Studies

Our experiment is the first to measure the above-listed variables in IV insertions on live humans, with the exception of axial insertion force. In [22], an IV needle was instrumented to measure axial insertion force during IV insertions. However, the ergonomics of the instrumented IV needle, as shown in Figure 2.2C, changed the practitioner’s grasp from the normal finger-grasp along the needle’s axis, seen with our device in Figures 2.2A-B, to a whole-handed grasp to the side of the needle. Whereas the finger-grasp allows for significant forward motion using only the fingers, the whole-handed grasp requires movement of the entire hand to insert the needle, changing the mechanics of the insertion. As such, we consider that we are the first to measure axial insertion force in a manner that is truly representative of normal IV insertion.

2.2.2 IV Needle and Sensor Details

2.2.2.1 IV Needle Choice and Modifications

Every insertion was performed with a Jelco 20-gauge, 1.25inch IV catheter (Smith Medical PN: 4056). We selected the Jelco-brand because its circular base allowed our robot to hold it easily and precisely, whereas all other needles that we surveyed had non-circular bases that were difficult to hold. We originally intended to allow the practitioners to select the gauge and length of the needle (within the Jelco-brand) for each patient as they would in a hospital, but they predicted that the 20-gauge, 1.25inch size would be their selection for normal, adult patients. Using a single needle size simplified the experimental process because it required only a single set of needle motion capture calibration devices, as described below in section 2.2.2.6.4.

We made two modifications to the IV needles: shortening the base and adding a vent hole. As we described in section 2.2.1, we replaced the IV needle’s base with an aluminum “dummy-finger-hold” to enable the force-sensing. We trimmed the needle base length to approximately 2mm using a sterilized, fresh razor blade, taking great care to deburr the cut edge so that the base would have the proper press-fit with the Delrin puck that held the needle. Because the Delrin puck closed the needle’s base to air, we drilled a 1.34mm vent-hole into the conical surface of the flash-chamber to allow air to escape the needle and blood-flash to occur. The only practical issue from this vent-hole was that it allowed a small amount of blood to drip from

the needle after blood-flash, whereas a normal IV needle would trap the blood via an air-permeable, blood-impermeable “flash-plug” that fits into the base of the needle. However, the practitioners found that the minor blood leakage from the vent-hole did not interfere with the insertion.

Because we press-fit the IV needle’s shortened base into the Delrin puck, we needed a custom tool to remove the needle without damaging the Flexiforce sensor or bending the needle. Other than the base, which was held by the puck, there was no structure on the needle large enough for fingers to gain a sufficient grip to remove the needle. Grabbing the needle shaft directly bent the needle and contaminated its sterility. Our solution was to modify a pair of needle-nose pliers to grasp the needle’s large-diameter, plastic flange (the structure that sat flush against the Delrin puck). We milled the jaws of the pliers to match the thickness of the needle flange and milled a circular hole in the middle of the jaws that was slightly smaller in diameter than the flange so that the flange’s perimeter was compressed very slightly to produce a solid, uniform grasp.

2.2.2.2 ATI Nano17 6-Axis Force/Torque Sensor

Our primary force sensor was the 6-axis ATI Nano17. This sensor is extremely light-weight (9.07 grams) and compact (diameter of 17mm and height of 14.5mm). The Nano17 offers three specific calibrations with a trade-off between force/torque sensing range and resolution. To handle the torques we expected and encountered, it was necessary to use the US-12-4 calibration, which provided the highest force/torque measurement range (12lbf for Fx/Fy, 17lbf for Fz, and 4in-lbf for Tx/Ty/Tz) at the lowest resolution (1/320lbf for Fx/Fy/Fz and 1/2000 for Tx/Ty/Tz). We were initially concerned that the sensor’s approximately 3mm-diameter signal cable would become tangled during the insertion or weigh on the practitioner’s hand, but these concerns were not realized. Also, the signal cable’s weight was not measured by the sensor because the cable exited the sensor on the side held by the hand and not the side that supported the IV needle. However, because the Nano17 supported the “Hook-And-Gap” structure, we needed to subtract the structure’s weight from the measurements to ensure that we were measuring only insertion forces and torques. To accomplish this, we moved the instrumented needle over the entire pitch and roll angle ranges to generate self-weight-vs-pose curves that were then used as gravity offsets.

2.2.2.3 Flexiforce 1-Axis Force Sensor

The Flexiforce sensor measured only axial insertion force. There are many models with different force-range, length, and temperature-range options, and we selected the 0-1lbf, 8inch, standard-temperature model (Tekscan PN: A201, distributed by Trossen Robotics, PN: S-20-1000-FS1LB). It consists of two adhesive-joined, outer layers of polyester film, two silver traces that coat the inner surfaces of the polyester, and pressure-sensitive ink in the middle. As pressure increases on the ink, its conductance increases linearly. When using the sensor as the input-resistor to an inverting-op-amp circuit, the circuit’s output voltage is proportional to the applied pressure. With a response time specified below 5 μ s, the sensor responds quite quickly to changes in pressure. As the force resolution depends on many variables and was not specified by the manufacturer, estimating it is difficult. However, when comparing the insertion force curves of the Nano17 and Flexiforce, the Flexiforce signal showed the same, minute perturbations in normal, axial loading.

The Flexiforce sensor is very light-weight (0.6g for the 8-inch model) and compact (0.208mm thick and

14mm wide overall, with a “sensing-area” diameter of 9.53mm). However, compared to the Nano17, it suffers from quite a few non-ideal performance issues, including repeatability ($\pm 2.5\%$ of full scale), linearity ($\pm 3\%$), hysteresis (below 4.5%), and drift (below 5%). We found that the effect of these issues was appreciably smaller than specified, with only the hysteresis being problematic. The sensors must be “conditioned”, or “exercised”, after periods of disuse exceeding a few hours to obtain accurate measurements. Conditioning involves cyclically-loading the sensor from 0 to 110% of its rated load for several minutes.

It should be noted that the Flexiforce sensor measures pressure applied over the entire sensing-area, not force. To enable force-sensing, we glued the center of the sensing-area in between a set of Delrin pucks to transfer forces, regardless of their geometric distributions, as pressure. One puck attached the sensor to the “Hook-And-Gap” structure, and the other puck attached the IV needle to the sensor, with the IV needle lightly press-fit into that puck. Per the manufacturer’s specifications, the puck covered 90% of the sensing-area and was centered to ensure that it was contained entirely within the sensing-area for accuracy. We used fine-sandpaper to texturize the bottom surface of the puck for enhanced adhesion with the glue (cyanoacrylate). The glue did not degrade the sensor’s polyester surface and provided a strong, stiff, and long-lasting bond. However, it also introduced a sensor-bias that had to be canceled during the calibration process.

We calibrated each sensor’s voltage-to-force curve using small, precision test-weights over the range 0-450g in 50g increments. To obtain accurate results, we calibrated the sensors only after they were installed/glued in the instrumented needle. There was some variability in the specific linear-fit constants between individual sensors, so we calibrated each sensor individually. As an example, one of the sensors had a voltage-vs-force curve that was described by: $Force(N) = 3.167 * V(volts) - 2.124N$, $R^2 = 0.991$. We discovered that lightly press-fitting the IV needle into the puck changed the calibrated offset(the b constant in $y = mx + b$), but did not change the slope (m). Unfortunately, we could not calibrate the sensor with the needle installed because the needle blocked the resting spot for the weights. As such, we had to adjust the calibrated offset after needle installation to ensure accurate results. The only weight that hung from the Flexiforce sensor was that of the distal Delrin puck and IV needle, and these were so light-weight that active gravity-subtraction based on the needle’s pose was not required for the Flexiforce.

We were initially concerned that attaching the IV needle to the main “Hook-And-Gap” structure through the Flexiforce sensor would introduce too much compliance or break the sensor. In fact, we found that the connection stiffness was very high for all force/torque directions until the sensor experienced sudden failure through delamination. When torques were applied to the needle about axes orthogonal to its axis (tending to bend the needle or pry it from its puck), the sensor’s two polyester halves would tend to be pulled apart. Fortunately, such delamination did not have a progressive effect on either the sensor’s values or its stiffness but would permit normal operation until sudden, easily-identified failure. When the sensor delaminated, the force readings would experience large variations, and the needle would bend severely to one side, making it impossible not to notice the failure. We experienced several of these failures in the course of our experiment but never during an active insertion, instead encountering them mainly during the “exercising” routine prior to an insertion. The reliability of the sensor until its sudden and obvious failure meant that such infrequent

failures were fixed quickly and did not affect the results of our study.

2.2.2.4 Intravenous Blood Pressure Transducer

After the catheter had been successfully inserted, we attached it to an IV line with an inline, disposable transducer that measured the intravenous blood pressure. Measurement of intravenous blood pressure resulted in the loss of only a few milliliters of blood that entered into the end of the IV tubing, so it was not dangerous for the volunteers. We measured the pressure first with the tourniquet remaining for approximately 1 minute and then without the tourniquet for another minute. Without the tourniquet, the intravenous blood pressure is generated simply by the height of the blood column above the hand such that if the patient crouches so that the heart and the hand vein are at the same height, the intravenous pressure measures at roughly zero. The tourniquet increases intravenous pressure because it closes the superficial veins located near the tourniquet, trapping blood inside the arm, but cannot close the deeper arteries that continue pumping blood into the arm. As reported in [44], there is a linear relationship between pressure applied externally by a congesting-cuff (essentially a blood pressure monitor used as a controllable tourniquet) and intravenous pressure.

We measured the intravenous blood pressure to test if it correlated with any of the other variables we measured, particularly axial force. If such a relationship existed, we would need to have our robot control the tourniquet tightness, and hence intravenous pressure, to provide for consistent insertion parameters. For example, in [64], it was shown that a vein’s circumferential and longitudinal stiffness increased with intravenous pressure, and we wondered if the increased stiffness would affect the insertion forces. If we observed a force “pop” with an absolute magnitude that depended on intravenous pressure, then our robot would need to apply a tourniquet/congesting-cuff pressure that set an intravenous pressure commensurate with the absolute force magnitudes expected by the force-“pop”-detection algorithm. We also planned to copy the measured pressures in our artificial vein phantoms to provide robotic test insertions that were as realistic as possible.

We used the inline intravenous pressure monitoring device that is standard at Stanford University Hospital (“Pressure Monitoring Kit with Disposable Transducer”, Hospira List No. 42585-05). This sensor output a differential voltage that we conditioned, amplified, and filtered via a Transducer Techniques TMO-1 Conditioner Module (2-Pole Butterworth, 220Hz bandwidth). To calibrate the sensor’s conditioned, analog signal to the line pressure, we set the IV bag at known heights above the sensor in a literal implementation of “inch water column” and further confirmed those pressures via a secondary device, the Extech HD700 Differential Pressure Manometer. There was some variability in the specific linear-fit constants between individual sensors, so we calibrated each sensor individually. As an example of the linear relationship, sensor number 32’s voltage-vs-pressure curve was: $pressure(mmHg) = 23.784 * V(volts) + 10.054mmHg, R^2 = 1$.

We are the first to measure intravenous pressure when using a standard, non-congesting-cuff tourniquet and first to test its relationships with other variables in IV insertion.

2.2.2.5 3-Axis Digital Accelerometer

To measure accelerations directly instead of relying on a noisy double-differentiation of the motion signals, we included a 3-axis digital accelerometer (Sparkfun PN: SEN-00758, using the STMicroelectronics LIS3LV02DQ IC, $\pm 2G$ range, 12-bit resolution at 640Hz). For both space-constraints and to minimize the




Figure 2.3: Motion Capture and Experimental Setup. A: Four infrared motion capture cameras were placed in a shallow arc facing the practitioner and patient chairs. B: An insertion in progress with the author watching to manually-code blood-flash, line/plane calibration, and active-insertion via touch-sensors, providing easier data analysis.

chances of the accelerometer’s wires becoming tangled with the patient’s or practitioner’s hands, we placed the accelerometer at the back end of the instrumented needle. The accelerometer’s Z-axis was concentric with the needle’s axis, and the accelerometer’s negative X-axis was aligned approximately with the needle’s “bevel-up” direction. We were concerned that the weight and stiffness of the five wires needed for power and communication with the computer would interfere with the insertions. To minimize any such effects, we used 42-gauge magnet wire (Antique Electronic Supply, PN: S-MW-42-D125) with an outer diameter of 0.0635mm (0.0025inch). The wires were so thin that they could not be felt when holding the IV needle, and we experienced no problems with the wires’ performance.

2.2.2.6 Motion Capture System for Measuring the 6D Needle Motions

To measure the 6D needle motions (X, Y, X, roll, pitch, and yaw), we used a passive, retroreflective-marker-based motion capture system (Motion Analysis, Eagle-4 cameras and Cortex 1.1.4 software). We found non-optical alternatives to be inadequate, with magnetic trackers too imprecise and mechanical-linkage trackers too intrusive. We chose a passive, instead of active, motion capture (“Mo-Cap”) system because it allowed for tracking with smaller, lighter-weight markers that did not require power or wires. The only disadvantage to using a passive system was that it required considerable post-processing effort on our part to help the analysis software disambiguate and re-identify markers that left and re-entered the cameras’ views.

The system worked by shining infrared light on the scene and triangulating the reflected light from the retroreflective markers, which were roughly spherical and had an approximate mass of 0.1g and diameter of 4.56mm. Because of the system’s reliance on reflected light, we initially experienced issues with the computer detecting reflections from shiny portions of the instrumented needle as false, unknown markers. To fix this, we placed matte tape over shiny portions of the structure, where practical, as shown in Figure 2.1A.

To verify that our calculated 6D motions were correct, we attached the instrumented needle to a Sensable Omni device, which senses all 6-axes of pose, and compared the motions recorded by the Omni to those generated by the motion capture system and our analysis code. The Omni and calculated values matched so closely that their respective curves were indistinguishable when plotted together in the same figure.

2.2.2.6.1 Motion Capture Camera and Marker Placement

We placed four motion capture cameras in a shallow arc partially surrounding the practitioner and patient chairs and looking down onto the patient's hand, as shown in Figure 2.3A. Although only three cameras were necessary for triangulation, the fourth camera was useful as a substitute in case another camera's view was blocked. We placed all of the cameras to one side of the line between patient and practitioner to allow for the author to stand on the opposite side and observe the insertion without blocking any of the cameras. This decision required that we place all of the Mo-Cap retroreflective markers on the side of the needle facing the cameras. The practitioner's handedness (right or left) then determined which stool he or she would sit on to allow the cameras to see the retroreflective markers on the side of the instrumented IV needle. Because the cameras maintained the same global position despite which chair the practitioner used, we had to keep track of which stool was used and account for it when creating a hand-frame coordinate system for the motion capture data to obtain consistent velocity directions between left-handed and right-handed practitioner insertions.

Although only three Mo-Cap markers were necessary to determine the pose of the instrumented needle, we included a fourth marker outside of the triangle formed by the other three markers to facilitate disambiguation. This was crucial during post-processing when the author needed to help the analysis software re-label markers that had temporarily been hidden. The only ergonomic consideration in the markers' placement was that the markers' support structure limited the range of the needle's roll angle. By positioning the "bevel-up" direction of the needle carefully, we could guarantee that the needle could roll ± 22.5 degrees from the "bevel-up" direction being aligned with the hand normal before the cameras were unable to see the markers or the support structure collided with the hand. We were confident that this roll range would be more than sufficient because we had observed minimal roll during insertions performed with normal IV needles prior to the experiment. Ultimately, the practitioners did not experience problems from this limitation, and the roll range that we measured in the course of the experiment confirmed our pre-experiment observation. The practitioners reported no ergonomic problems with the "on-needle" motion capture measurements.

2.2.2.6.2 Motion Capture System Calibration

The motion capture system required calibration before each insertion to provide accurate measurements. This consisted of two steps: Wand calibration and frame calibration. In wand calibration, a special wand with three Mo-Cap markers in a known configuration was moved slowly over the entire workspace and many orientations. The system then calculated the relative frame transformations between the cameras to a high degree of precision. After each calibration, the software displayed the mean and standard deviation for marker position error. Averaging across all trials, the mean marker position error was $\mu = 0.075 \pm 0.003\text{mm}$, and the standard deviation was $\sigma = 0.053 \pm 0.018\text{mm}$.

Frame calibration was used to establish the motion capture system's global coordinate system. This

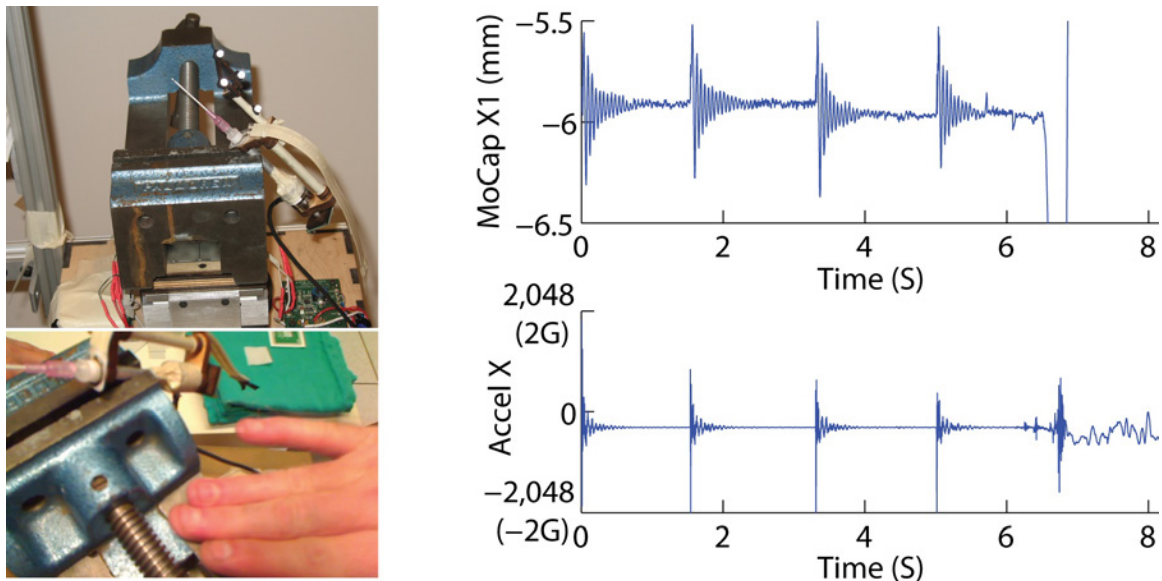


Figure 2.4: Synchronization of Motion Capture Data With Other Sensors Via “Impulse Alignment”. Repeatedly slapping a vice that held the instrumented IV needle produced peaks/impulses in the Mo-Cap marker coordinates and accelerometer data that, once aligned, allowed us to synchronize the data from their respective computers and form a global timer.

consisted of placing a special retroreflective jig where the patient’s hand would normally be placed and instructing the motion capture software to establish the coordinate system. Note that this global coordinate system was near where the hand was placed but that for each insertion we also established another, local hand-frame coordinate system specific to the patient’s hand/vein anatomy and insertion site location.

2.2.2.6.3 Synchronization of Motion Capture Data With Other Sensors

Because the motion capture system required its own, dedicated computer, we had to synchronize the time-stamps of the motion capture data with those of the other sensors recorded on a separate computer. To accomplish this, we held the instrumented needle in a vice that was placed on the hand-plate and slapped the vice four times to produce impulses in the accelerometer and Mo-Cap data. We calculated time-stamp-offsets that aligned the impulses from the two sensors and set the first impulse as the zero-time for our global timer. Figure 2.4 shows this process in detail.

The Mo-Cap data and other sensors were also recorded at different speeds. There was a trade-off between spatial resolution, Field-of-View, and frame rate for the Mo-Cap system, and we found that the optimal balance for our experiment yielded a maximum frame rate of 200Hz. Because we recorded all non-motion-capture data at 1000Hz, we interpolated the calculated needle poses using Linear Interpolation for translation and Spherical Linear Interpolation for rotation.

2.2.2.6.4 Calculating Needle Pose From Marker Positions

We developed custom hardware and software for calculating the needle’s pose based on the MoCap markers’ positions. The first step was to construct a local needle coordinate system by examining the positions of the three markers (M2, M3, and M4) that formed a triad and performing the following operations:

1. Calculating the origin as the centroid of M2, M3, and M4.
2. Calculating the the X-axis as the normalized vector between M2 and M3.
3. Calculating a vector V2 as the normalized vector between M2 and M4.
4. Calculating the Z-axis as the normalized cross-product X-axis x V2.
5. Calculating the Y-axis as the normalized cross-product Z-axis x X-axis.

Our next step was to calibrate the axis and tip position of the instrumented needle in the local needle frame, repeating this calibration every time that we replaced a component that would affect the current calibration. We used two special devices for this calibration. As shown in Figure 2.5A, the first device consisted of a spherical bearing and location jig that rotated the needle precisely about its tip. As such, the motion of each marker was on the surface of a sphere with its center at the tip of the needle. We performed a RANSAC sphere-fit for each marker's position-set and then used the average calculated center as the needle tip expressed in global coordinates. Iterating through the motion data, we constructed a local needle frame and calculated the needle tip position in terms of the local coordinates, taking the mean of all these values as our calibrated, local needle tip position. Our second device consisted of a set of ball bearings and a location jig that allowed the needle to be rotated precisely about its axis such that the motion of each marker was contained in a plane with a normal that was parallel to the needle's axis. We fit a plane to the position-set for each marker and then used the average plane normal as the needle axis expressed in global coordinates. Iterating through the motion data, we constructed a local needle frame and calculated the needle axis in terms of the local coordinates, taking the mean of all these values as our calibrated, local needle axis. To analyze new Mo-Cap data for an insertion, we constructed the updated local needle frame and transformed our calibrated, local needle axis vector and tip position into global coordinates.

2.2.2.6.5 Establishing the Hand Frame

To measure angles, distances, and velocities relative to the actual hand and vein geometry, we needed to establish a hand/vein coordinate system. This process entailed:

1. Setting the origin as the needle tip's position at the moment-of-skin-touch (KE1),
2. Fitting a plane that was tangent to the hand surface at the insertion point,
3. Setting the Z-axis as the normal of this tangent/XY plane,
4. Fitting a line to the vein's overall direction,
5. Calculating the X-axis as the line-fit normalized and projected onto the tangent plane, and
6. Calculating the Y-axis as the normalized cross-product Z-axis x X-axis.

To fit the tangent plane (Item 2), we instructed the practitioners to move the needle's tip over the surface of the hand in the neighborhood of their intended insertion site, giving us the general topography over an area that was approximately 30mm x 30mm. To prevent accidentally cutting the patient with the needle tip during this process, the practitioners kept the needle "bevel-up" and moved the needle backwards or sideways, but never forwards in the insertion direction. Although practitioners were generally able to keep the needle in light contact with the skin to give a fairly accurate hand topology, they would sometimes hover the needle over the surface of the hand, thereby offsetting the mapped topology from the actual anatomy. However, the

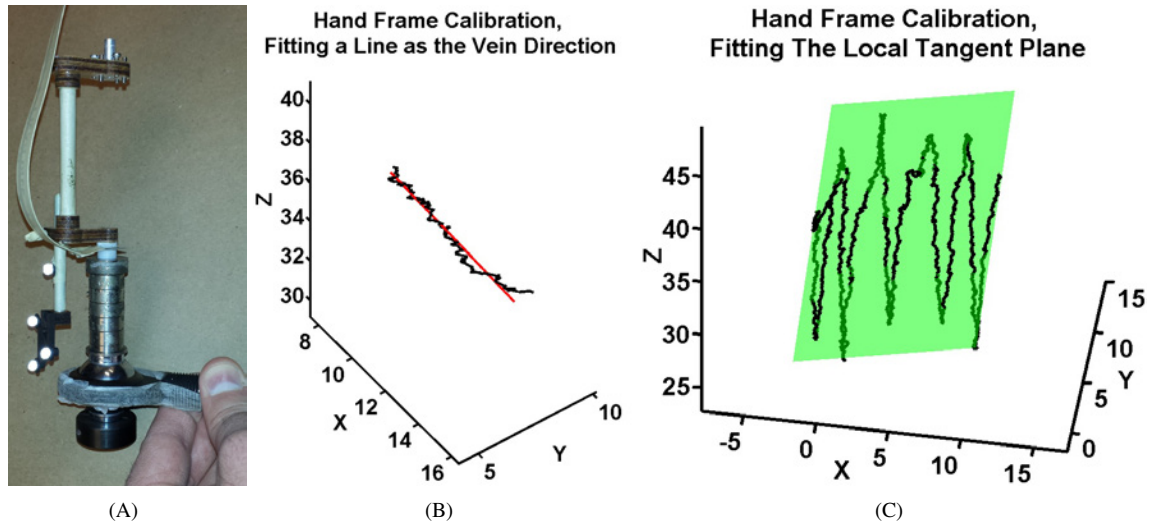


Figure 2.5: Calibrating the IV Needle’s Local Motion Capture Frame and Establishing the Hand Frame. A: We used two custom devices for calibrating the IV needle’s geometry with respect to the local frame formed by its motion capture markers. The device shown here uses a spherical bearing and location jig to rotate the needle precisely about its tip such that we can perform a RANSAC sphere-fit to establish the needle tip with respect to its markers. The other device consists of a set of ball bearings and a location jig that allow the needle to be rotated precisely about its axis such that we can calibrate the needle axis with respect to its markers. B: The practitioners were asked to drag the needle tip along the vein they wished to cannulate, and we fit a line to this data to establish the vein direction. C: The practitioners were asked to drag the needle tip across the hand near the insertion site, and we fit a plane to this data to establish the local tangent plane. We constructed a local hand frame using the line fit, plane fit, and needle tip position at the start of the insertion.

orientation of the plane-fit was only minimally affected by this, and we set the plane’s height above the hand based on the precise moment-of-skin-touch position (origin) confirmed by the force sensor. As described in section 2.2.2.7, the author manually-coded the timing of this process via a dedicated touch sensor to automate hand-plane generation. Figure 2.5C shows an example of this data and corresponding plane-fit.

To fit a line to the vein’s direction (Item 4), we instructed the practitioners to move the needle’s tip along the length of the vein that contained their intended insertion site in a manner identical to the plane-fitting process. Again, the author manually-coded this process via a dedicated touch sensor to enable automatic line-fitting. Figure 2.5B shows an example of this data and corresponding line-fit.

Calculating needle poses via our axis and sphere calibrations assumes that the needle is not bending, such that the tip and axis are fixed with respect to the Mo-Cap markers. In reality, the needle tip is able to bend under non-axial loading, and this is a potential source of error in our Mo-Cap tracking. However, in reviewing videos of the insertions, it appeared that departures from our straight-needle model were infrequent and likely did not have an appreciable impact on our results, although it is impossible to quantify the impact with the sensors available. The video is low resolution, monocular/2D such that it cannot calculate depth or out-of-plane bending, and does not include high-contrast, known-geometries near the needle that could provide precise scale. We could calculate the bending from the measured forces/torques, but we do not know their geometric distribution along the length of the needle, given the complex, unknown loading occurring

under the skin. For future clinical experiments, we plan to include some method of measuring the needle's bending to account for this possible error.

2.2.2.7 Touch-Sensors for Manually-Coding Sections of the Insertion

Because the insertions were long and complicated, we used four touch sensors for manually-coding sections of the insertions, including 1. hand-frame tangent-plane-fit calibration, 2. hand-frame vein-direction-line-fit calibration, 3. active insertion, and 4. blood-flash detection. As shown in Figure 2.3B, the author watched each insertion from overhead and pressed the touch-sensors during these four key sections. We could then quickly find the portions of data that were manually-coded during post-processing and analysis without having to examine data from the entire recording session, which included many minutes of non-active-insertion data. In Figures 2.8, 2.11A-B, and 2.12, the purple line shows the manual-coding of the active-insertion, and the red line shows the manual-coding of the blood-flash. It should be noted that the timing of those curves with respect to the data-curves from the instrumented needle's sensors was limited in precision and accuracy by the author's speed at observing and manually-coding the variables. For instance, the needle's moment-of-skin-touch often occurred slightly before the active-insertion flag was raised manually. As such, the manually-coded flags were used only as rough guides for finding portions of the insertion.

2.2.2.8 Video Cameras

We recorded each insertion with two video cameras placed far enough above the patient's hand as to avoid interfering with the practitioner's hands or view. Both cameras fed frames at 30Hz to the main data-recording program, which time-stamped and saved the frames. One camera recorded only visible light, and the other camera recorded only infrared light via a filter that blocked visible light frequencies. Because the motion capture system used infrared light to track the retroreflective markers, our NIR transilluminator would have interfered with the motion capture system. As such, the transilluminator was only used to examine patient anatomy before or after insertion, but not during insertion. However, the infrared camera was still able to see the veins in much better vein/non-vein contrast than the visible-light camera using just the infrared component of the room's overhead fluorescent lights and an incandescent work lamp. Figure 2.7 shows examples of both video feeds and their usefulness in identifying penetration-events.

2.2.3 Experimental Procedure

Each practitioner was asked to follow the same overall experimental procedure but to insert the IV needle in his or her own method. The overall procedure consisted of the following steps:

1. The volunteer was given informed consent and asked a series of demographic questions.
2. The tourniquet was applied to facilitate vein-detection.
3. The vein target was selected by the practitioner through standard visual and tactile methods.
4. Alcohol was used to clean the insertion site.
5. Lidocaine was injected to numb the insertion site.
6. The tourniquet was removed and the lidocaine allowed to start working for 5-10 minutes, as needed.
7. The motion capture system was calibrated.
8. The tourniquet was reapplied.

9. The insertion site was cleaned again with alcohol.
10. The “impulse alignment” synchronization of Mo-Cap and non-Mo-Cap sensors was performed.
11. The practitioner performed the tangent-plane-fit and vein-direction-line-fit processes.
12. The practitioner performed the IV insertion in his or her own method.
13. The intravenous pressure transducer was attached to the catheter and observed for 1 minute.
14. The tourniquet was removed and the baseline pressure observed for another minute.
15. The catheter was removed and the insertion site cleaned/bandaged.
16. The practitioner was debriefed and asked a series of questions about the insertion.

Although the active-insertion time was quite brief, the preparatory and calibration processes were significantly longer, requiring each volunteer to stay for a total of 45-60 minutes.

The decision to inject lidocaine was made for patient comfort and to match the conditions of our future robotic insertions in which we planned to administer some form of lidocaine. However, we had doubts as to whether it would affect the results or non-robotic, general applicability of the study. Lidocaine is a known vasoconstrictor, and there was some concern that smaller veins would change the insertion mechanics, strategies, or success rate. Further, lidocaine-injection is not a universally-adopted part of IV insertion but is typically left to the discretion of practitioners. We injected 0.1-0.3mL of lidocaine (1%-concentration), buffered in a 10:1 ratio with sodium bicarbonate (8.4%-concentration) to reduce the initial sting, using a 30-gauge needle and a 1mL syringe. Injection was used instead of a needle-free, topical patch because the injection works in minutes as opposed to up to an hour for the patch.

2.3 Results

To frame our specific results properly, we must first state some broader results that will be explained further in later sections, including:

1. Once the needle had been inserted and the catheter was being inserted, or “threaded”, into the vein, the practitioner’s fingers/hands on the catheter corrupted the force signals and blocked the cameras’ views (needed to help parse the data). Therefore we were only able to analyze up until catheter insertion.
2. The Flexiforce sensor was found to be inadequate for measuring force accurately on its own, so we were unable to include the “Simple-needle” trials in our results because the Flexiforce was the lone force sensor in those trials. A few exceptions were that we were able to include the procedural observations, intravenous pressure measurements, and clinical failure observations from those trials as they were not measured at penetration-events that required the force measurement for identification. Unless otherwise specified, all of the following values and figures will refer only to “Nano-needle” trials. We also relied solely upon the Nano17 force/torque data in our parsing of the “Nano-needle” trials.
3. In total, we performed 63 insertions. Unable to consider insertions without usable, complete force data (due to initial practice/debugging, clinical-failure, technical-failure, or “Simple-needle” type), we analyzed 10 “Nano-needle” trials (5 males and 5 females, age = 32 ± 7 years). Unless noted otherwise, all statistics in the following sections will refer to those 10 trials only.

4. Our clinical failure rate was 7.4% (2 out of 27 insertions) for the “Simple-needle” and 16.7% (6 out of 36 insertions) for the “Nano-needle”, confirming our practitioners’ qualitative assessment that the “Simple-needle” was somewhat easier to use and closer in ergonomics to a normal IV needle than was the “Nano-needle”. However, the “Nano-needle” failure rate of 16.7% is comparable to that experienced with one of the catheter types in [22] (8% for the PTFE-Venflon catheter and 16% for the PUR-Insyte catheter) and below clinical failure rates (28% for normal adults, higher for children [56] [38] [55] [46]). Ideally, our practitioners would have inserted normal, unaltered IV needles into the same patients as a control for clinical failure rate, but the added injury to patients did not seem justified at the time. Observing the same practitioners post-experiment to establish a control failure rate would not work because the conditions of the experiment could have impacted the practitioners’ performance even with a normal needle, due to the unusual setting (such as sitting instead of standing during an insertion) and stress of being recorded. Section 2.3.3 gives the details of each failure.

In the following sections, we will provide all of the specific results, along with the 19 “Key Discoveries” that describe those results more broadly.

2.3.1 The 27 Steps of Human IV Insertion

We observed the following 27 steps of IV insertion across all five practitioners and all insertions.

1. Gathering and opening all supplies for IV insertion,
2. Applying a tourniquet to the upper arm or forearm,
3. Bracing the hand against a stiff, stable structure (such as the hand-plate),
4. Cleaning the dorsum of the hand with alcohol,
5. Drying the dorsum of the hand,
6. Applying analgesic (in our case, a lidocaine injection),
7. Waiting for the analgesic to take effect,
8. Tapping the dorsum of the hand to stimulate vein vasodilation,
9. Palpating and observing the hand veins visually to determine the best target vein,
10. Searching the target vein for valves visually, as well as through palpation,
11. Bending the fingers downward,
12. Applying traction towards the wrist with the non-dominant hand,
13. Inserting the IV needle into the skin,
14. Inserting the IV needle into the vein,
15. Detecting the force “pop(s)” and visually observing the blood-flash,
16. Pausing the needle’s insertion to evaluate continued blood-return,
17. Advancing the needle a short distance beyond where the last “pop” and blood-flash occurred,
18. Advancing the catheter into the vein while holding the needle stationary,
19. Dealing with any valves that have been encountered,
20. Applying pressure to close the vein with one hand while withdrawing the needle with the other,
21. Withdrawing the needle fully while keeping the catheter stationary in the vein,

22. Attaching the IV bag to the Luer lock of the catheter,
23. Cleaning the hand of any spilled blood,
24. Securing the catheter to the dorsum of the hand with tape or other means,
25. Releasing pressure on the insertion site,
26. Removing the arm tourniquet, and
27. Disposing of the needle in a “sharps” container and other supplies in a Biohazard container.

Items 12 and 13 could theoretically be combined as one, fluid motion without stopping or slowing in between, but we always observed them as two, distinct steps. Item 19 is the only step that was not observed universally but was patient-dependent. We included it on the list because it is a mandatory step if the insertion encounters a valve and is not left to the discretion of the practitioner (like the choice of wearing gloves).

We made three, key discoveries about the step-by-step process of IV insertion:

- Key Discovery #1: IV insertion consists of 27 steps, each of which is necessary for safety and success.
- Key Discovery #2: IV insertion, as currently performed clinically, is an inherently 2-handed procedure.
- Key Discovery #3: One of the most important and difficult insertion steps is closing the vein with the non-dominant hand to prevent bleeding during needle withdrawal and Luer Lock attachment (Item 20).

An implication of Key Discovery #1 is that other prior research surrounding robotic phlebotomy [47] [58] [69] [51] [60] [59] [32] [39] [27] and IV insertion [16], has over-simplified the problem. With the exception of [12], each prior effort addressed only some subset of three automated steps (sometimes with manual help from the human operator to transition between steps), including vein-detection, insertion of the needle in a straight line, and halting the needle upon detection of the force “pop”. In [12], the system also automatically applied an analgesic patch and picked up/disposed of the IV needle to prevent accidental needle sticks or blood contamination. Our work is the first to develop and address a complete, step-by-step task-list of IV insertion that could allow for automation of the full procedure.

An implication of Key Discovery #2 is that our robot must have two, separate arms. Some of the items (12, 20, 22, and 24) absolutely require two arms, some items (18 and 19) require two arms in difficult cases, and many of the remaining items are typically performed with two hands for convenience but could theoretically be achieved, laboriously, with one hand. Our work is the first to establish the need for two arms to complete a full IV insertion and the first to provide them in a robotic IV system.

An implication of Key Discovery #3 is that small details of the step-by-step process can be very important and are sometimes overlooked by practitioners in their process descriptions. As shown in Figure 2.6, even an experienced practitioner can use insufficient pressure to close the vein during needle withdrawal, causing blood to pour from the catheter before it is attached to the Luer Lock of the IV line. This is traumatic for the patient and presents a blood-borne pathogen risk for the practitioner. We asked many practitioners to describe the step-by-step process of IV insertion, and not a single one mentioned this critical step.

2.3.2 Quantitative Results and “Key Discoveries”

We will discuss our quantitative results in the following sections, with most of the specific, numerical results contained in the captions of Figures 2.8- 2.12 and the broader “Key Discoveries” in the main text.

2.3.2.1 How We Parsed the Data

We originally envisioned that parsing the data to identify “pops” would require examination of just the axial force, as shown in Figure 2.8A. In our robotic insertions, to be described further in Chapter 4, this was exactly the case. However, in our human-practitioner insertions, the many non-penetration-event motions (including re-orientation of the needle, position adjustments, and unsuccessful penetration attempts) introduced false, “pop”-shaped waveforms in the force data that made it too confusing to parse on its own. Fortunately, the other sensors, both on and off the instrumented needle, supplemented the force data to allow for parsing. The following is a description of how we parsed the penetration-events, in order of their occurrence:

1. Geometric order-of-operations required that the key events occur in the following order: moment-of-skin-touch (KE1), skin needle “pop” (KE2), skin catheter “pop” (KE3), vein needle “pop” (KE4), and vein catheter “pop” (KE5), with blood-flash occurring at least after the vein needle “pop” and potentially after the vein catheter “pop” if delayed. We used the end of the vein catheter “pop”, when the axial force had returned to a resting level, as the end of needle insertion (KE6). As we were unable to analyze the catheter’s threading into the vein, KE6 was the final time for our analysis.
2. The active-insertion flag and video feed provided the time-span of actual insertion that we needed to examine, as the data included many more minutes of non-insertion preparation and calibration.
3. The moment-of-skin-touch was readily-identifiable in the axial force plot by finding the time at which the force began to ramp up from its resting level.
4. The video feed provided a fairly precise time estimate for when the skin needle “pop” occurred because the needle bevel would disappear beneath the skin. Sometimes this “pop” happened quickly in between frames, as shown in Figure 2.7A, and sometimes it happened more slowly over the course of several frames, as shown in Figure 2.7B. Often the skin would relax somewhat after the skin needle “pop”, as shown in Figure 2.7A, because the skin would jump backwards from its pre-penetration position to which it had been stretched by the needle tip to its new position resting against the catheter tip. Using the video’s time-stamp for the “pop” as a starting point for our search, we took the nearest “pop”-shaped waveform in the axial force plot as the actual penetration-event.
5. Identifying the skin catheter “pop” was more involved than for the skin needle “pop”. We could sometimes see it as a distinct event in the video feed with further skin relaxation, but it was less frequent and harder to determine than for the skin needle “pop”. The skin catheter “pop” generally happened at a higher needle speed than for the skin needle “pop” and experienced longer needle distance traveled immediately after it because the needle was free to slide underneath the skin. After the skin needle “pop”, the needle had only been able to move the short distance between the needle tip and the catheter tip before the skin blocked the catheter. A further contribution from the Mo-Cap data was that the practitioners would pause briefly and lower the pitch in between the skin and vein penetrations, allowing us to check the pitch angle to tell what time the skin catheter “pop” preceded. In 90% of the trials, the axial force magnitude of the skin needle “pop” was smaller than for the skin catheter “pop”, and this narrowed our search for the catheter “pop” to nearby, subsequent waveforms that were larger.




Figure 2.6: A: As the needle is withdrawn, leaving just the catheter inside the vein, one hand must apply pressure to close the vein so that blood does not escape from the catheter before it is attached to an IV line. B: An example of insufficient pressure on the vein during needle withdrawal allowing blood to escape. The patient was visibly scared by this.

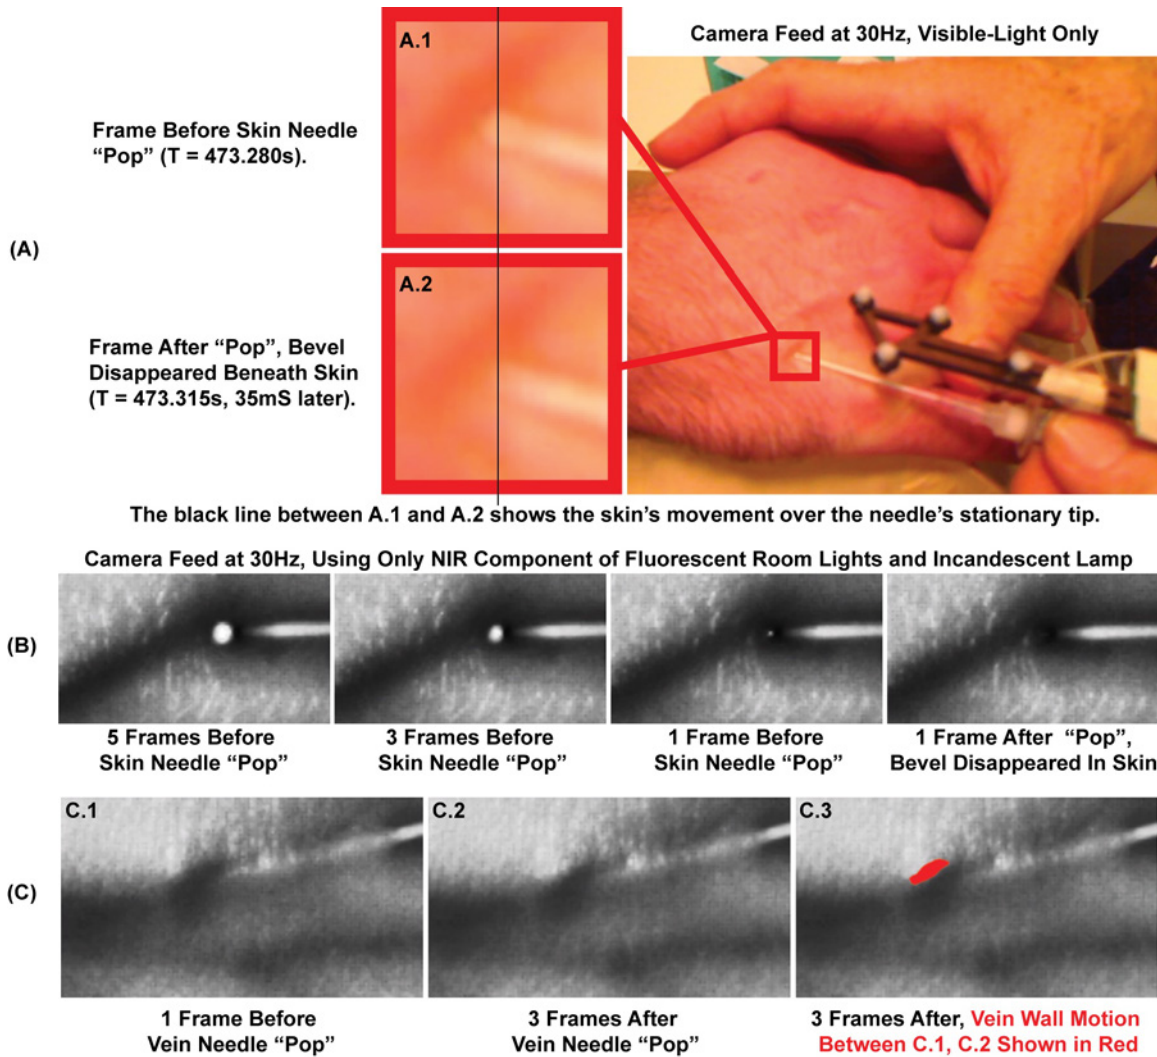


Figure 2.7: Using Video to Identify "Pops" and Prove Tissue Acceleration During "Pops". A-B: Examples of skin needle "pop". C: Example of vein needle "pop". A, C: In these examples, the needle remained stationary while the tissue moved/accelerated. This helps to explain why our needle-mounted accelerometer did not always detect the "pops".

6. We used the blood-flash, seen both in the blood-flash flag and the video feed, as a starting time from which we searched backwards to find the nearest, preceding “pop”-shaped waveform candidate for vein needle “pop”. Searching forwards in time would not have worked because the time in between the skin catheter “pop” and vein needle “pop” contained “pop”-shaped waveforms that were not actual penetration-events, due to the practitioner adjusting the needle’s position and orientation before attempting vein penetration. The infrared video feed was sometimes useful for identifying the vein needle “pop”, as shown in Figure 2.7C, because the vein wall could be seen suddenly relaxing from one position to another while the needle remained essentially stationary. Another way of disambiguating the vein needle “pop” from the middle, false “pops” that arose from the practitioner’s needle adjustments was to check whether the needle was inserting forwards (in the direction of insertion) or backwards/sideways. The middle adjustments often involved moving the needle backwards or to the side, and the penetration-events always occurred during motion in the forward direction.
7. The vein catheter “pop” was the last penetration-event by order-of-operations. The practitioners would typically pause between inserting the needle and beginning to slide the catheter into the vein, and this pause meant that the vein catheter “pop” was usually the last “pop”-shaped waveform during needle insertion, making it easier to identify than the penetration-events in the middle of the insertion. Further, the active-insertion flag and video feed provided time-stamps for this pause from which we could search backwards. Sometimes the blood-flash would be delayed from the vein force “pops” (usually due to a valve), making it more difficult to identify the vein needle and catheter “pops” because they both occurred before the blood-flash. In this case, their relative magnitudes helped identify them because in 90% of the trials, the vein needle “pop” was larger than the vein catheter “pop”.

We made two, key discoveries about parsing the data to identify the penetration-events:

- Key Discovery #4: Identifying the penetration-events is possible for a human, but it requires more sensors than simply the axial force. We also needed the blood-flash flag, video feeds, and sometimes the active-insertion flag and motion data.
- Key Discovery #5: The video camera is one of the most important tools for identifying the penetration-events. Sometimes it can be used as a stand-alone sensor for detecting “pops”, as in the skin and vein needle “pops” shown in Figure 2.7, and sometimes it must be used in conjunction with other sensors.

The following ideas would improve data parsing via video for future clinical experiments and our robot:

1. We could increase the visibility of the needle’s bevel so that its disappearance beneath the skin is more obvious, perhaps by shining light that will reflect from it or dipping the bevel in fluorescent dye.
2. To make visual detection of relative motion between the needle and tissue easier, we could mark a grid of rings on the catheter that would show the length remaining outside of the skin and possibly the vein.
3. To detect skin stretch better, we could apply a grid of permanent marker dots to the skin.
4. To increase vein/non-vein contrast in the infrared camera and facilitate detection of the vein needle and catheter “pops”, we could include narrow-spectrum bandpass filters on the motion capture cameras/IR-illuminators and our NIR transillumination system so that we could use both simultaneously.

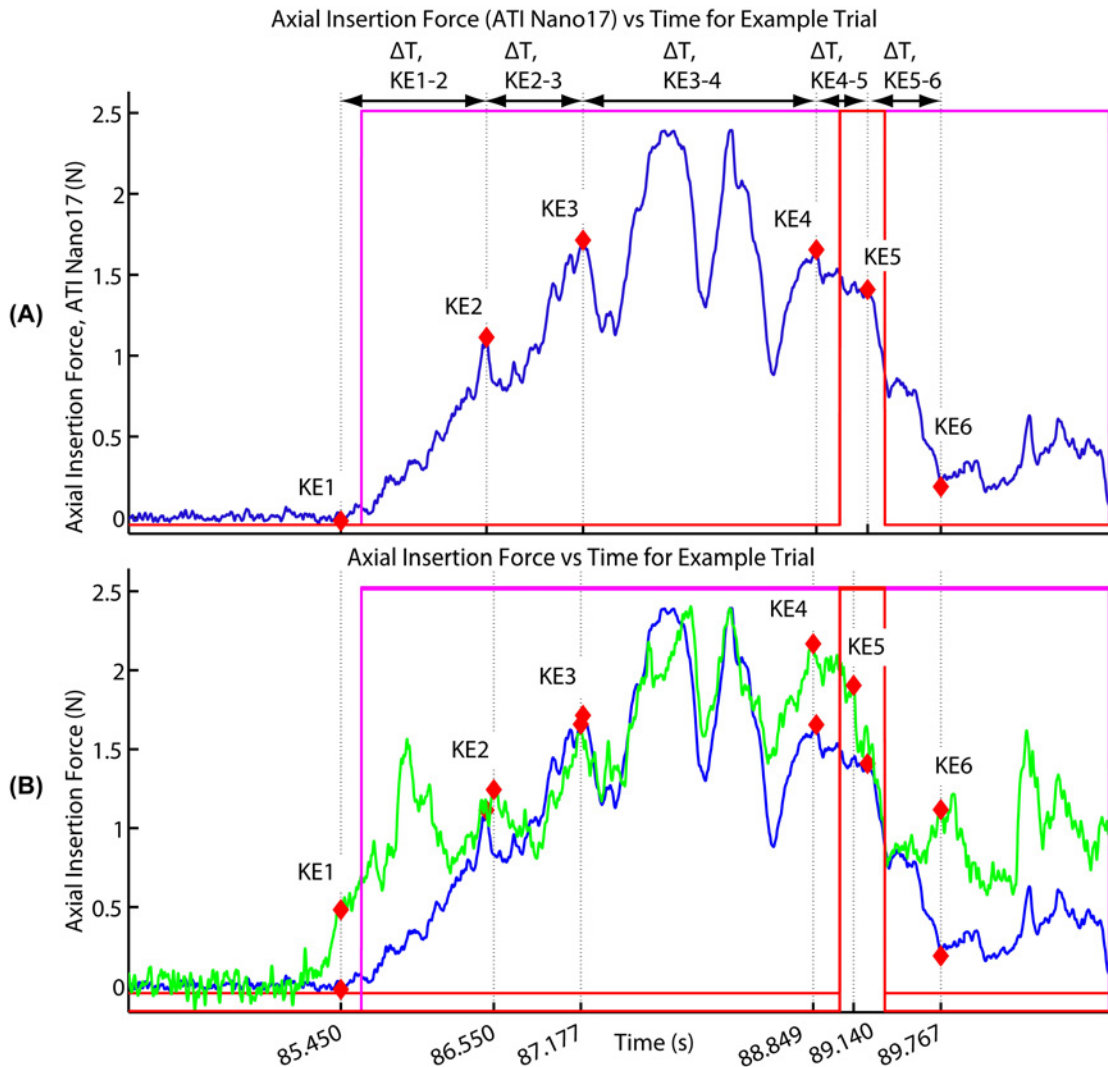


Figure 2.8: Values on this Figure’s axes are for this example trial only. Statistics below were computed across all trials. Red line = blood-flash flag, purple = active-insertion flag, blue = ATI Nano17 (“Nano-Fz”), green = Flexiforce (“FF-Force”).

A: Axial Insertion Force (ATI Nano17) vs Time.

KE1: moment-of-skin-touch Nano-Fz = $-0.041 \pm 0.101\text{N}$. KE2: skin needle “pop” Nano-Fz = $1.008 \pm 0.254\text{N}$.
 KE3: skin catheter “pop” Nano-Fz = $1.357 \pm 0.281\text{N}$. KE4: vein needle “pop” Nano-Fz = $1.978 \pm 0.618\text{N}$.
 KE5: vein catheter “pop” Nano-Fz = $1.462 \pm 0.497\text{N}$. KE6: end-of-needle-insertion Nano-Fz = $-0.021 \pm 0.199\text{N}$.
 $\Delta T, \text{KE1-2} = 1.317 \pm 0.510\text{s}$. $\Delta T, \text{KE2-3} = 0.826 \pm 0.564\text{s}$. $\Delta T, \text{KE3-4} = 7.774 \pm 8.230\text{s}$ ([min,max] = [1.680, 24.123]).
 $\Delta T, \text{KE4-5} = 0.377 \pm 0.258\text{s}$. $\Delta T, \text{KE5-6} = 0.713 \pm 0.510\text{s}$.

B: Axial Insertion Force (ATI Nano17 and Flexiforce Sensor) vs Time.

%Diff. = percent difference between sensors, using the Nano17 as ground truth. ΔT = detection time delay between sensors.
 Corr-R, Corr-P = R, P coefficients for correlation between vectors of the sensors’ values at key moments.
 KE1: FF-Force = $-0.025 \pm 0.433\text{N}$, %Diff. = $1179.188 \pm 5627.448\%$. Absolute Diff. = $0.015 \pm 0.452\text{N}$.
 KE2: FF-Force = $0.929 \pm 0.827\text{N}$, %Diff. = $15.142 \pm 140.073\%$. $\Delta T = -0.039 \pm 0.109\text{s}$. Corr-R = -0.529 , Corr-P = 0.116 .
 KE3: FF-Force = $1.118 \pm 0.639\text{N}$, %Diff. = $-14.367 \pm 59.905\%$. $\Delta T = -0.035 \pm 0.106\text{s}$. Corr-R = 0.170 , Corr-P = 0.638 .
 KE4: FF-Force = $1.809 \pm 0.988\text{N}$, %Diff. = $-10.318 \pm 31.908\%$. $\Delta T = 0.006 \pm 0.017\text{s}$. Corr-R = 0.799 , Corr-P = 0.006 .
 KE5: FF-Force = $1.402 \pm 0.854\text{N}$, %Diff. = $-2.710 \pm 41.886\%$. $\Delta T = 0.059 \pm 0.224\text{s}$. Corr-R = 0.589 , Corr-P = 0.073 .
 KE6: FF-Force = $-0.130 \pm 0.775\text{N}$, %Diff. = $1358.203 \pm 3890.700\%$. Absolute Diff. = $-0.109 \pm 0.708\text{N}$.
 Average correlation between sensor signal curves over entire insertion: $R = 0.879 \pm 0.049$, $P < 0.0001$.

2.3.2.2 Needle Forces, Torques

We made three, key discoveries about needle insertion forces/torques:

- Key Discovery #6: Identifying the “pops” requires axial force but not the non-axial force, torque axes. Non-axial axes might yield other useful information, such as vein motion in a controlled environment like robotic insertion, but they do not help “pop” detection.
- Key Discovery #7: We were always able to identify two “pops” (needle and catheter) per tissue-layer (skin and vein), for a total of four “pops” per insertion.
- Key Discovery #8: The Flexiforce sensor measured axial force fairly accurately as long as torques on the sensor were low. This makes the sensor potentially usable for robotic insertions but unusable for insertions performed by human practitioners because they cannot control the moment loading.

When moments were applied to the Flexiforce sensor, they registered as axial force (at least partially), shifting the force-curve to an extent that it could not be trusted for either absolute values or even overall shape of the waveform. The Flexiforce signal also experienced a time-delay due in part to the analog noise-filtering circuitry, but this could be reduced by further tuning of the circuit’s filter-frequency. As shown in Figure 2.8B, portions of the ATI Nano17 and Flexiforce curves aligned closely, but other portions showed considerable differences in both magnitude and shape. The Flexiforce could be suitable as an inexpensive, disposable sensor on a robot because the robot could control the needle pose to ensure low moments. However, it is not suitable for use by human practitioners who are free to apply arbitrary loads to the needle.

As shown in Figure 2.8, the axial force magnitudes of the four penetration-events were as follows: KE2: skin needle “pop” force = $1.008 \pm 0.254\text{N}$, KE3: skin catheter “pop” force = $1.357 \pm 0.281\text{N}$, KE4: vein needle “pop” force = $1.978 \pm 0.618\text{N}$, and KE5: vein catheter “pop” force = $1.462 \pm 0.497\text{N}$. While these normative values are helpful for establishing the force levels that will inform our robot’s mechanical and sensing design, there is also the question of whether their individual magnitudes can be used to identify them during the data parsing process. We found statistically-significant differences in the axial force means between skin and vein penetration, specifically between 1. skin needle and vein needle “pops” ($P = 0.0001$) and 2. skin catheter and vein needle “pops” ($P = 0.016$). This difference was almost established between skin needle and vein catheter “pops”, but the P-value (0.114) was slightly too high for confirmation. We were extremely close to establishing a statistically-significant difference in the axial force means between needle and catheter penetration of the vein ($P = 0.058$). Practically speaking, our parsing process would only be helped by finding statistically-significant differences between adjacent “pops” in the order-of-operations, so it is most important that we have established, or are near to establishing, statistically-significant differences in the skin catheter/vein needle ($P = 0.016$) and vein needle/vein catheter ($P = 0.058$) “pop” pairs.

2.3.2.2.1 Comparison of Axial Force Results to Other Studies

In [22], the authors compared the axial insertion forces for two different IV catheters, the PUR-Insyte and PTFE-Venflon. Their force-vs-time curves and statistics for latex and lamb skin penetrations showed separate needle and catheter “pops” per tissue-layer (skin and vein), but their actual IV insertions into live human volunteers showed only a single “pop” per layer, referred to simply as “skin” and “vein”. For insertions on

human volunteers, they established a statistically-significant difference ($P < 0.001$) in mean force between the two types of catheter at both skin and vein penetration-events. However, when examining each catheter individually, they did not find a statistically-significant difference in force levels between skin and vein penetration. Their specific force values for insertion into human volunteers were:

- skin penetration (no disambiguation between needle or catheter “pops”)
 - PUR-Insyte catheter: $3.9 \pm 1.6\text{N}$ ($N = 35$).
Their mean is 187.4/286.9% higher than our skin penetration means (needle/catheter).
 - PTFE-Venflon catheter: $2.7 \pm 0.8\text{N}$ ($N = 35$).
Their mean is 99.0/167.9% higher than our skin penetration means (needle/catheter).
- vein penetration (no disambiguation between needle or catheter “pops”)
 - PUR-Insyte catheter: $3.5 \pm 1.2\text{N}$ ($N = 31$).
Their mean is 77.0/139.4% higher than our vein penetration means (needle/catheter).
 - PTFE-Venflon catheter: $2.4 \pm 0.9\text{N}$ ($N = 34$).
Their mean is 21.3/64.2 % higher than our vein penetration means (needle/catheter).

Their results differ from ours in that they documented only a single “pop” per tissue layer, found no statistically-significant difference between skin and vein penetration-event forces, and measured forces that had a higher magnitude and variance than ours. Possible explanations for the differences include:

- Different catheter. Their results showed a significant difference in penetration forces between two catheter types (PUR-Insyte forces were 44.44-45.83% larger than PTFE-Venflon forces), and we suspect that our use of yet another, different catheter contributed heavily not only to the differences in force values, but also in the number of “pops” per layer and difference between skin and vein forces.
- Different force-sensing ergonomics. As discussed in section 2.2.1.1, our instrumented needle retained the on-axis, finger-grip ergonomics of a normal IV needle, whereas their instrumented needle required a whole-handed grip to the side of the needle’s axis, thereby possibly changing the insertion mechanics.
- No cancellation of sensor self-weight. They list a possible maximum error of 0.24N due to “skew forces of the transducer’s own weight and by its rotation when in function”, which leads us to believe that they did not dynamically remove the self-weight component of the transducer and needle (based on pitch angle) from their force signal, as we did.

2.3.2.3 Intravenous Blood Pressure

We made two, key discoveries about the intravenous blood pressure:

- Key Discovery #9: There can be statistically-significant differences between the tourniqueted intravenous blood pressures for tourniquets tied by different practitioners.
- Key Discovery #10: We did not find statistically-significant correlation between intravenous blood pressure and other measured variables, including the axial force magnitudes of “pops”. We found weak correlation ($R = 0.501$, $P = 0.140$) between the intravenous blood pressure and axial insertion force at the vein catheter “pop”, but the P-value was too high to be strictly considered as significant. However, the P-value is small enough that we feel further study is warranted.

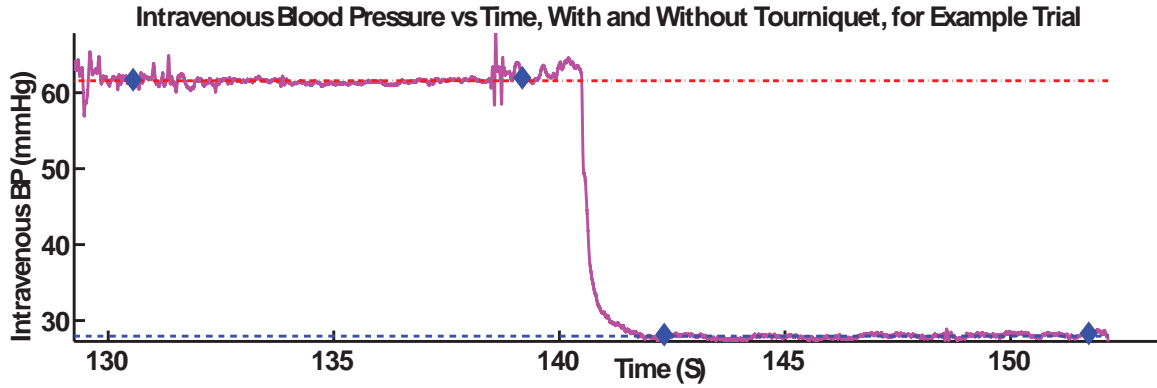


Figure 2.9: Intravenous Blood Pressure vs Time.

Values on this Figure’s axes are for this example trial only. Statistics below were computed across all trials.

Red dashed line = mean BP with tourniquet. Blue dashed line = mean BP without tourniquet.

Across all trials: BP with tourniquet = 50.917 ± 3.705 mmHg, BP without tourniquet = 23.281 ± 2.828 mmHg.

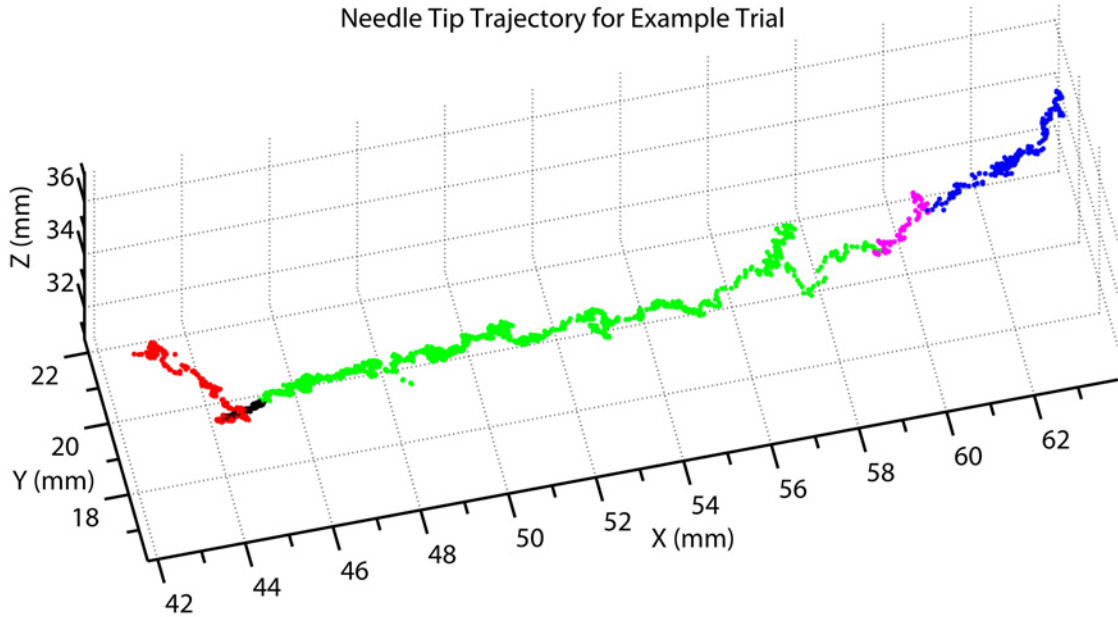


Figure 2.10: Needle Tip Trajectory.

Values on this Figure’s axes are for this example trial only. Statistics below were computed across all trials.

Blue = moment-of-skin-touch (KE1) to skin needle “pop” (KE2). Purple = skin needle “pop” (KE2) to skin catheter “pop” (KE3).

Green = skin catheter “pop” (KE3) to vein needle “pop” (KE4). Black = vein needle “pop” (KE4) to vein catheter “pop” (KE5).

Red = vein catheter “pop” (KE5) to end-of-needle-insertion (KE6).

Tip distance between moment-of-skin-touch and skin needle “pop” (KE1 to KE2) = 3.294 ± 1.268 mm.

Tip distance between skin needle and skin catheter “pops” (KE2 to KE3) = 1.730 ± 1.158 mm.

Tip distance between skin catheter and vein needle “pops” (KE3 to KE4) = 9.152 ± 5.419 mm.

Tip distance between vein needle and vein catheter “pops” (KE4 to KE5) = 1.109 ± 1.007 mm.

Tip distance between vein catheter “pop” and end-of-needle-insertion (KE5 to KE6) = 1.113 ± 1.798 mm.

Tip distance traveled over the entire needle insertion (KE1 to KE6) = 16.397 ± 4.564 mm.

Across all trials and all practitioners, we found a statistically-significant difference ($P < 0.0001$) between the mean, tourniqueted intravenous blood pressure of $51.380 \pm 13.712\text{mmHg}$ ($N = 47$) and the normal, resting pressure (without a tourniquet) of $25.580 \pm 3.483\text{mmHg}$ ($N = 49$). Figure 2.9 shows an example of this difference in pressures with and without a tourniquet. There was not a statistically-significant difference in normal, resting intravenous pressure (without a tourniquet) between patients in the the different practitioner groups. However, we measured a statistically-significant difference ($P = 0.004$) for tourniqueted intravenous pressure between tourniquets applied by Practitioner-1 ($56.269 \pm 12.927\text{mmHg}$, $N = 21$) and Practitioner-4 ($39.751 \pm 12.937\text{mmHg}$, $N = 12$) and a similar difference ($P = 0.046$) between Practitioner-4 ($39.751 \pm 12.937\text{mmHg}$, $N = 12$) and Practitioner-5 ($55.680 \pm 9.112\text{mmHg}$, $N = 8$).

2.3.2.4 Needle Distance and Speed

We made one, key discovery about the needle’s trajectory in terms of distance:

- Key Discovery #11: The needle travels a considerable distance underneath the skin before penetrating the vein. As shown in Figure 2.10, the distance traveled between the skin catheter and vein needle “pops” is the longest distance traveled (by far) for any of the intervals between key events (KE1-KE6).

The needle tip’s velocity vector was roughly in the direction of the vein at all times, with only minor, accidental deviations. As such, we will henceforth refer to the needle tip’s velocity as simply “needle speed”.

We made two, key discoveries about needle speed:

- Key Discovery #12: IV insertion is much slower than we had expected. Averaging across all trials, the maximum speed within an insertion was $4.746 \pm 0.876\text{mm/s}$, and the highest speed at one of the penetration-events was $3.029 \pm 1.099\text{mm/s}$ at the vein catheter “pop”.
- Key Discovery #13: IV insertion is not performed as a single, steady motion, but as a series of motions with pauses and different speeds in between.
- Key Discovery #14: Axial insertion force correlates moderately with insertion speed for most of the needle insertion but not at the key penetration-events.

Setting our robotic insertions at the same, low speeds that we measured for human insertions would:

1. Increase safety due to lower kinetic energy and more time to detect/react to problems, and
2. Make the performance requirements for the robot easier to achieve because our system could use a slower frame-rate for the NIR stereo camera and less powerful motors to decelerate from higher speeds.

As shown in Figure 2.11, there is statistically-significant, moderate correlation (averaged across all trials, $R = 0.630 \pm 0.273$, $P = 2\text{E-}7 \pm 6\text{E-}7$) between the axial insertion force and needle speed over the duration of the insertion. However, if we look at correlation at just the key penetration-events instead of over the entire insertion, we find no such correlation. Of the penetration-events, only the vein catheter “pop” showed correlation values ($R = 0.463$, $P = 0.178$) that even approached the level of statistical-significance. The moderate correlation when examining the entire duration of the insertion makes sense because the tissue is viscoelastic, and much of the needle (not catheter) insertion is comprised of pushing into, but not cutting/penetrating, the tissue. It also makes sense that there exists no, or little, correlation at key penetration-events because those events are comprised of cutting, which uses a different force-model than viscoelasticity.

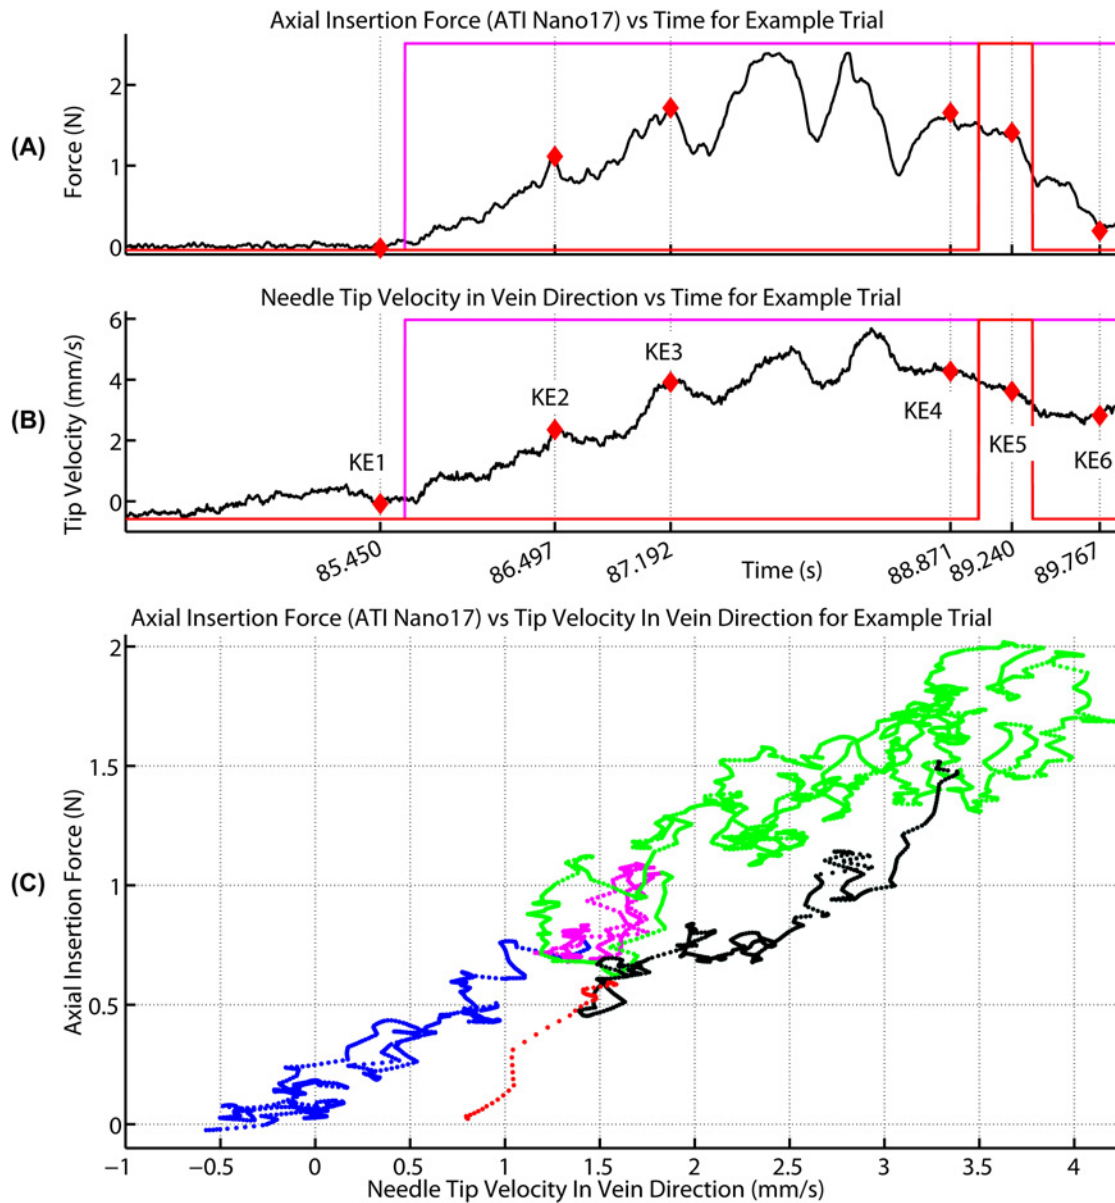


Figure 2.11: We use the term “velocity” to refer to the needle tip’s speed along the vein direction only. Values on this Figure’s axes are for this example trial only. Statistics below were computed across all trials.

A-B: Needle Tip Velocity and Axial Insertion Force (ATI Nano17) vs Time.

Red line = blood-flash state coded manually. Purple line = active-insertion state coded manually.

Correlation between axial insertion force and tip velocity for this example trial: $R = 0.876$, $P < 0.0001$.

KE1: moment-of-skin-touch speed = -0.882 ± 0.847 mm/s. KE2: skin needle “pop” speed = 1.713 ± 0.778 mm/s.

KE3: skin catheter “pop” speed = 2.424 ± 0.876 mm/s. KE4: vein needle “pop” speed = 2.899 ± 1.375 mm/s.

KE5: vein catheter “pop” speed = 3.029 ± 1.099 mm/s. KE6: end-of-needle-insertion speed = 2.105 ± 1.359 mm/s.

C: Axial Insertion Force (ATI Nano17) vs Needle Tip Velocity.

Blue = moment-of-skin-touch (KE1) to skin needle “pop” (KE2). Purple = skin needle “pop” (KE2) to skin catheter “pop” (KE3).

Green = skin catheter “pop” (KE3) to vein needle “pop” (KE4). Black = vein needle “pop” (KE4) to vein catheter “pop” (KE5).

Red = vein catheter “pop” (KE5) to end-of-needle-insertion (KE6).

Correlation for this example trial: $R = 0.930$, $P < 0.0001$. Average across all trials: $R = 0.630 \pm 0.273$, $P = 2E-7 \pm 6E-7$.

2.3.2.5 Needle Angles

We made three, key discoveries about needle angles:

- Key Discovery #15: Roll and yaw do not change appreciably during insertions. Roll keeps the needle approximately “bevel up”, and yaw keeps the needle aligned with the vein direction.
- Key Discovery #16: Pitch starts out high so that the needle does not slip from its initial skin contact point and is lowered after skin penetration to reduce the chance of penetrating the vein’s back wall.
- Key Discovery #17: We found no statistically-significant correlation between pitch angle and axial insertion force at the key penetration-events, with the strongest correlation occurring at the vein catheter “pop” ($R = 0.500$ $P = 0.141$). While the P-value was not low enough to confirm statistical-significance, we feel that it is low enough to warrant further investigation.

We came to these conclusions by computing statistics about the angles over the entire insertion, not just at the key penetration events. The following is a summary for all three angles:

1. Roll

- (a) The average range (max value - min) observed within insertions was $18.902 \pm 16.073\text{deg}$.
- (b) The average standard deviation observed within insertions was $4.477 \pm 2.724\text{deg}$.
- (c) We observed no characteristic roll-vs-time profile across different insertions.

2. Yaw

- (a) The average range (max value - min) observed within insertions was $10.337 \pm 4.120\text{deg}$.
- (b) The average standard deviation observed within insertions was $2.776 \pm 1.261\text{deg}$.
- (c) We observed no characteristic yaw-vs-time profile across different insertions.

3. Pitch

- (a) The average range (max value - min) observed within insertions was $19.121 \pm 6.070\text{deg}$.
- (b) The average standard deviation observed within insertions was $5.576 \pm 1.901\text{deg}$.
- (c) Pitch drops from an initial, high insertion angle for skin penetration (skin needle “pop” pitch = $17.316 \pm 6.095\text{deg}$, skin catheter “pop” pitch = $14.199 \pm 8.353\text{deg}$) to a lower insertion angle for vein penetration (vein needle “pop” pitch = $2.686 \pm 4.891\text{deg}$, vein catheter “pop” pitch = $2.427 \pm 5.206\text{deg}$), as shown in Figure 2.12A. While there is a statistically-significant difference between the mean pitch used for skin and vein penetration ($P < 0.001$ for all permutations), there is no difference in pitch between needle and catheter penetrations within each tissue layer.

2.3.2.6 Accelerometer

We made two, key discoveries from the accelerometer:

- Key Discovery #18: The accelerometer was not suitable for penetration-event detection, having both high false-positive and false-negative detection rates, with examples and statistics provided in Figure 2.12B. For the events that were detected, they were visible only in the accelerometer axis aligned with the needle axis (in our case, the Z-axis).
- Key Discovery #19: The accelerometer measures the pitch angle very accurately. As shown in Figure 2.12, the pitch-angle-vs-time and accelerationZ-vs-time curves are almost identical in overall shape.

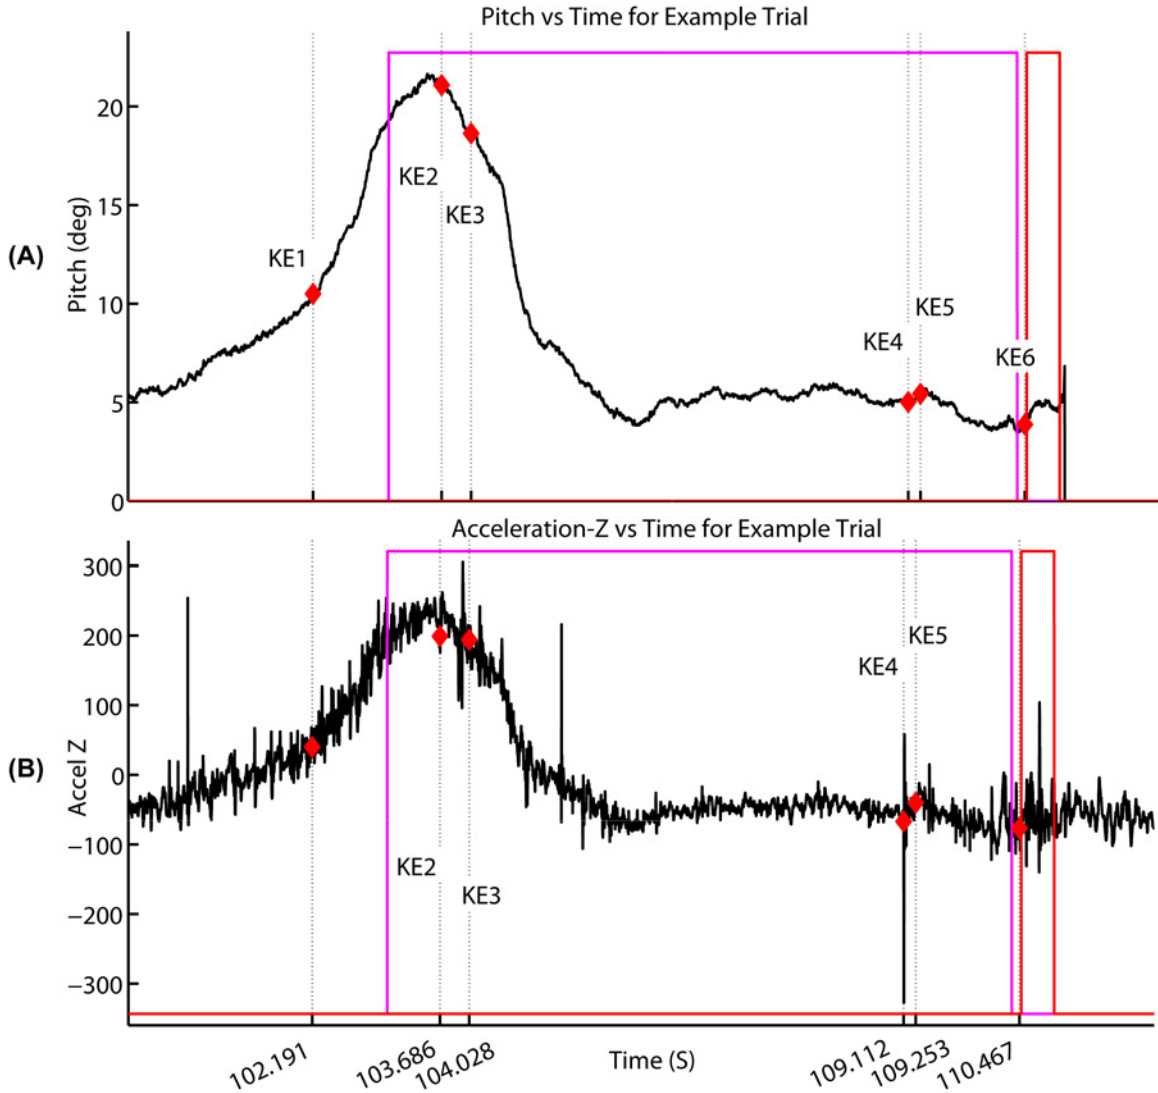


Figure 2.12: Values on this Figure’s axes are for this example trial only. Statistics below were computed across all trials. Red line = blood-flash state coded manually. Purple line = active-insertion state coded manually.

A: Pitch Angle vs Time.

KE1: moment-of-skin-touch pitch = 13.205 ± 7.451 deg. KE2: skin needle “pop” pitch = 17.316 ± 6.095 deg.
 KE3: skin catheter “pop” pitch = 14.199 ± 8.353 deg. KE4: vein needle “pop” pitch = 2.686 ± 4.891 deg.
 KE5: vein catheter “pop” pitch = 2.427 ± 5.206 deg. KE6: end-of-needle-insertion pitch = 2.565 ± 4.323 deg.

B: Acceleration Along Needle Axis (“Accel Z”) vs Time.

For the penetration-events that were detected, they were visible only in the accelerometer axis aligned with the needle axis. Acceleration was not considered for KE1, KE6 because they are not penetration-events.

KE2: skin needle “pop” accelerometer detection rate = 40%. KE3: skin catheter “pop” accelerometer detection rate = 20%.
 KE4: vein needle “pop” accelerometer detection rate = 50%. KE5: vein catheter “pop” accelerometer detection rate = 20%.
 Number of false-positive “pop” detections per insertion: 1.9 ± 1.8 “pops”.

A-B: The accelerometer gives us an accurate measurement of pitch, as shown by the strong resemblance between the curves.

The false-positive rate is due to the fact that the needle's motion was not smooth, and many of the sudden accelerations were unrelated to actual penetration. In our later testing of the accelerometer on robotic insertions, to be described in Chapter 4, the accelerometer no longer experienced false-positives because the robot inserted the needle in a fluid, continuous motion.

Our best explanation as to the false-negative rate is that it is often the tissue, and not the needle, that is often undergoing the bulk of the acceleration during the penetration-events. Figures 2.7A, C show examples of both a skin and vein needle "pop" wherein the needle remains essentially stationary while the tissue jumps from its compressed, pre-penetration position to its more relaxed, post-penetration position.

Measuring the pitch-angle with the accelerometer removes the need for the use of motion capture to measure pitch. If we switched to an accelerometer board that also contains a magnetic-compass (to allow for singularity-free measurement of all angles), we could remove all but a single motion capture marker. This single marker would serve as the origin of our local frame and would be placed as the last component behind the instrumented needle along its axis. This would greatly simplify the motion capture post-processing and pose-calculation process, allowing for real-time pose-tracking of the needle in future clinical experiments.

2.3.2.7 Variable Dependence on Age, Sex

We examined the effects of patient age and sex on all of our measured variables, finding results that approached or achieved statistical-significance in the following variables:

1. Intravenous blood pressure. There was a statistically-significant difference ($P = 0.016$) in normal, resting intravenous pressure (without a tourniquet), between male ($24.792 \pm 3.420\text{mmHg}$, $N = 35$) and female patients ($27.568 \pm 2.896\text{mmHg}$, $N = 12$), but not for tourniqueted pressure.
2. Axial insertion force.
 - (a) Correlation between age and axial force at skin catheter "pop": $R = 0.480$, $P = 0.160$.
 - (b) Correlation between age and axial force at vein catheter "pop": $R = 0.510$, $P = 0.132$.
 - (c) Hypothesis that axial force at skin catheter "pop" is different for men, women: $P = 0.185$.
3. Clinical failure rate. Examining all "Simple-needle" and "Nano-needle" trials, we found a difference in clinical failure rates between men (failure rate = $4/43$, or 9.3%) and women (failure rate = $4/20$, or 20.0%), but the difference was not statistically-significant ($P = 0.242$).

The P-values for the age/sex correlations with axial force are not low enough to be statistically-significant, but they are low enough that we feel they warrant further investigation, particularly in light of suggestions by others researchers about such relationships. In [22], the authors established statistically-significant, moderate correlation between higher age and lower penetration force into the skin ($R = 0.47$, $P < 0.004$) and vein ($R = 0.50$, $P < 0.002$) for one type of catheter (but not another type). Although their correlation runs in the opposite direction (lower forces for older patients) from the tentative relationship we found, it is still confirmation that age effects insertion force. They also found that men experienced higher insertion forces than women (median difference = $0.54\text{-}0.60\text{N}$ for skin and $0.19\text{-}0.92\text{N}$ for vein), but the difference was not statistically-significant. If in the future we confirm such relationships, then the robot could be programmed with penetration-event detection parameters that take age and sex into account.

2.3.3 Observed Clinical Failure Modes and Recovery Techniques

In the course of 63 insertions total (including all “Simple-needle” and “Nano-needle” insertions), we experienced 8 instances of clinical failure. The following list, which will henceforth be known as the “Clinical-Failure-List”, details the type and frequency of specific failure modes:

1. Two failures from the catheter being stuck outside of the vein wall (unable to slide off of the needle that had already penetrated the vein),
2. Three failures from back-wall penetrations not involving a valve (two of these due likely to vein rolling and one due likely to venospasm).
3. One failure from a valve blocking the catheter’s tip from sliding deeper into the vein after successful initial entry past the vein’s wall, and
4. Two failures from back-wall penetration due to the needle encountering a valve.

There were 2 failed insertions for the “Simple-needle” (failure rate = $2/27$, or 7.4%), and both failed in the mode of Item 2, one likely due to vein rolling and one likely due to venospasm. There were 6 failed insertions for the “Nano-needle” (failure rate = $6/36$, or 16.7%), and these failures spanned the modes of Items 1-4, with the Item 2 failure due likely to vein rolling.

We observed the practitioners’ techniques to recover from these failures and save the insertion, and we describe them below. Since we observed small numbers of each, further study would be required to understand all of their mechanics and how they would influence our robotic insertions. However, we include their description because for many of these cases this is the first time that they have been described in the literature.

2.3.3.1 Failure Mode: Catheter Stuck Outside Vein Wall

In Item 1 of the the “Clinical-Failure-List”, the catheter is stuck outside of the vein wall after the needle’s bevel/tip has pierced the vein, usually because the catheter tip is damaged so that it snags on the vein wall. In this scenario, the practitioner sees the initial blood-flash but might not feel the “pop”. We observed two solutions for this scenario, including:

1. Pushing hard on the catheter until it enters the vein, and
2. Spinning the catheter about the needle while pushing the catheter forward into the vein. Sometimes the interference between the catheter’s tip and the vein wall at one particular spot prevents the catheter from entering, and spinning the catheter can shift that spot sufficiently to allow entry.

Since the needle acts as a guidewire for the catheter’s entry into the vein, great care must be taken not to pull the needle out of the vein while trying to force the catheter into the vein. The robot would be better able to keep the needle stationary while manipulating the catheter than a human practitioner.

2.3.3.2 Failure Mode: Back-Wall Penetration

In Item 2 of the the “Clinical-Failure-List”, the needle has penetrated the front-wall of the vein, into the lumen, and through the back-wall such that the entire vein is skewered by the needle. Such an insertion may be saved by pulling the catheter back slowly until the tip is again in the lumen and then pushing the catheter forward at a slightly different pitch angle until it threads normally into the vein. We observed two instances of using this method to save insertions with back-wall penetration. However, the insertion may still fail, as

we observed in one case, because leakage from the hole in the back-wall can cause the surrounding tissue to swell considerably with blood or IV fluids. We would not use this strategy for the robot because the chance of success is outweighed by the risk of developing a massive bruise, or permanent tissue damage in the case of leaking certain medications. Such a bruise can often be prevented or mitigated by aborting the insertion and applying prolonged pressure to the insertion site to allow the blood to clot around the leaking hole.

2.3.3.3 Failure Mode: Valves Causing Blockage, Back-Wall Penetration

In Item 3 of the “Clinical-Failure-List”, the catheter is blocked by an internal valve that fails the insertion either by blocking blood-flow into the catheter’s tip or by blocking the catheter from entering sufficiently far into the vein to prevent the catheter from dislodging easily. In Item 4, the needle tip encounters the valve, and the struggle to push past it leads to a back-wall penetration. While valves were encountered frequently (15/63 insertions, or 23.81%), only three of these encounters (20%) caused failure because practitioners have the following techniques for handling valves:

1. Pulling the catheter back slowly until blood return increases because the catheter’s tip is no longer pushed against the valve (solves Item 3). This must be done carefully because it can result in pulling the catheter out of the vein and losing the insertion. Further, enough of the catheter must be threaded into the vein that it will withstand several days of patient movement without dislodging.
2. Injecting saline via the catheter to open the valve while pushing the catheter through (solves Item 3).
3. Pushing hard past the valve (solves Item 4). This must be done carefully so as to avoid back-wall penetration of the vein and may be a safer strategy for the robot than a human practitioner. Using its force sensor, the robot could detect when the needle penetrates beyond the valve and stop the needle before it penetrates the back-wall.
4. Pushing while spinning the needle to drill past the valve (solves Item 4). The robot would be able to spin the needle much faster and with greater control than a human.

The practitioner suspects that a valve has been encountered when the needle or catheter, once confirmed in the vein by the initial blood-flash and “pop”, cannot be inserted further into the vein or if there is not strong blood return in the catheter after needle withdrawal. However, the following two scenarios mimic this latter case, and they require different solutions:

1. Back-wall penetration. The practitioner felt the “pop” and saw the initial blood-flash during penetration of the front vein wall but then pushed the needle all the way through the vein such that the needle did not have access to the vein’s blood supply to provide strong, sustained blood return. As stated in section 2.3.3.2, saving a back-wall penetration is sometimes possible but not advisable.
2. Tissue blocking the needle bore. Sometimes a piece of tissue blocks the bore of the needle such that blood cannot flow strongly past it. To solve this, the practitioner withdraws the needle from the catheter while keeping the catheter in the vein, checks for and clears the blockage, and pushes the needle back into the catheter/vein. The robot could detect such blockage indirectly by watching the catheter’s hub for blood return after removing the needle. If the needle was blocked, then sterile, compressed air could be used to remove the blockage while the needle is outside of the vein.

2.3.3.4 Possible Failure Mode: Venospasm and Disappearing Veins

Some patients have sensitive veins that constrict when aggravated (venospasm), sometimes shrinking to the extent that they are no longer visible or palpable. We encountered venospasm in 11.11% of insertions (7 out of 63), and only 14.29% (1 out of 7) of those insertions failed. In the single failed case, a back-wall penetration was the ultimate cause of failure, and it is unclear whether venospasm caused that back-wall penetration. The practitioner felt a “pop” but did not see blood-flash during initial insertion. After determining that a back-wall penetration had occurred, the practitioner pulled the needle backwards out of the vein and experienced a blood-flash during this withdrawal. Our theory is that venospasm caused the vein to collapse such that the front and back walls were penetrated as a single wall with no blood in between them to enter the needle and cause an initial blood-flash. As the needle was withdrawn, friction between the exiting catheter and the front vein wall helped pull open the vein’s lumen, allowing blood to enter the needle for blood-flash.

Even if venospasm does not lead to outright failure, it can still increase the difficulty of the insertion. In two of the venospasm cases, veins that had been initially visible became completely invisible, and in an additional case, the veins became significantly harder to see. The robot would still be able to see a constricted vein as long as some blood remained in its lumen, but it is unclear whether the vein would be large enough for successful cannulation based on the needle size and the NIR stereo camera’s depth resolution. The robot should be able to detect constriction by continually tracking the vein’s diameter with the NIR stereo camera, and the robot could refuse to insert on veins that shrink below its needle-size or resolution limit.

2.3.4 Technical Explanations for Conventional Wisdom on Bifurcations, Traction

The conventional wisdom of many practitioners is that 1. inserting at a venous bifurcation is easier than inserting along the midsection of a single vein [52], in part because the bifurcations tend to roll less, and 2. applying traction to the dorsal hand skin also reduces vein rolling, making insertions easier. In the course of observing many IV insertions and practicing insertions with our robot, we were able to help explain the feasibility of such conventional wisdom from a technical perspective.

We developed three explanations for why targeting bifurcations could make insertions easier, including:

1. Larger cross-sectional area for easier targeting. As shown in Figures 2.13A-B, bifurcations provide a larger cross-sectional area than does a single segment of vein with the same internal diameter. The diameter and wall-thickness of this, particular bifurcation were measured using a 2D-ultrasound on the author’s veins, and we set the internal-branching angle of our model to 45 degrees for symmetry. For this specific geometry, the bifurcation’s area was larger than that of the corresponding single section of vein by a factor of 2.6. In our teleoperated insertions using HaemoBot, described in Chapter 4, the larger targeting area allowed for needle/vein alignment that was significantly faster and more reliable.
2. Providing a “skin valley” that guides the IV needle and decreases initial needle slippage. The skin droops somewhat between the branches of a bifurcation, providing a valley in which the IV needle can rest. This valley helps guide the needle into the vein and allows the needle tip to approach the bifurcation’s inner surface from a lower position and the skin from a higher angle, thereby reducing the chance of slipping from the intended insertion spot. Figure 2.13C shows this process.




Figure 2.13: Technical Explanations for Choosing Bifurcations and Applying Skin Traction. A-B: Bifurcations provide a larger target cross-sectional area than does a single segment of vein with the same internal diameter. Light, transparent blue = vein wall. Dark purple = vein lumen/blood. Green curve/area = target envelope for a single segment (offset from the lumen’s edge to compensate for the IV catheter’s diameter). Red curve/area = target envelope for a bifurcation. C: The skin droops somewhat between the branches of a bifurcation, providing a valley in which the IV needle can rest. This valley helps guide the needle into the vein and allows the needle tip to approach the bifurcation’s inner surface from a lower position and the skin from a higher angle, thereby reducing the chance of slipping from the intended insertion spot. The left-most image displays the bifurcation with the skin removed. The middle image shows the needle on top of the base/distal branch of the bifurcation such that it has no guidance from the “skin-valley”. The right-most image shows the needle depressed into the bifurcation’s “skin valley” with the needle’s tip gaining traction against the inner surface of the bifurcation. D: When veins are modeled as a network of 1D springs (axial stiffness “k” only and no transverse stiffness), we see that bifurcations provide 2D stiffness that resists all planar motions, including sideways motion along the X-axis, whereas a single section of vein provides only 1D stiffness that allows the vein to deflect sideways easily (note the left-most column of zeros in its stiffness matrix). E: Applying skin traction removes slack from, or pre-stretches, the skin and vein network to prevent such motion from occurring during the actual needle insertion when it would produce a moving target. It also increases the skin’s stiffness due to nonlinear elasticity, thereby decreasing elastic displacement (separate from the initial slack) due to the needle’s pre-penetration insertion force.

3. 2D instead of 1D stiffness. When veins are modeled as a network of 1D springs (assuming that their axial stiffness is much greater than transverse stiffness), bifurcations have a 2D stiffness, whereas single sections of vein have only a 1D stiffness. Figure 2.13D shows the spring models and corresponding 2D stiffness matrices for both cases. In our robotic insertions, to be described in Chapter 4, we found that our latex bifurcation moved sideways under the load of the IV needle far less than did a comparably sized, single section of latex tubing, confirming our 2D-vs-1D stiffness theory.

Applying skin traction removes slack from, or pre-stretches, the skin and vein network to prevent such motion from occurring during the needle insertion when it would produce a moving target. As shown in Figure 2.13E, applying traction can move the skin a distance several times the diameter of a vein. Further, due to the nonlinear elasticity of skin [17], applying traction increases the skin's stiffness, thereby decreasing elastic displacement (separate from the initial slack) due to the needle's pre-penetration insertion force.

2.4 Discussion and Thesis Contributions

We made five main, novel contributions in Chapter 2:

1. Empirically-determined a first-of-its-kind, complete list of human IV insertion steps that provides a framework for the robot to perform the whole process, not just a small portion.
2. Developed a novel methodology for parsing IV insertion data to identify the four penetration-events, or "pops", from a real IV insertion, proving that a human can parse insertions and providing a framework for robotic parsing.
3. Empirically-measured normative values for force, motion, and intravenous pressure overall and at key moments in real IV insertions on live patients using a uniquely-ergonomic, instrumented IV needle. These specific values provided 19 "Key Discoveries" that explain properties of IV insertion more broadly (such as non-axial forces/torques being unnecessary and insertion being slow and interrupted).
4. Observed and analyzed clinical failure modes and recovery techniques, many of these never before documented or explained, that could be used by the robot to prevent and recover from failures.
5. Developed three novel, technical explanations for why bifurcations can be easier insertion targets than single, non-branching sections of vein and two explanations for applying skin traction.

Given the complexity of this experiment, it is tempting to desire complicated or far-ranging conclusions. We would like to understand every aspect of every variable and to establish the "prototypical insertion", or how, on average, every single insertion could be expected to proceed. However, it is important to recognize the limitations of what claims we can make conclusively from the data and what desired claims would require another, different experiment, to substantiate them. In this experiment, we discovered that completely understanding the relationships between certain variables (such as force and speed) and generating a model for a "prototypical insertion" are impossible using human practitioners because of issues with variable isolation and varying techniques/strategies that are somewhat opaque to us. Although this experiment did not answer all of our questions, it laid the framework for an additional clinical study to be performed with the robot that will provide the necessary control and repeatability to provide the missing answers.

Chapter 3: Development of Vein-Imaging Hardware and Software

3.1 Motivation and Overview

Recall that one of the main difficulties in IV insertion is locating the vein. In this chapter, we solve this by proving that a robot can autonomously locate a vein and compute a 6D needle pose for insertion. To accomplish this, we

- developed a custom, NIR stereo camera and LED Grid, shown in Figure 3.1, that generated uniquely accurate, high-resolution, high-contrast, and 3D vein point-clouds, shown in Figure 3.14, and
- developed a novel algorithm for locating, characterizing, and tracking venous bifurcations in a 2D image, as shown in Figure 3.19, and an algorithm for generating a 6D needle pose for those detected bifurcations, as shown in Figure 3.21.

3.1.1 The Clinical Problem

The two standard methods of vein-finding are visual-detection with the naked-eyed and palpation (feeling with the finger tips). However, with the exceptions of non-obese, fair-skinned, adult males and geriatric patients, one or both of those methods are often inadequate in many demographics, including:

- Pediatric. The veins are often too small and underneath a thicker layer of fat to be seen or felt.
- Adult Female. The veins are often too small to be seen, but can usually be felt.
- Obese. The veins are often underneath a thicker layer of fat such that they cannot be seen or felt.
- Dark-skinned. The veins often cannot be seen due to the skin's darker pigmentation.

There is no inherent sensing problem with a practitioner detecting the veins visually, other than a partial lack of depth perception, but it only works for a moderately-small subset of patients. However, localizing a vein through palpation is possibly the most problematic of all vein-sensing methods as it suffers from:

1. spatial imprecision (both in terms of semi-planar location and depth),
2. an inability to get real-time, updated information during insertion in case of vein movement/rolling,
3. the “soda straw effect” in which the finger only samples a small subset of the venous network such that it cannot easily compare different insertion sites or locate larger macrostructures like bifurcations (similar to a camera with a very narrow Field-of-View),
4. the “observer effect” in which touching the vein to find its location pushes the vein into a different, sensed location that only lasts while the finger's pressure is on the vein,
5. the conflicting needs to feel the veins without gloves for increased sensitivity and to wear gloves for sterility/infection-control concerns, and
6. the difficulty for human practitioners of mapping and utilizing non-visual sensory data.

Some of the problems described in this list, which will henceforth be referred to as the “Vein-Sensing-Problems-List”, are shared by other vein-sensing technologies that will now be discussed as alternatives.

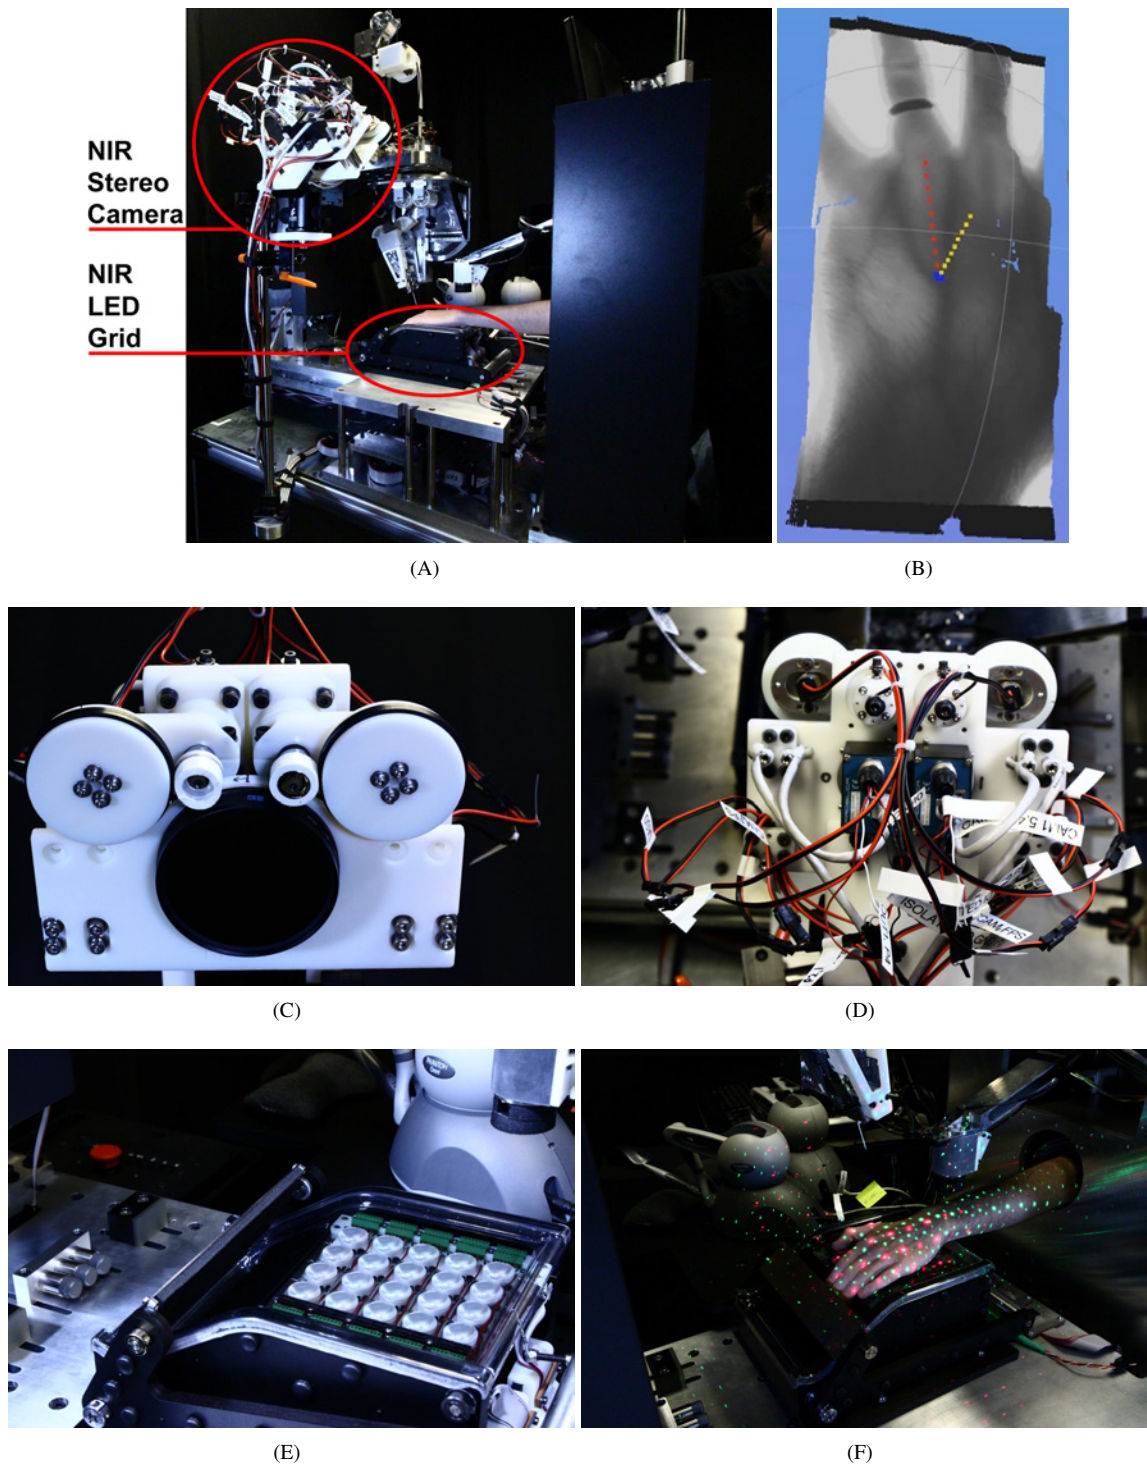


Figure 3.1: NIR Vein-Imaging System. A: Imaging system on HaemoBot. B: 3D NIR point-cloud with automatically detected bifurcation and calculated 6D needle pose (shown as red line). C-D: Custom, NIR stereo camera. E-F: Custom, NIR LED Grid. F: Stereo camera's texture projection for improving the point-cloud's density.

3.1.2 Vein Sensing Modalities and Prior Work

In designing our robot’s vein-sensing capabilities, we considered the standard, human approaches to see if they could be improved via robotics, as well as additional technologies that are typically reserved for more difficult insertions. These approaches included:

- Visible Light (“Naked-Eye”) Viewing. Human eyes are remarkably capable in the visible light spectrum such that a robot would not offer an advantage here, and visible light viewing works for only a small subset of patients.
- Palpation. Although robotic palpation can improve or solve problems 1, 5, and 6 of the “Vein-Sensing-Problems-List”, it cannot solve problems 2, 3, and 4. Nevertheless, several research groups have experimented with palpation-based vein-detection for blood-drawing robots [68] [69] [60] [39], with uniformly disappointing success rates and sensing precision. Further, patients reported discomfort/pain from the palpation process.
- 2D-Ultrasound. In this technology, a small transducer is placed in contact with the tissue (with an acoustic transmission gel in between), and sound waves are projected into the tissue and used to visualize a 2D, wedge-shaped slice of the tissue directly beneath the transducer. Figure 3.4A shows a 2D-ultrasound view of one of the author’s hand veins as an example. Some practitioners use 2D-ultrasound, typically when vision and palpation fail, because it works across all patient demographics, gives a precise depth estimate, disambiguates between veins and arteries/nerves, and helps to align the needle with the vein using a special jig. It should be noted that the 2D image does not show the IV needle with respect to the vein before or during insertion until the moment that the needle’s tip penetrates the 2D imaging-plane of the transducer.

Several research groups have incorporated 2D-ultrasound into their blood-drawing [53] [27] [12] and IV robots [16], but only one robot [16] used 2D-ultrasound as the primary vein-detection method. All of the other robots used Near Infrared (NIR) imaging as the primary vein-detection method and 2D-ultrasound to provide depth measurement and confirmation of sufficient blood-flow for sampling. Such use of ultrasound as an auxiliary sensor that confirms a suitable insertion location but uses NIR imaging for pre-insertion vein-detection and real-time guidance during insertion seems reasonable. However, if relied on as the only method of vein-sensing, 2D-ultrasound suffers from problems 1, 2, 3, and 5 of the “Vein-Sensing-Problems-List”, with problems 1 and 2 being present because the transducer is providing precise, real-time spatial measurements of the slice directly underneath the transducer and not where the needle tip is actually interacting with the skin and vein over the majority of the insertion. It is possible that the use of 3D-ultrasound, which provides a 3D, wedge-shaped sensing volume underneath the probe, could increase the proportion of the insertion for which the vein and needle tip can be visualized simultaneously, but it would still be unable to visualize the needle entering the skin and potentially part of the vein.

- 1D Doppler Ultrasound/“Smart Needle”. The “Smart Needle” [65] uses Doppler ultrasound to estimate fluid velocity in front of the probe along a single axis. This technology suffers from all of the same

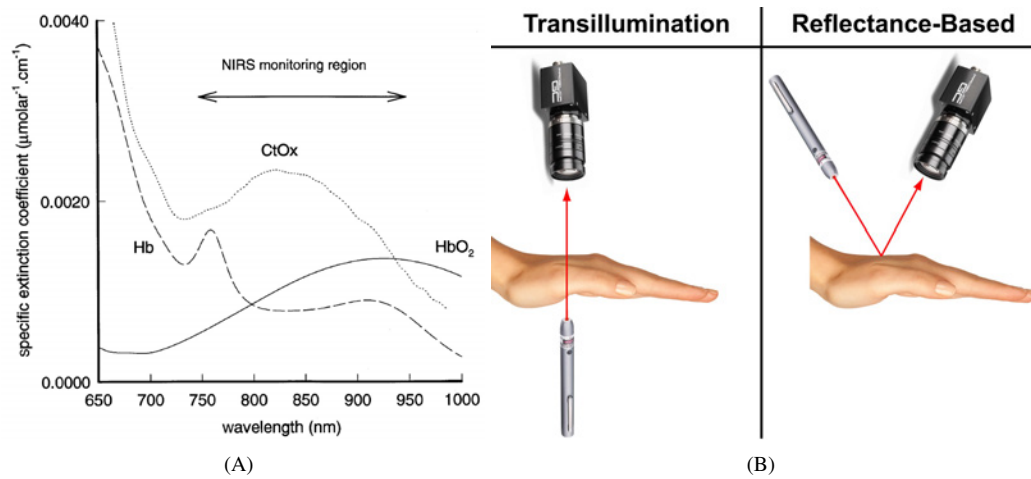


Figure 3.2: A: NIR light absorption-coefficient (“specific extinction coefficient”) versus wavelength for blood. Hb is deoxygenated/venous blood, and HbO₂ is oxygenated/arterial blood. Our imaging system works near the 760nm-peak in the Hb curve to maximize the vein’s contrast with respect to all other tissues, including arteries. Reprinted from [18] with permission. B: Transillumination vs reflectance-based camera/light placement configurations. Reflectance-based systems struggle with shadows and false-vein detection, so our system uses transillumination.

problems as 2D/3D-ultrasound but with greater severity due to its uniaxial operation. Further, obtaining sufficient venous blood-flow in a tourniqueted arm may render it unusable.

- Thermal Paper/Temperature-Sensitive Films. Films that change in color according to local temperature variations can be placed on the hand to show the venous network [45]. However, in addition to suffering from some of the same issues as palpation (items 1, 2, and 5), the needle must penetrate through the film, which could damage the delicate catheter tip.
- Near Infrared (NIR) imaging. The NIR spectrum of light shows veins clearly as dark structures on a light background of the hand, as seen in Figure 3.1B, even in patients with small veins, fat hands, or dark skin for whom visible light cannot show the veins. Further, the veins can be mapped in 3D by using one NIR camera with structured light or two NIR cameras as a stereo pair. As such, NIR imaging is the most common vein-sensing method in other researchers’ work on robotic blood-drawing [48] [49] [50] [51] [59] [32] [53] [27] [12]. There are also many commercially-available vein-finding devices that employ NIR imaging, and we provide a detailed analysis of them in Chapter 5. Stereo NIR solves every problem in the “Vein-Sensing-Problems-List” because it is able to work for all patient demographics, precise, able to provide depth information, real-time, wide in its Field-of-View, non-contact, and capable of visual confirmation by a human practitioner.

3.1.3 Using Near Infrared (NIR) Light For Vein-Visualization

3.1.3.1 Explanation of Why Near Infrared Light Shows Veins

Light in the wavelength range 600 to 1300nm (often referred to as the “Near Infrared (NIR) window”, “optical window”, or “therapeutic window”) interacts with tissues in the hand in special ways [3] [66] [35]

[63] that allow the veins to be seen in high contrast to the surrounding tissues.

Only 4-7% of light is reflected from the skin's surface, with the remaining 93-96% of light penetrating into the skin and possibly the tissue beneath. Once light enters the tissue, a combination of scattering (bouncing in a different direction) and absorption determines how deeply it penetrates. In the visible spectra with wavelengths lower than those in the NIR window, most of the light is quickly absorbed by the haemoglobin and melanin in the skin so that the light does not penetrate deeply. Similarly, in the infrared spectra with wavelengths higher than those in the NIR window, most of the light is quickly absorbed by water. However, in the NIR window, absorption levels by all tissues are much lower such that forward-scattering dominates over absorption [13] [31]. As such, the light is able to penetrate much deeper and scatter over considerably longer distances before being absorbed or exiting the tissue on the opposite side (having forward-scattered through the entire tissue layer).

The hand is composed of many tissues with different tissue-and-wavelength-dependent absorption coefficients, and even small differences in those coefficients are magnified by the longer optical path provided by scattering. What allows us to "see" the veins is that the absorption coefficient for haemoglobin in blood is higher than that of surrounding tissues [18] [41], so the blood absorbs more light and appears as dark on a light background. The absorption coefficients also differ between the deoxygenated haemoglobin (Hb) in the veins and oxygenated haemoglobin (HbO₂) in the arteries. Figure 3.2A shows the absorption-coefficient (also called the "specific extinction coefficient") versus wavelength curves for Hb and HbO₂, including the wavelength at which their absorption is equal (806nm, or the "isobetric point") and the wavelength which maximizes venous/arterial-blood contrast (the Hb-curve peak at 760nm). Accordingly, we designed our vision hardware to use light near this wavelength of 760nm to maximize vein contrast with respect to all other tissues, including arteries.

It should be noted that there are two commercially-available vein-finding devices, the Phillips WeeSight and Translite VeinLite-LED, which use red (approximately 629nm) visible-light for imaging veins. Such red light operates in the same manner just described for NIR imaging, albeit less efficiently, and can be seen by the human eye without assistance. However, it provides greatly-reduced vein/non-vein contrast and requires a significantly brighter light source (and corresponding increase in eye-safety risk) to achieve the same level of brightness as with NIR imaging (to the point that the room must be completely dark for the veins to be seen). In an assistive device that attempts to let the human practitioner retain his/her natural vision, red light could make sense for very thin hands that require more shallow tissue penetration, but in a robotic system that inherently requires a camera, the choice of NIR light is far better.

3.1.3.2 Transillumination vs Reflectance-Based Camera/Light Configurations

As shown in 3.2B, there are two main camera/light configurations for NIR imaging: transillumination and reflectance-based (sometimes referred to as "Epi-Illumination"). In transillumination, the light-source is underneath the palm of the hand and shines light through the hand to an NIR camera on the dorsal side. In the reflectance-based configuration, the light-source and camera are both on the dorsal side of the hand, and the light-source bounces light off of the hand's surface and into the camera.

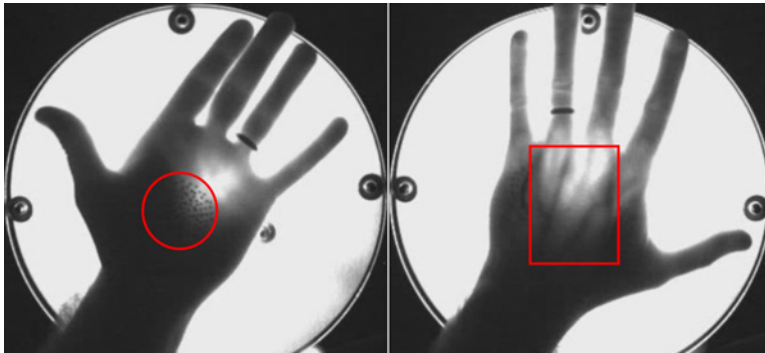


Figure 3.3: Example of Light Scattering Hiding Deeper Structures/Arteries. Permanent marker dots and veins (both highlighted in red) can only be seen when on the surface of the hand closest to the camera (when they are superficial). As the hand is transilluminated from the opposite side and the dots/veins become “deeper” structures, they disappear from the image. Dots shown in left image only, veins in right.

All of the prior efforts on NIR vein-detection described in section 3.1.2, as well as four out of five of the NIR-based, commercially-available vein-finders described in Chapter 5, utilized a reflectance-based configuration. However, there are two considerable problems with the reflectance-based configuration:

1. In the reflectance-based configuration, any anatomical structure that casts a shadow onto the hand (including hair, skin/fat creases, and the valleys in between metacarpal bones) is falsely-detected as a vein. In Chapter 5, we show examples of false detection using two commercially-available vein-finders.
2. The robot’s mechanism could move in between the light source and the hand, thereby potentially blocking the light from detecting veins or creating shadows that are falsely-detected as veins.

3.1.3.3 How Light Scattering in Transillumination Hides Arteries

As discussed in section 3.1.3.1, one of the benefits of selecting the wavelength of 760nm, as opposed to more common NIR-LED wavelengths like 808nm or 850nm, is maximizing the contrast between veins and arteries. We wanted to ensure that our images showed only veins so that the system did not accidentally select an artery for cannulation. However, when we tested 808nm and 850nm LEDs, we obtained the same vein patterns as we did at 760nm without seeing even traces of arteries.

We determined that the amount of scattering light undergoes as it passes through the hand hides deeper structures, like arteries, from being seen. If we imagine a theoretical, collimated light source (all parallel rays) beaming into the hand without scattering, then the section of light underneath the artery would be absorbed and the light to its sides transmitted such that the artery would appear as a dark shadow on a light background. In actuality, the light that is transmitted beyond the sides of the artery eventually scatters/bounces in front of the artery, illuminating what would otherwise have been a shadow without scattering. Although this same behavior exists for imaging both veins and arteries, it is only witnessed for arteries because they are appreciably deeper to the surface of the hand such that there is sufficient tissue distance after the artery in which the light can bounce into its shadow. For veins, there is not enough distance between the vein and the surface of the hand for the light to bounce in front of the vein.

To test our theory, we used a permanent marker to draw NIR-blocking black dots on the palm of the hand and then transilluminated from first the palm and then the dorsum. As shown in Figure 3.3, the dots/veins could only be seen when they were on the surface of the hand closest to the camera (superficial), with the dots/veins that were closer to the light-source (deeper) hidden. In the left image, the black dots in the middle




Figure 3.4: NIR Light Shows Venous Blood But Not Vein Walls. A: 2D-ultrasound images of the author’s hand veins taken with a Sonosite Turbo Ultrasound with L25 linear probe. B-C: Veins disappearing in NIR image (from AccuVein AV400) when blood is squeezed out, meaning that NIR light shows only the venous blood and not the vein walls. Images used with permission of AccuVein Inc.

of the palm can be seen (circled in red), but they do not appear in the right image. In the right image, the veins in the middle of the hand (highlighted in a red box) can be seen, but they do not appear in the left image.

The practical implication of this phenomenon is that we do not have to take wavelength into account when attempting to hide the appearance of arteries because the hand’s scattering will accomplish this task. As such, we can pick whatever NIR wavelength will best show the veins within the practical design limitations of our system. Although 760nm is theoretically the best wavelength for providing high vein contrast, other frequencies, such as 808nm or 850nm, could be used if necessary for practical reasons (such as LED availability and performance characteristics) without showing arteries.

3.1.3.4 NIR Light Shows Only The Vein’s Interior

A basic question that, until now, has not been addressed is this: What portion of a vein does NIR imaging show? In all other prior work related to NIR vein-imaging, the term “vein” has been used without specifying if that term refers simply to the venous blood inside the vein or the entire vessel, including the walls. This distinction matters considerably in determining how we target the vessel. If the walls are seen in the NIR image, then we must determine how to estimate the wall thickness and subtract it from the image so that we are able to target the inner lumen only.

We posit that NIR imaging shows only the blood within the lumen of the vein and does not show the vein’s walls, based on two key pieces of evidence. Our most basic proof is that measurements of the inner diameter of the author’s hand veins taken with a 2D-ultrasound, as shown in Figure 3.4A, closely match measurements taken in the 3D point-clouds shown in Figures 3.1 and 3.14. Measuring at three different vein branches, we found an average difference between the “outer” diameter in the point-cloud and the inner diameter in the ultrasound images of $0.04 \pm 0.03\text{mm}$ (or $1.23 \pm 0.95\%$ of the diameter). Since this difference in diametral measurements is considerably smaller than the average wall thickness of $0.68 \pm 0.12\text{mm}$ that we measured with the ultrasound, we conclude that the “outer” diameter measured in the point-cloud is really the inner diameter of the vein.

Our second piece of evidence is that collapsed/blood-free veins are not shown in NIR images. Recall that veins often contain one-way valves that prevent blood from flowing away from the heart. If a vein is occluded upstream of a valve (further from the heart) with one finger while another finger swipes along the vein in the downstream direction past a valve (in the direction of blood flow), the blood will be squeezed past the valve and unable to flow backwards to refill the collapsed section of vein in between the occluding-finger and the valve. Even after the swiping action is completed, this section of vein will remain collapsed until the occluding finger is removed. This scenario is shown in Figure 3.4 using a commercially-available, reflectance-based NIR vein-viewing device, the AccuVein AV400. Figure 3.4A shows a large vein before swiping that disappears in Figure 3.4B after the swiping action has occurred. If NIR imaging showed the vein wall in addition to the venous blood, then the vein would not disappear when collapsed/drained of blood.

3.2 Computer Vision Hardware

As shown in Figure 3.1A, our custom imaging hardware consists of two subassemblies: an NIR stereo camera and NIR LED Grid that transilluminates the hand. The following sections provide details of each.

3.2.1 NIR Stereo Camera

3.2.1.1 Commercially-Available 3D Depth Cameras and Why We Built Our Own

Our three main requirements for a 3D depth camera were that it 1. be compatible with NIR imaging, 2. have sufficient depth-resolution, and 3. provide a sufficient frame-rate to be practical. We searched for a commercially-available stereo or non-stereo 3D depth camera that met our design needs but were unsuccessful. Of the non-stereo systems we surveyed (including the Microsoft Kinect, Velodyne LIDAR, and IPA-based Dynast camera), we found that they all used some form of structured infrared light that was projected onto the scene to generate the 3D structure that was then painted with the visible image. Such projected infrared light would have interfered with our NIR imaging of the veins unless the manufacturer gave us deep, firmware-level access to the projector's controls so that we were able to selectively disable it.

Of the stereo camera options (the best being from Minoru, Videre, and Focus Robotics), we found that:

1. Most stereo cameras are meant to work in visible light, but some can be adapted for NIR imaging by removing the IR-cut/blocking filters. While the Videre and Focus Robotics cameras had removable lenses that allowed for the IR-cut filters to be removed, the Minoru camera did not. However, even if the IR-cut filters can be removed, if the imaging-chip's NIR-spectrum-sensitivity is low (which is often the case for color cameras), the resulting images may be noisy and low quality.
2. No stereo camera had a sufficiently small baseline (distance between the two, individual cameras) to provide the required depth-resolution, with the smallest being 60mm-baseline cameras from Minoru, Videre, and Focus Robotics. The specifics of this requirement will be discussed further below.

Having been unable to find a commercially-available system that met our needs, we decided to design and build our own NIR stereo camera, as shown in Figures 3.1C-D.




Figure 3.5: Setting Stereo Depth-Resolution Based on Hand Vein Inner Diameter.

A: The IV needle is centered in the vein's lumen. "R" represents the \pm stereo depth-resolution that we must achieve to penetrate the lumen fully without intersecting the wall. B: Penetrating at the edge of the lumen with the catheter entirely inside. C: Penetrating at the edge of the lumen with the needle entirely inside and expecting the catheter/vein collision to resolve itself due to wall elasticity.

3.2.1.2 Required Depth-Resolution

Although the camera's planar (XY) resolution is inherently higher than its stereo depth-resolution, we set the depth-resolution as the limiting factor of precision in all three dimensions (XYZ) to be conservative and safely cover all possible camera-hand configurations wherein the depth/Z-axis might be tilted such that it was no longer normal to the hand's surface.

3.2.1.2.1 Hand Vein Inner Diameter

To set the depth-resolution requirement for our camera, we needed to know the size of the veins we were targeting. Fortunately, multiple clinical studies have measured hand vein diameter and are in rough agreement in their results. It should be noted that in the following references, hand vein diameter refers to the inner diameter only and is often referred to by the abbreviation "HVD". To remind us that we are discussing the inner diameter only, we will henceforth adopt the abbreviation "HVID" for Hand Vein Inner Diameter.

Before we discuss the specific clinical HVID measurements and the corresponding depth-resolutions they would require, we must consider the maximum alignment error between the vein and IV needle that would permit successful cannulation. The 20-gauge Jelco IV needle that we used both for the clinical experiment in Chapter 2, as well as for HaemoBot, has a needle outer diameter of 0.7mm and a catheter outer diameter of 1.15mm. The most conservative view is that the needle must puncture close enough to the vein's center that the catheter would slide against the edge of the lumen during insertion, as shown in Figure 3.5B. Therefore, the required stereo depth-resolution "R" is half of the gap between the catheter and the wall of the lumen, as shown in Figure 3.5A, or $R = (HVID - 1.15mm)/2$. However, given the elasticity of the vein wall, it is possible to imagine that penetrating the vein further from the center such that only the needle is fully within the lumen, with part of the catheter's tip intersecting the vein wall, would still result in successful cannulation, as shown in Figure 3.5C. This rests on the supposition that the needle would make a well-defined hole that the catheter would follow, pushing the wall out of the way in the process. In this case, the less-conservative, required depth-resolution would be $R = (HVID - 0.7mm)/2$. Although we believe that this scenario would result in successful cannulation, we cannot accept it as fact without clinical proof.

In the following list, henceforth referred to as the "HVID-clinical-measurements-list", we state the measured HVID mean for each study along with the corresponding depth-resolutions it would require:

1. In [44], HVID was found to have a mean value of $2.12 \pm 0.67\text{mm}$ ($N = 41$).
(Conservative $R = 0.485\text{mm}$, less-conservative $R = 0.710\text{mm}$)
2. In [2], HVID was found to have a mean value of $1.79 \pm 0.36\text{mm}$ ($N = 25$).
(Conservative $R = 0.320\text{mm}$, less-conservative $R = 0.545\text{mm}$)
3. In [67], the authors dissected 90 cadaver specimens, finding a
 - 3.1. mean HVID of $0.9 \pm 0.3\text{mm}$ in the most superficial hand veins
(Conservative $R = -0.125\text{mm}$, less-conservative $R = 0.100\text{mm}$)
 - 3.2. mean HVID of $1.4 \pm 0.5\text{mm}$ in the deeper veins.
(Conservative $R = 0.125\text{mm}$, less-conservative $R = 0.350\text{mm}$)
 - 3.3. More specifically, they found:
 - i. a perforating branch in the first metacarpal space (present in 100% of patients) with a mean HVID of $1.8 \pm 0.3\text{mm}$.
(Conservative $R = 0.325\text{mm}$, less-conservative $R = 0.550\text{mm}$)
 - ii. a specific set of three dorsal metacarpal veins (present in 93.3% of patients) with a mean HVID of $1.6 \pm 0.2\text{mm}$.
(Conservative $R = 0.225\text{mm}$, less-conservative $R = 0.450\text{mm}$)
4. For reference, we measured the author's HVID with a 2D-ultrasound, finding a mean of $2.8 \pm 0.6\text{mm}$.
(Conservative $R = 0.825\text{mm}$, less-conservative $R = 1.050\text{mm}$)

The above measurements were made under normal conditions and specific tourniquet tightness. In practice, HVID is dependent upon many variables. Certain medications [2] [28] and hypertension [28] can shrink the veins from their normal, expected size (vasoconstriction). Certain medications [34] and local warming [44] can dilate the veins (vasodilation). The tourniquet tightness (or congesting-cuff pressure if a blood-pressure cuff is used as a tourniquet) also has a large effect on vein diameter. In [44], a linear relationship between congesting-cuff pressure and intravenous pressure was established (with a slope of approximately $1\text{mmHg}/1\text{mmHg}$), and a non-linear relationship between congesting-cuff pressure and HVID was established, with an increase in cuff-pressure from 45mmHg to 60mmHg leading to a 23-33% increase in HVID. For pressures greater than 40mmHg , the relationship was approximated as linear with a slope of 0.1mm diametral increase per 5mmHg cuff-pressure increase. That is, tying the tourniquet tighter increases the vein's internal pressure, which makes the veins larger.

For our purposes, we assumed that the pharmacological and temperature effects would be less common than the tourniquet-pressure effects, so we concentrated solely on the latter in our HVID estimates. In Chapter 2, we measured internal venous blood pressure to be $51.380 \pm 13.712\text{mmHg}$ ($N = 47$), tourniqueted, and $25.580 \pm 3.483\text{mmHg}$ ($N = 49$) without a tourniquet. Given the $0.1\text{mm}/5\text{mmHg}$ slope established in [44], we would expect that the diameters reported in studies 1 [44] and 2 [2] of the "HVID-clinical-measurements-list", wherein they used a 45mmHg congesting-cuff pressure, would have been larger, on average, by $0.128 \pm 0.274\text{mm}$ under our testing conditions. This relationship between cuff pressure and HVID is compelling motivation for robotic control of the congesting-cuff to ensure that we encounter veins in the size range for

which we designed our stereo camera's depth-resolution.

3.2.1.2.2 Calculating the Depth-Resolution

To meet the clinical resolution requirements, we designed our stereo camera to have a depth-resolution of 0.184mm, calculated using two standard equations [10]: 1. depth-resolution $R = delZ = (Z^2)*delD/(b*f)$ and 2. quantized disparity (in mm) $delD = N * W/S$, with the parameters

1. Depth between camera and hand (normal distance) $Z = 667\text{mm}$. This distance was set by the ergonomics and space constraints of HaemoBot that required the stereo camera to be placed in a certain position and orientation with respect to the hand.
2. Lens focal length $f = 12.926\text{mm}$. For the required camera-hand configuration, we found empirically that lenses with a focal length of approximately 12mm were optimal for filling the camera's Field-of-View with the hand without clipping too much of the hand or leaving too much of the image unused.
3. Camera binning (degree of combining adjacent pixels) $N = 1$. This was an internal camera parameter, and setting it to 1 was the natural decision for maximizing depth-resolution.
4. Stereo-matching pixel interpolation factor $S = 16$. This was hard-coded in the stereo-matching package, part of the Robot Operating System (ROS), that we used.
5. Pixel width $W = 3.45\mu\text{m}$. This parameter is the physical width of the monocular camera's pixels, and we aimed to minimize it by selecting a high-resolution camera.
6. Camera baseline $b = 40.372\text{mm}$ (40mm nominally). This was the parameter over which we had the most independent control for setting the depth-resolution. We calculated 40mm as the approximate baseline that we needed to achieved our desired depth-resolution and then selected suitable monocular cameras that could achieve this baseline.

Recalling the average HVID values from the "HVID-clinical-measurements-list", our camera's depth-resolution of 0.184mm was able to meet the conservative requirements for all but items 3.1 and 3.2, although it met the less-conservative requirement for item 3.2. However, since our camera met the conservative requirements for other veins in the same study (items 3.3.i and 3.3.ii) that were universally or near-universally present in patients, we state that our camera met the depth-resolution requirements set by all three studies, as well as our self-testing.

3.2.1.3 Monocular Camera

We selected a pair of Prosilica GC2450C monocular cameras for their fairly unique combination of favorable optical and electrical properties (described below) with the ability to be spaced with a baseline of only 40mm. Although meant for color imaging, their easily-removable IR-cut filters and good NIR-spectrum-sensitivity allowed for high-quality NIR images. We chose color cameras instead of black-and-white cameras that would have better NIR properties so that they could see a green, dummy-needle calibration tool for vision-to-robot calibration during installation. The cameras' global shutter and ability to be triggered using an external electrical pulse allowed for synchronized imaging of the left and right images with minimal motion blur. The ability to set the gain and exposure manually, unlike most web-cameras or other inexpensive cameras, allowed us to tune the system for optimal vein/non-vein contrast. The ability to disable autofocus

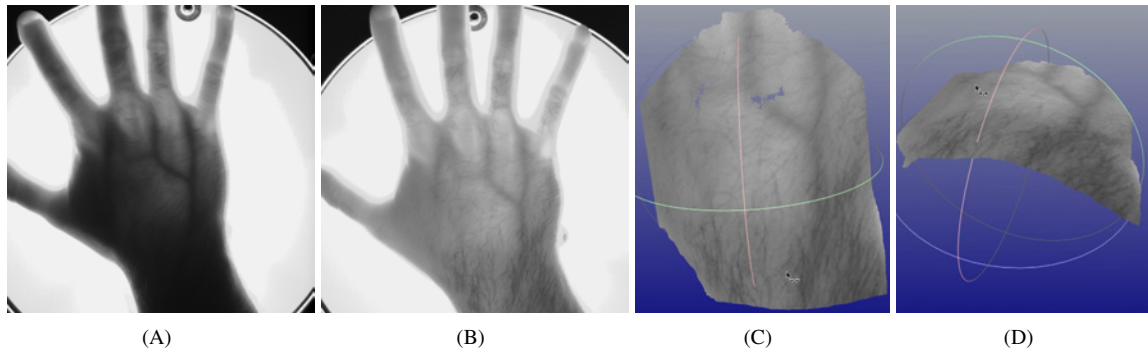


Figure 3.6: The Effects of Overhead, Fluorescent Room Lights on NIR Image/Point-Cloud Quality. Hand transilluminated by a single LED (no LED Grid). A: Imaging in a dark room. B: Imaging in a room with overhead, fluorescent lighting. C-D: 3D NIR point-clouds of veins with overhead, fluorescent lighting. Vein/non-vein contrast was still high.

was critical for obtaining a stereo calibration that persisted throughout the procedure and after electrical shutdown of the cameras. We used a high-resolution, 12mm (nominal) focal-length lens (Kowa PN LM12JC) that zoomed-in on the hand to fill most, but not all, of the image. With this lens, our depth of field was approximately 30mm, which allowed for the hand to remain in focus at any position on the imaging plate. We further set a digital region-of-interest (ROI) in the cameras' firmware that cropped the image to include just the hand and allowed for a higher frame-rate than with a full-frame image. To reduce the effects of ambient room lighting on our system, we included a single, large longpass filter (Zykkor IR720, 77mm, 720nm longpass) that spanned both cameras simultaneously, ensuring nearly-identical filtering in both cameras (as opposed to trying to match two, individual filters, one on each camera).

3.2.1.4 Texture Projection

Stereo vision requires visual texture on the objects being imaged, and it is sometimes necessary to project additional texture onto objects that do not have enough themselves to be seen. For this reason, we included two laser projectors on the stereo camera: one that projected red light and one that projected green light. This texture projection facilitated the imaging of hands, as shown in Figure 3.7, as well as nearby robot hardware like the imaging plate, without actually painting texture onto the hands or hardware.

The projectors consisted of inexpensive, novelty laser-pointers with dot-texture-projecting heads that could spin. Spinning the coarse dot-projection gave us similar coverage to a denser, stationary dot-projection and worked well as long as the speed was low enough to prevent significant motion blur in the images. We controlled the laser's ON-OFF state via a relay connected to its battery compartment and motorized the projection-heads via a friction-drive, as shown in Figures 3.1C-D, that provided velocity-control ($\pm 0-255$).

The red light was sufficiently close in wavelength to our NIR light-source that it was visible in the NIR images with the longpass filter installed and could provide texture for the transilluminated hand. Figure 3.7 shows the extent to which adding projected texture was able to improve the density of the point-cloud. We found that the need for extra texture depended on the thickness of the hand and amount of hair. For thinner hands, our light-source was able to brighten the hand and show more of its natural texture such that additional




Figure 3.7: The Effects of Stereo Camera Laser Texture Projection. Added texture can improve the density of the point-cloud for certain hands. Without added texture (A) the point-cloud is much sparser than with texture (B).

texture was unnecessary. For thicker hands that could not be brightened as much, the natural texture was too dim on its own to be sufficient. Further, hairy hands typically required less additional texture.

The green light was invisible to our stereo camera with the bandpass filters installed, but we included it in case we needed to add a visible-light camera in the future that would use texture projection without interfering with our NIR stereo camera.

It should be noted that our bifurcation-finding algorithm was able to function properly in images that contained this projected texture, but the algorithm’s parameters had to be tuned much more carefully, making the algorithm less robust. Specifically, increasing the ability of the image-preprocessing to blur/remove the high-frequency pattern of the texture also removed some of the smaller, finer vein details. In the future, we plan to experiment with blinking the texture projection so that every-other frame results in a texturized, dense 3D point-cloud, and every-other frame results in a clean, texture-free image that can be used for the bifurcation-detection algorithm. This would reduce our frame-rate by 50% but allow for vein-detection that is significantly more robust.

Although we had moderate success with our custom texture projection system, we later found better success with the integral, NIR texture projector on the Microsoft Kinect, as it provided a dimmer, higher-density grid that did not require rotation to cover the entire hand. Figures 3.21C-D show examples of the 3D point-cloud generated by our camera while pointing the Kinect’s texture projector at the hand. However, fully integrating the Kinect’s texture projector with our system would have required more control over the pattern’s appearance and operation than could be obtained without dissecting and re-wiring the projector.

3.2.1.5 Stereo Camera Controller Electronics and Code

As shown in Figure 3.8A, the camera’s controller electronics consisted of:

1. a “Baby Orangutan B-328” microprocessor/motor-controller board (Pololu PN: 1220) that drove the cameras’ external trigger-line at a configurable frame-rate (FPS), relays to control power to the laser pointers, and motors that rotated the texture projectors. This microprocessor communicated with




Figure 3.8: Controller Electronics for A: Stereo Camera and B: LED Grid.

HaemoBot’s main processing computer over an asynchronous USB-to-serial stream according to a custom communication protocol that we developed.

2. an ethernet switch for connecting the monocular cameras to the processing computer.

We used built-in Prosilica camera-handlers and stereo-processing packages in the Robot Operating System (ROS) to collect images from the cameras and generate the 3D point-clouds. Only minimal changes to these packages were required, with the largest being that we increased the stereo-processor’s maximum disparity from 128 to 256. While using ROS enabled us to develop the stereo camera very quickly, its speed was insufficient for real-time insertion, providing a maximum frame-rate of 2Hz for point-clouds with a depth-resolution of 0.184mm and approximately 1.3 million points. We were able to increase the speed of the system to 5Hz by raising the stereo camera binning to 4 and were able to increase the speed to 8-10Hz by further reducing the stereo correlation window size (allowing more noise). However, for a binning value of 4, the depth-resolution was 0.736mm, which was below the resolution we required.

To maintain the required resolution and also obtain higher speed, we would need to switch the stereo-processing algorithm from running in the computer to running on a chip. Deploying a stereo-processing algorithm on a Field Programmable Gate Array (FPGA) for “On-Chip-Processing” at rates above 30Hz is common-place for auto manufacturers who use stereo vision for collision avoidance, and it is simply a matter of time and cost (not technical feasibility) to produce the FPGA. However, if we increased the frame-rate too high, we would eventually need a shorter camera exposure time, which would require us to increase the brightness of our light-source to maintain the same image brightness.

3.2.1.6 Verifying Correct 3D Geometry from Stereo Camera

Initially the topography of the point-clouds, with their many dimples and valleys, seemed unnatural and like a possible mistake, as shown in Figures 3.9A-B. To verify that our stereo camera was representing objects with the correct size and shape, we used the following three techniques:

1. We imaged a 3D-printed pyramid, shown in Figure 3.9C, that we manually texturized using a fine-tipped permanent marker. The camera produced a 3D point-cloud that was extremely accurate to the

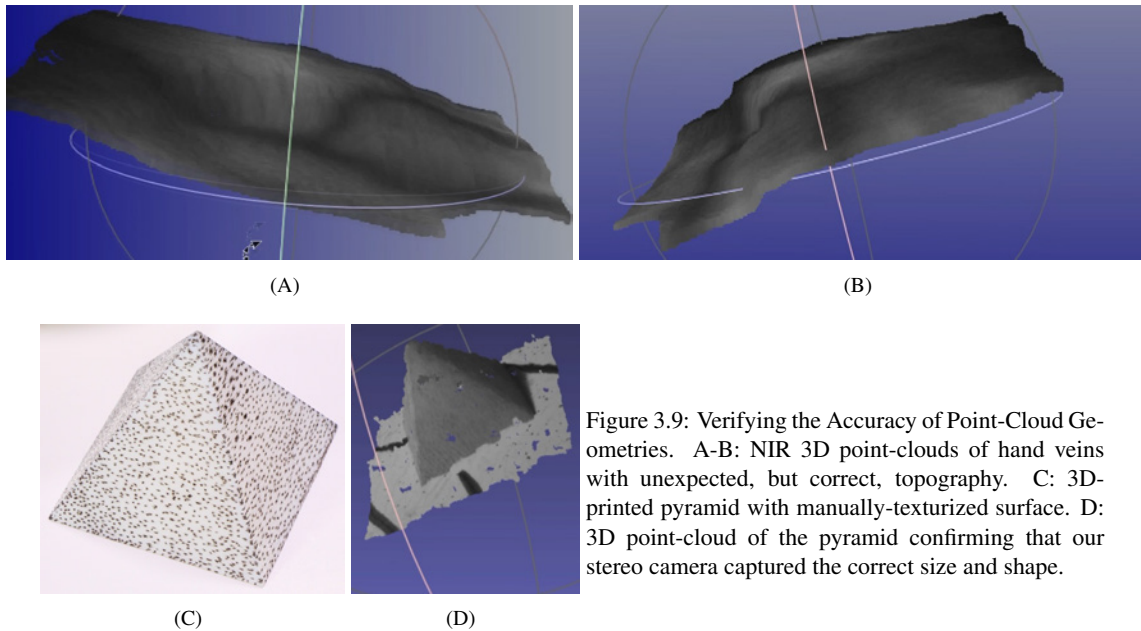


Figure 3.9: Verifying the Accuracy of Point-Cloud Geometries. A-B: NIR 3D point-clouds of hand veins with unexpected, but correct, topography. C: 3D-printed pyramid with manually-texturized surface. D: 3D point-cloud of the pyramid confirming that our stereo camera captured the correct size and shape.

pyramid's dimensions and shape, as shown in Figure 3.9D.

2. We included a 2D calibration grid on the hand and measured the distance between tick-marks in the point-cloud, finding that they matched the expected dimensions almost exactly.
3. We compared the vein diameters measured in our point-clouds to those measured using a standard, 2D-ultrasound and found that the values matched to within a few percent.

After verifying that the unexpected hand/vein topography in our point-clouds was real, we set about explaining it. The two main causes were the valleys in between the metacarpal bones and skin wrinkling, the effects of which were much more pronounced with the hand flattened against the imaging-plate than with the hand formed in a fist. The only practical problem of skin/bone valleys or creases is that they cannot be fully seen by the stereo camera, resulting sometimes in what look like tears or rips in the NIR point-cloud.

3.2.2 NIR LED Grid

The NIR LED grid, as shown in Figures 3.1E-F, served as our transillumination light source. It consisted of a 5 x 4 grid of 753.5nm LEDs (26.5mm in between rows/columns and the grid's long axis aligned with the wrist axis), each of which had 8-bit (0-255) controllable brightness. Using a grid of LEDs, as opposed to a single LED, provided several benefits, including:

1. The LED grid can transilluminate all parts of the hand with sufficient brightness regardless of the hand size or placement. Figure 3.10 shows transillumination of the hand using a single LED, with the LED moving left to right underneath the hand in each image. The tissue directly above the LED is lit brightly (with the veins visible), but the tissue far from the LED is too dim for veins to be seen, producing a narrow effective-Field-of-View. By using a grid of LEDs that is larger than even a large, adult hand, we can brighten the entire hand regardless of where it is placed so that we see the entire

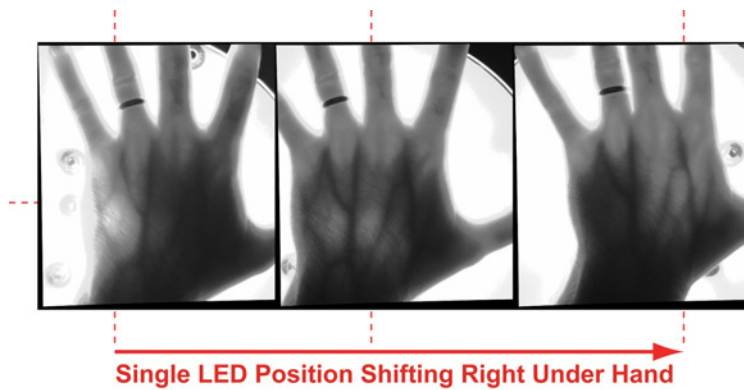


Figure 3.10: A single NIR LED is moved underneath the hand from left to right in each image, shifting the small, transilluminated region that shows veins with it and leaving areas not directly underneath the LED too dim to show veins. Showing only a small portion of the overall venous network limits insertion options, and precise placement of the patient's hand on a single LED is difficult.

venous network instead of a small subset of veins. For infant hands, we can turn off all but one LED underneath the hand without worrying about placing the hand precisely in one spot, although some easy, manual adjustment of the hand to place it roughly on top of one of the LEDs may be necessary.

2. The grid allows us to tune the amount of light to the local thickness of tissue above it to obtain an optimal image. Too much light in a thin section of tissue (e.g. near the hand edge or finger webbing) causes camera saturation/pixel-blooming, and too little light in a thick section of the hand results in an overly-dim image and unseen veins. By setting each LED's brightness according to the brightness and vein/non-vein contrast of the image portion directly above the LED, we can obtain better overall vein/non-vein contrast in the image without camera saturation.
3. The grid allows us to improve vein/non-vein contrast by illuminating the veins from all sides and orientations, in a manner similar to how a microscope's ring-light works.

3.2.2.1 LED Grid Mechanical and Optical Design

Figure 3.11A shows the components of each LED subassembly, which included:

- Aluminum Base Plate. All LED subassemblies bolted to a single, aluminum base plate that provided mechanical support, as well as passive cooling.
- LED. We used the LED Engin 740nm Metal-Core PCB (Manufacturer PN: LZ1-00R300, distributed by Mouser Electronics as PN: 897-LZ110R300). Running at its maximum continuous current, the LED's junction temperature rose such that the peak wavelength shifted to 753.5nm, which is very close to the desired 760nm. The LEDs can produce significantly higher optical power if pulsed at higher currents, but we decided to run them at their maximum continuous current to ensure safety if an electrical or software error railed the LED signal to its maximum for an extended period. The Metal-Core PCB (MCPCB) mounting removed the need for complicated soldering of the actual LED to its PCB and allowed for easy attachment to the main, aluminum base plate via two M2 screws with nylon washers and Arctic Silver 5 thermal compound between the MCPCB and the base plate. The ceramic casing of the LED served as the alignment feature for installing the LEDs in a precise grid.
- Lens. To capture and use more of the LED's light, we attached a Ledil FCN10786-LE1-D lens (distributed by Mouser as PN: 928-FA10786LE1D). Without the lens, not enough of the light transilluminated the hand for the veins to be seen clearly. This polycarbonate lens served as a collimator, focusing




Figure 3.11: LED Subassembly Details.

lens, and diffuser all in one package. Although scattering in the hand diffuses the light sufficiently for vein-imaging, we wanted an integral diffuser in the lens to protect the stereo camera from focused light that could damage the camera without the hand present. The LED/lens combination produced a maximum of 108mW optical power at the palm of the hand, with the optical power decreasing non-linearly to 5mW at 213mm (8.39 inches) from the hand-plate.

- **Lens Holders.** The lens came with an integral, ABS lens holder that is typically mounted to the MCPCB via adhesive. However, we designed an external lens holder using laser-cut delrin that used a light press-fit to mount the integral lens holder and lens on the MCPCB such that we could disassemble it easily. This external lens holder contained alignment features that allowed us to install the external lens holder, integral lens holder, and LED in a precise grid on the aluminum base plate, as well as two M1.6 tapped holes for screws that facilitated the LED wire management system.
- **LED Wire Management.** We wrapped thin, braided Dyneema around the screws of each external lens holder in multiple columns such that the Dyneema spanned across all of the LED wires, keeping them neatly in place. To route the wires through tight bends without breaking them, we used “Hi-Flex” wires (E-Z-Hook, 28 AWG, 19/40 conductor strand, Digikey PN: 461-1113-ND).
- **The LEDs generated enough heat that they required active cooling.** As such, we attached a fan (Sunon PN: MC25060V1-000U-A99, distributed by Digikey as PN: 259-1552-ND) and aluminum heat sink (Advanced Thermal Solutions Inc. PN: ATS-61270D-C1-R0, distributed by Digikey as PN: ATS1416-ND) to the bottom of the aluminum base plate beneath each LED, with Artic Silver 5 thermal compound to facilitate heat conduction. We selected this specific combination of heat sink and fan because they were the largest components that fit together in the 26.5mm-grid spacing. This cooling method allowed for 5 minutes of continuous imaging at maximum brightness for all 20 LEDs before the hand-plate became uncomfortably warm, requiring a 10 minute break to cool. In the future, we plan to replace the individual fans with thermoelectric cooling or several larger fans that could cool more effectively.

We considered using lasers instead of LEDs to obtain grid spacing less than the 26.5mm we achieved with the LEDs, but there were many practical issues with the lasers. We were unable to find a laser that was




Figure 3.12: A: With cross-polarization set to its minimum and all LEDs set to 1.6% of their maximum brightness, saturation/pixel-blooming in areas not under the hand is a problem. B: With cross-polarization set to its maximum, the majority of the stray light is removed.

affordable, approximately 100mW, TTL-dimmable, approximately-760nm, and sufficiently short. Further, lasers that came close to meeting those requirements became too hot to keep near the hand and presented a blinding hazard to both the human operator and the stereo camera.

3.2.2.2 Using Cross Polarization to Reduce Stray Light.

Placing the LEDs underneath a plastic hand-plate (with a gap of 4mm) allowed some light to reflect off of the hand-plate, bounce internally within the LED grid, and escape to the side of the hand (what we refer to as “stray light”), even for LEDs directly underneath the hand. In Figure 3.3, we show transillumination by a single LED underneath a flat hand-plate, with the entire background saturated white. This problem was exacerbated in the LED grid, with the stray light from multiple LEDs saturating the stereo camera.

We used cross-polarization to solve this problem, placing linear polarizers on both the LED grid and the stereo camera with a relative angle of 90 degrees between their polarizing axes. Because the hand scatters light significantly, light that enters the hand polarized exits the hand unpolarized. Therefore, the useful light that transilluminates the hand and provides vein data was depolarized and able to pass through the stereo camera’s polarizing filter with significantly less attenuation than the polarized stray light that did not enter the hand. Figure 3.12 shows the reduction of stray light due to cross-polarization.

For the LED grid, we adhered a 0.2mm-thick, adhesive-backed linear polarizing film (AFlash Photonics PN: PFSC-AD-SM) to the top of the hand-plate such that the polarizer was the last medium through which the light traveled before entering the hand. This film is marketed as “high-contrast”, with a maximum transmission rate of 43.4% (for light aligned with the polarizer) and a crossed transmission rate of 0.0078% (with the stereo camera/LED grid polarizer axes rotated 90 degrees relative to each other for maximum cross polarization). We placed a single, 77mm linear polarizer (Edmund Optics PN: 64901) in front of both monocular cameras to ensure an identical amount of cross-polarization in each of the cameras. Had we used two smaller, separate filters, it would have been extremely difficult to rotate both of the filters to the same relative angle with respect to the LED grid’s polarizing filter, resulting in different amounts of cross-polarization that would have made stereo-matching and point-cloud generation difficult.

The main detraction of cross-polarization is that it reduced the brightness of the hand, or the absolute magnitude of our useful signal, by approximately half. However, it reduced the magnitude of the stray light, or noise, much further, so that we experienced a drastic increase in signal-to-noise ratio. Ultimately, we were

still able to achieve the necessary brightness by running the LEDs closer to their maximum continuous limit than required without the polarizers.

We discovered that the linear polarizing filter must be the last medium through which the light travels before it enters the hand to prevent premature depolarization. That is, the filter must be adhered on top of the plastic hand-plate instead of beneath it, or else its polarizing effect will be canceled by the hand-plate. Although we do not fully understand this mechanism, it appears that the internal stresses and uneven thickness of the plastic caused this depolarization.

3.2.2.3 LED Grid Electrical Design

As shown in Figure 3.8B, the LED grid controller electronics consisted of a custom PCB which housed ten “Baby Orangutan B-328” microprocessor/motor-controller boards (Pololu PN: 1220), each of which drove two LEDs via its PWM motor-control circuitry. To ensure that we could not drive the LEDs above their rated maximum continuous current accidentally, we added a fixed, high-wattage current-limiting-resistor in line with each LED such that a PWM duty cycle of 100% corresponded to 100% of its rated current. We tested the commanded PWM duty cycle versus optical power of the LED/lens subassembly and found the relationship to be extremely linear. To keep the circuitry properly cooled, we placed two powerful fans to blow directly on all components. The microprocessors communicated with HaemoBot’s main processing computer over an asynchronous USB-to-serial stream according to a custom communication protocol that we developed. A “dead-man” timer shut off the LEDs if there was a gap in the received serial brightness commands of longer than 30mS.

A particular design challenge was routing the LED and fan wires from the controller electronics in the base of HaemoBot to the actual LED grid in the middle, requiring us to use four 20-conductor, “Medium Flex” ribbon cables (28-AWG, 19/40 conductor strands, Digikey PN: HF20G-100-ND). We needed a ribbon cable with extra flexibility to navigate the tortuous path through the robot, as well as the many sharp, 90-degree-folds we included inside of the LED grid to route portions of the ribbon cable to rows of LEDs. At over 4 meters in length, each of the ribbon cables contributed a significant amount of resistance to the LEDs’ circuitry, and we had to account for this extra resistance when selecting the current-limiting-resistor.

3.2.2.4 LED Grid Manual Brightness-Setting GUI

We provided a Graphical User Interface (GUI) for manually setting the brightness levels on each of the LEDs during development and debugging of the vision hardware and automated brightness-setting algorithm. As shown in Figure 3.13A, the GUI allowed for the user to click on an individual LED or set of LEDs and use a slider to control the brightness. The background color, as well as text in the middle of the LED’s cell, reflected the current brightness of each LED.

3.2.2.5 LED Grid Automated Brightness-Setting Algorithm

To obtain an optimal image and enable full-autonomy in the imaging system, we developed an algorithm for setting all twenty LED brightness values automatically. The goal of the algorithm was to maximize the brightness of each LED without causing camera saturation or “pixel-blooming” in the hand or wrist portion of the image.

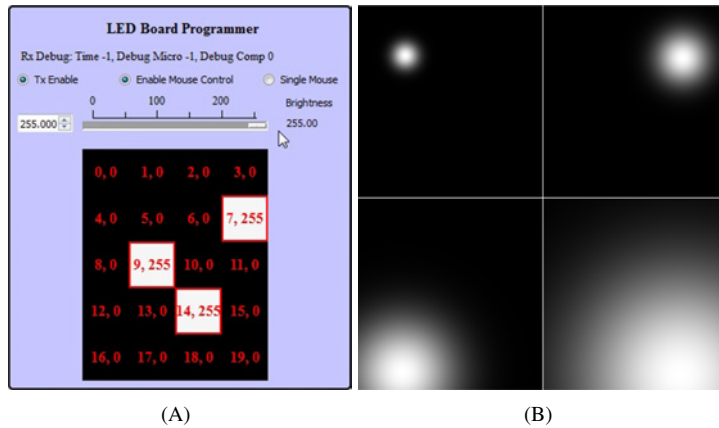


Figure 3.13: LED Grid Software Control. A: Manual Brightness-Setting GUI. B: Automatic brightness-setting, Point Spread Function (PSF) Examples. Top Left: LED 0, $\sigma = 42$ pixels, Top Right: LED 3, $\sigma = 85$ pixels, Bottom Left: LED 16, $\sigma = 170$ pixels, Bottom Right: LED 19, $\sigma = 339$ pixels.

Our algorithm was a modified Binary Search that used two additional tests before the standard search to:

1. Protect the camera from prolonged saturation/blinding by testing a very low brightness level (4%) first (instead of the standard 50% starting value) and
2. Minimize the search time based on our observation that most LEDs needed to be either fully off or fully on (0% or 100%). Trying these 0% and 100% states before performing the standard binary search improved search time.

Describing it in a pseudo-code format, the algorithm proceeded as follows:

For each of the 20 LEDs

```

{
  Set the Nth LED at a low brightness (4%).
  If it saturates the camera, part of the LED must be directly exposed to the camera.
    Permanently turn off the LED (0%). Start testing the next LED.
  Set the LED at maximum brightness (100%).
  If the camera is not saturated
    Keep this maximum brightness (100%). Start testing the next LED.
  Else
    Perform a binary search to find the maximum brightness that does not saturate the camera.
}

```

Our method for determining whether an LED had saturated the camera consisted of:

1. Thresholding the image to look only at pixels above a certain allowable brightness (near 255).
2. Weighting each offending pixel according to a Point Spread Function (PSF), with pixels closer to the LED's center receiving a higher penalty. Lacking the expensive equipment to measure our LED/lens subassembly's PSF empirically (and the manufacturer also lacking it), we approximated the PSF using a standard, 2D-Gaussian function. Figure 3.13B shows examples of the 2D-Gaussian PSF image masks for different LED positions and standard deviation (σ) values. Implementing this weighting-system required us to calibrate the position of each LED within the left and right images.

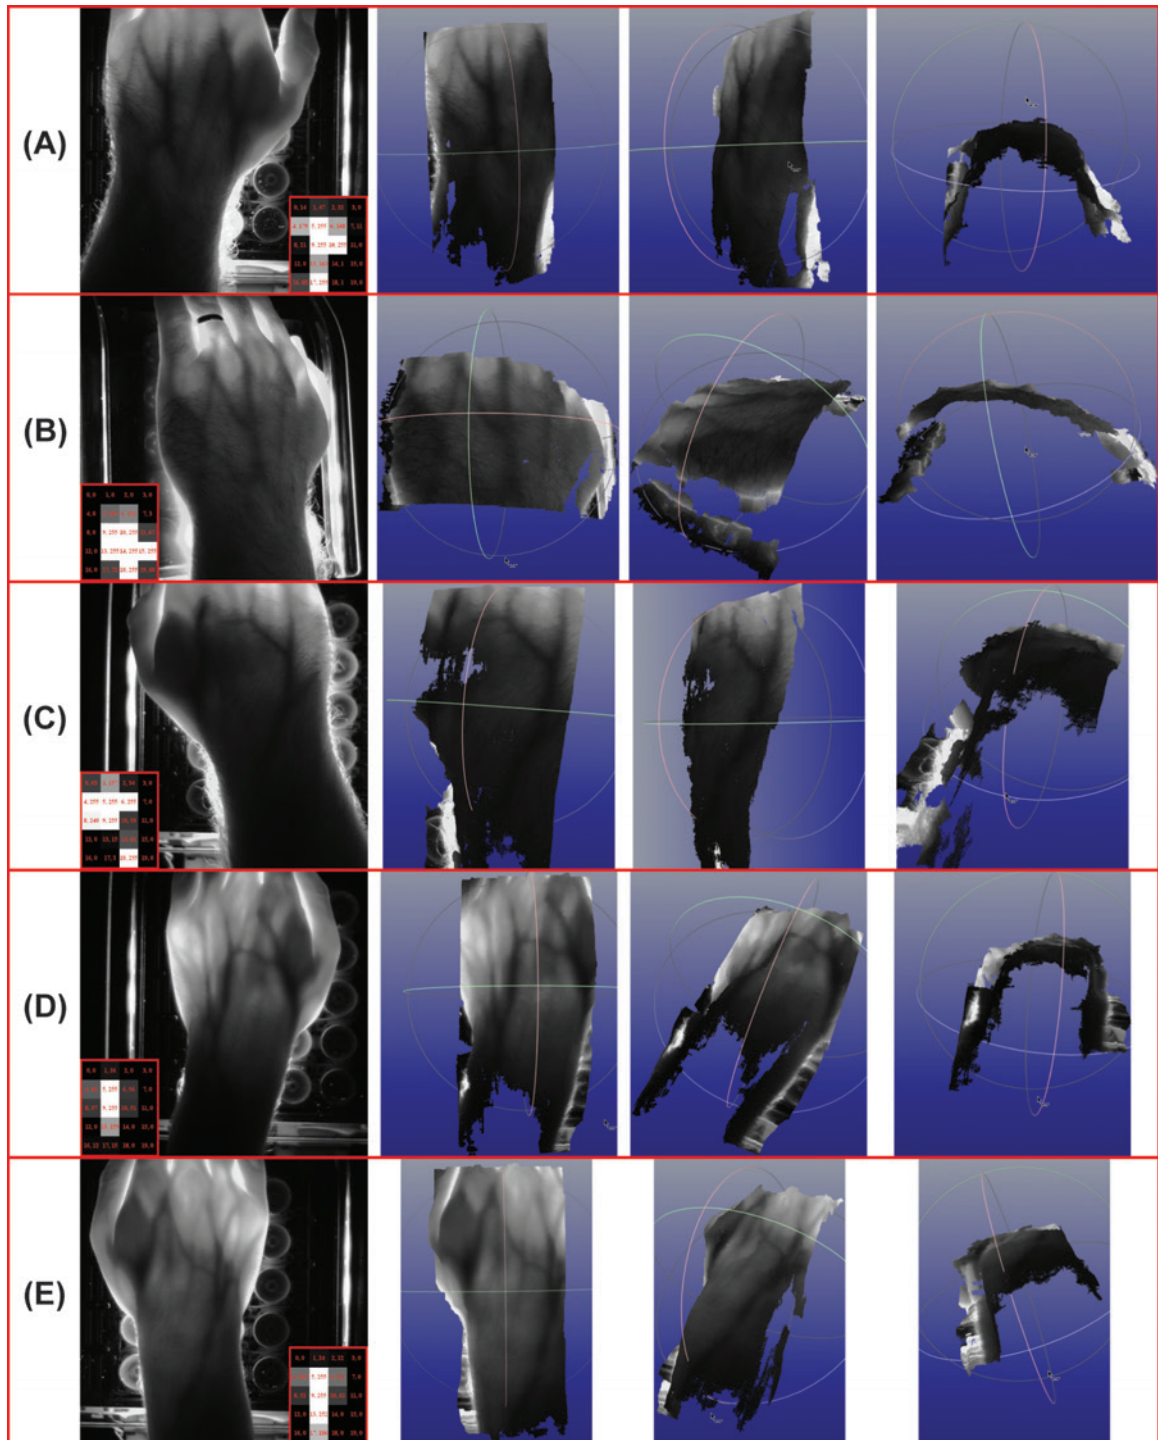


Figure 3.14: LED Grid, Example Point-Clouds Generated By The Automated Brightness-Setting Algorithm.
 A-B: Author's left hand. A: Stereo camera placed directly over the hand with its depth/Z-axis aligned with the hand normal for testing. B: Stereo camera installed on HaemoBot, above and to the side of the hand such that the camera's depth/Z-axis and hand normal were misaligned.

3. Summing all thresholded, weighted pixels and comparing the sum to a preset maximum which was determined heuristically (according to what “looked blinding” to the human programmer.)

Figure 3.14 shows the results of the Automated Brightness-Setting Algorithm for four different hands. In each subfigure, the left-most image shows one of the 2D stereo-pair images with the automatically-generated brightness levels displayed in one of the bottom corners. The three right-most images show the 3D point-clouds from various perspectives. Figures 3.14A, B, and C show the author’s large, thick hands, and Figures 3.14D-E show a volunteer’s smaller, thinner hands which resulted in brighter overall point-clouds.

Figures 3.14A and B show the author’s left hand from two different camera configurations. In Figure 3.14A, the stereo camera was placed directly above the LED grid in an off-robot testing setup, with the camera’s depth/Z axis aligned with the LED grid and hand normal. In Figure 3.14B, the entire vision system was installed on HaemoBot with the stereo camera placed above and to the side of the LED grid, such that the camera’s depth/Z axis was angled with respect to the LED grid and hand normal. We were still able to obtain good results even with the camera placed far to the side of the hand.

Our only issue with these results was the brightness of the point-clouds for thicker hands, as shown in Figures 3.14A, B, and C. We added cross-polarization, which transmits 43.4% of our light, after designing the LED grid electronics to limit the LEDs to their rated maximum continuous current, and hence brightness, instead of the considerably brighter pulsed-operation. Future work includes building a fail-safe, pulsed-operation driver for the LEDs so that we have more light for improved imaging of thick hands.

3.3 Computer Vision Software

We developed novel computer vision algorithms to 1. find a venous bifurcation and 2. calculate a 6D needle pose for insertion. Bifurcation-detection was performed on a 2D image from the stereo pair, and the results were then projected into the 3D point-cloud to calculate the 6D needle pose. Some of the following text and figures regarding the 2D bifurcation-detection were originally published in [6] (© 2010 IEEE), but we describe the algorithm more completely, along with the 6D pose calculation, in the following sections.

3.3.1 2D Bifurcation-Detection Algorithm

Given a 2D NIR image, the bifurcation-detection algorithm detects all bifurcations, calculates the planar insertion position and angle for each, and recommends the best candidate for insertion, as shown in Figures 3.19A-C. The following sections detail the novel approaches we developed to achieve this. We first detail a method for locating the wrist veins as a reliable starting point for vessel tracking. We then present a new annular tracking window that examines the macrostructure of the vein, allowing it to 1. track along a vein’s path regardless of noise, changing size, or erratic edges and 2. detect bifurcations during the tracking process. We describe a method for centering an initial bifurcation estimate on the true center of the bifurcation by using a set of concentric annular windows. We detail methods for error-checking and finding the actual position and angle of insertion in between the branches of the bifurcation. We discuss how continuous usage of the centering-algorithm allows for tracking of bifurcations as they move. Finally, we discuss the experimental results of applying our bifurcation-detection algorithm to a set of sample images.

3.3.1.1 Related Work

Bifurcation-detection has been studied widely for images of retinal arteries. In [4], the entire image is searched using the assumption that bifurcations will occur at areas of high variation in Sobel edge direction. Similarly, in [62] and [9] a coarse grid of seed-points that looks for local gray-scale minima between oppositely-signed edges is applied across the entire image. Recursive tracing proceeds from these seed-points by moving in the direction that best fits the vein model of a gray-scale minimum between antiparallel edges. However, searching the entire image can generate many false positives, especially if the image is noisy or has artifacts. In the case of our particular vein images, the variation in gray-scale intensity and contrast of the veins lead to poor edge detection such that the edges cannot be followed reliably. An alternative to searching the entire image is to trace along the arteries from a known landmark, hopefully reducing the number of false-positives. In [23], Canny edge detection is performed, and tracing occurs along the detected edges, starting from the optic disk. In [37], a steerable Gaussian filter is used to trace from the optic disk. This tracing assumes thin, well-defined vessels with consistent cross-sectional profiles. Unfortunately, the vessels in our images exhibit cross-sectional intensity profiles of varying size, noisiness, and degree of saturation such that looking at their cross-sectional profile does not give much reliable information.

3.3.1.2 Image Enhancement

The veins are often difficult to identify in the raw images due to low contrast and fading of some vein sections, as seen in Figure 3.15A. For this reason, we enhance the images to increase the vein/non-vein contrast and normalize the appearance of the veins so that all sections appear similar. We use Laplace of Gaussian (LoG) filtering ($\sigma = 8$, window width of 5σ , determined empirically) to accentuate the vein macrostructure over the background and histogram equalization to improve the overall contrast in the image. Figure 3.15A shows the raw image on the left and the enhanced, or “preprocessed”, version of the raw image on the right.

3.3.1.3 Identifying Wrist Veins

Our bifurcation finder operates by finding the most prominent veins at the base of the wrist and tracking along those veins with a window that looks for bifurcations. The wrist veins are composed of the basilic, cephalic, and their branches, such as the accessory cephalic vein. The desired bifurcations occur in the dorsal metacarpal veins in the top of the hand. Whereas the venous network in the hand varies greatly between individuals, the location of wrist veins is fairly consistent, providing an ideal starting location for tracking. As in [23] and [37], we track along the veins from a known location instead of searching the entire image because a bifurcation that we detect while tracking along a vein is more likely to be real than a bifurcation that we detect at a random point in the image. To identify the wrist veins, we take a horizontal sample strip (height = 8 pixels, determined empirically) across the bottom of the image and examine the smoothed histogram of gray-scale intensity, as shown in Figure 3.15B. After collapsing the strip to a single pixel-width height by taking the median in the vertical direction, we apply Gaussian filtering in the horizontal direction and hysteresis-thresholding to mitigate the noise in the histogram. Since veins now appear as bright against a dark background in the preprocessed image, they are identified on the histogram as local maxima. Each identified wrist vein becomes a starting point for tracking. We use the histogram to size each tracking window

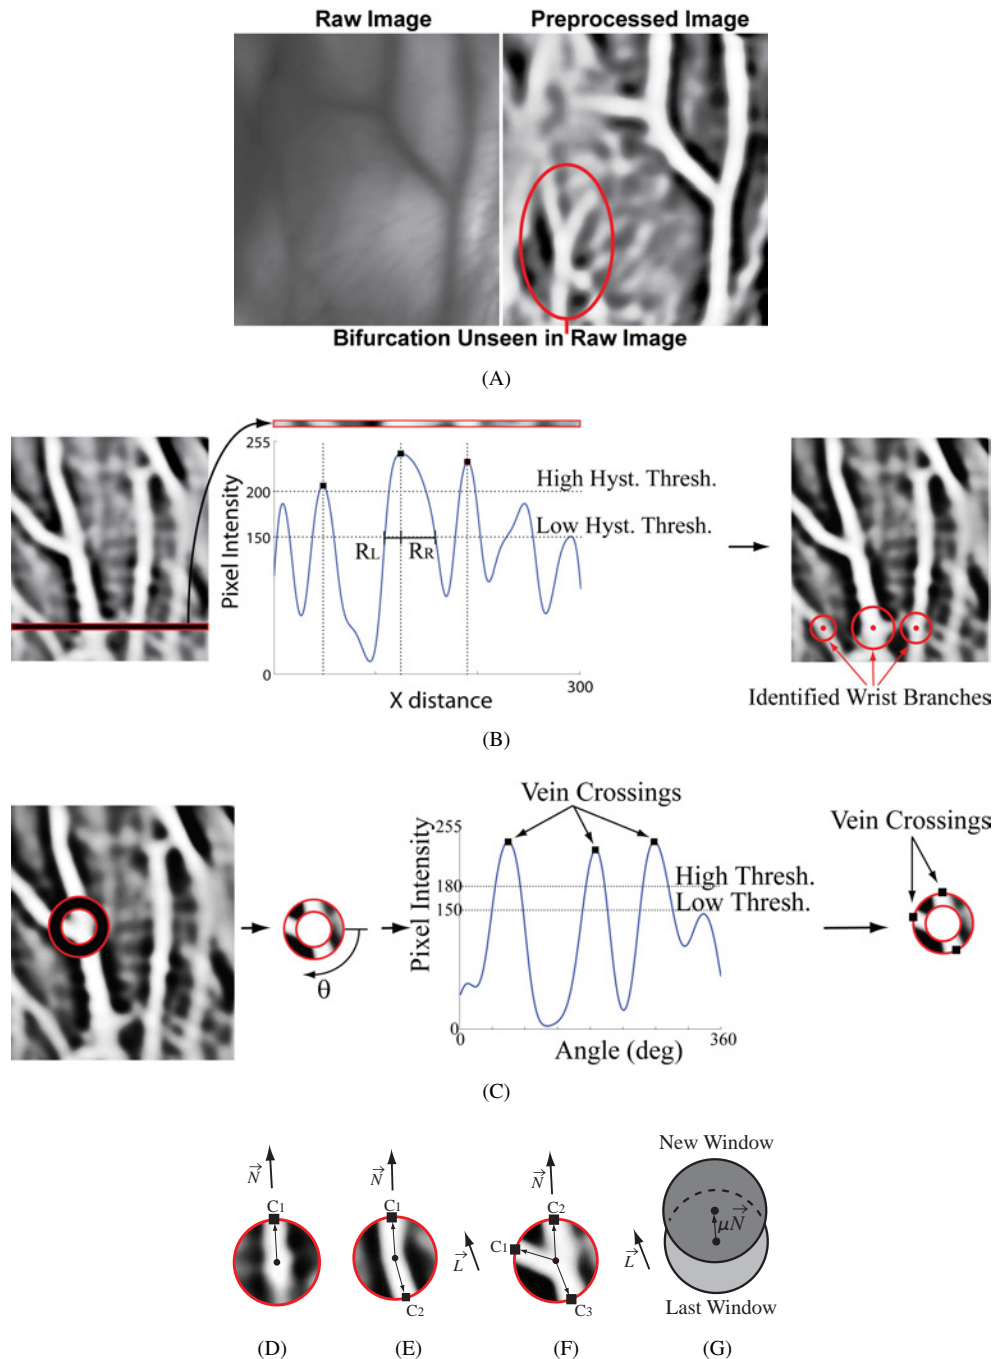


Figure 3.15: 2D Bifurcation-Detection: Image Preprocessing, Wrist Vein Detection, and Annular Bifurcation-Detection and Tracking Window. A: Raw image (left) and preprocessed image (right). B: Identifying wrist branches with histogram at bottom of image. C: Polar histogram of annular tracking window. D: 1 vein crossing represents a faded, single vein. E: 2 vein crossings represent a solid, single vein. F: 3 vein crossings represent a bifurcation. G: Incremental movement of tracking window.

to be twice the width of the wrist vein that it will track. The tracking window is larger than the vein being tracked so that the tracking window is never entirely inside the vein.

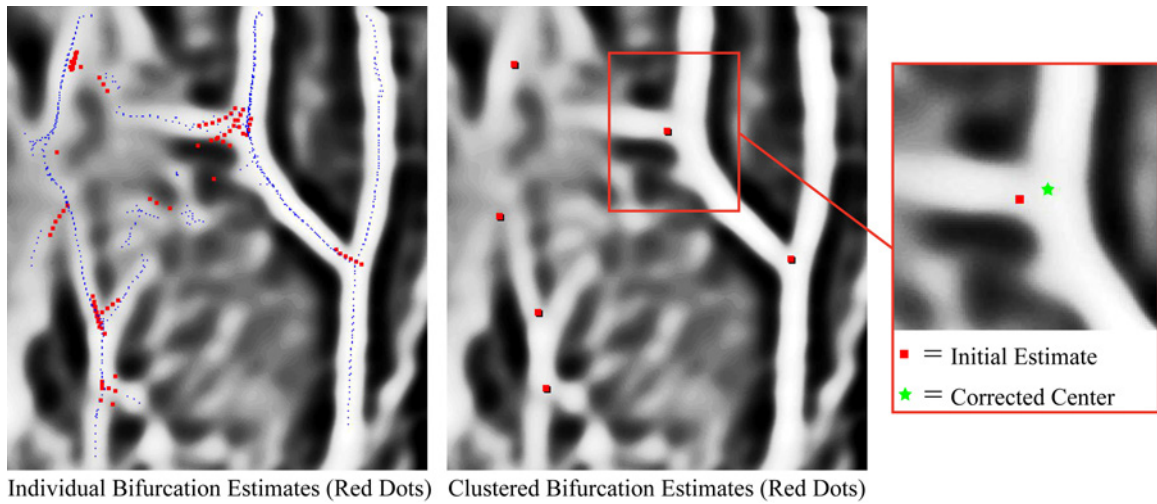
3.3.1.4 Tracking Along Veins and Detecting Bifurcations

Tracking along the wrist veins employs a steerable window that uses information about the macrostructure of the vein section contained within the window to determine how to move along the vein. The tracking window consists of an annulus (thickness = 8 pixels, determined empirically) that is centered on a section of the vein. A polar histogram of gray-scale intensity in the annular sample describes the shape of the vein section contained within the window. As before, the veins are identified in the smoothed histogram as local maxima. After collapsing the annulus to a single pixel-width thickness by taking the median in the radial direction, we apply Gaussian smoothing in the angular direction and hysteresis-thresholding to reduce noise in the histogram. Figure 3.15C shows the tracking window in detail.

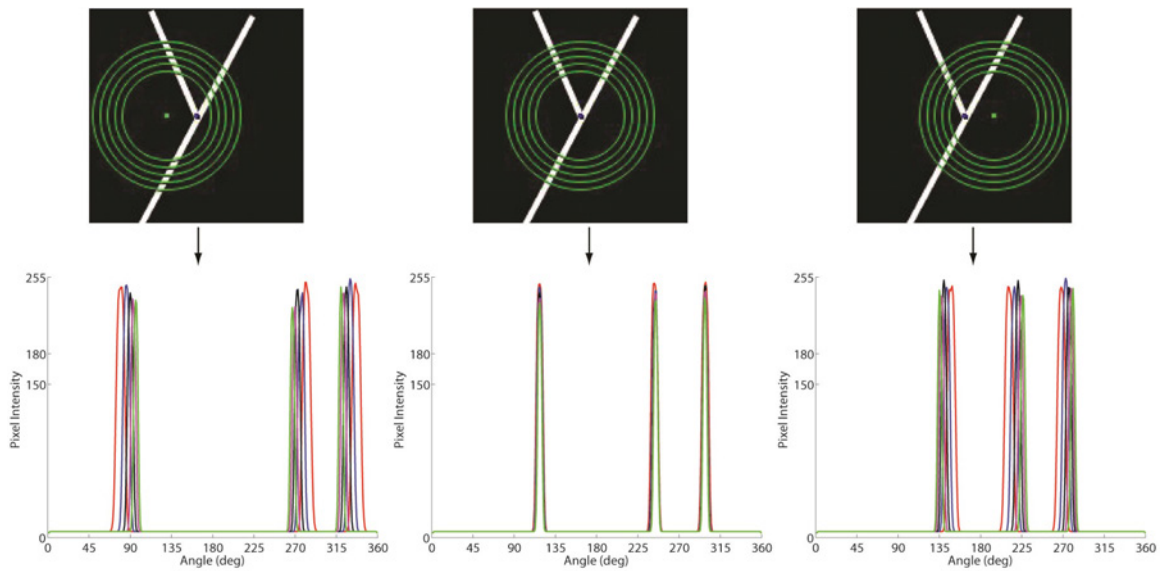
Figures 3.15D- 3.15G show how the number of vein crossings detected in the tracking window conveys information about the section of vein within the window and how the tracking window should move along the vein. We assume that one vein crossing (C_1) indicates that the tracking window is centered on a faded, single section of vein, two crossings (C_1, C_2) indicate that the tracking window is centered on a solid, single section of vein, and three vein crossings (C_1, C_2, C_3) indicate that the tracking window is centered on a vein near a bifurcation. In all three cases, the tracking window steps a small amount μ (5 pixels, determined empirically) in the new direction \vec{N} that is computed as the vector between the window center and the vein crossing most closely aligned with the previous direction of movement \vec{L} . The degenerate cases are zero vein crossings, which indicates a lack of information, and more than three crossings, which indicates noise. In either degenerate case, the tracking window steps in the last known direction of movement. In the case of a detected bifurcation, the tracking window continues tracking along the branch most closely aligned with the last known direction of movement until it reaches an edge of the image or stalls, whereupon it returns to track along the alternate (third) branch of the bifurcation. Figure 3.16A shows an example of vein tracking and initial bifurcation detection. Due to the thickness of the veins, there are many points in the vicinity of a single bifurcation that show three vein crossings, or a detected bifurcation. Clustering all of the bifurcation estimates based on distance from each other reliably thins this mass of estimates to one estimate per actual bifurcation detected, as shown in the middle image of Figure 3.16A.

3.3.1.5 Centering the Bifurcation Estimate


The initial bifurcation estimate often does not coincide with the true center of the bifurcation due to the asymmetric distribution of initial guesses that were eventually clustered. Thus, we need an algorithm for finding the true center of the bifurcation given a moderately-close initial estimate. We have devised a solution to this problem by looking at the correlation between polar histograms in a set of concentric annular samples taken at different diameters, as shown in Figure 3.16B. These samples are resized versions of the tracking window described above. If we position this set of annuli at each pixel in the vicinity of a bifurcation with three branches and examine the correlation, we find that the global maximum of the correlation corresponds to the true center of the bifurcation. Since there are more than two samples, we compute the correlation as the



(A)



(B)



(C)

Figure 3.16: 2D Bifurcation-Detection: First-Pass Bifurcation Detection, Clustering, and Centering Details. A: In the left image, small red dots show un-clustered initial bifurcation estimates, and small black dots show the path of the tracking window. In the middle image, large red dots show clustered bifurcation estimates. In the right image, we show an example of the centering algorithm working on one of those clustered estimates. B: Concentric annular samples near a bifurcation and their polar histograms. C: Correlation in the set of annular samples near the bifurcation from Subfigure B.

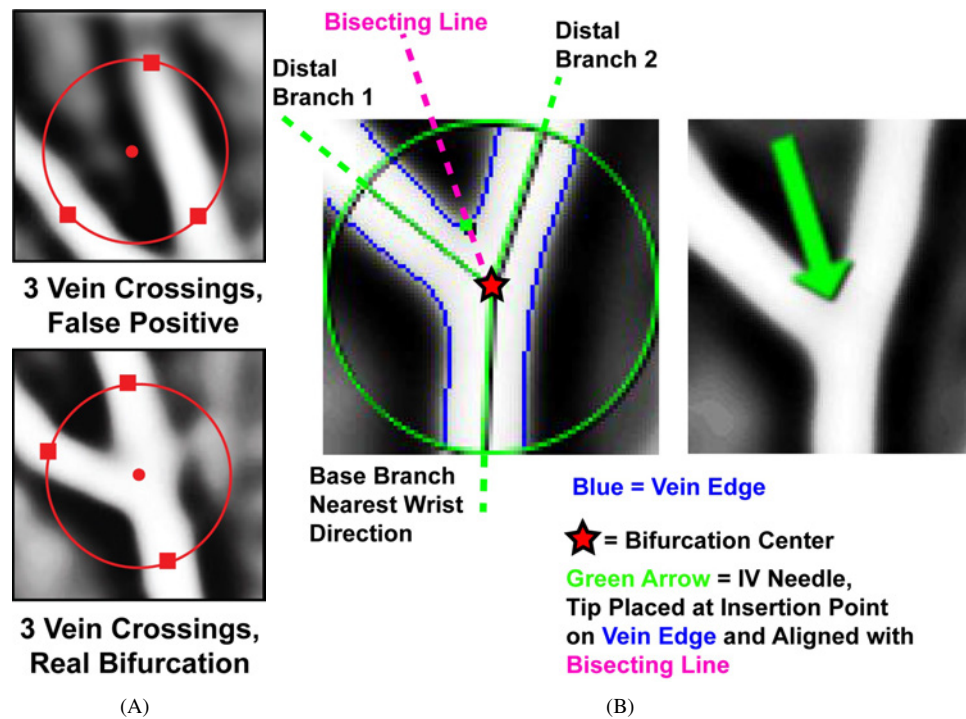


Figure 3.17: 2D Bifurcation-Detection: Error Checking and Calculating Insertion Point/Angle. A: Error-checking for false-positives. B: Calculating insertion point/angle along the edge and bisecting line of the bifurcation's distal branches.

sum of correlations between adjacent rings. Figure 3.16C shows the correlation plot for the example bifurcation in Figure 3.16B. The right-most image in Figure 3.16A shows an example of the centering algorithm working on a real bifurcation. Note that the initial estimate is appreciably off-center but that the centered estimate appears directly at the center of the bifurcation.

3.3.1.6 Error-Checking

We employ two simple methods of error-checking to detect false-positive bifurcations. In a fashion similar to the centering algorithm, we examine concentric annular samples of each bifurcation over 35 different diameters and count the number of diameters that elicit the characteristic three vein crossings of a true bifurcation. If our estimate is a true bifurcation, then it looks like a bifurcation over many scales, whereas false-positives appear like bifurcations over only a few diameters. In practice, false-positives almost uniformly return a low number of counts (below 10), whereas true bifurcations return a much higher number (typically above 25). A final, simple error-check is to ensure that each bifurcation is centered on a bright spot (vein). Figure 3.17A shows examples of both error-checks. Note that the false positive is centered on a dark spot (non-vein) and only shows the three vein crossings for the particular diameter shown, whereas the true bifurcation is centered on a bright spot and shows the three vein crossings over most diameters.




Figure 3.18: 2D Bifurcation-Detection: Tracking with the Centering Algorithm. The red circle represents the center of the bifurcation, and the purple arrow represents the calculated IV needle placement. Applying our centering algorithm continuously allows us to track a bifurcation undergoing translation, rotation, skew, dilation, or loss-of-focus.

3.3.1.7 Calculating the Insertion Point

The desired insertion point is not the center of the bifurcation but along the edge in between the two branches of the bifurcation furthest from the wrist, as shown in Figure 3.17B. To find this point, we use Canny edge detection to compute the edges of the veins in the vicinity of the bifurcation. This edge detection uses locally determined threshold values because the intensity of the veins varies somewhat over the image, precluding the usage of global threshold values. Once the local vein edges are found, we sample along the line that bisects the two distal branches of the bifurcation until we intersect with a vein edge. This intersection at the edge of the vein is the desired insertion point, and the bisecting line is the desired direction for the IV needle. It is acceptable for the bisecting line and needle not to align exactly with the branch nearest the wrist because the needle will not enter the vein far enough to hit the back wall, and the catheter is sufficiently flexible to deflect safely away from the back wall and into the lumen of the base vein branch.

3.3.1.8 Weighting Multiple Bifurcations

If multiple bifurcations are detected in the same image, we must be able to decide which is optimal for insertion. Towards this end, we examine each bifurcation and compute a weighted score based on a variety of desirable traits. These traits include the correlation found during the centering process, orientation with respect to the wrist direction, distance from the center of the hand, and average gray-scale intensity in the vicinity of the bifurcation. The correlation and intensity metrics address our confidence in the authenticity of the bifurcation. The orientation and distance metrics address the practical concerns that we wish to insert the needle in the direction of the wrist and near the center of the hand, as is standard practice among practitioners. While these metrics make sense from an engineering perspective, it will be necessary in the future to determine additional clinical criteria used by practitioners to select amongst bifurcations and implement the same weighting in our algorithm.

3.3.1.9 Tracking the Bifurcation During Hand/Vein Movement

Because our centering algorithm can find the true center of a bifurcation based on an off-center initial estimate, applying it continuously to a detected bifurcation allows us to track that bifurcation as it moves.




Figure 3.19: 2D Bifurcation-Detection: Examples of Successes and Failures. A-C: Examples of correctly identified and characterized bifurcations. The arrows show the desired position and orientation of the IV needle. D: Failure modes. Point A is a false negative, point B is a correctly identified but misplaced bifurcation, and point C is a false positive that arose because two separate veins appeared to be a bifurcation.

Such tracking is necessary in the event that the vein moves under the pressure of the needle or the patient intentionally moves his or her hand. Our centering algorithm can track a moving bifurcation in translation, rotation, skew, dilation, or loss-of-focus (the image gets blurry), some of which are shown in Figure 3.18. Future work includes increasing the speed of this tracking from a few frames-per-second to real-time.

3.3.1.10 Experimental Methods and Results

We developed our bifurcation-detection algorithm on a set of thirteen images that provided for iterative refinement. We determined algorithm parameters such as LoG filter σ , window step size μ , and annular thickness empirically in these 13 images. To test the general applicability of the algorithm, we applied it to fifty images that were unseen during development. The test population was 72% male, had an age range of 22-56 ($\mu = 31.0$, $\sigma = 9.1$), and was 72% caucasian and 28% asian. To compare the performance of our algorithm with that of human practitioners, we provided four medical doctors with the same images and asked them to identify all bifurcations, without preference to size or desirability. Our algorithm successfully detected and characterized 32.4% of the bifurcations identified by the doctors with a false positive rate of 11.6%. However, our algorithm correctly found and characterized at least one bifurcation in 82.6% of the images that contained bifurcations, which comprised 92% of all images (4 images did not contain any bifurcations). Assuming that it is clinically sufficient to find only one (instead of every) bifurcation per hand for the robot to insert an IV, our vein-finding algorithm would allow us to insert on the majority ($82.6\% * 92\% = 76\%$) of individuals. Figures 3.19A-C show examples of successfully detected and characterized bifurcations, and Figure 3.19D shows several failure cases. Although the algorithm was not tested on children, we expect future pediatric tests to be successful due to the algorithm's ability to auto-size the tracking window to differently-sized veins. It should be noted that the large width of the tracking window creates margins at the edges of the image that are unsearchable by the algorithm, whereas humans can search the entire image. In fact, many of the bifurcations that the doctors identified and our algorithm did not were near the edges of the image/hand where the algorithm could not search. These "missed" bifurcations are not as important as others because it would be inconvenient to insert an IV at the edge of the hand.

The experimental results have revealed several insights as to the limitations of our algorithm. The algorithm has particular difficulty in detecting closely-spaced bifurcations because they appear as a single, noisy bifurcation rather than distinct bifurcations. A limitation of the centering algorithm is the assumption that bifurcations consist of 3 straight branches. In some cases, there are more than 3 branches, or the branches are sufficiently curved to break the centering process. A further problem is that each tracking window maintains a constant size as it tracks along a wrist vein, even though the veins periodically change diameter. This can lead to either too much noise for an oversized tracking window or too little information for an undersized tracking window. Dynamic resizing of the tracking window is a potential solution to this problem.

Our algorithm required approximately 1 second/bifurcation for initial detection/characterization and was able to track the bifurcations at roughly 3 frames per second. Although this initial planning time is sufficiently short, future work includes increasing the speed of the tracking to real-time rates (above 30 Hz).

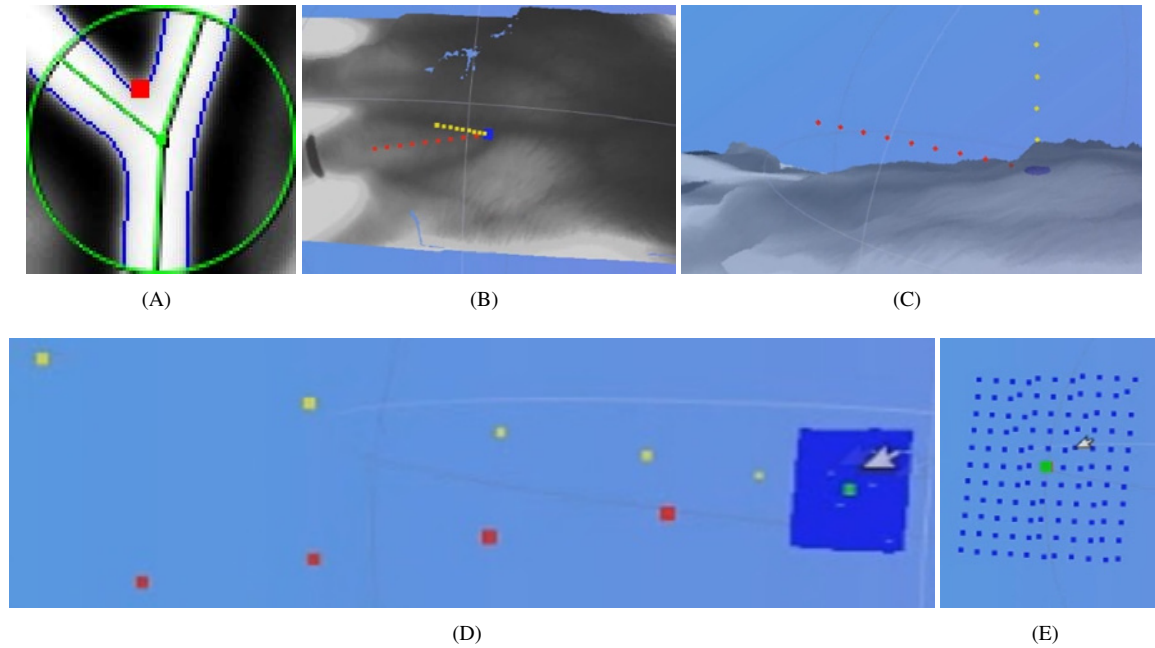


Figure 3.20: 6D Needle Pose Generation: Finding Neighbors in Point-Cloud and Fitting a Tangent Plane. A: Bifurcation detected and characterized by the 2D bifurcation-detection algorithm. The red square shows the neighborhood pixels around the insertion point that will be looked up in the 3D point-cloud and used to fit a local tangent plane. B-E: Generating the tangent plane and the 6D needle pose. Blue = insertion neighborhood points for fitting a local tangent plane via RANSAC. Green = insertion point from 2D bifurcation-detection algorithm projected into the point-cloud. Yellow = tangent plane normal vector. Red = needle insertion vector.

3.3.2 6D Needle Pose Generation Algorithm

We now use the planar insertion position and angle found in the previous section, along with the 3D point-cloud, to generate a 6D needle pose (X , Y , Z , roll, pitch, and yaw) that can be fed to the robot for automatic insertion. Figures 3.21A-C show successful examples of this process for different hands, with the red line showing the 6D needle pose, and Figure 3.21D shows a representative failure mode.

The 6D needle pose generation process is far simpler than the 2D bifurcation-detection algorithm and is composed mainly of simple vector algebra that is too basic to warrant inclusion in this text. Describing the process generally, it entails:

1. Picking neighboring pixels around the original insertion point in the 2D image and looking-up the corresponding insertion and neighborhood pixels as 3D points in the point-cloud. The red square in Figure 3.20A shows the local, 2D neighborhood around the original insertion point shown as a green dot in Figure 3.17B. In Figures 3.20B-E, the 3D neighborhood appears as blue points. We now have XYZ coordinates for the original insertion point.
2. Fitting a local tangent plane to the 3D neighboring points via RANSAC with the constraint that the plane intersects the original insertion point. In Figures 3.20B-D, the yellow line represents the plane's normal vector. The needle bevel uses the plane normal as "up", thereby setting the needle's roll angle.

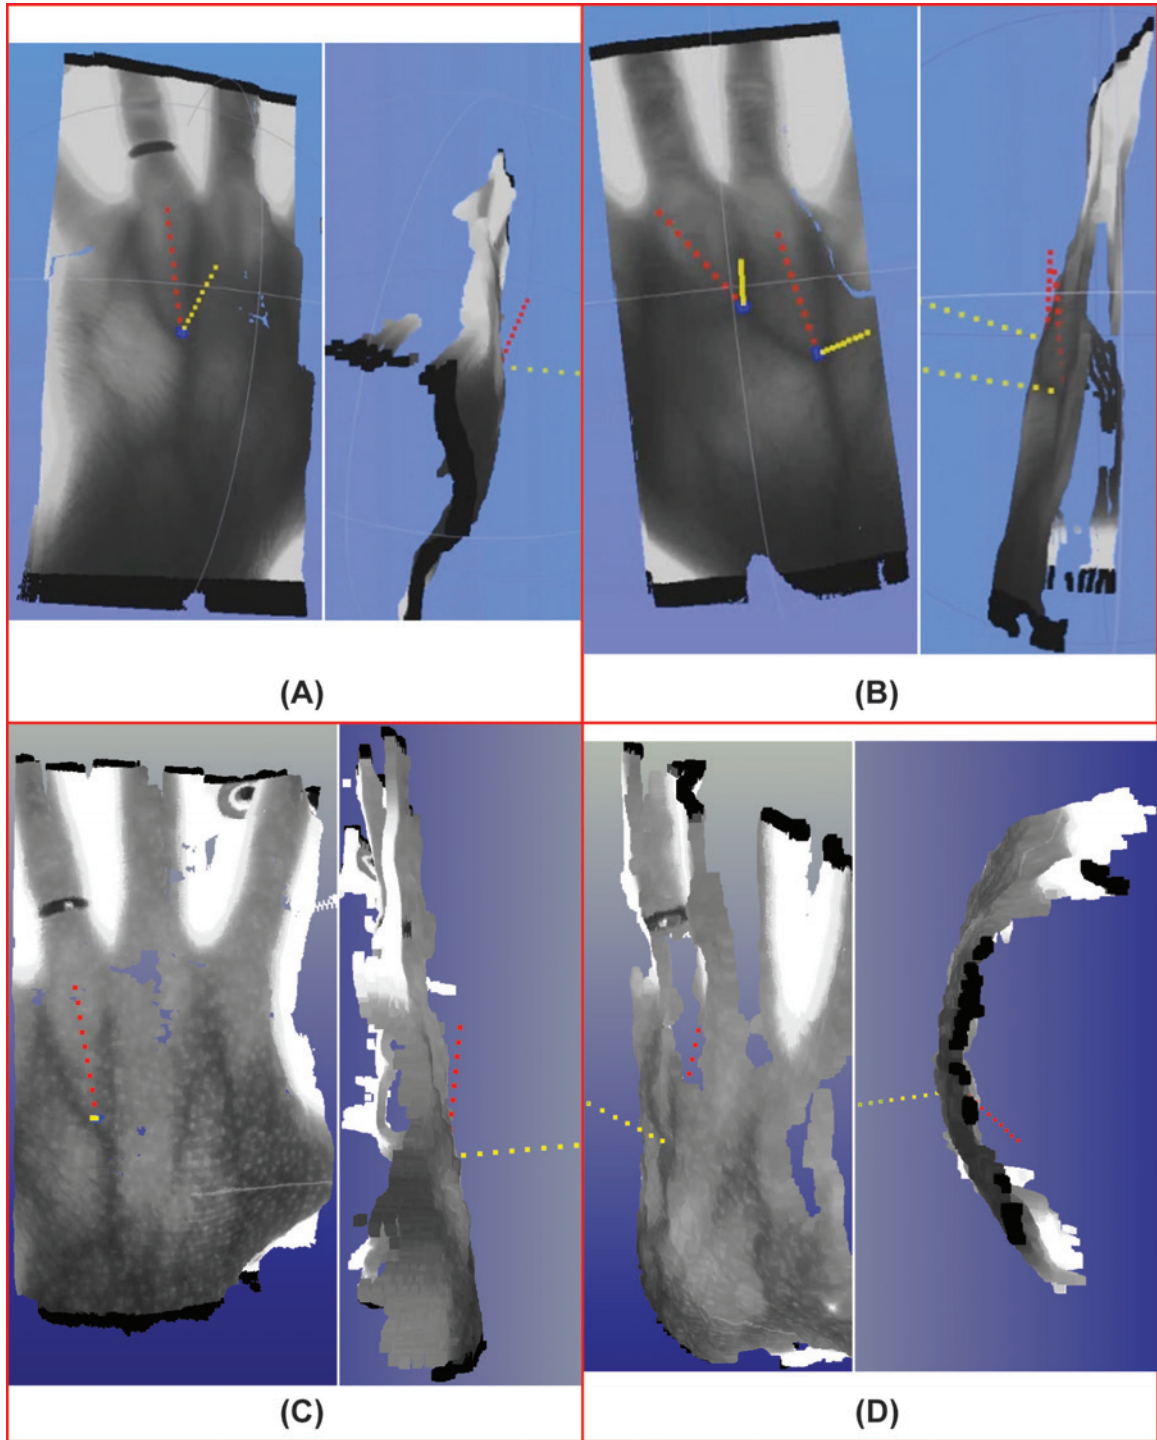


Figure 3.21: 6D Needle Pose Generation: Results. Yellow line = tangent plane normal vector, and red line = needle insertion vector. A-C: Success. D: Failure (the tangent plane was fit incorrectly such that the IV needle was placed within the hand). A-B: Without texture projection. C-D: With Microsoft Kinect texture projection.

3. Projecting the 2D insertion vector (the green arrow in Figures 3.17B and 3.19A-D) onto the tangent plane and then rotating it up off the plane/hand by the desired pitch angle (shown as 10 degrees in all of our examples). We now have the yaw and pitch angles.

There were two main issues with the pose-generation process:

1. Speed. The 6D needle pose generation required approximately 6-9 seconds/bifurcation, which is not problematic for initial planning but is prohibitively slow for real-time tracking. Future work includes modifying the algorithm's approach to tracking to increase its speed.
2. Isolated instances of incorrectly-generated tangent planes. Sometimes the size of the neighborhood around the insertion point was too small such that the RANSAC fitting of the plane was skewed by local anatomical features, such as dimples and valleys, that were not representative of the overall angle and taper of the hand. Figure 3.21D shows an example of this, wherein the plane was fit incorrectly, resulting in a needle pose that placed the entire IV needle within the hand instead of above it. Future work includes performing the tangent plane-fit over several different neighborhood scales to prevent local anatomical features from skewing the results.

3.4 Discussion and Thesis Contributions

In this chapter, our three main contributions were:

- Proving that a robot can locate a vein and compute a 6D needle pose autonomously.
- Developing novel NIR transillumination and stereo camera hardware, as shown in Figure 3.1, capable of generating uniquely accurate, high-resolution, high-contrast, and three-dimensional hand/vein point-clouds, as shown in Figure 3.14.
- Developing novel computer-vision algorithms for locating, characterizing, and tracking venous bifurcations in a 2D image, as shown in Figure 3.19, and then generating a 6D needle pose in the 3D point-cloud, as shown in Figure 3.21.

Future work includes increasing the brightness of the LED grid for imaging thicker hands, increasing the speed of the stereo camera, bifurcation-tracking, and 6D needle pose generation to real-time (greater than 30 Hz), and adding noise-reduction algorithms for smoothing the 3D point-cloud.

Chapter 4: Development of HaemoBot, a Robotic System for Teleoperated or Autonomous Insertion

4.1 Motivation and Overview

Recall that one of the primary difficulties of insertion is moving a small, 2-part needle/catheter into a vein that is not much bigger than the needle, making the margin for error small. As we observed in Chapter 2, the practitioner must sense the penetration-events and carefully coordinate the motions of needle and catheter with those events for successful insertion. Delays in sensing or reacting to the penetration-events can lead to back-wall penetration as the practitioner pushes the needle too far. The difficulty of manipulating many degrees-of-freedom can also lead to failure, particularly with respect to holding the needle still as a guidewire for catheter insertion, maintaining proper traction while inserting the catheter, overshooting when attempting to force the needle past a valve, and rolling the needle about its axis to drill past valves. In this chapter, we developed a highly sensitive and dexterous robot, HaemoBot, that was able to sense and react to the penetration-events quickly and with precise, coordinated motion from the many degrees-of-freedom. Figures 4.1- 4.2 show HaemoBot from different vantage points.

The driving question behind our development process was: Can we show a robot performing all of the mechanical steps of IV insertion autonomously using the sensory cues that we observed with human practitioners on realistic hand, arm, and vein geometries? This question generated many broad, system-level requirements, as well as specific design requirements. When compared to other researchers' efforts towards developing phlebotomy and IV robots, HaemoBot made three main, unique design contributions, including:

1. Performing an insertion with two arms,
2. Moving the catheter independently of the needle, and
3. Implementing inherent, mechanical safety.

In the following sections, we will discuss these three, main contributions in detail, as well as the design elements that supported and enabled those contributions. In the interest of brevity, we will only discuss those enabling, technical details that were critical, unique, or not obvious.

4.1.1 Unique Contribution 1: Performing an Insertion with Two Arms

We are the first researchers to develop a 2-armed robot that can perform all of the 1-handed and 2-handed motions that we observed human practitioners performing during insertions in Chapter 2, including all of the preparatory (cleaning, tourniquetting, and restraining), assistive (applying traction and clamping the vein/catheter at the end of the insertion), and actual needle-insertion steps. Further, we are the first to use a complete insertion task-list that was developed through controlled, clinical observations of real IV insertions.


We refer to HaemoBot's two arms as "NeedleBot" and "AssistoBot". NeedleBot controlled the motion of

the needle (XYZ, roll, pitch, and yaw) and catheter (prismatic insertion and 2 grab-and-release DOFs). We included full, 6-DOF motion control for the needle to target different hand and vein placements/geometries without requiring the human operator to pre-set any joint positions manually. Assistobot performed all of the tasks typically performed by the non-dominant hand, including:


1. Restraining the hand during insertion,
2. Restraining the arm and applying tourniquet pressure,
3. Cleaning the insertion site with alcohol,
4. Rubbing and tapping the veins to promote vasodilation,
5. Applying skin traction during insertion, and
6. Clamping the vein during needle withdrawal to prevent catheter dislodgement or blood loss.

Although HaemoBot performed more steps in the complete IV procedure than any other robot, it did not perform all of the steps that we identified in Chapter 2. Of the “27 Steps of Human IV Insertion”, HaemoBot performed all but the following steps:

- Gathering/opening IV supplies (Step 1) and disposing of them (Step 27). Although HaemoBot applied the alcohol and tourniquet automatically, it required a human operator to install and remove the IV needle and IV line. These are easy tasks for the human but very difficult for the robot.
- Applying analgesic (Steps 6-7). Lidocaine injection presented safety and technical concerns beyond the scope of this proof-of-concept prototype. We considered applying a topical analgesic cream or patch, but the 30-60 minutes required for the numbing effect to work was too long for this step to be performed with the patient’s arm inside of the robot. Such a step should be performed by the human operator an hour prior to robotic insertion.
- Palpating the veins (Step 9). Given the many problems related to palpation that were discussed in Chapter 3, we used only our NIR vision system for vein-finding.
- Searching for valves (Step 10). We have not yet developed a sufficiently reliable method of locating valves with the robot, and this remains a focus for future work.
- Attaching the IV bag’s Luer Lock to the catheter (Step 22). The mechanical complexity and safety concerns for this task outweigh the benefits of its autonomous completion. We consider the human operator’s inspection of the completed insertion to be a necessary safety-check before attaching the IV bag and delivering fluids, and requiring the human operator to attach the Luer Lock mechanically forces this inspection. Further, attachment of the Luer Lock is not a difficult or error-prone task for humans that requires robotic help.
- Cleaning the hand of spilled blood (Step 23). The mechanical complexity of this step would make it difficult for the robot with little gain for the human.
- Securing the catheter to the hand (Step 24). This step would be quite helpful for the robot to complete because it would reduce the risk of catheter dislodgement before or during the human operator’s inspection and attachment of the Luer Lock. However, due to the mechanical complexity required to perform this step safely and reliably, we kept its completion as a topic for future work.




(A)




(B)

Figure 4.1: HaemoBot, Overall View, Picture Set 1.



(A)



(B)

Figure 4.2: HaemoBot, Overall View, Picture Set 2.

4.1.2 Unique Contribution 2: Moving the Catheter Independently of the Needle

We are the first researchers to treat the needle and catheter as two, separate parts and to provide the additional degrees-of-freedom and more complicated insertion strategy that are necessary to move the needle and catheter separately. For the insertion to be fully autonomous, the robot must be able to 1. hold onto the catheter during insertion, 2. slide the catheter off of the needle and into the vein while holding the needle still, and 3. let go of the catheter. We included a dedicated mechanism for inserting the catheter and two DOFs that acted together to hold/release the catheter. Further, these catheter-control DOFs were able to perform the grab, move, and release processes repeatedly during an insertion to allow HaemoBot to implement some of the Failure Mode Recovery Techniques for back-wall penetrations and valves observed in Chapter 2.

4.1.3 Unique Contribution 3: Implementing Inherent, Mechanical Safety

We are the first researchers to develop an IV robot that implements a comprehensive mechanical design for safety. Although there are additional safety precautions that would be required for FDA-approval and clinical deployment, we solved the main safety issues, including:

1. Backdrivability of all joints (except catheter grab-and-release) to allow for quick needle/hand removal,
2. Passive counterbalancing of all joints affected by gravity,
3. Force-limited restraint of the hand and arm,
4. Redundant position/angle-sensing through incremental and absolute encoders,
5. Collision-sensing through an accelerometer on the Needle Driver,
6. “Deadman timer” and global/specific Emergency-switches to mitigate computer and electrical failures,
7. Laser-targeting through the needle to facilitate safe vein/needle alignment during teleoperation, and
8. A novel, mechanical Remote Center of Motion (RCM) design that is immune to power loss and has many benefits over existing designs.

We are the first to employ laser-targeting through the IV needle in general or a mechanical RCM for setting the needle pose in a phlebotomy/IV robot.

4.2 HaemoBot, the Final Design

4.2.1 Overall Design

HaemoBot required enormous complexity to enable full, autonomous IV insertion. The following list gives a broad overview of this complexity, and the remaining sections are devoted to describing the critical aspects of the design in greater detail:

- 19 mechanical DOFs.
 - NeedleBot had 9 mechanical DOFs:
 - Needle X, Y, Z, roll, pitch, and yaw + catheter insertion and two grab-and-release DOFs.
 - AssistoBot had 8 mechanical DOFs:
 - X, Y, Z, blood-cuff inflation + deflation, alcohol spraying, air spraying, and hand restraint.
 - Stereo camera had 2 mechanical DOFs for spinning the laser texture projectors.

- 23 optical DOFs: 20 LEDs in grid (8-bit control each), 1 laser on Needle Driver, and 2 lasers on camera.
- 1,674 separate part designs:
 - 27% of parts designed by us for custom fabrication: 313 laser-cut, 94 machined (separate-designs), 38 3D-printed, 4 CNC sheet-metal bending, and 1 vacuum-formed.
 - Remainder of parts from commercial vendors such as McMaster-Carr, Misumi, and Stock Drive.
- 13,306 parts total, including duplicates (like the same screw installed 100 times).
- 4 Data Acquisition Cards, 6 custom PCBs, and 8 Atmega microprocessor/motor-controller boards.
- 20,000 lines code: 9,000 lines for vein-finding and 11,000 lines for non-vision robot controls.
- Footprint with patient chair was 2063mm x 1858mm (along patient's arm axis) x 2105mm tall.
- Weighed 589 kg / 1298 lbf / 0.589 metric ton.

All of this complexity required a large, heavy robot. To keep it transportable, we required its size and weight to be within the limits of standard hospital doorways and elevators while mostly-assembled (only the haptic device support plate had to be removed to fit through the doorways). To ensure that the robot could insert on real humans (not just dummy, latex tubes), we required that the robot be able to accommodate adult arms of different handedness and length with the patient's arm inserted into the robot beyond the elbow.

4.2.2 NeedleBot

We decoupled translation and rotation, using a prismatic XYZ system to carry the RCM and Needle Driver that contained the rotational DOFs. This decoupling provided greater design flexibility for both the translation and rotation systems, allowed us to use a mechanical RCM for yaw and pitch as a safety measure, and allowed us to re-calibrate the robot for a different needle length by making adjustments to the RCM only instead of involving the translation stages or control-code as well.

4.2.2.1 Translational XYZ

We used Misumi LX45 ball-screw-driven translation stages for the X, Y, and Z axes. Their backlash was on the order of 5 μ m, and their 10mm-lead allowed for backdriving under approximately 8lbf, which was achievable using only a single hand. The effective workspace for the XYZ system was 182mm (X-axis) x 168mm (Y-axis) x 121mm (Z-axis). We included a much larger workspace than required anatomically (more like 100mm x 100mm x 40mm) to provide flexibility in the robot's calibration, maintenance, and operation.

We drove each stage using a Maxon RE-50 DC brushed motor that attached to the ball screw via flexible coupler, and the motors were powered through Copely Junus-20amp amplifiers. We picked these motors and amplifiers to allow deceleration of all axes from a maximum speed of 50mm/s to rest in a distance of 0.1mm. We used US Digital 1250-CPR incremental encoders attached to the back motor shaft to sense position. In the future, we plan to add high-precision stringpots to each axis as redundant, absolute position-sensing for safety. We attached a 3-axis, digital accelerometer to the yaw plate that measured the XYZ-acceleration directly (instead of double-differentiating the encoder signals) to provide a clean acceleration measurement that could be used to reduce positioning overshoot via feed-back control and feed-forward dynamic compensation.

We placed a pair of redundant end-of-travel sensors at the ends of each translation axis. We used an optical slot sensor and moving, stainless steel flag that would trip first, followed by a magnetic sensor and

moving magnet that would trip 5mm later in case the optical sensor failed. Both of these sensors fed into the safety-microprocessor-subassembly to generate a directional, stop-motion command that was wired to the end-of-travel command lines in the motor amplifiers. We included soft end-stops 8mm past the magnetic sensors to soften impact and prevent damage to the ball screws in case all sensors failed to prevent a head crash. A final, hard end-stop was located 5mm behind the soft end-stops as an absolute stopping point.

The X and Y axes were homed via an optical slot sensor and a homing plate with a step in its middle such that half of the plate would trip the sensor, and half of it would not. This allowed the robot to know on startup which direction it needed to travel to find the homing position based on the initial state of the homing sensor. We homed the Z-axis using the bottom end-of-travel optical sensor for convenience.

We used non-contact sensors for homing and end-of-travel detection to ensure that we would not break the sensors through inadvertent collision, as can happen easily with a contact/touch sensor. These sensors all experienced some hysteresis, in that the position which tripped the sensor was different depending on the direction from which the flag/magnet approached the sensor. Fortunately, these different tripping-positions were reliable, so we could offset for them in the robot's control code.

4.2.2.1.1 Z-Axis

The Z-axis was the most complicated of the translation axes because it required 1. careful routing of many wires, 2. fabrication of the main, square support tube, and 3. a complicated counterbalance mechanism.

Because there were many wires that needed to move through the Z-axis to distal DOFs, we included a wire carrier, as shown in Figure 4.3A, to ensure neat, compact, and safe wire routing. The carrier was custom-made by Gortrac to contain the unusual feature of a middle inflection point (flipping from convex to concave). We used a second, smaller wire carrier on the opposite side to route the signal cable of the Needle Driver's ATI-Mini40 force/torque sensor because 1. the sensor's fixed-length signal cable was not long enough to route through the main wire carrier to its signal-processing box, and 2. the sensor's signal cable was sufficiently delicate that it would not have survived being surrounded by so many other wires.

Fabricating the main, square support tube was challenging. To keep the structure light-weight and simple, we used a stock, square aluminum tube and machined/squared it precisely on all four, long faces. This allowed us to bolt the Z-axis translation stage flat against the outer surface without bending the tube or the stage. To provide a secure, serviceable, and non-galling means of bolting parts to the aluminum tube's thin walls (5.1mm), we designed stainless steel, flanged plugs that fit into larger holes in the square support tube and contained a long, tapped hole for mounting the other parts. The plugs inserted into the tube's wall from the inside, and the large flange prevented them from being pulled through the wall. The plugs were held in place by four smaller screws between the flange and wall, but they experienced only shear, not axial, forces.

4.2.2.1.1.1 Z-Axis, Spring Counterbalance System

To achieve a Z-axis that was easily backdriveable without falling down under the weight of its payload (the RCM), we needed to counterbalance it. However, the RCM was sufficiently heavy that using a mass counterbalance would have added too much mass for the XY stages to support and accelerate quickly. For this reason, we used a spring counterbalance, as shown in Figures 4.3B-C.

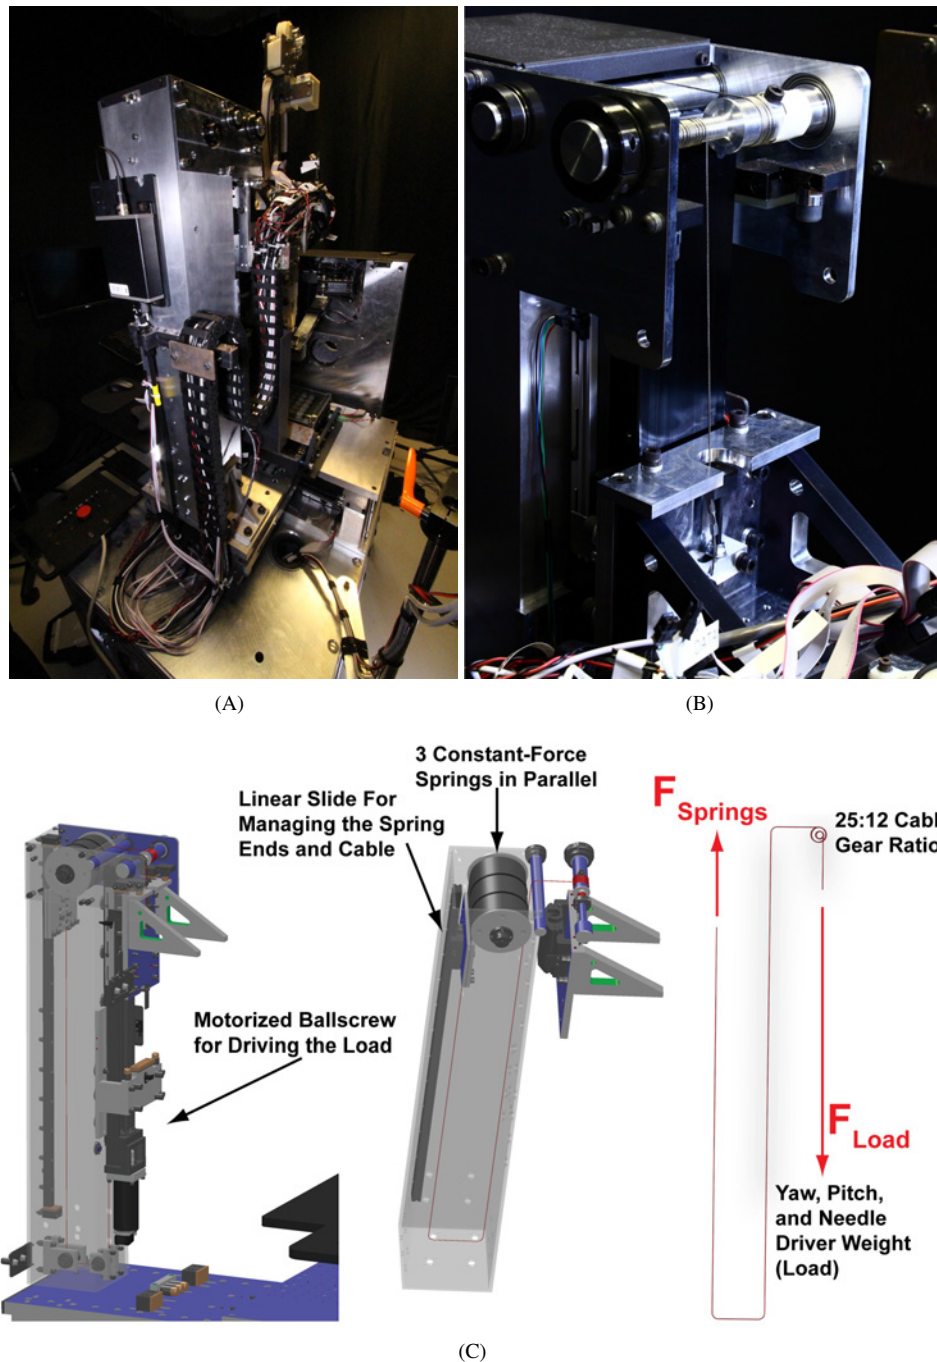


Figure 4.3: Z-Axis Details. A: The main wire carrier routes wires safely and neatly through the Z-axis to the RCM and Needle Driver. B-C: The Z-axis counterbalance uses springs to minimize the mass that must be moved by the XY stages. Three constant-force springs hide within the square support tube (for safety) and pull in parallel on a cable that routes out of the tube to a cable-reduction that increases lifting force. The cable-reduction pulls on the yaw plate, thereby supporting the weight of the RCM and Needle Driver. The springs bolt to the moving head of a linear slide, the rail of which is attached to the square support tube, thereby summing their forces and constraining their motion neatly.


Because high-force springs are a safety risk in the event of sudden failure, we enclosed them entirely within the square support tube with a steel, sheet-metal cap screwed onto the top of the tube. We used three constant-force springs in parallel to generate the counterbalancing force. We considered using gas springs instead of constant-force springs because the gas springs have a more gradual failure process (no sudden snapping), but we ultimately chose constant-force springs because their force-vs-distance profile is flatter in the middle of their travel. We used three springs in parallel instead of a single, higher-force spring because 1. we could not fit a bigger spring in the square support tube, 2. parallel springs allowed for a single spring failure without complete counterbalance-failure, and 3. we obtained better cycle-life from the smaller springs. To increase the springs' counterbalancing force, we included a 25:12 (2.083:1) cable-reduction right above the lifting cable's attachment point to the yaw plate, as shown up close in Figure 4.3B.

The linear slide that bolted to the back wall of the square support tube was critical for the counterbalance's operation. The ends of the three springs attached to a plate that rode on the linear slide's moving head, and the counterbalance cable (Sava Cable PN: 2047) also attached to that moving head. Beyond providing a method of summing the three spring forces and transmitting them to the counterbalance cable, the linear slide forced the springs to unfurl in a clean, controlled manner without twisting or rubbing on each other. We initially experienced trouble with bolting the linear slide against the unmachined, raw inner surface of the square support tube because the tube's inner surface warped the slide and caused it to bind when the slide was bolted down tightly at every mounting hole. We solved this by reducing the number of bolts attaching the linear slide to the tube's inner surface, thereby allowing the slide to keep its shape instead of bending with the tube.

4.2.2.2 Rotation, Yaw

The yaw DOF rotated the RCM about the vertical axis, or azimuth. The bottom, rotating structure that housed the pitch mechanism was bolted to the bottom of the hollow, stainless steel yaw shaft such that the bottom pitch-structure and yaw shaft rotated together. The yaw shaft inserted through a matched-pair of angular contact bearings that were housed in a central recess in the aluminum yaw plate, and the shaft was held in place by a 50mm shaft collar on top of the bearings. To remove all play from the bearings, we trapped and preloaded them using a special preloading-plate that bolted to the yaw plate over the bearings.


We drove the yaw axis via a 72:14 (5.143:1) timing-belt-reduction, as shown in Figure 4.4A, and this reduction provided better torque and angular resolution for the same motor/encoder. Because this was a small reduction, the yaw DOF was still easily backdriveable. A timing belt pulley was located concentrically by the yaw shaft and bolted to the bottom pitch-structure such that all three rotated as a single unit. We used a Maxon RE-30 DC brushed motor and the same type of motor encoder and amplifier that were used on the XYZ motors. The motor was fixed above the yaw plate and connected via flexible shaft coupler to a drive-shaft that was mounted on bearings in the yaw plate. The drive-shaft was fixed to a timing belt pinion that ultimately drove the yaw shaft through the timing belt. In this manner, we transmitted torque from the motor to the yaw shaft while protecting the motor from the radial loads of the belt. The belt was tightened by moving adjustable flat-belt idlers into the sides of the belt. A pair of hard end-stops limited the yaw DOF to a range of $\pm 109.7^\circ$ from the central angle that aligned the needle with the patient's arm axis, thereby



(A)


(B)

Figure 4.4: Yaw Details. A: Timing belt transmission. B: The Needle Driver wires translate and twist through the hollow yaw shaft. The shaft liner (white) sits against the shaft’s inner surface, isolating the pitch mechanism wires from the translating/twisting Needle Driver wires. The homing process uses an optical slot sensor and stepped-optical-flag.

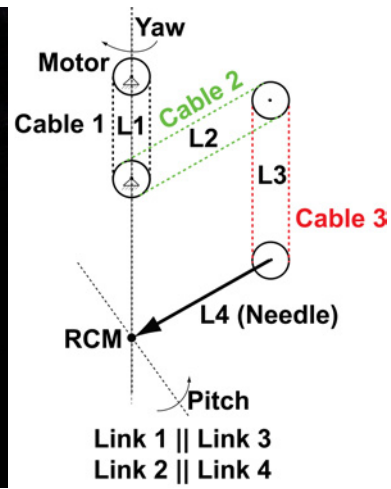


(A)

(B)



(C)



(D)

Figure 4.5: RCM, Earlier Designs/Prototypes. A: Our “U-bracket-RCM” prototype. B: Our “Parallelogram-RCM” prototype. C-D: Our earlier prototype of the single NeedleBot arm used the “Jensen-RCM” design.

preventing it from rotating too far in either direction and crashing the back end of the Needle Driver into the “AssistoBot-Arm-Box-Assembly” (the structure which housed the BP-cuff for the patient’s arm). We included a hole-and-slot pinning system in the yaw and pitch-structure plates such that an 8mm dowel pin could be inserted between the plates to fix yaw at its zero-angle during calibration/transportation.

We homed the yaw axis by rotating a stainless steel flag attached to the yaw shaft into a stationary slot sensor, as shown in Figure 4.4B. The stepped-shape of the flag tripped the sensor over only half of the yaw range, allowing the robot to know which direction to rotate for homing based on the sensor’s initial value.

We opted not to counterbalance the yaw DOF because gravity would only apply torques to the vertical yaw axis if the robot was tilted from being level with the floor. A human operator could not tilt the robot by accident because a special pallet jack was required to lift/tilt the robot. Adjustable feet allowed for easy, precise leveling of the overall robot until the yaw-axis was perfectly vertical and did not require counterbalancing.

The hardest design problem for the yaw DOF was routing the pitch and Needle Driver wires through the hollow yaw shaft, the details of which will be discussed in section 4.2.2.3.1.1.

4.2.2.3 Rotation, Pitch: Novel Remote Center of Motion (RCM) Design

We developed a novel, Remote Center of Motion (RCM) design for the pitch mechanism. In our previous prototypes of HaemoBot leading up to this final design, we fabricated and tested most of the standard RCM designs. Dissatisfied with each of them in some respect, we developed our own design that addressed the key problems that we had experienced with the other designs.

Our first, and simplest, RCM prototype is shown in Figure 4.5A, and we refer to this as the “U-bracket-RCM” design. The RCM consisted of two upside-down “U”-frames (an outer and an inner) that were connected at the bottom by bearings on either side of the “U”-shape. The yaw DOF rotated the outer frame about a vertical axis that intersected with the horizontal pitch axis about which the inner frame rotated. This design was extremely precise, simple to build, and required no calibration. However, it contained pinch-points and required that the inner-frame’s structure extend below the RCM origin. As such, the hand would need to rest entirely within the inner-frame’s horizontal space to prevent collisions, making for a very large inner-frame.

The “Parallelogram-RCM” is arguably the most famous RCM design, and our prototype of it is shown in Figure 4.5B. It used a rigid-linkage, parallelogram mechanism to force the final link (needle) to rotate at the same angle as the proximal/first rigid link. It was extremely precise, easy to fabricate and install, and required no calibration. The main detractors were its 1. limited range-of-motion/inability to rotate through the inflection point and 2. pinch-points created by the rigid-linkages. The limited range-of-motion was an inherent design limitation due to the parallel linkages colliding, and we wanted a larger range to provide clinical flexibility. The pinch-points were also an inherent design limitation due to the rigid-linkages. We pinched ourselves on several occasions, and this seemed to be a needless safety hazard, especially once the mechanism was fabricated out of a stronger, sharper material like metal.

The best RCM design that we tried before developing our novel design is shown in Figures 4.5C-D, and we refer to it as the “Jensen-RCM” after its primary inventor [33]. It used three cables to maintain parallelism between two sets of links, the last of which was the Needle Driver, acting like a cable-driven version of

the “Parallelogram-RCM” wherein cables replaced the rigid-linkages that maintained parallelism. The main advantages were its high range-of-motion through the vertical inflection point (unlike in the “Parallelogram-RCM” design) and lack of support structure near the hand (unlike in the “U-bracket-RCM” design). Unfortunately, the three cable-transmissions made fabrication, calibration, and maintenance extremely difficult.

4.2.2.3.1 HaemoBot RCM Operation Explained

Our novel RCM design relies on a set of orthogonal, linear slides and a single cable-transmission, as shown in Figures 4.6A-B. Figures 4.6C-F show the RCM moving over its entire yaw and pitch ranges. The green sphere, which was part of the RCM-calibration-needle-dummy and centered where the needle’s tip would be placed, was stationary throughout the RCM’s motion. Our design has multiple benefits, including that it 1. is easy to fabricate and calibrate, 2. contains only one cable, 3. has a high range-of-motion through its inflection point, 4. is easily gravity-counterbalanced, 5. contains no easily-accessible pinch-points, 6. contains no support structure that could collide with the hand, and 7. uses a novel, dedicated wire management system that prevents wire tangling. We will now briefly describe our RCM’s operation using terminology that matches the annotations in Figures 4.6A-B, and the full details can be found in the patent application [7].

The RCM mechanism was driven by a Maxon RE-30 DC brushed motor that was offset from and bolted to the back-pitch-plate. The motor attached via flexible shaft coupler to a drive shaft, or pinion, that was supported by a pair of ball bearings in the back-pitch-plate. The pinion drove the pitch-capstan via a 16:190 (1:11.875) cable-reduction, thereby transmitting the motor’s torque while isolating it from the radial loads of the cable-reduction. As shown in Figure 4.7C, the pitch-capstan contained integral cable-routing paths that were milled into it, and we used a screw-tensioning mechanism to tension the cable (Sava Cable PN: 1024). Homing of the axis used the same optical slot sensor and stepped-optical-flag design as for the yaw DOF. The pitch-capstan clamped to the top-pulley-shaft that inserted through the arc-shaped recess in the back-pitch-plate. The vertical-support-bar was supported by and pivoted about this top-pulley-shaft via a pair of ball bearings. At the end of the top-pulley-shaft, on the other side of the vertical-support-bar, a top-threaded-pulley was attached via set-screws.

The horizontal-linear-slide attached to the bottom of the back-pitch-plate, and a vertical-linear-slide attached to the moving head of the horizontal-linear-slide via a stainless steel slide-connecting-plate. The vertical-support-bar connected to the moving head of the vertical-linear-slide. Both linear slides were located precisely during installation by the slides’ integral locating holes and faces. Due to its connection to the linear slides, the vertical-support-bar could translate but not rotate with respect to the slides. Due to its pivot connection to the top-pulley-shaft, the vertical-support-bar moved in a circle centered on the RCM origin. A bottom-threaded-pulley connected rigidly to the Needle Driver on one side and was supported by a pair of bearings at the bottom of the vertical-support-bar on the other side. Because the vertical-support-bar translated in a circular arc, the bottom-threaded-pulley also translated in a circular arc. As the pitch-capstan rotated, the radial vector between its central pitch-axis and the top-pulley-shaft rotated with respect to the vertical-support-bar, which always remained vertical. This relative angle was transmitted via a 1:1 cable-transmission from the top-threaded-pulley to the bottom-threaded-pulley and Needle Driver. A turnbuckle

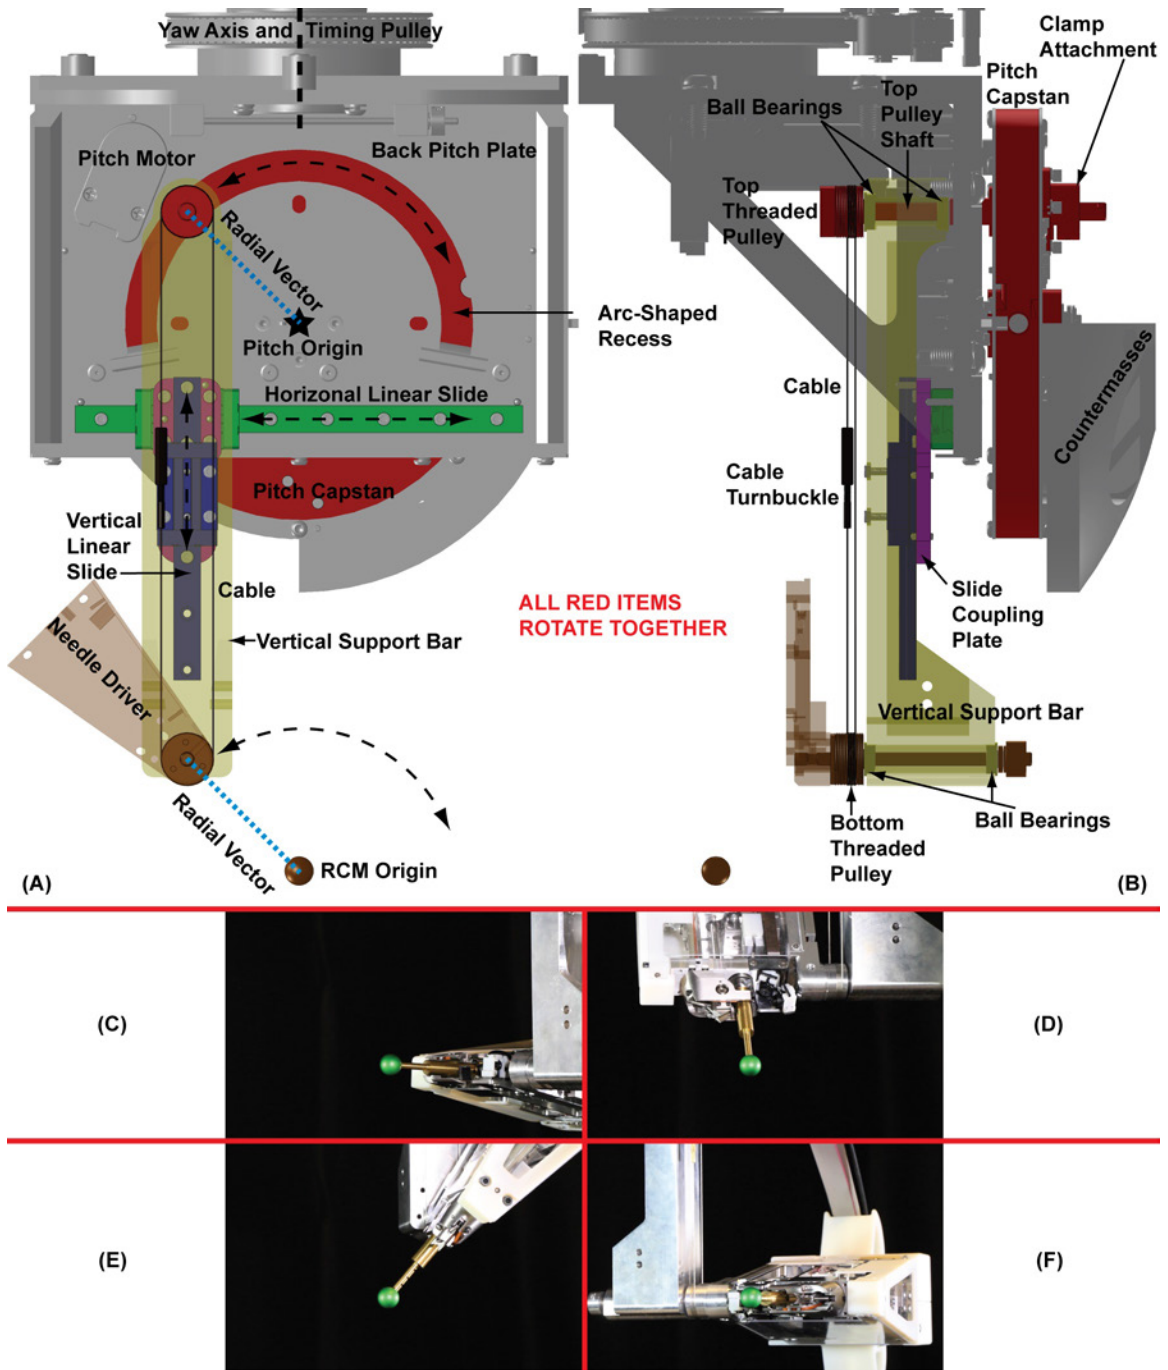


Figure 4.6: RCM, Novel Design and Operation Explained. A-B: Color-coded, annotated CAD rendering of RCM mechanism. C-F: Moving the RCM over its full yaw and pitch ranges with the IV needle replaced by the RCM-calibration-needle-dummy (green sphere tip) for optical measurement of accuracy. C: At one set of yaw, pitch motion limits (Needle Driver upright). F: At the opposite set of motion limits (Needle Driver upside-down). D-E: The Needle Driver travels through the pitch DOF's vertical inflection point between these two Subfigures.

was used to tension the cable-transmission on one side of the cable.

The pitch DOF's range was ± 92.7 deg from its vertical inflection point, extending the needle axis to 2.7deg below the horizontal plane in both the pointing-left and pointing-right configurations (Needle Driver upright or upside-down). Figure 4.7A shows the Needle Driver at its motion limits, as well as the vertical inflection point in between them.

We mass-counterbalanced the pitch DOF using a combination of fixed and sliding counterweights on the pitch-capstan. Splitting the counterbalancing into fixed and sliding components allowed for safer, easier adjustment of the counterbalancing force than adjusting a single, larger counterweight. As shown in Figures 4.7C-D, the fixed counterweight was a larger, wedge-shaped stainless steel part, and the sliding counterweight was a smaller, 18.5-alloy tungsten part that fit within a recess in the fixed counterweight. We contoured the fixed counterweight to stay within the motion envelope of the other RCM parts, and this required complex, extensive machining for which tungsten would have been unsuitable. As such, stainless steel was the best alternative. The geometry of the sliding counterweight was sufficiently simple that machining it from 18.5-alloy tungsten was straight-forward and provided a much higher density than for stainless steel.

4.2.2.3.1.1 Novel Wire Management System for RCM

We designed a novel wire management system for our RCM that routed the Needle's Driver's electrical wires from the back of the Needle Driver, through the RCM, and to the yaw plate where they were restrained. The wires inserted through the hollow yaw shaft so that they were aligned with the Needle Driver's back exit and did not interfere with yaw rotation by wrapping externally around the RCM. As the RCM's pitch mechanism rotated downward, the Needle Driver wires had to travel a longer distance from the hollow yaw shaft. Our system automatically unspooled that extra length of wire while keeping light tension on it to prevent slack, with the tension low enough that it did not affect the pitch DOF's gravity-counterbalancing. Had we instead clamped the wires at the exit of the yaw shaft such that there was a fixed length of wire that accommodated the maximum required length, the wires would have become tangled with the Needle Driver as it rotated upwards towards its vertical inflection point where a very short length of wire was needed. Figure 4.7A shows the drastic difference in wire length needed between the pitch mechanism's inflection point and bottom limits.

The Needle Driver's wires were comprised of two, 20-conductor, ribbon-cables (1mm-pitch, 0.9mm thick, "HiFlex" braiding to prevent fatigue failures) and an approximately 3mm-diameter ATI force sensor cable. To eliminate the need for a separate cable, we spliced the Needle Driver's USB cable into one of the ribbon-cables before it exited the Needle Driver. The ATI force sensor's cable contained too many sensitive signals for us to repeat this splicing process, so we kept it as a separate cable. The wire management system required great flexibility in the wires for proper operation, and this precluded us from mechanically joining the cables such that they would have behaved as one, stiffer cable. Even inserting the cables through a loose, braided jacket increased the stiffness too much (and also snagged on all of our routing pulleys). The one exception was that we included a single, loose piece of electrical tape around the cables near the middle of their spoolable length because the ATI cable was separating from and tangling with the ribbon-cables at that point. Other than at that, specific location, we did not experience issues with the cables separating from each other.




Figure 4.7: RCM, Pitch Design Details. A: The pitch DOF rotates the Needle Driver from one motion limit, through its vertical inflection point (middle image), to its opposite motion limit. B: The Needle Driver wires are guided into and along the axis of the hollow yaw shaft by the wedge-shaped features on the Needle Driver's back cover and the sets of pulleys at the entrance to the yaw shaft, always providing a large bending radius to prevent fatigue failure. C: Homing is accomplished via an optical slot sensor and stepped-optical-flag. Cable tensioning is accomplished via integral cable-routing paths in the pitch-capstan and a screw-tensioning mechanism. D: The pitch DOF is driven by a cable-reduction and counterbalanced via an outer, fixed countermass and inner, slidable countermass.

Our wire management system routed the wires in a special manner at three main locations: the Needle Driver's exit, the entrance to the hollow yaw shaft, and at the tensioning-pulley system on top of the yaw plate. We will now discuss the design details of each location.

Without a special design, the wires would have wrapped around the back edge of the Needle Driver and been pulled in a straight line to the hollow yaw shaft, thereby bending over such a small radius that they would have experienced fatigue failures. To provide a safe bending radius, we included two flanged, wedge-shaped features on the Needle Driver's back, 3D-printed cover around which the wires bent on their way to the yaw shaft, as shown in Figures 4.7A-B, 4.6F, and 4.8D. Their symmetry and placement allowed for the Needle Driver to rotate through its vertical inflection point with the wires transitioning seamlessly from bending around one of the wedges to the other.

Similarly, we needed to bend the wires into the hollow yaw shaft along its central axis at a safe bending radius. As shown in Figures 4.7B and 4.8A-B, we included a pair of large, 3D-printed pulleys at the entrance to the yaw shaft, and the wires passed between the pulleys to enter the hollow shaft. We experienced an issue with the wires translating along the pulleys' axial direction and rubbing on the inner surface of the yaw shaft, so we included a set of rotating shafts (orthogonal to the main pulley axes) that constrained this walking motion and changed the contact mechanism from sliding to rolling.

The wires exited the hollow yaw shaft and routed over the "wire-tensioning-pulley", as shown in Figures 4.8C-E and 4.9. This pulley was lifted upwards by a spring counterbalancing mechanism, thereby lifting and tensioning the wires as well. The pulley was mounted on the moving head of a vertical, linear slide that constrained its motion and was positioned so that the wires left the pulley along the yaw axis. The pulley included flanges and a top cover that prevented the wires from falling off of the pulley if the wires went slack. After routing over the wire-tensioning-pulley, the wires were gently strain-relieved on the yaw plate, again being careful to provide a safe bending radius before the strain-relief.

The spring counterbalance consisted of a spring pullbox (Stock Drive PN: A3Z55-C12, "spring power reel") that pulled on the wire-tensioning-pulley through a cable-transmission/reduction. We used the pullbox because of its low cost, compact size, convenient cable termination, and near-constant force-vs-distance profile. To increase the pullbox's force, we included a 51.5:12 (4.267:1) cable-reduction between the pullbox cable and lifting cable, and the details of these cable paths are shown in Figures 4.8E-F and 4.9A. To allow for changes in the Needle Driver components that would affect the fixed spring counterbalancing, we included a variable mass, or "shot-bay", on the linear slide's head, as well. The shot-bay was filled with tungsten shot (Tungsten Heavy Powder, Inc., Tungsten Shot #9, 17.0g/cc) that was double-wrapped in nitrile gloves as an extra precaution against spilling in the event of a catastrophic failure.

The wires twisted within the hollow yaw shaft as the yaw DOF rotated, as shown in Figure 4.8D, because their strain-relief was fixed on the yaw plate while their distal end rotated about the yaw axis with the Needle Driver. Because the distance over which this twisting occurred was relatively long, there was little torsional stress applied to the wires. However, the twisting meant that no other wires, like those for the pitch mechanism, could fit between the flat sides of the Needle Driver wires and the yaw shaft's inner surface without




Figure 4.8: RCM, Wire Management System, Picture Set 1. A-B: Two, orthogonal sets of pulleys (large/tan and small/gray) route the Needle Driver wires into the yaw shaft at a safe bend radius and without rubbing against the shaft. C-F: A spring-counterbalance automatically spools/unspools the Needle Driver wires by pulling them upwards to remove slack, thereby preventing them from tangling with the Needle Driver as it rotates through the pitch DOF's vertical inflection point. D: The Needle Driver wires twist within the hollow yaw shaft as the yaw DOF rotates. E: The wire-tensioning-pulley includes flanges and a top cover that prevent the wires from falling off of the pulley. F: A cable-reduction between the pullbox cable and lifting cable increases the pullbox's counterbalancing force.





Figure 4.9: RCM, Wire Management System, Picture Set 2. A: The wire-tensioning-pulley mounts to the moving head of a vertical, linear slide that constrains its motion and is positioned so that the wires leave the pulley along the yaw axis. The shot-bay contains tungsten shot and serves as an adjustable mass to accommodate changing counterbalancing requirements on the Needle Driver for the same spring/pullbox. The Needle Driver wires attach to the yaw plate at the bottom of the linear slide after curving under a gentle radius. B: A considerable length of wire unspools into the RCM as the Needle Driver rotates to its bottom motion limit (left image, “Down”).

special precautions to prevent collisions. Our solution was to insert a 3D-printed liner along the inner surface of the yaw shaft that isolated/protected the pitch mechanism wires in two access spaces between the shaft's inner surface and liner, as shown in Figures 4.8A-B. This liner was composed of two, separate halves such that it could be closed around the Needle Driver's wires without disassembling anything.

4.2.2.3.2 HaemoBot RCM Calibration Explained

4.2.2.3.2.1 Measuring the IV Needles

The accuracy of our RCM was predicated on the IV needles having consistent dimensions such that the RCM could be calibrated for a single needle-length without requiring adjustments for individual needles. We found high consistency in our measurements of 33 needles, the dimensions of which are shown in Figure 4.10A.

4.2.2.3.2.2 Pitch Angle and Needle Length Calibration

Setting the needle axis of the Needle Driver to point radially towards the RCM origin required special calibration. To accomplish this, we first dowel-pinned the pitch-capstan to the back-pitch-plate and the Needle Driver to the RCM's vertical-support-bar in one of their bottom motion-limit configurations, as shown in Figures 4.10B-C. We then installed the cable-transmission between the top-threaded-pulley and bottom-threaded-pulley, tightening it in place via the turnbuckle. Using a single turnbuckle produced uneven tension between the two cable sides that needed to be equalized for the calibration to succeed. Because the top-pulley-shaft was clamped to the pitch-capstan, we were able to loosen and wiggle it during the cable tensioning process to equalize tension in the two cable sides. Once the tension was sufficiently high and equalized, we clamped the top-pulley-shaft rigidly and removed the dowel pins to complete the calibration process.

We used two special mechanisms to set the needle tip precisely at the RCM's origin: a Needle Length Adjustment Mechanism and a RCM-calibration-needle-dummy. The Needle Driver connected to the RCM structure via a split plate with sides that could translate with respect to each other such that their relative separation could be adjusted. This allowed for the needle tip to be placed at the RCM origin even with different needle lengths and errors from stacked machining tolerances. Figure 4.10C shows this mechanism, with the fixed side shown in pink and the translating side shown in yellow. The details of this mechanism will be discussed later in section 4.2.2.4.9 after more context has been provided.

The RCM-calibration-needle-dummy provided a strong, tightly-toleranced, and color-segmented surrogate, or "dummy", for the IV needle that allowed for easier, more reliable calibration via the Needle Length Adjustment Mechanism than could be obtained using an actual IV needle. It consisted of a green-anodized, aluminum sphere that was centered within 0.001in of where the needle tip theoretically would be and a brass stem/support with the same diameter as the IV needle base such that it could be held by the needle-holder, to be discussed in section 4.2.2.4.2. The sphere was made a different color from any surrounding parts to allow for optical measurement of the RCM's accuracy, to be discussed in section 4.2.2.3.2.3. Figures 4.6C-F and 4.10C show how clearly the green sphere stood out against the surrounding parts. To calibrate the RCM and Needle Driver for the correct needle length, they were first pinned as for pitch angle calibration, as shown in Figure 4.10B. A tightly-toleranced needle-length-fixture was bolted to the side of the vertical-support-bar, as shown in Figure 4.10C. With the RCM-calibration-needle-dummy inserted into the Needle Driver in place




Figure 4.10: RCM, Pitch Angle and Needle Length Calibration.

A: We measured important dimensions on 33 new Jelco 20-gauge, 1.25in IV catheters (Smith Medical 4056).

$D_{nb} = 5.838 \pm 0.017\text{mm}$. $D_n = 0.700 \pm 0.013\text{mm}$. $D_c = 1.150 \pm 0.025\text{mm}$.

$D_{ct1} = 6.373 \pm 0.011\text{mm}$. $D_{ct2} = 5.491 \pm 0.021\text{mm}$. $L_{ncb} = 6.735 \pm 0.023\text{mm}$.

$L_{ct} = 9.437 \pm 0.024\text{mm}$. $L_n = 59.167 \pm 0.034\text{mm}$.

B: Setting the Needle Driver's needle axis to point radially towards the RCM origin requires 1. dowel-pinning the pitch-capstan to the back-pitch-plate and the Needle Driver to the RCM's vertical-support-bar in one of their bottom motion-limit configurations and 2. installing the cable-transmission and tightening it in place via the turnbuckle. C: Setting the needle tip precisely at the RCM's origin requires adjusting the separation between two plates on the Needle Driver so as to translate the RCM-calibration-needle-dummy until it contacts the needle-length-fixture.

of an actual IV needle, the Needle Length Adjustment Mechanism on the Needle Driver was adjusted until the green sphere contacted the needle-length-fixture (the face of which was offset from the RCM's origin by the sphere's radius). To detect that contact, we used the continuity-testing function of a Digital Multimeter (DMM), placing one lead on the needle-length-fixture and the other lead on the RCM-calibration-needle-dummy. The DMM beeped when first contact was made because the circuit was completed. Because the Needle Length Adjustment Mechanism used a fine-pitch adjustment screw to change the needle tip's position, we were able to make extremely fine adjustments very slowly and carefully.

4.2.2.3.2.3 Accuracy Analysis of the RCM

The RCM-calibration-needle-dummy allowed for optical measurement of the RCM's accuracy, or the size of the needle tip's motion envelope (which theoretically should have been a point). The spherical tip appeared the same in all orientations, providing a constant geometry that we could track accurately, whereas the needle bevel looked different from different vantage points such that tracking its positional change would have been inaccurate. Optical measurement allowed for non-invasive measurement (without introducing "observer error"), whereas the contact force of a mechanical dial gauge would have caused deflections and inaccurate measurements. The RCM-calibration-needle-dummy could also be used for calculating the frame transformation between NeedleBot and the stereo camera (with its longpass filter removed) by translating NeedleBot along its XYZ axes and comparing the transformed vectors in the stereo camera's frame.

To measure the RCM's accuracy optically, we first recorded video of the RCM via a macro lens while moving the yaw and pitch over their entire ranges at constant velocities. The macro lens magnified the sphere to fill the majority of the video frame, thereby providing a higher spatial resolution. We attempted to light the sphere uniformly and without specular reflections to produce more robust color-thresholding, but this process requires further improvement in the future. We then performed a frame-by-frame image analysis wherein we 1. created a black-and-white mask based on the sphere's green color (RGB-color-thresholding) and 2. fit a circle (yielding a center location and radius estimate) to the masked image using the MATLAB function "imfindcircles" (a Circular Hough Transform). By dividing the sphere's known, high-precision radius (0.250 ± 0.001 in) by the mean, estimated radius from the Hough Transform, we obtained a pixels-to-mm gain that gave our error physical dimensions. We took the median circle position as the center of our motion envelope and then calculated the average distance from that position as 0.694 ± 0.303 mm. That is, the majority of the motion envelope was contained within a circle with a diameter of approximately 2mm. Such a large error exceeds what we anticipated, and we have investigated it to discover the following explanations:

1. Since the needle length was very consistent ($L_n = 59.167 \pm 0.034$ mm), its variation could not explain the large errors we observed.
2. Almost all of the error occurred when changing the pitch angle, not yaw.
3. During our calibration, the external-ring-gear on the back of the needle-holder loosened, thereby allowing the needle-holder to wiggle in its supporting ball bearings. This play, combined with the cantilevered mass of the RCM-calibration-needle-dummy, caused the dummy to droop in the direction of gravity, which was confirmed in the error-vs-pitch-angle curve. In the future, we must ensure

proper bearing preloading and perhaps switch to a lighter-weight design for the RCM-calibration-needle-dummy which more closely matches the mass of the IV needle, which is minimal.

4. As the pitch angle lowered, the Needle Driver's mass was cantilevered at a greater distance, thereby applying greater torque to the RCM's cable-transmission. The cable stretched, thereby allowing the needle tip/sphere to move away from the RCM's origin. We also observed that cable loosening over the course of many experiments caused similar positional error. We plan to improve this by 1. increasing the cable and pulley diameters such that the stress in the cable (and hence elastic deformation/stretch) is reduced and by 2. implementing a lock-nut design on the cable's turnbuckle that prevents it from loosening the cable over time.
5. Occasional stiction in the RCM wire management system caused the Needle Driver to be pulled against its wires, creating much larger torques than usually experienced and stretching the cable-transmission, as described in Item 4. Including white lithium grease on the vertical, linear slide should eliminate or mitigate this issue.
6. When the sphere's attachment to its brass stem was visible in the camera's frame, the sphere appeared to have a section missing such that the Circular Hough Transform was less accurate. We noted that the positional error was greater for those frames. This stem attachment is an inherent limitation, but we plan to investigate methods of compensating for it in the analysis code.

As some of these fixes require remaking/modifying parts, we have yet to test the extent to which addressing these issues would affect the measured accuracy/error. However, we are optimistic that we can greatly reduce the error to a practical level that will not have a clinical impact.

4.2.2.4 Needle Driver

The Needle Driver housed the 1. needle, 2. laser-targeting, roll, catheter insertion, and catheter grab-and-release DOFs, and 3. insertion sensors (accelerometer and 6-axis force/torque sensor). Except for its back wire-exit, the entire Needle Driver fit within a ± 10 deg wedge, centered on the needle, such that the needle's pitch could be lowered to within 10deg of the hand before the Needle Driver's structure made contact.

4.2.2.4.1 Laser-Targeting Through the IV Needle

We developed and patented [8] a method of shining a laser directly through the IV needle to facilitate safe vein/needle alignment, especially during teleoperation when the practitioner, and not the robot's computer vision algorithm, was responsible for that alignment. Without laser-targeting, the needle tip must be close to the skin's surface to provide an accurate visual estimate of the insertion point, increasing the chances of the needle accidentally cutting the patient during the alignment process. Keeping the needle tip at a larger, safer distance from the skin during alignment provides a less accurate visual estimate, increasing the chances of missing the targeted vessel. Our solution was to shine a red laser down the bore of the needle to show precisely where the needle would insert while keeping the needle at a safe distance from the skin, as shown in Figure 4.11A. Although any wavelength of light could be used, we chose red so that both the human teleoperator and the robot's NIR vision system could see the targeting-laser. There were two, novel design components that enabled this laser-through-needle targeting, including:

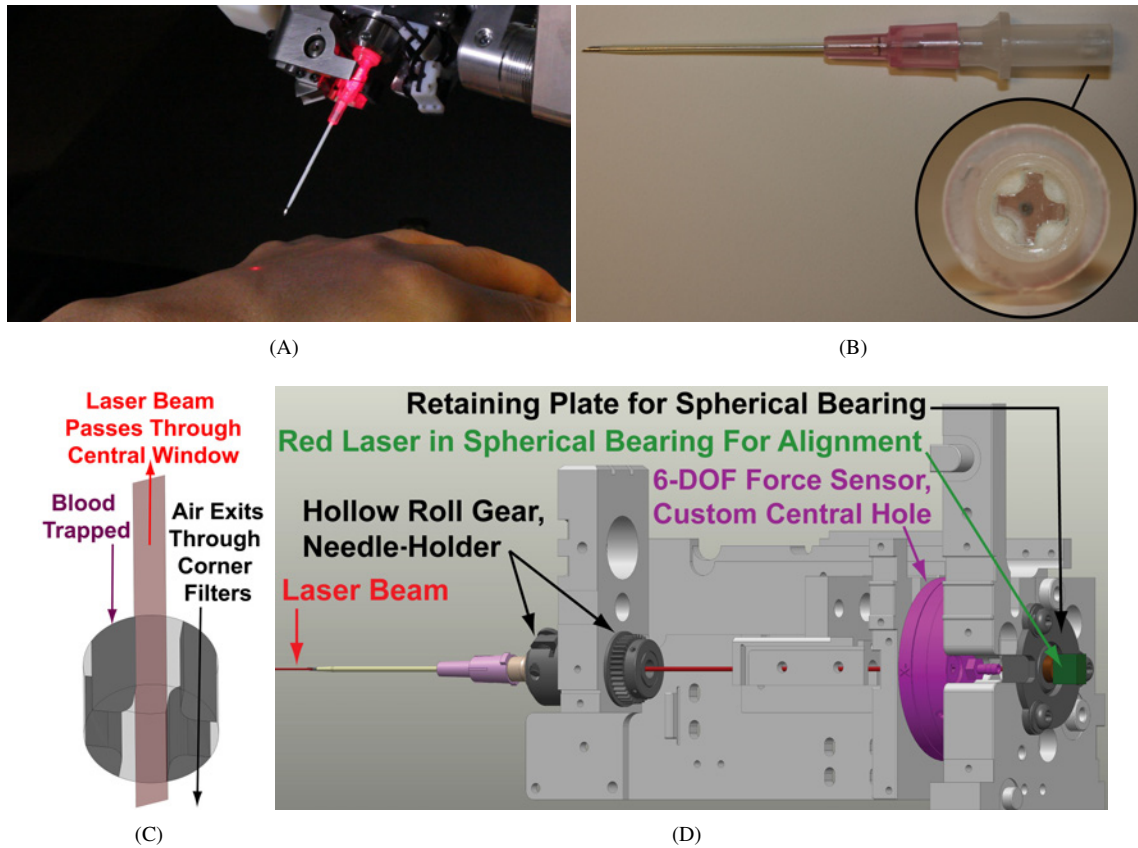


Figure 4.11: A: We patented [8] a method of shining a laser directly through the IV needle to allow vein/needle alignment from a safe distance, especially during teleoperation. B-C: Novel, transparent flashplug/filter that allows the laser to pass, traps blood, and allows air to escape. D: Implementation on the Needle Driver. The catheter insertion DOF (hidden here) also allows the targeting-laser to pass.

1. A transparent flashplug/filter design that maintained the needle's normal operation and sterility and
2. A method of aligning the laser beam with the needle bore both in translation and rotation.

Standard IV needles use an air-permeable, blood-impermeable filter, or flashplug, to allow air to escape the needle as it is displaced by blood entering and remaining trapped in the needle during the blood-flash. However, this filter is opaque, thereby preventing a laser from passing through the needle bore. We developed a transparent flashplug/filter that allowed the laser to pass through the needle without interfering with the normal air-escape and blood-entrapment processes. Our filter included a transparent center and an off-center, air-permeable/blood-impermeable material that allowed air through but not blood. Figure 4.11C demonstrates how the filter works, and Figure 4.11B shows our filter installed in an IV needle. Although the central window requires material in the middle where the laser beam passes, where and how that central window is supported is quite flexible. The window can extend any number of prongs of various geometries to the inner surface of the needle base for support but does not require them. We used a 4-prong design because it could be lightly press-fit into the needle base without the need for adhesive, and its symmetry best preserved the circular shape of the needle base after slight deformation. We used acrylic as the central window and cotton

(plucked from a cotton-ball) as the filter material in between the prongs. The cotton's blood-impermeability was verified with a small blood sample obtained from the author's finger via a blood-glucose-testing lancet.

Relying on precise machining tolerances to ensure laser/needle alignment would be expensive and difficult. Further, we used inexpensive lasers which varied widely in the concentricity between their beam and cylindrical housing. To enable the use of machined parts and lasers with poor tolerances, we devised methods for rotational/translational, post-installation alignment.

The laser was mounted in a spherical bearing that allowed for rotational alignment. As the beam was aligned more closely with the needle's bore, the flash-chamber and targeting-dot through the laser glowed brighter. Using this visual feedback and separating the alignment process into its X and Y components allowed for alignment in a matter of 1-2 minutes. The spherical bearing contained a nylon sleeve between the inner and outer races such that friction maintained the laser's orientation after the alignment process. Due to the spherical bearing's size and need to sit along the needle axis, it could only be placed at the back end of the Needle Driver. As a result, we were forced to design the 6-axis force/torque sensor, catheter insertion DOF mechanism, roll DOF mechanism, needle-holder, and Needle Driver's main base plate to allow for the laser to pass through or around them, as shown in Figure 4.11D. For translational alignment, we used a laser with a beam diameter that was significantly larger than the needle's bore such that the bore would still be contained entirely within the laser beam even with some translational misalignment.

4.2.2.4.2 Needle-Holder

The needle-holder held the IV needle by its base. It was supported by two ball bearings that allowed for infinite rotation about the needle axis, and it was hollow to allow the targeting-laser to pass. The front end of the needle-holder clamped onto the needle base using an integral slit-clamp that was closed by an M3 torx-screw, as shown in Figures 4.12A-B. The hinge side of the clamp contained only a 0.2mm gap such that the clamp could not be damaged by opening it wider than the gap. Due to the IV needles' consistency in base diameter, we were able to set the clamp's diameter once to produce a light press-fit between the clamp and needle that held the needle securely but also allowed for the needle to be changed easily by hand without needing to adjust the clamp screw. As shown in Figure 4.12A, the back end of the needle-holder located and connected to the roll DOF's external-ring-gear via set screws that pressed into two, perpendicular flats.

4.2.2.4.3 Roll DOF

The roll DOF served to 1. align the needle "bevel-up" with the insertion site's surface normal and 2. drill past valves, as described in Chapter 2 as a Failure Mode Recovery Technique. Whereas the first purpose required only a single revolution at low speed and torque, the drilling purpose required infinite rotation at higher speed and torque. The roll mechanism needed to be hollow to allow the targeting-laser to pass, so we placed the motor to the side of the needle-holder and connected the two with a gear-reduction. As shown in Figures 4.12A-C, the roll mechanism was comprised of 1. an aluminum motor-mount that clamped/held 2. a Maxon RE-10 DC brushed motor with a 4:1 gearhead, the output shaft of which was epoxied to 3. a 10-tooth pinion that drove 4. a 36-tooth hollow, external-ring-gear (0.4mm module) that was attached via two, perpendicular set screws to 5. the needle-holder.




Figure 4.12: Needle-Holder and Roll DOF Details. Both were hollow to permit the targeting-laser to pass.

The external-ring-gear was machined from gear-stock to include the set screws and integral bearing-inner-race-spacer. Beyond connecting the motor to the needle-holder, the 3.6:1 gear-reduction improved the roll torque, which was needed for drilling through valves, and angular resolution, which was needed because the motor had to be paired with the factory-assembled, stock 256-CPR Maxon encoder. Without a precise torque requirement/specification for drilling valves, we set the torque high enough that we were unable to stop the roll DOF by gripping the needle's shaft with our fingers. One disadvantage of the gear reduction was that its backlash caused limit-cycling that required considerable fine-tuning of the control gains to resolve.

The aluminum motor-mount clamped lightly around the motor's circumference. The motor-mount was bolted to the Needle Driver's base plate via two screws that went through slots such that the motor assembly could be translated closer to or further from the external-ring-gear, using an integral, machined locating-surface in the base plate as an alignment guide. This distance-adjustment mechanism allowed us to reduce backlash by positioning the gears as close as possible before binding without relying on machining tolerances.

Routing the motor/encoder wires was challenging because the wires were in a short, 10-conductor FFC (Flat Flex Cable) that was the only option from the manufacturer and not modifiable. The FFC only reached partway between the roll motor and the Needle Driver's main PCB located at the back, so we placed a breakout-PCB partway along the Needle Driver, as shown in Figure 4.13A, to route the FFC signals from both the roll and catheter insertion motors/encoders into a 20-conductor ribbon cable that ultimately exited

the Needle Driver. As shown in Figure 4.12C, we designed a 3D-printed part that clamped onto the roll motor and routed the FFC safely from the roll motor to the breakout-PCB. This part routed the FFC through a tortuous path that took up the FFC's excess length, bent it gently underneath the Needle Driver's 10-degree-wedge-profile, and twisted/translated it to reach the receptacle on the breakout-PCB.

4.2.2.4.4 Catheter Insertion DOF

The catheter insertion DOF translated the catheter entirely off of the needle's shaft (42mm) and into the vein. Figure 4.13A shows the catheter insertion DOF fully retracted and extended, and Figure 4.14 labels all of the important mechanism components. The drive transmission consisted of the following:



1. An aluminum motor-mount that clamped/held
2. A Maxon RE-10 DC brushed motor with an integral 64:1-reduction gearhead that connected to
3. A 2-3mm, bore-changing flexible shaft coupler (SDP-SI PN: S50HS9M1P02P03) that connected to
4. A 3mm drive-shaft supported by a pair of ball bearings with a set screw flat for attaching to
5. A 14-tooth, 2mm-pitch (GT2) timing belt pulley (SDP-SI PN: A6A51-014DF0304) that connected via
6. A 49-groove/98mm-long, 2mm-pitch (GT2) timing belt (SDP-SI PN: A6R51M049030) to
7. An identical timing belt pulley mounted via set screw to a 16-start, 0.375in-lead, 84%-efficiency, 3/16-in-diameter lead screw (McMaster-Carr PN: 6642K47, Haydon-Kerk PN: LSSSSR-018-0375-FY18) that was supported at both ends by ball bearings and translated
8. A matching lead screw nut (McMaster-Carr PN: 6642K57) that was bolted to
9. An aluminum carriage-block that attached to and translated
10. The catheter insertion rod, which translated through a 9mm-ID bushing/spline to move
11. The catheter grab-and-release DOFs' distal-contact-mechanism that held/released the catheter.

The aluminum motor-mount clamped lightly around the motor's circumference. The motor-mount was bolted to the Needle Driver's base plate via two screws that went through slots such that the motor assembly could be translated closer to or further from the lead screw, using an integral, machined locating-surface in the base plate as an alignment guide. This distance-adjustment mechanism was necessary for installing and setting the proper tension on the timing belt without relying on tight machining tolerances. The motor-mount also housed the flexible shaft coupler, bearings, and drive-shaft to the first timing belt pulley so that the motor's gearhead was isolated from the significant radial loads exerted by the timing belt.


The gearhead provided the required insertion force at a reasonable speed and also improved the positional resolution afforded by the stock incremental encoder, which was limited to 256-CPR and was the only option available for the motor. The positioning of the motor allowed the motor/encoder FFC to bend gently backwards to the receptacle on the breakout-PCB without an additional wire-routing mechanism.

The timing belt transmission was necessary for space issues. We could only fit the motor on one side of the Needle Driver's center-line and the lead screw on the other side, so we needed the timing belt to connect the two while keeping the needle axis hollow to allow the targeting-laser to pass.


We selected the lead screw for its small diameter, backdrivability, and low backlash. We machined the ends of the lead screw stock so that they could be supported by 1/8-in ball bearings, with one end receiving a



(A)




(B)




(C)

Figure 4.13: Needle Driver, Catheter Insertion DOF. A: Top image, fully retracted. Bottom image, fully extended. B-C: The stringpot serves as a redundant, absolute position sensor for added safety.



(A)



(B)

Figure 4.14: Needle Driver, Catheter Insertion DOF, Design Details. B: The catheter insertion rod is highlighted in red and the spline in green. The spline prevents rotation of the catheter insertion rod and everything connected to it (including the catheter grab-and-release DOFs). The spline's two prongs protrude into the same slots in the catheter insertion rod that route the Teflon tubes/nitinol push-pull wires for the catheter grab-and-release DOFs. Top, left image: The catheter insertion rod is shown as opaque. Bottom, left image: The catheter insertion rod is shown as transparent to reveal the Teflon tube routing in the bottom, curved slot.

retaining ring slot and the other end receiving a set screw flat for attaching the timing belt pulley. We modified the lead screw nut to include recesses for the mounting screws to be flush with the nut's flange.

The carriage-block, catheter insertion rod, and catheter grab-and-release DOFs' mechanism all translated together as a single unit. The carriage-block's front end clamped to and moved the catheter insertion rod, which supported and moved the catheter grab-and-release DOFs' distal-contact-mechanism that actually held the catheter. The carriage-block's back end attached to and moved the catheter grab-and-release DOFs' proximal-drive-assembly. Homing of the incremental encoder was accomplished via an optical flag that was attached to the carriage-block and tripped a slot sensor at the back end-of-travel. We used an M3 x 0.25mm fine-pitch adjustment-screw with a delrin lock-nut for vibration resistance as an adjustable end-stop to set the back motion-limit for the carriage-block. We required an adjustable end-stop so that we could leave a small gap (0.1mm) between the catheter hub and needle flange to prevent rubbing during rolling/drilling motions.

4.2.2.4.4.1 Catheter Insertion DOF, Spline Details

There were several design features that were critical to the catheter insertion rod's successful motion. Due to the tight space constraints near the tip of the Needle Driver, we were forced to use a plain bushing to provide the catheter insertion rod's linear motion. We used oil-impregnated, SAE-bronze for the bushing to reduce friction and reamed the bushing precisely to minimize play between the bushing and rod.

In our first prototype with only the bushing (and lead screw) constraining the linear motion, the machining tolerances allowed for the catheter insertion rod to rotate ± 15 degrees from its zero-angle. This also tilted the catheter grab-and-release mechanism that held the catheter, thereby bending the catheter/needle and causing the grab-and-release mechanism to collide with the hand. Adding a laser-cut, delrin spline after the bronze bushing solved this problem completely. The spline contained two prongs that protruded into slots in the catheter insertion rod to prevent rotation. These slots were originally designed for routing the Teflon tubes that were part of the catheter grab-and-release DOFs' transmission. The Teflon tubes extended from the proximal-drive-assembly on the back end of the catheter insertion rod, through the slots in the catheter insertion rod, to the distal-contact-mechanism that held/released the catheter. We used delrin for its lubricity, durability, and high-precision tolerances during laser-cutting. Figure 4.14B shows the catheter insertion rod, spline, Teflon tubes, and distal-contact-mechanism. Figures 4.16A-C show the catheter grab-and-release DOFs' distal-contact-mechanism and Teflon-tube-routing of the nitinol push-pull wires that actuate them.

4.2.2.4.4.2 Catheter Insertion DOF, Stringpot

As shown in Figures 4.13B-C, we used a string potentiometer, or "stringpot", as a redundant, absolute position sensor on the catheter insertion DOF as a safety precaution. The stringpot bolted to a fixed part of the Needle Driver and measured the distance that a steel cable was pulled out of the stringpot by the catheter insertion DOF's translating mechanism. The stringpot spooled the steel cable internally via a constant-force spring around a precision potentiometer that output an analog voltage proportional to the distance of cable that had been pulled out of the stringpot. For the catheter insertion DOF's full range of motion, the stringpot's voltage output range was 0.478-4.342volts. This signal was sufficiently large and strong to allow for routing of the raw, analog signal over 12ft from the stringpot to the DAQ card that read the signal without appreciable

noise/resolution-loss. We mounted the stringpot's housing on the sensing-side of the force sensor to prevent the stringpot's spring force from appearing as part of the insertion force.

The free end of the stringpot's steel cable attached to the catheter grab-and-release DOFs' proximal-drive-assembly such that the cable formed the hypotenuse of a triangle with one vertex at the stringpot's cable-exit, one vertex on the proximal-drive-assembly, and the base of the triangle parallel to the Needle Driver's central/needle axis. Figure 4.13B shows this triangle clearly. The Pythagorean Theorem allowed us to calculate the length of the triangle's base as a starting length offset + catheter insertion distance moved.

There were several, key design aspects to ensuring an accurate measurement via the stringpot, including:

1. The cable's free end needed to rotate so that the cable could follow the hypotenuse direction fully,
2. The cable needed to be pulled so that its resting-crimp, which sat against the stringpot's cable-exit at zero distance, did not bottom-out when the catheter insertion DOF was fully retracted, and
3. The cable needed to be attached to the pivot such that we could install, adjust, service, and uninstall it without permanently damaging the cable.

To meet all of these requirements, we attached the stringpot cable to a shoulder-bolt pivot on the side of the catheter grab-and-release DOFs' proximal-drive-assembly using a custom assembly that clamped the cable over a long, curved distance without damaging the cable. The clamp consisted of a 3D-printed part with a curved groove that guided the cable and a laser-cut, stainless steel clamp-plate that trapped the cable in the groove. We placed silicone tubing around the clamping screws to provide fine control over the amount of pressure applied to the cable. With the catheter insertion DOF fully retracted, we pulled the stringpot's cable gently through the semi-tightened clamp assembly until the cable's resting-crimp was extended a small distance from the stringpot's cable-exit and then tightened the clamp on the cable to fix that position.

4.2.2.4.4.3 Catheter Insertion DOF, Electrical Wiring Design

We needed to route fifteen wires (eleven sensor signals and four, 1-amp motor control lines) from the main PCB in the back of the Needle Driver to the catheter grab-and-release DOFs' proximal-drive-assembly that rode on the back of the catheter insertion rod. To accomplish this, we used a 15-conductor, 1.27mm-pitch, high-flexibility ribbon-cable that bent over itself in a "U"-shape, with its stationary end mounted to the main PCB, and its moving end mounted to a breakout-PCB on the proximal-drive-assembly. As the catheter insertion DOF extended, the inflection point of the "U"-shape rolled behind the moving proximal-drive-assembly, as shown in Figure 4.13A. The main detractor of this design was that it introduced extraneous, non-insertion forces on the 6-axis force/torque sensor because the ribbon-cable's stationary end was on the mounting-side of the force sensor, and the moving end was on the tool/sensing-side of the force sensor. However, this force was repeatable such that a look-up table of extraneous force versus catheter insertion distance could be used to cancel its effect. Additionally, this extraneous force was constant while the catheter insertion DOF was stationary during needle insertion, and that is when we were most concerned with using the insertion force to identify the penetration-events. We were able to use the insertion force at the moment right before needle insertion began as an offset to cancel out this constant, extraneous force even without a look-up table.

4.2.2.4.5 Catheter Grab-and-Release DOFs

The catheter grab-and-release DOFs held the catheter rigidly during insertion and released it at the end. It was comprised of a rotating “swing-arm” that served as a location jig for the catheter’s hub and a “nitinol-loop” that actively pulled the hub against the swing-arm. Figure 4.15 demonstrates the two-step process of grabbing the catheter, and the release process simply works in reverse. We used two DOFs for grab-and-release to allow the robot to release the catheter and retract the catheter insertion DOF (and hence the swing-arm and nitinol-loop) backwards from the released catheter without snagging on and dislodging the catheter. The nitinol-loop provided the actual holding force for the catheter, but it required a location jig against which to pull the catheter to prevent bending the catheter and needle sideways. The location jig needed to match the catheter hub’s conical shape, and this shape prevented the catheter from sliding axially past the location jig. As such, we turned the location jig into the rotating swing-arm such that it could serve as a location jig when in its fixturing-position and then rotate away from the catheter’s hub to allow the catheter to slide axially past the mechanism without snagging. This design required that the swing-arm be rotated into its fixturing-position before the nitinol-loop was tightened and the reverse order of operations during release.

The swing-arm and nitinol-loop comprised the “distal-contact-mechanism” that actually contacted the catheter, and the distal-contact-mechanism was driven via two nitinol push-pull wires that were moved by the “proximal-drive-assembly”. Figures 4.16A-C show the design details of the “distal-contact-mechanism”, and Figure 4.16D details the “proximal-drive-assembly”. Recall from section 4.2.2.4.4 and Figure 4.14 that the catheter insertion rod translated the catheter for insertion and carried both the distal-contact-mechanism and proximal-drive-assembly with it. Since the distal-contact-mechanism and proximal-drive-assembly were located on opposite sides of the 9mm bronze bushing and delrin spline through which the catheter insertion rod translated, we required a method of mechanically connecting the two that could fit within the catheter insertion rod. We accomplished this by recessing Teflon tubes in top and bottom slots in the catheter insertion rod and routing nitinol push-pull wires through the Teflon tubes. We used Teflon tubing (McMaster-Carr PN: 5762K12, 0.020in ID x 1/16-in OD) for its lubricity and nitinol (Memry Corporation PN: N1SS16A, 0.016in diameter, “amber oxide” finish) for its superelasticity such that it could move cyclically through the tight, tortuous paths without failing.

4.2.2.4.5.1 Distal-Contact-Mechanism

We designed the distal-contact-mechanism to fit within the Needle Driver’s ± 10 deg wedge-shaped profile so that the needle could be rotated to within 10deg of the hand without the mechanism colliding with or digging into the hand. This shape is evident in the design of the “distal-housing”, which was the single part that supported the swing-arm and nitinol-loop.

The swing-arm’s catheter-contacting surface was conical to match the shape of the catheter hub, and it also contained a central notch to allow clearance for the evenly-spaced ridges on the catheter hub, as well as a back notch to accommodate the hub’s flange. To set the open and fixturing-positions, a dowel pin was positioned within the swing-arm’s arc-shaped recess, and the swing-arm was always pulled into one of those two contacts with the pin. The swing-arm’s nitinol push-pull wire exited its Teflon sheath behind the swing-arm’s pivot,

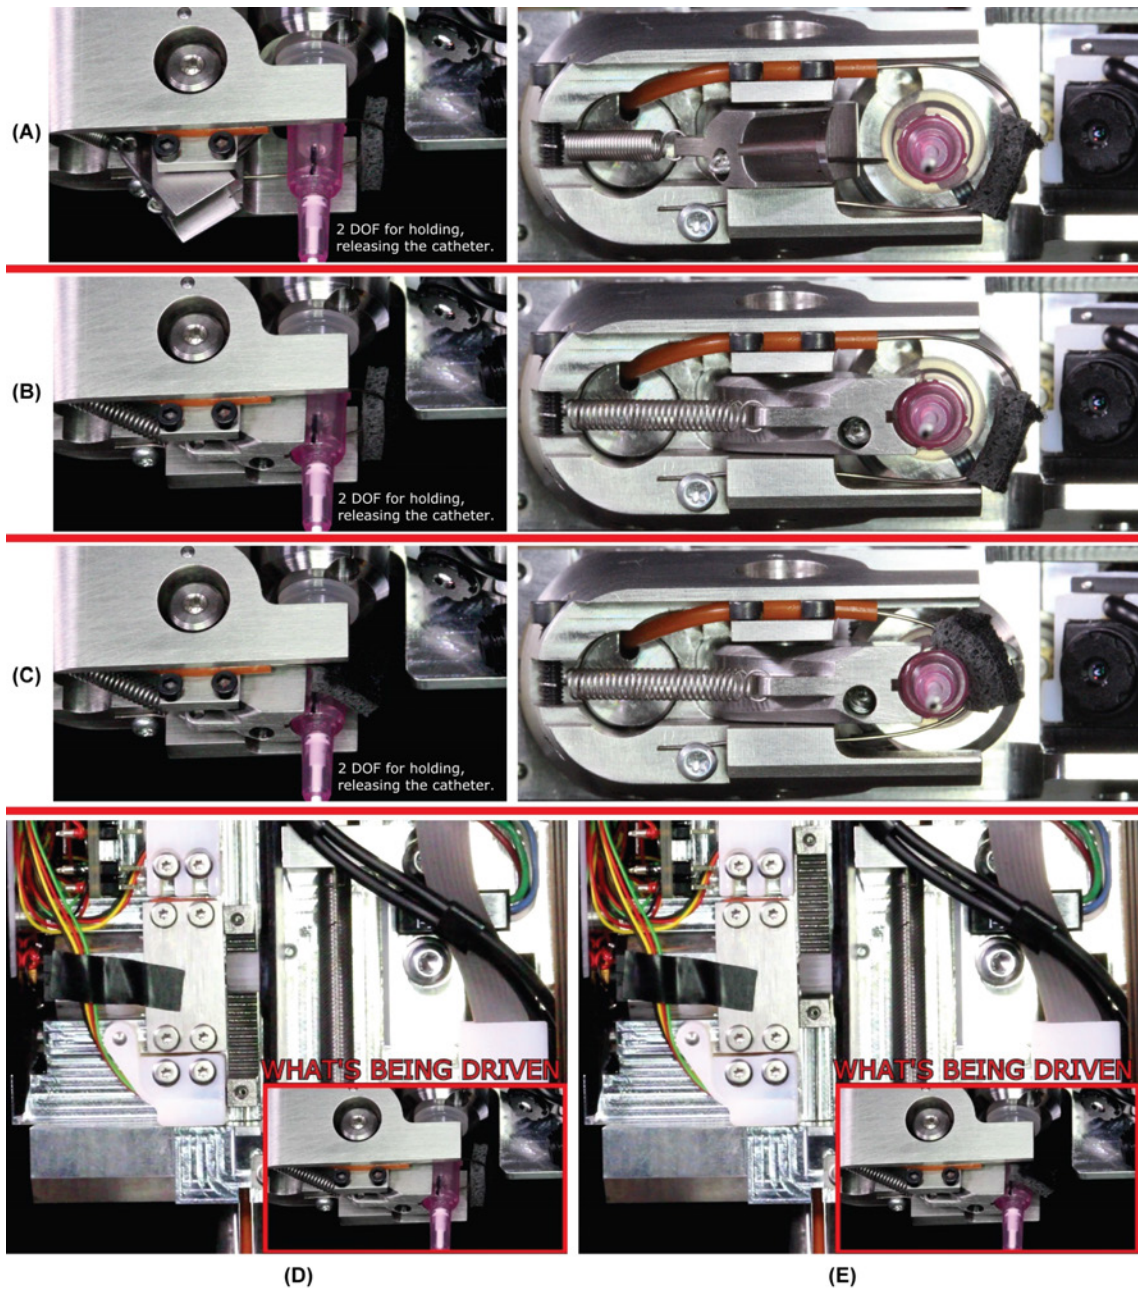


Figure 4.15: Needle Driver, Catheter Grab-and-Release DOFs, Demonstration. A: The catheter grab-and-release swing-arm and nitinol-loop are both open such that nothing is touching or restraining the catheter. B: The swing-arm closes to its fixturing-position to locate the catheter. This step must be completed before the nitinol-loop is closed. The swing-arm's conical profile and central slot match the shape of the catheter hub exactly. The crimp terminating the swing-arm's nitinol push-pull wire sticks out of the swing-arm in Subfigure A but recedes into its spring hole in Subfigure B as it compresses the inline spring. C: The nitinol-loop closes around the catheter, trapping it against the swing-arm. The textured, flexible pad conforms to the catheter's hub and provides high traction. A-C: The release process is simply the reverse of the process shown here (Subfigure C, B, A). D-E: A demonstration of the rack-and-pinion drive mechanism closing the nitinol-loop. The nitinol-loop is open until the rack moves backward, pulling on the nitinol-loop to close it.

routed around a back, cylindrical portion of the swing-arm, through a small hole, and into a spring chamber where it terminated. The wire inserted through a small compression spring (Lee Spring PN: CI012BC03M) before it was crimped such that the spring provided series, tensile compliance. This compliance allowed the nitinol-wire to be pulled a little too far past when the swing-arm was at its fixturing-position without damaging its rack-and-pinion driver in the proximal-drive-assembly. An extension spring (Lee Spring PN: EI011C03S) attached the swing-arm to the distal-housing to help open the swing-arm, thereby reducing the pushing force required from the nitinol wire that could have caused it to buckle.

The nitinol-loop trapped, or “lassoed”, the catheter against the swing-arm by pulling a piece of flexible, textured material against the unsupported side of the catheter. We used a portion of a mousepad for this material, as it was the perfect stiffness, included texture that increased traction, and had a boundary between the cloth backing and foam body through which we could insert the nitinol wire easily. This pad conformed to the unusual shape of the catheter’s hub to provide a strong, positive lock with the hub’s geometry, and it provided series compliance to protect the rack-and-pinion driver in the proximal-drive-assembly. We were careful to tune the nitinol-loop’s force against the catheter such that it was strong enough to provide sufficient traction but not so strong that it bent the catheter’s hub against the inner needle.

4.2.2.4.5.2 Proximal-Drive-Assembly

The proximal-drive-assembly moved each nitinol push-pull wire via symmetric transmissions on the assembly’s top and bottom. Each transmission consisted of a small gearmotor that rotated a 14-tooth/72-pitch delrin pinion that translated a custom, stainless steel rack gear that connected to and moved the nitinol push-pull wire. The gearmotors were not backdriveable due to their high gearing ratio (298:1), which was necessary to provide sufficient push-pull force in a compact space. However, this was not a safety concern because all other NeedleBot DOFs, including catheter insertion, were backdriveable, such that the entire catheter could be retracted or NeedleBot moved backwards and away from the hand without needing to release the catheter.

The gearmotor (Robot Marketplace PN: 0-KG12T50-KG54) contained a 3mm shaft that was D-shaped to transmit large torques to the pinion. We laser-cut the pinion (based on the 3D model of Stock Drive PN: S1G87Z-072S14) out of delrin and press-fit it onto the gearmotor shaft. The combination of the plastic pinion being weaker than the metal rack and the stress-concentration at the edges of the pinion’s D-shaped bore ensured that the pinion would be the part that broke in the event of a transmission overload. We intentionally designed this failure location because it was the only location on the transmission that could be fixed inexpensively and easily without the need for disassembly/re-calibration.

The racks were 3D-printed using Direct Metal Laser Sintering (DMLS) of 316L stainless steel (51 μ m XY resolution and 20 μ m-thick layers). Each rack sat in a rectangular slot in the assembly’s aluminum base, and the pinion trapped the rack in the slot. To decrease stiction and wear, we included shallow trenches at the corners of the slots that we filled with white lithium grease. A central, bottom through-hole in the rack was reamed for a tight slip-fit with thin-walled, stainless steel tubing (McMaster-Carr PN: 8988K49, 0.036in OD x 0.020in ID) that provided a smooth, tightly-toleranced bore for the nitinol wire. Two nylon-tipped, M2 set screws threaded vertically into the front and back of the rack and through two notches in the tubing to clamp




Figure 4.16: A-C: The catheter grab-and-release DOFs’ distal-contact-mechanism that contacts the catheter. The orange, Teflon tubes route inner nitinol push-pull wires that actuate the swing-arm and nitinol-loop. C: The recessed, inline spring provides tensile compliance for the swing-arm’s nitinol wire. The dowel pin positioned within the arc-shaped recess sets the opened and closed/fixturing-position angles. An extension spring attaches between the swing-arm and distal-housing to help open the swing-arm and reduce the risk of the nitinol buckling. D: The proximal-drive-assembly pushes/pulls the nitinol wires to actuate the swing-arm and nitinol-loop.

the nitinol wire in place against the tubing/rack, thereby providing a strong, adjustable, and non-marring attachment between the rack and nitinol wire. We provided access holes in adjacent Needle Driver parts such that we could access the set screws and adjust where on the nitinol wires the racks attached without needing to disassemble anything (with the exception of removing a single cover). This adjustment was critical for setting the correct boolean positions of the swing-arm and nitinol-loop.

Because the nitinol wires were naturally curved, their free ends that exited the racks (away from the Teflon tubes) needed to be constrained so as to prevent scratching adjacent components. To accomplish this, we included laser-cut, delrin plates at the back end of the assembly with small holes through which the free ends inserted and were terminated with the plastic insulation from stripped ribbon-cable. The swing-arm's driving rack had a throw, or distance between boolean positions, of 4.8mm, and the nitinol-loop's driving rack had a throw of 10.8mm. We left a minimum distance gap of 2.725mm between the end of the racks and the entrance of the Teflon tubes (when the racks were closest to the tubes) to ensure that any misalignments between the rack's nitinol axis and the Teflon tube's bore would not result in jamming or buckling. At the racks' maximum distance away from the entrance of the Teflon tubes, the unsupported section of nitinol was 13.525mm in length, which was short enough that the wires did not buckle when pushed on by the racks.

Because the gearmotor was very small, attaching a position encoder was impractical. It was also unnecessary because the swing-arm and nitinol-loop required boolean, not continuous, position control. To sense each rack's ends-of-travel, we included a pair of small, analog reflective sensors (Digikey PN: HOA1397-032) with the IR emitter and detector located next to each other on the same sensor face. When the rack's end reached approximately halfway between the emitter and detector, the sensor registered as "blocked". To allow for adjustable filtering and threshold-setting of what analog value (and hence exact distance) constituted "blocked", we fed the sensor's analog voltage directly to the Needle Driver's microprocessor for digital processing. We mounted each sensor in a 3D-printed, plastic housing that attached to the main assembly via screws-in-slots so that we could adjust the exact end-of-travel locations. The housings also protected the sensor's exposed, through-hole leads and wire-wrap connections from shorting on the assembly's aluminum base plate. We included a 3D-printed-plastic, opaque "light guard" in front of the entire proximal-drive-assembly to provide more constant background lighting for the reflective sensors. Without the light guard, the sensors' emitted light reflected from adjacent, shiny parts and provided spurious end-of-travel signals as the assembly translated through the Needle Driver along with the catheter insertion DOF.

As described in section 4.2.2.4.4.3, the electrical wires for the motors and sensors connected to the Needle Driver's main, microprocessor-PCB via a 15-wire ribbon-cable that moved with the catheter insertion DOF. This ribbon-cable attached to a custom breakout-PCB on the proximal-drive-assembly that routed the ribbon-cable's signals to square, vertical header pins. We used wire-wrap (30-gauge AWG) to connect each of these headers to the individual motor and sensor leads, as shown in Figures 4.13B-C, because it was thin enough to fit the large number of wires into a small space and maintained its shape after installation so that the wires would not snag on nearby structures during translation of the catheter insertion DOF.




Figure 4.17: Needle Driver, Force/Torque Sensor. The roll and catheter-control DOFs are cantilevered from the sensor.

4.2.2.4.6 Force/Torque Sensor

We proved in Chapter 2 that non-axial force/torque axes are unnecessary for identifying the penetration-events. However, these other axes can be useful for detecting vein rolling, hand movement, collisions between the Needle Driver and hand/AssistoBot, and the success of valve drilling (by examining the resistance torque). As such, we included an ATI-Mini40 6-axis force/torque sensor that gave us the four axes we wanted (F_x , F_y , F_z , and T_z /Troll) with two additional torque axes. The manufacturer customized the sensor to meet our design needs by providing 1. an off-center, axial exit for the signal cable that directed the cable out the back of the Needle Driver with minimal bending and 2. a central, 6.35mm-diameter hole that allowed the targeting-laser to pass. Our sensor used the SI-20-1 calibration, which provided a 20N sensing-range for F_x , F_y (0.005N resolution), 60N for F_z (0.010N resolution), and 1NM for T_x , T_y , and T_z (0.000125 NM resolution).

As shown in Figure 4.17, we cantilevered the main Needle Driver base plate with the roll and catheter-control DOFs from the tool/sensing-side of the force sensor. The ATI-Mini40 was stiff enough to support the cantilevered mass without introducing any bandwidth limitations that were noticeable in the insertion force profiles. This placement allowed us to 1. use a larger, stronger sensor than could have fit nearer the needle and 2. measure roll torque on the needle as it drilled infinitely. The detraction was that we were also measuring the self-weight of the structure that changed with different values of pitch angle and catheter insertion (because the catheter insertion DOF carried the heavy catheter grab-and-release DOFs with it). This required calibration of the force sensor to remove the gravity component as a function of pitch and catheter insertion distance. The only parts that broke the mechanical isolation of the force sensor's mounting and tool/sensing-sides were the electrical wires from the cantilevered motors and sensors. To mitigate their effect on the force measurements, we strain-relieved the wires on both sides of the force sensor to provide a constant force/torque offset that was removed during the calibration process. We included 3D-printed covers on as many sides of the force sensor as possible to prevent debris from lodging in the sensor's thin midline-gap that separated the mounting and tool/sensing-sides, which would have caused inaccurate readings.

The analog force signals experienced moderate noise (despite their being differential signals), so we filtered them digitally using a 2nd-order Bessel Filter. We picked a Bessel Filter over other potential low-pass

filters because the Bessel filter preserves the shape of a waveform the best. We designed the filter in MATLAB for $\omega_0 = 30\text{rad/s}$, or $f = 4.775\text{Hz}$, such that frequencies up to ω_0 experienced a nearly-constant group delay $D = 1/\omega_0 = 33.3\text{ms}$. For all of our automatic insertions, we used a constant axial insertion speed $I_s = 10\text{mm/s}$, so the delay caused the robot to insert the needle an extra $I_s \cdot D = 0.333\text{mm}$ beyond when the robot should have detected the force “pop” and halted, which is minimal compared to the vein’s internal clearance for the IV needle at a low pitch angle. Had the needle been inserted at the fastest of the penetration-event insertion speeds that we measured in Chapter 2 ($3.029 \pm 1.099\text{mm/s}$ at the vein catheter “pop”), the needle would have moved only an additional $0.100 \pm 0.036\text{mm}$ due to the detection delay.

4.2.2.4.7 Accelerometer

Our discovery in Chapter 2 that an accelerometer is unreliable for detecting penetration-events during insertions performed by human practitioners made us question its inclusion in HaemoBot. However, we included it on the robot to compare and contrast its detection abilities for robotic insertions. We found that the detection rate was worse (essentially zero) for robotic insertions than for human insertions (20-50%, depending on which “pop”), likely because the robot had higher stiffness and mass than the instrumented IV needle such that the penetration force caused undetectably-low acceleration of the robot.

We did, however, find that the accelerometer was quite useful as a redundant, absolute encoder for the pitch DOF, which added extra safety, and also for leveling the Z-axis and robot overall with respect to an uneven floor. Because the stationary accelerometer displayed the Needle Driver’s gravity vector, we could adjust the Z-axis/robot-feet until the accelerometer’s gravity vector converged on fully-vertical. We also found that the accelerometer was quite good at detecting even minor collisions with the patient’s hand or other parts of the robot, such as AssistoBot, and in the future we plan to use this ability to issue “halt-insertion” commands that pause the insertion while the human operator investigates the cause of the collision.

We bolted the 3-axis, digital accelerometer (Sparkfun PN: SEN-09836) to the underside of the Needle Driver’s base plate, as shown in red in Figure 4.13C. This placement positioned the accelerometer as close to the IV needle as possible, with as few compliant parts in between the accelerometer and the needle as possible. To prevent electrical shorts between the accelerometer and aluminum base plate, we epoxied the accelerometer’s PCB to a non-conductive, 3D-printed-plastic mounting plate that was then bolted to the Needle Driver’s base plate. The accelerometer communicated with the Needle Driver’s microprocessor via I2C serial communication and streamed data from all three axes at 400Hz, 10-bit resolution, and $\pm 2\text{G}$ range.

4.2.2.4.8 Blood-Flash and First-Person-Shooter (FPS) Cameras

We included two small cameras (approximately $9 \times 9 \times 9\text{mm}$ cubes) on the Needle Driver: the First-Person-Shooter (FPS) camera which looked at the needle’s tip and the blood-flash camera which looked at the needle’s flash-chamber for blood-flash. Figures 4.18A-B show the placement of these cameras, Figures 4.18C-D show feeds from the FPS camera, and Figures 4.18E-F show an example blood-flash as viewed by the blood-flash camera. We mounted each camera with a single screw that allowed us to rotate the camera about its vertical axis with respect to the needle, and this allowed us to center the desired needle geometry in each camera’s view even for differently sized needles. The cameras did not have an integral means of




Figure 4.18: A-B: There are two small cameras (9x9x9mm cubes) on the Needle Driver. The First-Person-Shooter (FPS) camera points at the needle's tip, and the blood-flash camera points at the needle's flash-chamber. C-D: The FPS camera helps to align the needle/vein and shows the skin needle "pop" as the bevel disappears beneath the skin. D: The FPS camera can also detect needle bending, as shown here. E-F: The blood-flash camera shows the blood-flash to confirm success or failure of the insertion. Although these images are black-and-white due to a recording error, we used a color camera that shows the blood as red, making it much easier to detect the blood-flash. E: Blood-flash has not yet occurred. F: Blood-flash has occurred and is visible. The blood-flash shown here is from semi-transparent, artificial blood that is much harder to see in the needle than real blood.

attaching them to the Needle Driver, so we designed a 3D-printed housing that trapped the camera against a delrin mounting plate that could be bolted to the Needle Driver's base plate via a single M2 screw. Due to the small size of all parts, we bolted the 3D-printed housing to the delrin mounting plate via four M0.6 x 5mm screws (US Micro Screw PN: M06-50-M-SS-P), and the combination of the delrin's tight laser-cutting tolerances and lubricity enabled its successful hole-tapping. The housing also included wire-redirection and strain-relief features that routed the wires away from their origin at the back of the camera to exit the housing at its top so that the wires did not interfere with the nearby roll and catheter insertion motors.

The blood-flash camera observed the blood-flash to serve as ground-truth confirmation of successful insertion. When the robot detected the penetration-events via the faster axial insertion force signal, it would pause insertion until the blood-flash confirmed the penetration-events. We mounted a color camera (SuperCircuits PN: PC208) to the side of the IV needle such that it was pointed at the needle's flash-chamber. Figures 4.18E-F show an example of this blood-flash detection. Although we could have used a non-camera sensor to detect the blood-flash optically (such as a slot-sensor or single-pixel, RGB color sensor), we chose a video camera because its view could be fed to a human teleoperator for normal, visual interpretation (non-algorithmic/non-numerical) or to a computer-vision algorithm for automatic blood-flash detection. Tuning the camera's flash-detection algorithm to be immune to variations in blood color would also be easier than for a single-pixel sensor because the camera provides more information through its many pixels.


We included the FPS, NIR-compatible camera (SuperCircuits PN: PC206XP) to show the vein/needle tip interaction. The FPS view offered some critical benefits, including:

1. Helping the practitioner to aim straight at the vein during teleoperation,
2. Allowing the needle to be aligned "bevel-up" along the surface normal,
3. Allowing detection of needle bending, as shown in Figure 4.18D,
4. Allowing detection of some of the penetration-events (particularly the skin needle "pop"), and
5. Acting as a secondary blood-flash detector by observing blood-flash in the catheter's hub during catheter insertion and needle withdrawal.


For visible light, the FPS camera was extremely helpful. However, its vein/non-vein contrast was low when used with the NIR LED grid, perhaps due to poor NIR-spectral-sensitivity for our, specific wavelength. In the future, we plan to switch to another camera that yields higher vein/non-vein contrast. Figure 4.18C shows the FPS camera's view of a straight needle, and Figure 4.18D shows the needle bending.

4.2.2.4.9 Needle Length Adjustment Mechanism

For the RCM to work properly, the tip of the IV needle must be placed precisely at the RCM's origin. Relying on machining tolerances alone to position the tip would be inadequate because the tolerances accumulate quickly in the serial chain of parts connecting the Needle Driver's bottom-threaded-pulley to the IV needle. As such, we needed a precise, easy method of adjusting the Needle Driver's length to place the needle tip at the RCM's origin. Because the length of our IV needles was quite consistent, we only needed to perform this calibration once. However, the calibration could be repeated to accommodate a different IV model that was shorter or longer than the one for which we originally designed.




(A)



(B)

GREEN parts connected rigidly. **RED** parts connected rigidly.
Needle length is set by sliding **RED** against grounded **GREEN**.

Repelling, Adjustable Neodymium Magnets



(C)

Figure 4.19: Needle Driver, Needle Length Adjustment Mechanism.

As shown in Figure 4.19, our adjustment mechanism split the plate that connected the main body of the Needle Driver to the bottom-threaded-pulley into two parts that could translate with respect to each other. We refer to the stationary plate that attached to the pulley as the “grounded-plate” and the plate that moved with the rest of the Needle Driver as the “sliding-plate”. A preloaded, reciprocating-ball linear guide connected the two plates rigidly such that their separation in the direction of the needle’s axis could be adjusted while maintaining precise, relative pose in all other dimensions. The high stiffness of the linear guide and plates prevented any detectable deflection in other directions, even under high loading.

The separation distance between the plates was set by the magnetically spring-loaded contact between the adjustment screw on the grounded-plate and the stop-block on the sliding-plate. The Neodymium (NdFeB) magnets on the grounded-plate and sliding-plate tried to separate the plates until the adjustment screw contacted the stop-block, so the separation distance could be adjusted by turning the adjustment screw. We used an M3 x 0.25mm fine-pitch, ball-tip adjustment screw with a delrin lock-nut for vibration resistance. The adjustment screw was turned by inserting an allen-key through access holes that extended to the back end of the Needle Driver after loosening the lock-nut with two, custom wrenches. The directionality of the adjustment mechanism was such that the needle insertion force pushed the stop-block further into the adjustment screw (they were already in contact by default at the beginning of an insertion). As such, the Needle Length Adjustment Mechanism did not introduce any compliance into the insertion process. The magnetic spring-loading ensured that there was always contact between the stop-block and attachment screw, except for when the needle was pulled away/out of the needle-holder with a force that exceeded the magnetic spring-loading. The frictional force during needle withdrawal from the catheter was much smaller than the magnetic spring-loading, so there was no normal situation during the IV insertion wherein the attachment screw/stop-block contact was broken and the needle moved away from the RCM origin. The magnet on the sliding-plate was secured by two nylon-tipped set screws such that the magnet’s separation from the opposing, fixed magnet could be adjusted to account for a change in needle length or a change in magnetic spring-loading force.

4.2.2.4.10 Electronics and Communication

We designed three custom, 4-layer PCBs for the Needle Driver: a breakout-PCB that routed the roll and catheter insertion motor/encoder FFCs to the ribbon-cable that exited the Needle Driver, the breakout-PCB that routed signals from the moving ribbon-cable to the motors and sensors of the catheter grab-and-release proximal-drive-assembly, and the main/microprocessor-PCB (located at the back of the Needle Driver) that performed all of the sensor processing and motor control (with the exception of the roll and catheter insertion motors). The main-PCB housed a “Baby Orangutan B-328” microprocessor/motor-controller board (Pololu PN: 1220) that provided serial communication, motor control circuitry, and analog/digital I/O lines for our sensors. We designed the additional circuitry that controlled the laser-targeting and powered, conditioned, and routed the various sensors. Communication between the microprocessor and the robot’s main computer program was achieved via a bidirectional, asynchronous USB-to-serial stream and a custom communication protocol that we developed. We clocked the communication speed at 500Hz in each direction with a delay of approximately a few milliseconds. A “deadman timer” measured the time between received serial packets




Figure 4.20: AssistoBot's Main Subassemblies. A, B, C: The patient's arm inserts through the "AssistoBot-Arm-Box-Assembly", and the blood-cuff restrains/tourniquets it. A, C: The "AssistoBot-Hand-Plate-Assembly" restrains the hand, and the "AssistoBot-Arm" sits above and applies forces, alcohol, and compressed air to the hand. C-D: The IV line inserts with the arm into the robot to prevent tangling. The inserted catheter and attached IV line are taped to the hand/arm and withdrawn from the robot with the arm.

and stopped all motion in the event of silence longer than 30mS. We encased the main-PCB in a 3D-printed cover to protect it and its wires while leaving a transparent window for viewing the onboard debugging LEDs.

4.2.3 AssistoBot

AssistoBot performed all of the non-needle-insertion functions and served as the non-dominant hand when cooperating with NeedleBot during two-handed procedures (such as clamping the vein/catheter during needle withdrawal). It was comprised of three main subassemblies, including:

- The "AssistoBot-Arm-Box-Assembly" that restrained the arm and applied a tourniquet,
- The "AssistoBot-Hand-Plate-Assembly" that restrained the patient's hand over the LED grid, and
- The "AssistoBot-Arm" that applied forces, alcohol, and compressed air to the hand.

Figures 4.20A/C show all three subassemblies simultaneously, with and without a patient's arm inserted.

4.2.3.1 “AssistoBot-Arm-Box-Assembly”: Arm Restraint and Tourniquet

The “AssistoBot-Arm-Box-Assembly” contained a hole through which the patient’s arm was inserted, as shown in Figure 4.20C. This hole was lined with an inflatable blood-cuff that served both as an adjustable-pressure tourniquet and as an arm restraint to prevent the patient’s arm from being injured by premature retraction. We modified a blood-pressure monitoring device by adding a 10-bit pressure sensor and attaching our own electronics and code to control the the blood-cuff’s inflation pump and air release valve. The blood-cuff attached to the inside of the arm hole via industrial Velcro and could be removed for cleaning. Figures 4.20A-B show the difference in diameter between its fully deflated and inflated states, from which it is clear that the cuff could accommodate a wide range of arm sizes.

To allow the patient’s arm to be retracted from the arm hole with a catheter that was attached to an IV line/bag, the IV line needed to route through the arm hole as well. Otherwise, the line would become tied to the robot through the arm hole. We found that initially taping the IV line to the patient’s forearm so that the IV line was inserted through the arm hole with the patient’s arm provided for easy, reliable line-handling. The patient was able to feel if any part of the line became stuck on the structure and could make adjustments to untangle the line. Figure 4.20D shows how the IV line was taped to the patient’s arm before the procedure began and inserted along with the arm into the robot. Figure 4.20C shows how a successfully inserted catheter and attached IV line would be taped to the patient’s hand and arm for retraction.

4.2.3.2 “AssistoBot-Hand-Plate-Assembly”: Hand Restraint

The “AssistoBot-Hand-Plate-Assembly” performed many functions, including:

- Placing the hand in a known location that was reachable by both NeedleBot and AssistoBot,
- Supporting the hand over the LED grid via the “Hand-Plate”,
- Restraining the fingers to prevent the hand from flailing and becoming injured in the event of panic,
- Collecting any blood-contaminated fluids generated during insertion, and
- Bending the fingers down and away from the Needle Driver so as to
 - Apply gentle, natural traction to the hand’s dorsal skin through the finger-bending process,
 - Prevent collisions with the Needle Driver, and
 - Promote higher-quality vein imaging through improved contact with the hand.

4.2.3.2.1 “Hand-Plate”

There were many design elements that drove the geometry of the “Hand-Plate”. The main, top portion needed to be flat so that 1. the hand could be imaged properly by the LED grid below and 2. the hand would be parallel to the robot’s base. The distal end of the plate needed to slope downwards to help bend the fingers for the reasons listed above. We included a continuous ridge around the plate to prevent any fluids from escaping and a valley below the fingers for collecting any fluids that rolled down the slope. This ridge also served as the press-fit attachment mechanism that kept the “Hand-Plate” on the rest of the assembly. We sized the plate so that even a large hand could be stretched completely flat against the plate with the fingers together. If the plate was too small, the hand would be forced into a cup-shape with a concave palm, thereby producing low-quality NIR vein-images (likely due to light being lost to reflections between the plate and palm). For




Figure 4.21: “AssistoBot-Hand-Plate-Assembly”. The fingers are restrained using only the weight of the restraint-mass.

the best image quality, the palm needed to be in full contact with the plate. The “Hand-Plate” also needed to be made of an optically-clear material to allow the NIR transillumination to work through it. With such a complicated geometry and the need for an optically-clear material, we were forced to vacuum-form the plate out of plastic. After experimenting with different thicknesses and materials, we determined that 1/8-inch-thick PETG provided the best combination of formability, ductility, stiffness, durability and optical-quality (minimal bubbles and other impurities). Sterility of the “Hand-Plate” could be achieved by 1. changing the “Hand-Plate” in between patients, 2. including a thin, optically-clear slip-cover over the main, structural “Hand-Plate” that would be changed in between patients, or 3. having the robot sterilize a fixed “Hand-Plate” using the alcohol sprayer in the head of the “AssistoBot-Arm”.

4.2.3.2.2 Finger Restraint Mechanism

The finger restraint mechanism consisted of a weighted lever which rested on the fingers and could be lifted off by a motor that connected to the weight via a cable. We intentionally designed the mechanism so that the clamping force was due entirely to gravity and the motor could pull the weight off of the hand but not push on and potentially hurt it. In the event of motor failure, the human operator could easily lift the restraint

manually. Figures 4.21A-B show the finger restraint mechanism in action, and Figures 4.21C-D show the mechanism details. The mechanism consisted of three main parts:

- The main, “check-mark”-shaped lever which rotated about its end closest to the fingers,
- The restraint-mass at the wrist-end of the lever which provided the main clamping force, and
- The rotating finger-clamp-bar that attached to the main lever in its middle section and transmitted the weight of the restraint-mass onto the fingers.

The restraint-mass consisted of a heavy piece of stainless steel with attachment points for the mechanical lifting-cable and main lever. We placed the mass at the end of the lever to maximize its mechanical advantage.

We used a Dynamixel AX-12+ Servo Motor to lift the restraint-mass via a cable that routed over a spinning redirect shaft near the top of the assembly. To prevent excessive loads on the motor shaft’s front bearing, we connected the front and back ends of the motor shaft via a stainless steel standoff that bolted between a custom back-plate (laser-cut, 1/8-inch stainless steel) and a custom servo-head (sandwiched layers of laser-cut, 1/4-inch delrin and 1/16-inch stainless steel). Our custom servo-head provided flanges that kept the cable on the pulley during moments of cable-slack and an adjustable cable-termination that allowed us to set the rest/lift heights of the restraint-mass mechanically without changing the servo-head’s commanded positions.

The finger-clamp-bar connected to the main lever via pillow blocks so that it rotated passively to maximize contact with the fingers, regardless of size. It was composed of a top plate (1/16-inch, laser-cut, low-carbon steel) that was coated on the finger-side with 1/4-inch foam-padding for comfort. We used closed-cell foam so that it did not trap fluids in the foam’s cell structure. In the future, we envision having the alcohol sprayer in the head of the “AssistoBot-Arm” automatically spray the patient’s hand as it is inserted into the robot to sterilize the fingers before contacting the finger-clamp-bar. We could also add disposable, adhesive liners that cover the foam or an automatic chemical-drip system that could sterilize the foam directly.

4.2.3.3 “AssistoBot-Arm”: Applying Forces, Alcohol, and Air

The “AssistoBot-Arm” performed many functions, including:

- Spraying isopropyl alcohol (IPA) onto the insertion site to clean it,
- Rubbing the IPA into the hand to facilitate the cleaning process and vasodilate the veins,
- Spraying the insertion site with compressed air to dry it of residual IPA before insertion,
- Applying traction to the hand skin during insertion of the needle, and
- Clamping the catheter/vein to prevent catheter dislodgement or bleeding during needle withdrawal.

These tasks required 3-DOF, XYZ motion, a large dexterous workspace, a minimal structural footprint to prevent collisions with NeedleBot, and the ability to sense interaction-forces with the hand and catheter.

4.2.3.3.1 Kinematics and Drive Details

We used RRP-kinematics to achieve XYZ motion, with a planar RR-mechanism that provided XY motion riding on a vertical, prismatic Z-axis. The XYZ workspace was 117mm x 108mm x 96mm and allowed us to 1. reach any reasonable insertion site in the hand or wrist and to 2. raise the “AssistoBot-Arm” high above the “Hand-Plate” to prevent collisions with the hand while it was being inserted into the robot.

4.2.3.3.1.1 Planar XY Mechanism

As shown in Figures 4.22B and 4.23, we used a parallel mechanism for the planar, XY motion to increase in-plane stiffness and attached the two, parallel arms from opposite, vertical sides of their common-link to increase out-of-plane stiffness. The construction of the linkages was designed to maximize their bending-stiffness to weight ratio while providing easy fabrication. Each link was a “sandwich” of inner, laser-cut 1/4-in delrin plate(s) and outer, laser-cut 1/16-in steel plates. The layers were aligned with dowel pins and clamped together rigidly by screw-and-nut assemblies. The common-link that connected both of the parallel arms contained top and bottom “sandwiches” connected by an aluminum offset.

Both the top and bottom links contained an integral cable-capstan on its proximal end, and we connected the capstan via a 1:17 cable-reduction directly to a stepper motor’s 5mm shaft to minimize the number of parts. These particular stepper motors (Pololu PN: 1209) had shafts that were supported front-and-back by large-diameter, thick bearings that could take the radial loads of the cable-reduction without damaging the bearings. This reduction was strong enough to produce the desired end-effector force while keeping the rotational axes easily backdriveable in the event of power loss. When the steppers were driven via full-stepping for full torque and speed, their cable-magnified pull-out torque exceeded what could be backdriven by hand, which allowed us to drive the steppers as an open-loop position-source that did not require encoders. Our pinion-to-cable-diameter ratio was 19.7:1 (using Sava Cable PN: 2010), which was below the manufacturer’s recommended minimum ratio of 40:1 for maximizing fatigue life, but we experienced no fatigue or strength-degrading issues. We used the linkages’ “sandwich”-construction to incorporate integral cable-routing features, including outer flanges and depth-milled cable-paths in the inner Delrin layers that provided for clean, easy termination and tensioning of the cable. The cable terminated on a non-capstan portion of the linkage, routed through the top, depth-milled cable-path to the capstan portion, wrapped around the motor 6 times (to prevent slippage and to set the proper helix-angle of the cable), routed through the bottom, depth-milled cable-path, and finally terminated on a swing-arm that tensioned the cable by rotating and pulling on it.

We homed each capstan using the same slot-sensor/stepped-flag design that was used on the pitch and yaw DOFs. We could not fit end-of-travel sensors, but we included end-stops to limit the capstans’ rotation.

4.2.3.3.1.2 Final-Link Parallelism Mechanism

We hid the main structure of the “AssistoBot-Arm” inside of the “AssistoBot-Arm-Box-Assembly” structure to save space and prevent collisions with NeedleBot. We used a long, final-link to bridge most of the distance between the main structure of the “AssistoBot-Arm” and NeedleBot, and this final-link remained directly behind the head of “AssistoBot-Arm” and aligned with the patient’s arm-axis (robot’s Y-axis) at all times to prevent collisions. That is, the final link translated in XYZ but never changed its orientation. Figures 4.22B and 4.23B show the passive, timing-belt mechanism that achieved this parallelism constraint, and Figure 4.22C demonstrates its operation over the entire workspace.

The top and bottom capstan links rotated about a fixed, central shaft, to which we also attached our “ground” timing belt pulley halfway between the two capstans. We positioned the timing belts close to the XY mechanism’s center line (in between the top and bottom capstan links) to prevent asymmetrical belt forces

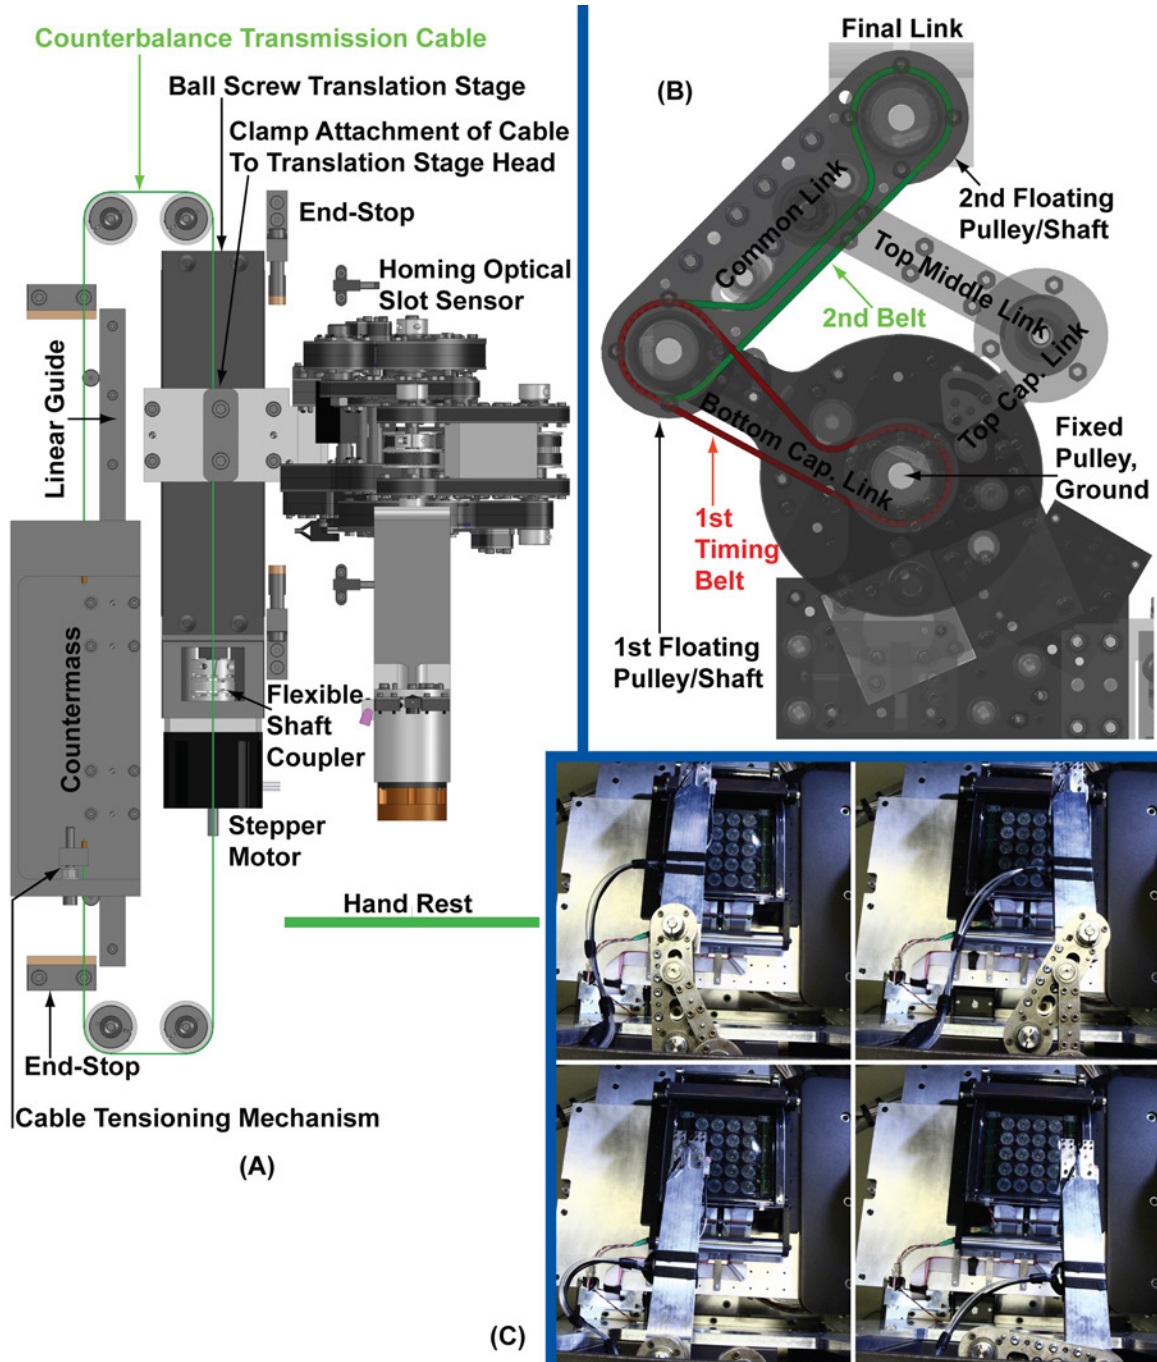


Figure 4.22: “AssistoBot-Arm”: Design Details of the Z-Axis and Final-Link Parallelism Mechanism. A: Z-axis details. B: Final-link parallelism mechanism, links and belts configuration. C: Demonstration of final-link parallelism mechanism. The final-link remains parallel to the patient’s arm and maximally distanced from the Needle Driver over the entire workspace. The special arrangement of belts and pulleys in Subfigure B is what enforces this parallelism.





Figure 4.23: "AssistoBot-Arm": Design Details of the Planar XY Mechanism.

from bending the arm out-of-plane. We connected the orientation of the “ground” pulley, or the “ground-angle”, to that of the final-link via two timing belts (Misumi PN: HTUN318S3M-60). The first timing belt transmitted the “ground-angle” to the first-floating-pulley/shaft that floated on the shaft that connected the bottom capstan link to the proximal-end of the common-link. The second timing belt transmitted the “ground-angle” from the first-floating-pulley/shaft to a second-floating-pulley/shaft at the distal-end of the common-link. This second-floating-pulley/shaft was connected rigidly to the final-link so that its angle was transmitted to the final-link, completing the transmission of the “ground-angle” to the final-link. The final-link attached to this shaft via a clamp such that we could adjust the angle of the final-link relative to it and then clamp it in place to hold that alignment. To set the final-link parallel to the robot’s Y-axis, we used three hole-and-slot dowel pin alignment mechanisms to pin the capstan links and the final-link in their zero-configurations while we tightened the clamp on the final-link. We used flat-belt idlers on movable-posts to tension the belts by sliding the idlers into the flat backside of the belts. Obtaining high tension on the belts was critical to providing an accurate, stiff parallelism mechanism.

4.2.3.3.1.3 Z-Axis Mechanism

As shown in Figure 4.22A, the Z-axis motion was provided by a ball-screw-driven translation stage (Misumi LX3010-series) that was connected via flexible shaft coupler to a stepper motor (Pololu PN:1209). The stepper motor’s detent-torque increased the stage’s backdriving force, but the stage was still easily backdriveable when unpowered. We homed the axis via an optical slot sensor at the top end-of-travel.

Because the weight of the planar XY mechanism that rode on the stage exceeded the stage’s backdriving-force, we included a counterbalancing mechanism. The counterbalance consisted of a stainless steel counter-mass that was constrained to move along a reciprocating-ball linear guide and was driven by a closed-loop mechanical cable (Sava Cable PN: 1024) that also connected to the translation stage head via an adjustable clamp. We machined the cable-tensioning system as an integral part of the counter-mass such that one crimped end of the cable routed through and pulled against the top ridge of the counter-mass, and the other crimped end routed through and was screw-tensioned against the bottom ridge.

4.2.3.3.2 Foam-Tip

We placed FDA-compliant foam at the tip of the “AssistoBot-Arm” as the only part of the arm that would contact the hand or catheter, as shown in Figures 4.24A/C. We used a pliable material to prevent damage to the patient or catheter and to maintain contact between the flat foam-tip and contoured hand/catheter. We chose a moderately-stiff, Polyethylene foam (McMaster-Carr PN: 8722K27, load of 5psi causing 25% deflection) so that we had enough compliance to maintain safe, reliable contact without deflecting so much as to prove ineffective at clamping the vein/catheter. We chose a closed-cell foam for sterility so that fluids would not soak into the foam’s cell-structure. We further envision using a disposable slip-cover that would be replaced in between patients or having the robot spray IPA onto the “Hand-Plate” and then scrubbing it with the foam-tip as a form of self-cleaning.

4.2.3.3.3 Force Sensor

We used an ATI-Mini40 6-axis force/torque sensor to measure and limit the interaction XYZ forces between

the foam-tip and the hand/catheter during several phases of insertion, including IPA-scrubbing, traction-pulling, and vein/catheter clamping. We used the Z-axis torque measurement during automated insertion by having AssistoBot remain clamped on the catheter/vein indefinitely until the human operator twisted the foam-tip about its Z-axis, thereby signaling that the IV line had been attached and that clamping could be relaxed. Since the force sensor came into close proximity to the hand and Needle Driver, we recessed it completely inside of a protective aluminum casing so that it could only be touched through the foam-tip attached to its bottom surface, as shown in Figure 4.24A. This casing saved the sensor during multiple collisions between AssistoBot and NeedleBot during debugging of the robot's control-code.

4.2.3.3.4 Air and IPA Sprayers

The head of the "AssistoBot-Arm" included nozzles for spraying isopropyl alcohol and compressed air onto the hand for pre-insertion sterilization and drying, as shown in Figures 4.24A/C. Figures 4.24E-F show the electromechanical spraying mechanisms for the compressed air and IPA, respectively. The reservoirs and electromechanical controls were located within the "AssistoBot-Arm-Box-Assembly" and were accessed by swinging the sheet-metal enclosure up and away from the assembly, as shown in Figure 4.24D. The enclosure covered all but the NeedleBot-facing side of the structure and was attached via two plastic, living hinges at the top, NeedleBot-facing edge so that it could swing upwards to reveal the spray assemblies. A laser-cut, stainless steel prop underneath the enclosure could be rotated upward to prop the enclosure open during maintenance or refilling operations inside of the assembly. When the enclosure was in its operational, down position, two knobs at the bottom of the assembly clamped it in place against the foam-lined edges of the structure to prevent rattling during an insertion.

4.2.3.3.4.1 IPA Sprayer

The isopropyl alcohol (IPA) sprayer consisted of a motorized squirt bottle that connected via thin tubing to a nozzle that rode on the head of the "AssistoBot-Arm". We motorized a squirt bottle because it was the simplest design and allowed for single sprays based on boolean position changes (pushing the lever to one side, then pulling it back to the other side to reset), instead of continuous, unidirectional motion like in a pump motor. In this manner, a servo motor that became stuck in a position/velocity would simply force the bottle's lever to one side without additional spraying. In contrast to this, a pump could become stuck at a position/velocity that continued to spray until the reservoir depleted, potentially soaking parts of the robot.

We used a Dynamixel AX-12+ Servo Motor to pull on a mechanical cable that routed around a rotating redirect-shaft and attached to the end of the squirt bottle's replaced lever. We replaced the original lever on the squirt bottle with a laser-cut 1/4-in plywood and 1/16-in stainless steel assembly that allowed us to attach the cable to the lever in an easy, adjustable manner. The squirt bottle did not need to be removed for refilling because the head could be unscrewed in-place and IPA poured into the reservoir easily without disassembling any other part of the mechanism. We removed the nozzle from the spray bottle and replaced it with a custom-machined, delrin plug with an outer diameter that press-fit into the squirt bottle's housing and an inner diameter that fit the tubing that routed to the final nozzle on the head of the "AssistoBot-Arm". We used thin, flexible tubing so that it could ride along the links of "AssistoBot-Arm" without kinking or

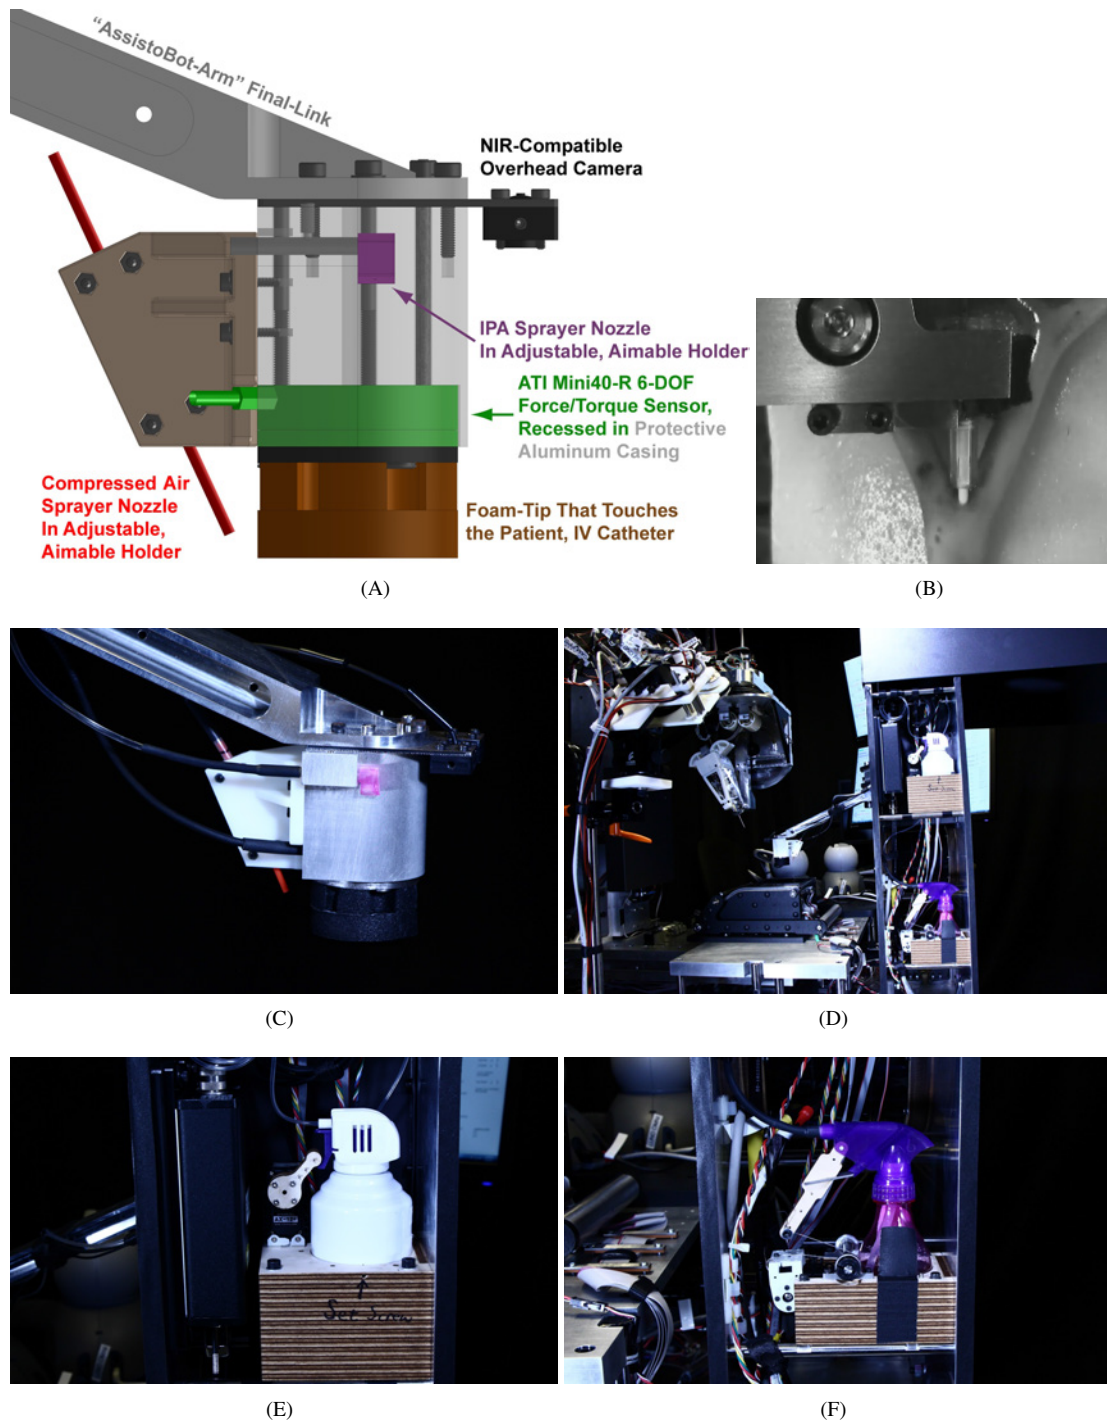


Figure 4.24: “AssistoBot-Arm”, Head Design Details. A, C: The head of the “AssistoBot-Arm” contains spraying nozzles for isopropyl alcohol and compressed air, as well as an overhead camera, recessed 6-DOF force/torque sensor, and foam-tip that contacts the hand/catheter. B: The overhead, NIR camera lets us watch the insertion from above and provides an estimate of catheter insertion depth beneath the skin. D: The alcohol and air spraying mechanisms are housed inside of the “AssistoBot-Arm-Box-Assembly” and can be accessed by swinging the hinged, sheet-metal enclosure upwards. E: Air spraying mechanism. F: Alcohol spraying mechanism.

providing a strong spring-force. We sprayed 99.8%-pure IPA (Techspray PN: 1610-G1) instead of the typical, lower-concentration IPA sprays or gels that are used clinically because we wanted to minimize water-content that could short electronics, minimize residue on the “Hand-Plate” that would distort the NIR vein images, and provide lower-viscosity for easier, wider-coverage spraying.

4.2.3.3.4.2 Air Sprayer

The air sprayer consisted of a motorized canister of electronics duster that connected via tubing to a nozzle that rode on the head of the “AssistoBot-Arm”. We motorized a canister of compressed air instead of pumping and pressurizing the air ourselves because it was the simplest, safest way of ensuring moisture-free air. Although the convenience of a self-generating, nondisposable air supply was appealing, the risks associated with the pressurized/pumped air were too great. Each disposable canister could be used for hundreds of insertions before needing to be replaced, and changing the canister was quick and easy because it connected to its housing via a single set screw and to the tubing via light press-fit. We chose the specific electronics duster (Techspray PN: 1671-10S) because it was completely anhydrous (100% HFC-134a), safe to use on powered electronics, nonflammable, safe on plastics, and left no residue.

We used a Dynamixel AX-12+ Servo Motor to rotate a custom swing-arm into the canister’s trigger. Because the swing-arm’s angle with respect to the trigger changed as the trigger was depressed, we made the tip of the swing-arm circular to provide smooth sliding between the two. We needed thin, flexible tubing to route from the canister inside of the “AssistoBot-Arm-Box-Assembly”, along the moving links of the “AssistoBot-Arm”, to its termination on the head. Because the air stream was quite narrow, we needed a method of aiming it at the insertion site. To achieve this, we developed a 3D-printed, adjustable clamp that held the short, rigid tube that came with the electronics duster by default (shown in red in Figures 4.24A/C), such that it aimed the air stream 20mm directly beneath the center of the foam-tip.

4.2.3.3.5 Overhead NIR Camera

We included a small camera on the head of the “AssistoBot-Arm” that provided an up-close, overhead view of the needle insertion. We used the same NIR-compatible camera on AssistoBot as we used for the Needle Driver’s FPS camera, described in section 4.2.2.4.8. The camera was held in a small gap at the head’s top, front edge by a nylon-tipped set screw that pushed the side of the camera against the opposite side of the gap to secure its pose without damaging it. Figures 4.24A/C show the camera’s mounting, and Figure 4.24B shows an example of how the camera viewed an insertion. The overhead camera served many critical uses, especially during teleoperated insertions, including:

- Aligning the IV needle with the vein,
- Coordinating the two arms (NeedleBot and AssistoBot) to prevent collisions,
- Aligning the foam-tip with the catheter hub to prevent crushing it into the hand during vein clamping,
- Identifying penetration-events by watching the needle bevel and tissue movement,
- Confirming blood-flash in the catheter hub during needle withdrawal, and
- Estimating the catheter’s depth of insertion beneath the skin.

4.2.3.3.6 Control Electronics and Communication

We controlled the stepper motors and blood-pressure cuff inflation/deflation locally via a microprocessor board (SparkFun PN: DEV-09152, “Arduino Mega Microprocessor Development Board”). The Arduino Mega drove the stepper motors through separate, dedicated controller boards (Pololu PN: 1201, “Stepper Motor Driver Carriers”) and the blood-cuff pump motor and air release solenoid via custom H-bridge circuitry. We wired an emergency switch directly to the enable-lines of all motor-driving circuitry as a safety precaution. Communication between the microprocessor and the robot’s main computer program was achieved via a bidirectional, asynchronous USB-to-serial stream and a custom communication protocol that we developed. We clocked the communication speed at 500Hz in each direction with a delay of approximately a few milliseconds. A “deadman timer” measured the time between received serial packets and stopped all motion in the event of silence longer than 30mS. We controlled the hand-restraint, IPA-sprayer, and air-sprayer servo motors via a USB-to-serial stream that followed the Dynamixel-command format.

4.2.4 Interacting With the Robot

4.2.4.1 Robot Controller GUI

We developed a QT-based Graphical User Interface, or “GUI”, that allowed for complete, manual control of the robot. The GUI included real-time plots of all important variables (such as forces, positions, and angles) and allowed for adjustment of all key parameters (such as motor control gains and “pop” force-detection thresholds) during live operation of the robot such that changes were implemented in real-time without having to stop and recompile the program. It also allowed for manual position control of all DOFs, which greatly facilitated debugging and insertion setup. It used a color-coding scheme for displaying errors and motor-states intuitively so that the human operator could determine the robot’s state with a quick glance at the two, main screens. Three, small LCD screens displayed the feeds from the AssistoBot overhead camera and Needle Driver blood-flash and FPS cameras below the main GUI screens, as shown in Figures 4.25A-B.

4.2.4.2 Teleoperation Inputs

We used two Sensable Omni devices as the teleoperation inputs to our two arms, NeedleBot and AssistoBot, as shown in Figure 4.25C. Each Omni device provided six sensed DOFs (XYZ, roll, pitch, and yaw), three force-feedback DOFs (XYZ), and two boolean buttons. For controlling NeedleBot, the Omni’s 6 sensed DOFs mapped intuitively to the needle’s 6-DOF pose, and we used the buttons to control the catheter insertion DOF (holding one button for constant-velocity insertion and the other button for retraction). For AssistoBot, we mapped the Omni’s XYZ input to the arm’s XYZ motion, ignored the rotational inputs, and mapped the buttons to the IPA and air sprayers. As we had used all of the buttons on the Omni devices, we included a 4-button foot-pedal to control the two catheter grab-and-release DOFs and to clutch into and out of teleoperated control for both NeedleBot and AssistoBot independently, as shown in Figure 4.25D. We left control of the hand-restraint and blood-cuff inflation to the robot controller GUI as these would only be changed at the start or end of insertion when the practitioner’s hands were free of the Omni devices. The robot controller GUI included a button that set the handedness (left or right) of the teleoperation by mapping the human operator’s dominant hand to control NeedleBot and the non-dominant hand to AssistoBot without actually moving the

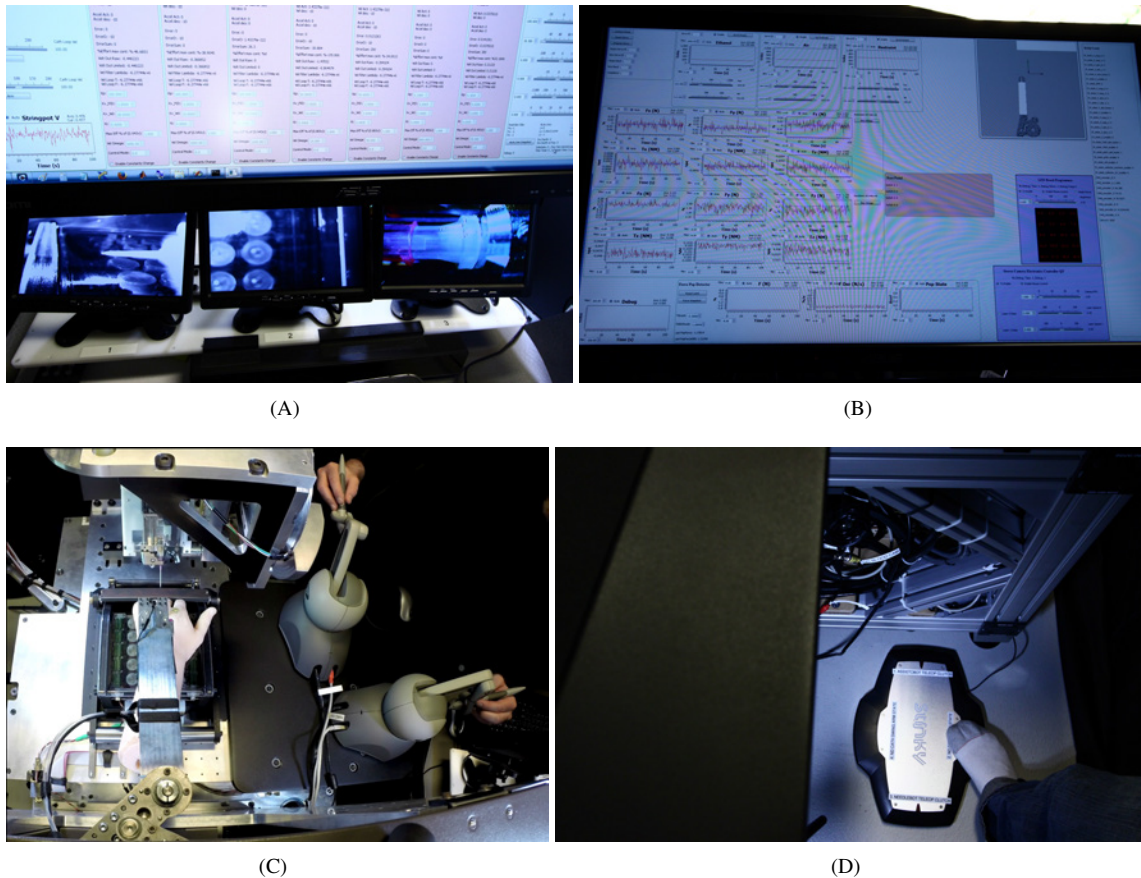


Figure 4.25: GUI and Teleoperated Control of the Robot. A-B: The Graphical User Interface (GUI) allows for complete, manual control of the robot. C-D: Two Sensable Omni haptic devices allow for teleoperation, one device per robot arm, along with a foot-pedal for controlling the catheter grab-and-release DOFs and clutching into and out of teleoperated control.

devices. The input devices had to be the same model to permit such software handedness-swapping, which is why we could not use a simpler, 3-DOF device to control AssistoBot's three, main DOFs.

4.2.5 Robot Base and Electronics

HaemoBot's base, which consisted of an aluminum frame and a top steel plate, contained most of the electronics, as shown in Figures 4.26A/C. We constructed the frame out of 80/20-brand aluminum extrusions as they provided for easy assembly, high stiffness/strength-to-weight ratios, and a flexible mounting scheme for adding components later in HaemoBot's development for which we had not initially allocated space. The frame was supported by four, adjustable-height feet that could be set individually so as to level the robot on an uneven floor. We designed the frame such that a narrow, hydraulic pallet jack (Big Joe Lift PN: BJL50-12, 20.5in wide x 48in long x 3in minimum height) could be inserted underneath the frame to lift the robot both for transportation and for leveling, as shown in Figure 4.26C. We included a single, 12mm-thick steel plate on top of the frame that tied NeedleBot, AssistoBot, and the frame together rigidly. The plate included





Figure 4.26: A: Most of the electronics are located in the robot's base. B: Emergency switches, both global (large, red button) and individual (row of small switches). C: Robot base and electronics, CAD model. The frame is constructed of 80/20 aluminum extrusions and a single, 12mm-thick, top steel plate that ties NeedleBot, AssistoBot, and the frame together rigidly and precisely. The light-weight covers attach to the frame via Velcro. Four adjustable feet allow for leveling of the robot on an uneven floor. All wires route from the electronics in the base, through the four access holes in the top plate, to the robot above. A narrow pallet jack (orange) lifts/transport the robot.

locating features that allowed NeedleBot and AssistoBot to be installed precisely with respect to each other. All wires were routed internally through the robot via four large, grommet-lined holes in the top plate to prevent wires from needing to be routed on the outer edges of the robot where they might be snagged. To protect the electronics, we attached covers to the frame's exterior that were laser-cut from 0.5-in-thick "Ultra Board" (polystyrene foam core with exterior layers of solid polystyrene, manufactured by United Industries). This material proved to be extremely durable, flat (good for laser-cutting), and sufficiently light-weight/stiff that it did not vibrate during the robot's operation. We attached the covers via Velcro inserts for the 80/20 extrusions and adhesive-backed, industrial Velcro stuck to the inner face of the covers.

In keeping with our flexible component placement theme, the primary electronics-mounting board was a 3/4-in-thick MDF board with a 40 x 40mm grid of M4-Teenuts, and all of our electronic subassemblies matched that 40mm grid so that they could be moved anywhere on the board as the design evolved. The robot's base contained many electronics subassemblies, including: the main processing computer, stereo camera driver, LED grid driver, Maxon RE-10 motor drivers, bank of Copely motor amplifiers with external inductors and power supplies, AssistoBot control circuitry, safety-microprocessor-subassembly, DAQ card breakout boards, analog filter for intravenous pressure transducer, Omni power supplies, switching power supplies that powered sensors and microprocessors, a "star ground" system, and many power strips.

We used switching, regulated power supplies to power all of our sensor and microprocessor boards because of their efficiency, accuracy, and safety features. However, we used unregulated, linear power supplies for driving the larger, Maxon motors and stepper motors (40V-supply PST-040-13-DP-F for the Maxon RE-30s, 70V-supply PST-075-10 for the Maxon RE-50s, and PowerVolt 24V-supply UPV-24Y610 for the stepper motors). We chose linear power supplies for the large motors because their inductive fields could handle sudden loading due to high motor accelerations, whereas switching supplies often interpreted sudden loading as a fault-condition that triggered safety-auto-shutdown. The one exception was in powering the Dynamixel servo motors, for which we used a switching power supply (9V-supply Meanwell S-150-9) so that we could adjust its voltage via the integral trim-pot to the manufacturer's recommended operating voltage of 9.6V.

We used three National Instruments Data Acquisition, or "DAQ", cards to provide the main computer's I/O lines. Specifically, we used the PCI-6723 (mostly analog-out lines for controlling the Maxon motors), PCI-6602 (mostly counter-lines for the encoders), and PCI-6224 (mostly analog-in and digital-in lines for the force, intravenous blood pressure, homing, and end-of-travel sensors). Because the Linux driver support for these cards was quite outdated (often resulting in unpredictable installation and performance problems), we used the Windows 7 operating system. Our main concern with this decision was whether the Windows scheduler would allow a sufficiently steady servo-loop rate for the robot to remain stable. However, we found that feeding one of the DAQ card's waveform-generation lines into a separate counter line provided a reliable, high resolution clock for measuring time between and scheduling servo-loop calls. We toggled a digital-out line every time the servo-loop was entered, and examination of this signal on an oscilloscope showed a fairly reliable, 2kHz servo-loop rate with only occasional blips where the loop was temporarily delayed due to a background program taking priority in the scheduler. Using this precisely measured servo-loop period instead of hard-coding the desired value allowed for our servo-loop velocity calculations to remain accurate even when these scheduler delays occurred intermittently. The analog cameras (FPS, blood-flash, and AssistoBot overhead) connected to the computer via a 4-channel, PCI-E frame grabber (Sensoray PN: 811).

We developed the emergency-stop-box and safety-microprocessor-subassembly to help provide safety limits on most of NeedleBot's motors, including those for XYZ, yaw, pitch, roll, and catheter insertion. Such limits were not necessary for the catheter grab-and-release DOFs because they were not mechanically dangerous. All of NeedleBot's safety-limited motors were driven by motor amplifiers that included an enable line, and all but the roll and catheter insertion motors also included directional stop-motion lines. The safety-microprocessor-subassembly logically combined a "deadman timer", global emergency-stop signal, and individual stop-signals for each motor and fed the results directly to each motor's enable line. The "deadman timer" was output by the control program's servo-loop to show that the program had not crashed/stalled. Each time the robot entered the servo-loop, it toggled a digital-output line, thereby producing a square wave. The safety-microprocessor-subassembly measured the time since the last toggle and issued an emergency-stop signal if it exceeded 5mS. It also fed the signals from the end-of-travel sensors (logically combined for the redundant sensors in the case of XYZ) to the directional stop-motion lines on the motor amplifiers to prevent further motion in that direction. AssistoBot was developed before the safety-microprocessor-subassembly

and could not accommodate individual motor or directional motion disabling. We did, however, include a single emergency-switch that could disable all of the stepper motors' enable lines at once.

As shown in Figure 4.26B, the emergency-stop-box consisted of a box that contained a large, global emergency-stop button that halted all motors and smaller, individual switches that halted only a single motor. Each of the switches fed into the safety-microprocessor-subassembly for processing and contained an LED that displayed its enable-state. This emergency-stop-box attached magnetically to HaemoBot's top, steel plate next to the human operator, but it could also be taken up to 25 feet from the robot in case a supervisor/trainer needed to take charge of safety while a separate operator performed the insertion. Integral handles allowed the box to be held securely, and a retractable lanyard allowed it to be suspended from the neck.

4.2.6 Controlling the Robot

HaemoBot could perform insertions under the direct control of a human operator, or it could insert autonomously. Human control was achievable via individual-motor-control using the GUI or full teleoperated control using the Omni devices and foot-pedal. Individual-motor-control was helpful for aligning the needle and vein while setting up an autonomous insertion because we could "jog" each axis using small increments. When we switched between control methods, we were careful to use the robot's values at the end of the previous control method as the initial-values for the next control method to provide a seamless transition.

Each of NeedleBot's motors (except for the catheter grab-and-release DOFs) was controlled via PID-controller, although only the catheter insertion DOF required a significant integration, or "I", term due to the stiction in its lead screw. Although we closed the digital-loop around position-control, we drove the motors electrically with current-control amplifiers that would allow for direct force-control in the future if we desired. Thus far, the only influence of force on the robot's motions was in setting the force levels when AssistoBot applied traction and clamped the catheter/vein and in using force "pop" detection to halt needle insertion and start the catheter insertion/release step-sequence.

4.2.6.1 Proof-of-Concept Autonomous Insertion

HaemoBot successfully inserted into an artificial, venous bifurcation using the axial insertion force to detect the "pop" and guide its motions and step-sequencing. As will be discussed in Chapter 6, the artificial vein was a nearly-exact replica of the author's venous network, and its force-profile had a similar shape and magnitude to the force "pops" observed clinically in Chapter 2. The vein resided on an exact, silicone replica of the author's hand such that we could test the robot's ergonomics and hand restraint system as well. Figure 4.28 provides a step-by-step demonstration of an autonomous insertion. To isolate variables during initial testing, we simplified the insertion by:

1. Aligning the needle with the vein manually instead of using the vision system's autonomous vein-finding and 6D-needle-pose-generation algorithms,
2. Inserting into a vein only (no skin) such that we only considered one set of tissue-layer force "pops",
3. Using only the catheter "pop" as the penetration-event without considering the needle "pop",
4. Inserting at a constant velocity, and
5. Inserting at a constant pitch angle instead of varying it as we observed clinically.

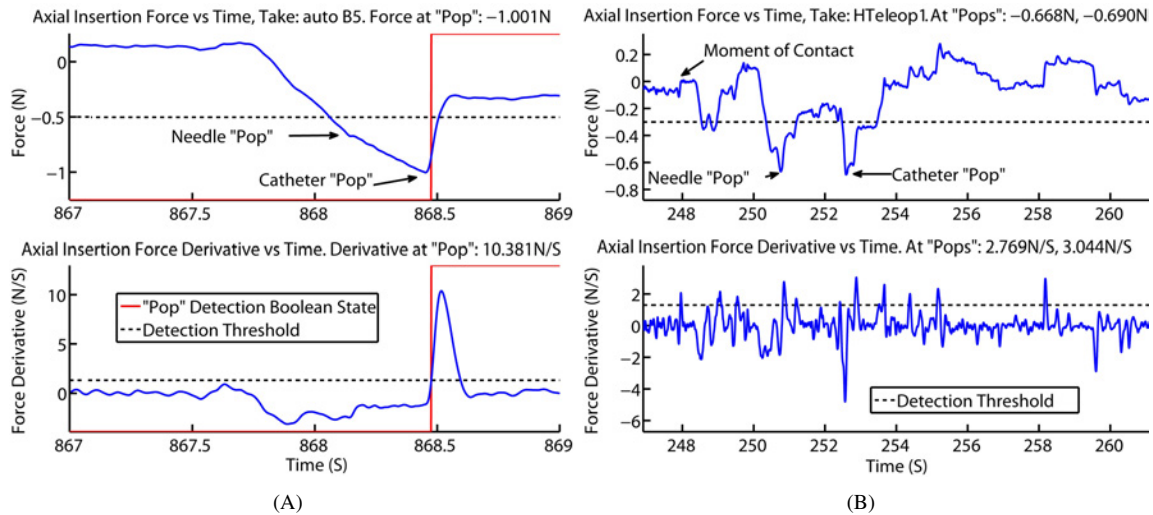


Figure 4.27: Examples of “Pop” Detection Using Axial Insertion Force.

The red curve displays the robot’s latched, “pop” detection state (1 = “pop” detected, transition from 0 to 1 is the moment of detection). The “pop” detection algorithm waits for the force and force time derivative to exceed heuristically-preset thresholds simultaneously, with the thresholds set directionally to ensure pushing, not pulling, on the needle and a force decrease at the “pop”. The force derivative requirement addresses the sudden nature of the force decrease at the “pop”. A: Autonomous insertion into our custom, latex bifurcation. B: Teleoperated insertion into the same bifurcation with “pop” detection disabled.

We chose an axial insertion speed of 10mm/s, which was higher than the maximum penetration-event speed observed in Chapter 2 (3.029 ± 1.099 mm/s for the vein catheter “pop”), because we became impatient waiting over many trials. Future work includes removing each of the testing simplifications and adjusting the control strategy to accommodate the more complex scenarios.

Our force “pop” detection algorithm consisted of waiting for the axial force and force time derivative to exceed preset thresholds simultaneously. The force derivative requirement addressed the sudden nature of the force decrease at the “pop”. For both of these requirements, the thresholds were set directionally such that the force had to be pointing in the correct direction (pushing, not pulling, on the needle) and decreasing towards zero at the “pop”. We set the detection thresholds heuristically based on our observations of the force-vs-time profiles for different insertion tests. To offset biasing forces, we subtracted the starting force level from the forces measured during the insertion. Figure 4.27A shows these force curves and successful “pop” detection during autonomous insertion.

As shown in Figures 4.18E-F, our blood-flash camera observed the blood-flash, but we had not yet implemented a computer vision algorithm to detect it automatically. This remains a topic for future work.

Performing the subset of human insertion steps described in section 4.1.1 required 33 individual steps for the robot. This did not include the steps of inflating/deflating the blood-cuff for arm restraint/tourniquetting, restraining the hand (since it was already installed), or visual vein-finding. The steps included:

1. “AssistoBot-Arm”: Move far from the hand during its insertion into the robot to prevent collision.
2. Needle Driver: Close the catheter grab-and-release swing-arm into its fixturing-position.
3. Needle Driver: Close the catheter grab-and-release nitinol-loop to hold the catheter.

4. NeedleBot: Back the needle far away from the hand to give AssistoBot room to prepare the hand.
5. “AssistoBot-Arm”: Move in the XY-plane a little left of the insertion site for IPA spraying/cleaning.
6. “AssistoBot-Arm”: Move downwards along the Z-axis to just above the insertion site.
7. AssistoBot IPA-sprayer: Close the IPA sprayer servo (start spraying).
8. AssistoBot IPA-sprayer: Open the IPA sprayer servo (stop spraying).
9. “AssistoBot-Arm”: Move in the XY-plane to right above the insertion site.
10. “AssistoBot-Arm”: Move downwards along the Z-axis until the force sensor reads above a threshold.
11. “AssistoBot-Arm”: Move/scrub in a circular motion in the XY-plane to clean and vasodilate.
12. “AssistoBot-Arm”: Move upwards along the Z-axis to just above the insertion site.
13. AssistoBot Air-sprayer: Close the air sprayer servo (start spraying).
14. AssistoBot Air-sprayer: Open the air sprayer servo (stop spraying).
15. “AssistoBot-Arm”: Move in the XY-plane to just behind the insertion site to prepare for traction.
16. “AssistoBot-Arm”: Move downwards along the Z-axis until the force sensor reads above a threshold.
17. “AssistoBot-Arm”: Move away from the fingertips along the patient’s arm-axis to apply skin traction.
18. NeedleBot: Insert the needle forward almost until the intended insertion site.
19. NeedleBot: Pause briefly to alert the human operator that needle insertion is imminent.
20. NeedleBot: Insert the needle forward at a constant speed and pose until the force “pop” is sensed.
21. NeedleBot: Pause briefly after the “pop” while the blood-flash is confirmed.
22. NeedleBot: Insert the needle forward 1mm beyond the position at which the “pop” was detected.
23. NeedleBot: Hold the needle stationary, and insert the catheter fully into the vein.
24. “AssistoBot-Arm”: Move upwards along the Z-axis to prepare for clamping the catheter/vein.
25. “AssistoBot-Arm”: Move in the XY-plane until the foam’s leading edge is just behind the insertion site.
26. “AssistoBot-Arm”: Move downwards until the force is above a threshold, clamping the catheter/vein.
27. Needle Driver: Open the catheter grab-and-release nitinol-loop to release the catheter.
28. Needle Driver: Open the catheter grab-and-release swing-arm away from the catheter for clearance.
29. Needle Driver: Retract the catheter grab-and-release distal-contact-mechanism away from the catheter.
30. NeedleBot: Withdraw the needle backwards away from the insertion site along the original trajectory.
31. NeedleBot: Stow the needle in a safe pose far away from the hand.
32. “AssistoBot-Arm”: Keep clamping the catheter until the human operator signals the IV line is attached.
33. “AssistoBot-Arm”: Move far away from the hand so that it can be withdrawn from the robot safely.

4.2.6.2 Proof-of-Concept Teleoperated Insertion

We successfully performed complete, teleoperated insertions into the same artificial, venous bifurcation as was used in the proof-of-concept autonomous insertions. Figure 4.27B shows the force and force-derivative-vs-time curves for a teleoperated insertion. The robot was not allowed to detect the “pop”, instead testing if the human teleoperator could make that determination. Teleoperation of the full insertion with both arms was difficult because of the many DOFs that had to be controlled simultaneously and the order-of-operations that had to be remembered for the catheter grab-and-release DOFs. As such, teleoperated insertion took longer

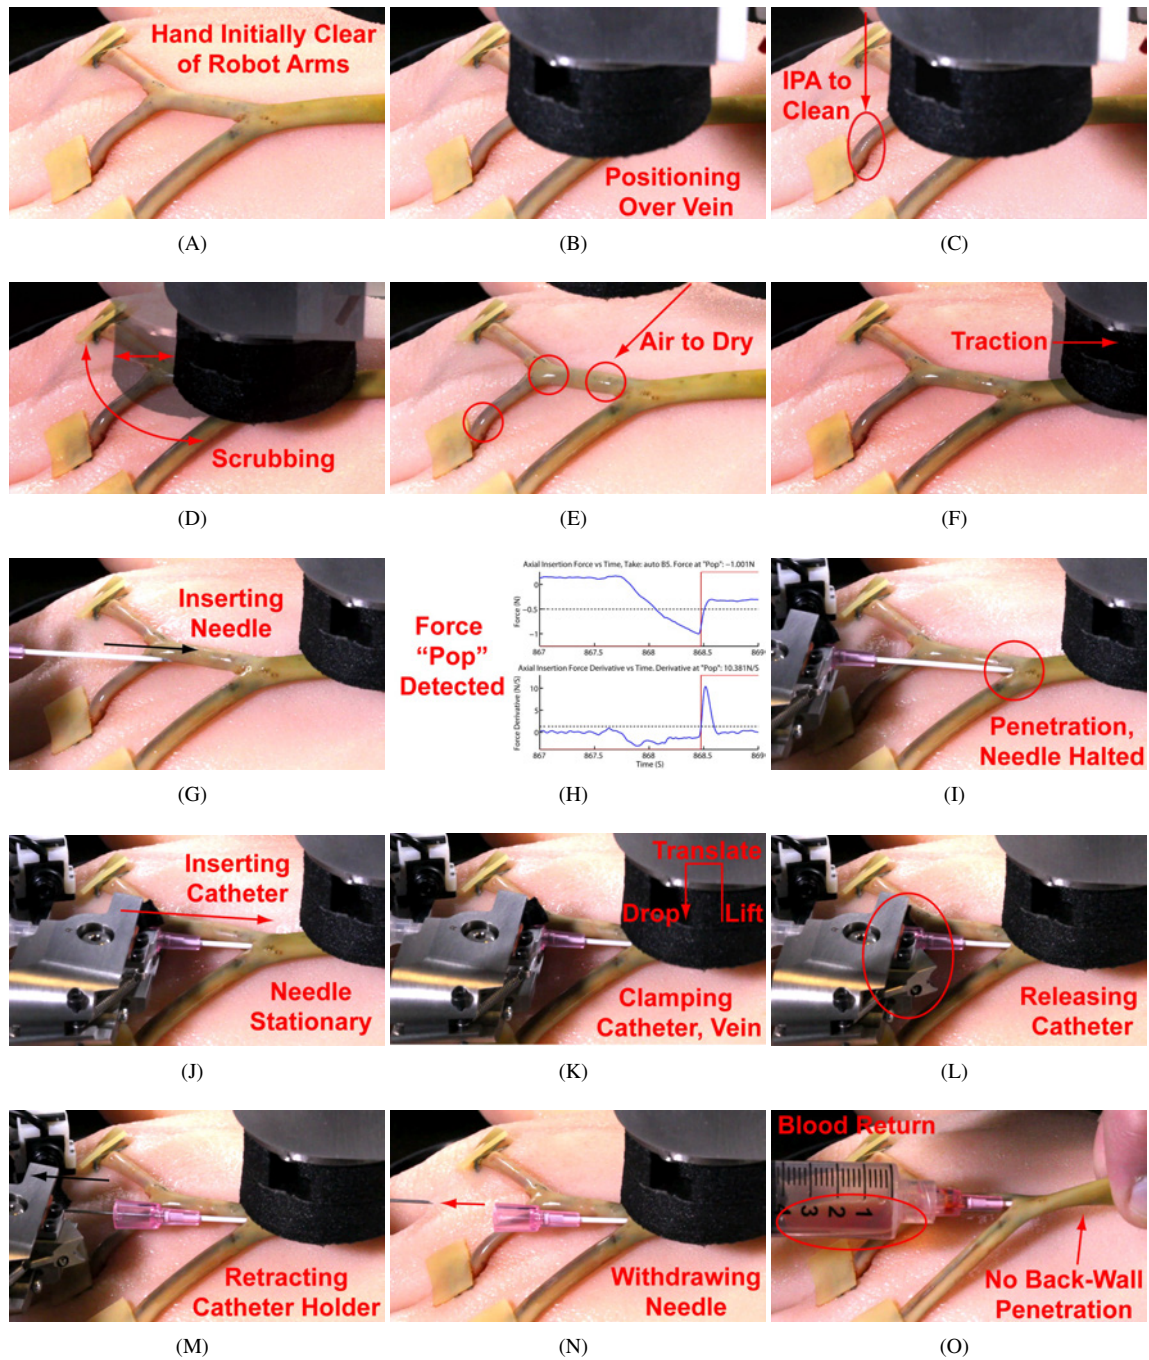


Figure 4.28: Proof-of-Concept Autonomous Insertion. A-B: AssistoBot moving to the insertion site. C: Spraying isopropyl alcohol (IPA) to disinfect the insertion site. D: Scrubbing IPA into the hand to clean the insertion site and vasodilate the vein. E: Spraying air to dry the insertion site. F: Applying skin traction to minimize skin/vein motion during insertion. G: Inserting the needle until a force "pop" is detected. H: The force "pop" is detected. I: The needle is halted upon the "pop". J: Inserting the plastic catheter into the vein while holding the needle stationary as a guidewire. K: AssistoBot clamping the catheter/vein to prevent bleeding or dislodgement of the catheter. L: Releasing the catheter. M: Retracting the catheter grab-and-release mechanism away from the catheter. N: Withdrawing the needle and stowing it in a safe position/pose. O: Attaching a syringe to the catheter and verifying blood return (success).

and had noisier, harder-to-parse force-vs-time curves than for autonomous insertion. We found that the best use for HaemoBot's teleoperation was aligning the needle with the vein to set up an autonomous insertion because this process was much faster via teleoperation than with individual-motor-control through the GUI.

Whereas the needle was always perfectly straight for autonomous insertions, it was often bent temporarily during the teleoperated insertions. This happened especially during the two-handed operation of catheter/vein clamping because our attention was devoted to clamping AssistoBot onto the catheter with the non-dominant hand, making it difficult to focus also on holding the needle still with the dominant hand. We ultimately solved this problem by temporarily clutching out of teleoperated control for whichever arm (AssistoBot or NeedleBot) was required to hold still such that the arm held its pose without our concentration or hand tremor being problematic. One-handed teleoperation was significantly easier than two-handed, and one-handed teleoperation was easier for NeedleBot than for AssistoBot because NeedleBot was fast enough to keep up with any hand motions without lagging behind like AssistoBot.

We experimented with applying force-feedback to the teleoperator's hands via the Omnis, both through feeding the force sensors' measurements directly to the Omnis and by adding a virtual spring between each Omni device and its controlled robot arm. The virtual spring provided the most assistance for AssistoBot because it improved the lag between the operator's hand and the actual robot arm, but this decreased the transparency of feeling the tissue/robot contact forces during traction pulling and catheter clamping. Because NeedleBot was fast enough that it did not really need a virtual spring to slow the operator's hand, it was more beneficial to feed the needle forces directly to the controlling Omni.

4.3 Discussion and Thesis Contributions

HaemoBot's design is unique from other researchers' efforts towards phlebotomy and IV robots in that it performs the insertion with two arms, moves the catheter independently of the needle, and implements inherent, mechanical safety. This unique design allowed us to prove for the first time that a robot can insert an IV on its own through our proof-of-concept autonomous insertions. Thinking more specifically, we made five main, novel contributions in this chapter:

1. Proved that a robot can autonomously perform all of the 1-handed and 2-handed mechanical steps of IV insertion, including independent manipulation of the catheter as a separate part from the needle,
2. Proved that a robot can perform these steps based on the force "pop" that we observed clinically,
3. Proved that we can do all of this on realistic hand, arm, and vein geometries (not just rubber tubes),
4. Developed a patented method of shining a targeting-laser through an IV needle while maintaining its sterility and normal functionality, and
5. Developed a patent-pending design for a Remote Center of Motion that has many superior fabrication, installation, and performance properties over existing RCM designs.

Chapter 5: Development of eye.V.ease: a Simple, Inexpensive Assistive Device for IV Insertion

5.1 Motivation and Overview

This chapter seeks to distill the best lessons from our previous work into a simple assistive device, “eye.V.ease”, to help a human practitioner insert IVs more easily and effectively.

Recall that the fundamental, driving question behind our development of HaemoBot was whether or not we could raise the IV insertion success rate of a robot beyond that of a human practitioner. Although we proved in Chapters 3 and 4 that it is technically possible for a robot to insert an IV autonomously, we did not address the non-technical questions of whether such a complicated, autonomous robot would be permitted, accepted, and used clinically, all of which would need to be answered affirmatively for us to confirm and implement robotic insertion as a superior alternative to human insertion.

HaemoBot was necessarily complicated and expensive because its purpose was to perform a complete IV insertion autonomously, as well as to serve as a research platform for exploring different clinical questions and insertion techniques. However, HaemoBot’s autonomy and enabling-complexity are actually problems from the perspective of commercial-deployment. When we spoke with practitioners and other robotics researchers about our system, they often raised concerns about its prospects for commercialization, including:

- FDA concerns over inherent safety, especially in view of autonomous operation,
- Practitioner misgivings about the system’s operational complexity and ability to replace them,
- Patient anxiety about a robotically-inserted needle, and
- Hospital questions about the benefit-to-cost ratio of such a system.

Our choices are to 1. keep our overall vision of HaemoBot as a complicated, autonomous technical solution that we force against the non-technical concerns of the humans who will use it, or to 2. change the form of our solution, using the best lessons from HaemoBot as inspiration, to be more practical and palatable to FDA-approvers, practitioners, patients, and hospital procurement decision makers. Since our primary objective is to make IV insertions more successful, we must change our solution to address both the technical and non-technical problems. As such, we now update our original, driving question to:

- Can we raise the IV insertion success rate of a human practitioner with a non-robotic solution that implements the best lessons we learned from the development of our robot? More specifically,
 - Can lessons we learned from developing the robot be used in a simple assistive device to make IV insertion easier for human practitioners, while keeping the device a minimal safety-risk for the FDA, simple for practitioners, non-threatening to patients, and inexpensive for hospitals?
 - Can we develop a tool to allow someone to insert an IV on themselves by simplifying IV insertion from a 2-handed to 1-handed procedure? This would partially address one of our original goals

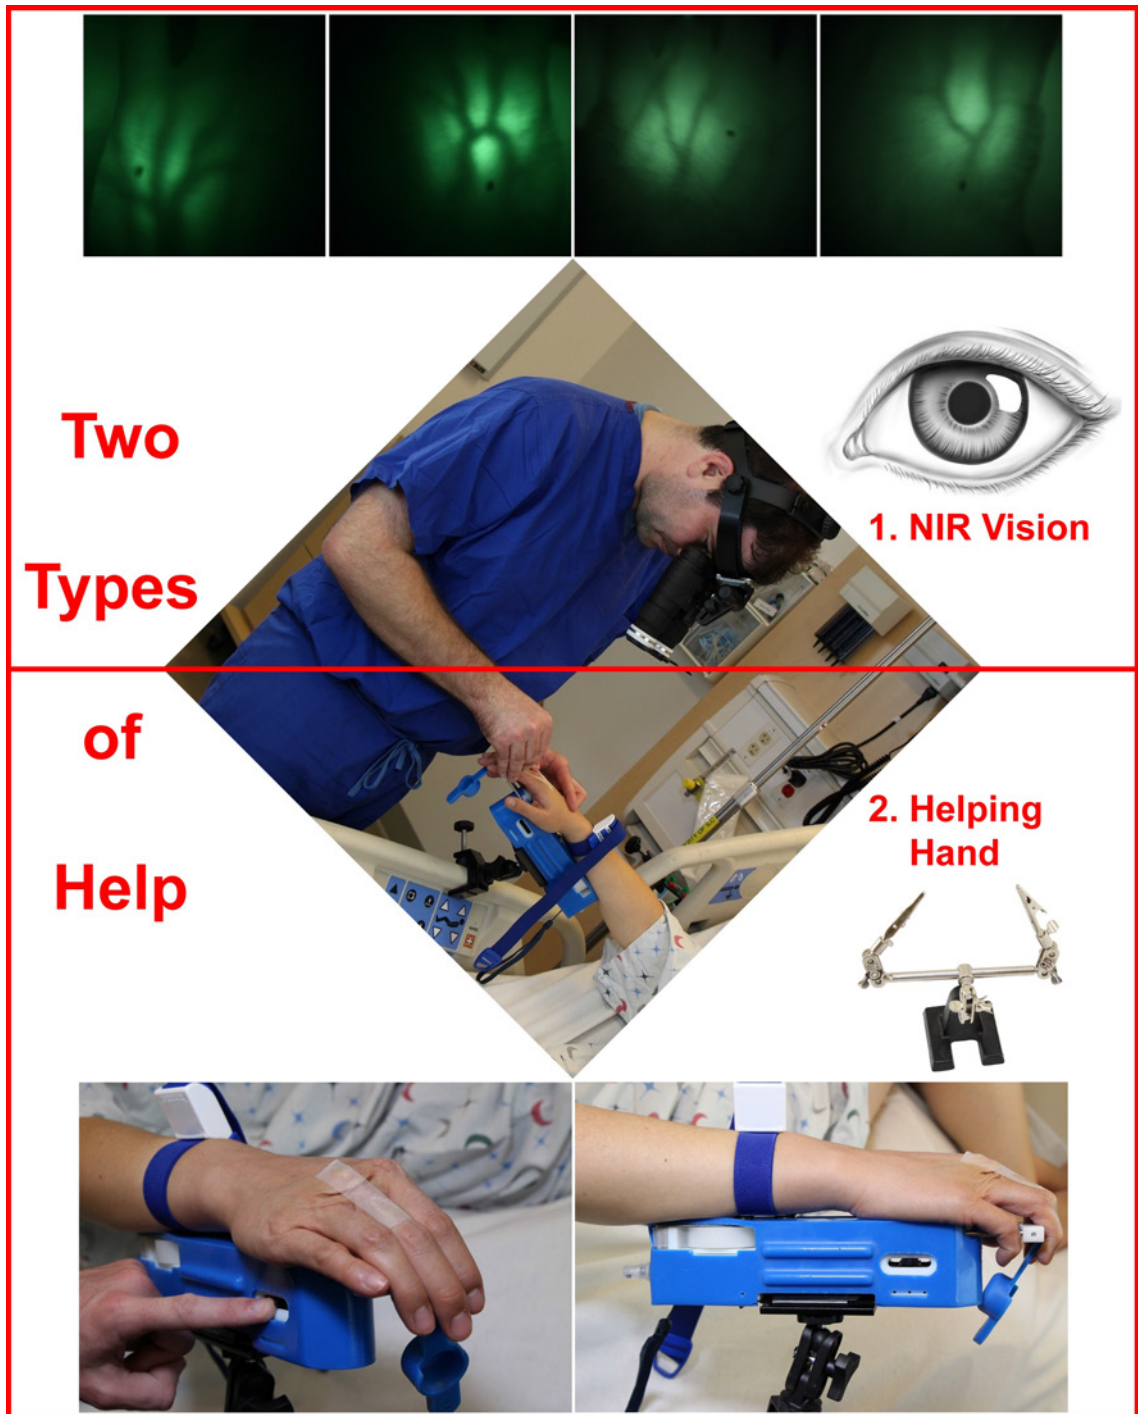


Figure 5.1: eye.V.ease, Two Types of Help: NIR Vision and “Helping Hand”. The NIR-Goggles provide reliable, stereo NIR vein-visualization with natural depth perception and without false-vein-detection. The special ergonomics of our device, integral supplies (tourniquet, alcohol sprayer, and medical tape), and pairing of our device with a special IV needle eliminate all of the non-dominant hand tasks during active insertion and most tasks prior to insertion (with the exception of a few setup and package-opening steps).

of enabling IV insertions in remote locations where a human practitioner could not be physically present. Although we originally imagined using HaemoBot to insert the IV autonomously or through teleoperation, we now wonder if we could empower patients to insert on themselves.

Having decided to build an assistive device that makes it easier for a human to insert on themselves or others, the next logical question is what type of assistance to provide. That is, what are the main lessons that we learned from our prior work on HaemoBot that we can include in our device to solve unmet clinical needs? From our clinical experiment in Chapter 2 and work on NIR vein-imaging in Chapter 3, we observed two main types of help our device could offer that would address currently unmet clinical needs, including:

1. Reliable, stereo NIR vein-visualization that is free of false-vein-detection and
2. A “helping hand” that performs most or all of the functions of the practitioner’s non-dominant hand.

Figure 5.1 shows these two methods of help, and the following sections discuss each in detail.

5.1.1 Goal 1: Reliable, Stereo NIR Vein-Visualization Without False-Vein-Detection

Although there are many commercially-available vein-viewing devices, as shown in Figure 5.2, we found that none of them met our following technical requirements for providing reliability, safety, and effectiveness:

1. Transillumination-based lighting to prevent false-vein-detection. Reflection-based systems often see local anatomical features that cast shadows (e.g. hair, the valley in between metacarpal bones near the knuckles, and skin/fat creases) as veins. Figures 5.2B-C show such false-vein-detection as compared against our device, shown in Figure 5.2A, which displays the true vein-pattern.
2. True 3D stereo display for depth perception and precise positioning of the needle in the vein.
3. Infrared, not visible light. As described in Chapter 3, NIR light penetrates tissue more effectively and provides higher vein/non-vein contrast than other frequencies of light. In limited cases (very thin hands in a dark room), visible light can be used to transilluminate the hand, but the veins appear much fainter than when viewed with NIR light. Figures 5.2H-I show such visible-light devices and the low vein/non-vein contrast they provide.
4. Shows only veins and not arteries. Showing both types of blood vessel is a cognitive load on practitioners as they must disambiguate between vessel types to prevent selecting an artery for cannulation. Figure 5.2G shows how difficult it is to disambiguate between vessel types when both are visible.

5.1.1.1 Comparison With Commercially-Available Vein-Finding Devices

Table 5.1 scores the main commercially-available vein-visualization devices on our design requirements. We will now discuss each of them in detail below:

5.1.1.1.1 AccuVein AV400

The AccuVein AV400 is a handheld/handsfree, reflectance-based vein-finder, about the size and aspect ratio of a medium-sized flashlight, that detects the vein pattern and projects it back onto the patient’s hand as a red image. The current model, the AV400, has a list price of \$6,000 USD, including the wheeled stand with an adjustable arm for handsfree operation. The operation of the device, as best as can be determined by [24], is this: a 740nm near infrared laser and a 640nm red laser are combined and fed into a MEMS scanner that directs the combined beam onto the patient’s target site in a raster-scan pattern. At each position of the beam,

Table 5.1: Comparison of Commercially-Available Vein-Finding Devices. Red cells show properties that violate our requirements for a vein-visualization system, and blue cells show uncertainty. Requirement 1: Transillumination-based lighting to prevent false-vein-detection. Requirement 2: True 3D stereo display for depth perception and precise positioning of the needle in the vein. Requirement 3: Infrared, not visible light to provide higher vein/non-vein contrast. Requirement 4: Shows only veins and not arteries to lower cognitive load and prevent selecting an artery for cannulation.

	Our Device	AccuVein	Christie Digital	VueTek	Evena	INRIS	TransLite	Philips
Model		AV400	FLEX	VeinSite	Eyes-On Glasses	Vascular Viewer	VeinLite LED	Wee Sight
Price	\$1339	\$6,000	\$4,950	\$5,995	TBD	\$14,900	\$469	\$140
Vessels Shown	Veins	Veins	Veins	Veins	Veins	Veins, Arteries	Veins	Veins
Reflection vs Transillumination	Trans.	Ref.	Ref.	Ref.	Ref.	Trans.	Trans.	Trans.
Detection Light	NIR	NIR	NIR	NIR	NIR	NIR	Visible	Visible
Display Method	3D Goggles	Projected Red Light	Projected Green Light	2D Screen Goggles	Semi-3D Glasses	2D Screen	Naked Eye	Naked Eye
Stereo Display	Yes	Yes	Yes	No	Unclear	No	Yes	Yes

a photodetector tuned to the 740nm light outputs an analog signal which is proportionate to the amount of reflected light from the 740nm beam and, hence, the beam's proximity to a vein. This raster-scan is repeated at 60 different power levels of the 740nm laser, with scans performed at higher power levels penetrating further into the tissue and giving depth information. At the end of the 60 scans, all of the images are combined into a single image that is projected onto the patient's anatomy via the 640nm laser. The raster-scanning is faster than the human eye can detect, so the projection appears to be a real-time video.

In [19], usage of the AccuVein AV400 was evaluated in 6 nursing units and reported to yield a 45% decrease in escalation calls (requiring a more-experienced practitioner to attempt the insertion again), with non-pediatric units experiencing most of the improvement. In our testing of the device, we found its only problem to be false-vein-detection. Figure 5.2B shows how the device falsely-detected a vein in the author's hand not seen with our own device, as shown in Figure 5.2A. The frequency of false-vein-detection was quite high when tested on a 7-month-old infant due to the many fat/skin creases that were detected as veins.

5.1.1.1.2 Christie Digital Flex, Luminetx VeinViewer

Christie Digital offers several vein-viewers that are similar to the AccuVein device in principle and operation (handsfree, NIR, and reflectance-based vein-detection with a visible-light projection of the detected veins onto the hand). Their device is a re-engineered version of those sold by Luminetx, which Christie Digital acquired. Their least expensive model is the Flex, which retails for \$4,950 USD and can be held by hand or mounted on a flexible arm for handsfree operation. The main differences from the AccuVein device are that Christie Digital's devices use 1. an LED/camera vein-detection system instead of a laser/photodetector system and 2. a compact image-projector that displays a green vein image on the patient's hand instead of a red laser-rastered image. Clinical evaluations of the Luminetx devices yielded mixed results, with one study showing no effect on success rates [11], other studies showing improved success rates [29] [36], and another study showing worse success rates [61]. Figure 5.2C shows the Luminetx system falsely-detecting a vein in the author's hand. Figure 5.2D shows the green projected vein image from the Christie Digital Flex device.




Figure 5.2: Demonstrations of Commercially-Available Vein-Finding Devices. A: Our device. B: AccuVein AV400. The falsely-detected vein (circled in black) was not seen by our device (Subfigure A). C: Luminetx VeinViewer. The falsely-detected vein (circled in red) was not seen by our device (Subfigure A). D: Christie Digital VeinViewer Flex. E-F: VueTek Veinsite, 2D Goggles. G: InfraRed Imaging Systems (INRIS) VascularViewer. Both veins and arteries visible. Image used with permission. H: TransLite VeinLite LED in a dark room. Vein circled in black. I: Philips Wee Sight Transilluminator in a dark room. Note that the veins appear less visible here due to camera limitations than in person.

5.1.1.1.3 VueTek Veinsite

The VueTek Veinsite is an NIR, reflectance-based vein-viewer in the form of a single-piece helmet (all elements contained within the helmet), and its list price is \$5,995 USD. We tested the device on several adults and found that it showed the hand, arm, leg, and foot veins very clearly, as shown in Figure 5.2F, but had some trouble disambiguating hair from veins. Further, wearing the device as a helmet was not challenging ergonomically. We tested it on a 7-month-old infant and found that it was able to show at least 1 vein in the infant's hand and foot, but the number and clarity of insertion sites were greatly diminished over what was experienced for the adults. A recent study [14] found that the Veinsite significantly increased the number of veins detected over unassisted, standard vein-detection techniques and was a useful planning tool.

Unfortunately, this device lacks the depth perception needed for use during active insertion because its single camera displays the same image in both eye-LCDs. We tried to trace the vein pattern on a patient's hand under the guidance of the device and found that making correct contact between the hand and marker was very difficult. For this reason, the device's manual states that the infrared-view is for planning-purposes only and should not be used as a live feed during insertion. The user must first look up into the LCDs to see the 2D, infrared-view of the veins for planning purposes and then look below the edge of the helmet to watch the patient directly, and unassisted, while performing the insertion, as shown in Figure 5.2E.

5.1.1.1.4 Evena Medical Eyes-On Glasses

Evena Medical offers multiple NIR, reflectance-based vein-finders, including the Owl, Sparrow, and Eyes-On Glasses. Whereas the Owl and Sparrow models display the detected veins on a 2D-screen, the Eyes-On Glasses are a heads-up-display (HUD) in the form of glasses. Each eye has its own display with the unmagnified, infrared feed in the center, and the user can see through the displays to the "real-world"/non-infrared view. The glasses themselves are roughly the size of a bulky pair of safety goggles, are able to fit over most glasses, and connect to a belt-mounted battery and computer (roughly the size of a large cellphone) via a wire that runs up the back of the neck. The glasses also offer the ability to overlay the feed from a special ultrasound, consult with another practitioner via built-in speakers and a microphone, and wirelessly-stream video to a monitor for training purposes and updating the patient's record.

The Eyes-On Glasses have two cameras, each of which feeds its infrared-view into a respective eye-LCD-display for a 3D, stereoscopic view. Looking through the glasses is similar to using bifocals in that seeing the infrared-view requires near-field focus, and looking at the real-world view past the LCDs requires far-field focus. That is, because the infrared and real-world/non-infrared views require the eye to focus differently, the views are not seen simultaneously, or overlaid, on each other. Instead, the user focuses on one view at a time. Beyond not being overlaid on each other, the two views are also not registered to each other, meaning that an insertion site will appear in different locations in the two views. Because the views are not registered, the placement of the glasses on the user's face is not critical. The glasses are intended to be used for real-time guidance during insertion but can also be used just for pre-insertion planning.

Evena products use multispectral-lighting, wherein four different wavelengths of infrared light (likely

including 730nm, 850nm, and 920nm [25] [26]) are flashed sequentially and the resulting four images combined to provide a more complete view of the veins. The additional wavelengths allow for imaging of patients for whom single-wavelength imaging was unable to see the veins, often due to factors such as skin pigmentation, fat, and hair. The wavelengths were “tuned” to hide arteries.

5.1.1.1.5 InfraRed Imaging Systems (INRIS) VascularViewer

The InfraRed Imaging Systems (INRIS) VascularViewer uses transillumination at 850nm to show veins and arteries, as seen in Figure 5.2G, so that it can assist with both venipuncture and arterial blood-gas (ABG) sticks. The system retails for \$14,900 USD, and each Disposable Light Source, which is replaced between patients, costs \$19. The device cannot provide depth perception because it displays the veins on a 2D screen, with earlier versions using a monocular-viewer instead. A clinical study [21] showed that the INRIS system reduced the number of sticks and the completion-time for peripheral ABG sticks, and the company advertises an unpublished study that showed decreased failure rate and completion-time for IV insertions.

5.1.1.1.6 TransLite VeinLite

The VeinLite LED, made by TransLite Inc., is a small, handheld device that uses a combination of orange and red LEDs to show the veins to the naked eye through “side-transillumination”. The adult-version retails for \$469 USD and the pediatric version for \$188 USD. We tested the VeinLite LED and found that it showed veins on the hands of both an adult male and female with the overhead lights on, although the device showed only a difficult-to-see, small segment of the overall venous network. The results of our testing on a 7-month-old infant were significantly better, as seen in Figure 5.2H, with the veins in the hands and feet seen more clearly than on the adults. Of all the vein-finding devices we tested, the VeinLite provided the second-best images of the infant’s veins, with the Philips Wee Sight providing the best.

The VeinLite’s method of side-transillumination is quite different from other transilluminators, such as our device. Whereas a traditional transilluminator beams its light through the tissue to be viewed on the opposite side, the VeinLite delivers its light from the same side from which the vein is viewed, employing a ring of 24 radially-and-downward-oriented LEDs that sit on the bottom of the device near the skin and angle the light into the skin at a 45 degree angle such that back-scattering of the light illuminates the vein [20] [42].

According to [43], the usage of orange light (580nm-620nm) is meant to allow for “enhanced imaging of superficial veins”, and the red light is included to image “deeper veins and darker skin”. The usage of visible light allows for practitioners to view the veins with the naked eye instead of needing hardware to visualize infrared light. However, the cost of this decision is 1. decreased vein/non-vein contrast due to a lower absorption coefficient of deoxygenated haemoglobin for visible versus NIR light and 2. the need for higher optical power to penetrate tissue of a given thickness compared to that required for NIR light. To avoid the need for darkening the room, the Veinlite includes a “light shield” over the transilluminated vein that blocks ambient light while allowing the practitioner to view the vein through a small opening [42].

5.1.1.1.7 Philips Wee Sight

The Philips Wee Sight Transilluminator is a small, handheld transilluminator that is intended primarily for neonates and very young infants. It emits red, LED light at 629nm (with the same benefits and detractions

described above for visible light), and it costs \$140 USD. It is intended to fit in the palm of a practitioner's hand with the infant's arm or leg on top of the device so that the practitioner's fingers can restrain both the arm/leg and the device with only one hand. We tested the device on two adults (one male and one female) and a 7-month-old, fair-skinned infant, as shown in Figure 5.2I. It showed veins on the thinner portions of the adult female's hands in a darkened room, but it was unable to show veins in any portions of the adult male's hands due to their thickness. However, the device delivered a clear, high-contrast view of the venous network on the infant's hands and a less-distinct, but still visible, view of the veins in the infant's feet in a darkened room. This device provided the clearest, most-comprehensive view of the infant's hand veins of any commercially-available device we tested. However, the fact that the device must be used in a dark room makes its usage somewhat inconvenient. Although we obtained excellent views of the infant's hand veins, some researchers question the clinical usefulness of this device. In [54], a 5.5% decrease in IV insertion success rate was found when using the Wee Sight instead of normal, unassisted vein-identification through ambient-light observation and palpation.

5.1.2 Goal 2: A “Helping Hand” to Replace the Practitioner’s Non-Dominant Hand

We have eliminated all of the non-dominant hand tasks during active insertion and most tasks prior to insertion (with the exception of a few setup and package-opening steps) through the special ergonomics of our device, the integral supplies (tourniquet, alcohol sprayer, and medical tape), and the pairing of our device with a special IV needle. The following list, which will henceforth be known as the “Non-Dominant-Hand-Task-List”, details these tasks and to what degree our device reduced their complexity:

1. Applying the tourniquet (often used to require both hands, now requires one hand).
2. Opening the alcohol wipe to clean the insertion site (used to require both hands, now requires one).
3. Restraining the hand/arm of a noncompliant child (used to require the non-dominant hand or an additional practitioner, now requires no hands after initial setup by the practitioner).
4. Bending the fingers downward to keep them out of the way and help with traction (used to require the non-dominant hand, now requires no hands).
5. Pulling on the hand skin for traction (used to require the non-dominant hand, now requires no hands after initial setup by the practitioner).
6. Ripping a piece of tape for securing the catheter (used to require both hands, now requires one hand).
7. Applying pressure to the vein to prevent blood from escaping before the IV line's Luer Lock is attached (used to require the non-dominant hand, now requires no hands when using a special needle).
8. Preventing the catheter from rotating as it is secured via tape or as the IV line's Luer Lock is attached (used to require the non-dominant hand, now requires no hands when using a special needle).

Items 7 and 8 were solved by using the Braun B. Introcan Safety 3 Closed IV Catheter (20ga x 1 1/4inch, distributed as BoundTree Medical PN: 1613-13020). Unlike in a traditional IV catheter, the hub of this “closed” catheter contains a split, flexible, and multiuse membrane that remains sealed against blood-leakage unless a needle, syringe tip, or Luer Lock tip pushes it open. This catheter also contains stabilizing “wings” on the sides of its bottom surface that prevent rotation as it is taped in place or the Luer Lock is attached.

Freeing the non-dominant hand from menial tasks allows the practitioner to use it for more advanced, dedicated, and effective help during the IV insertion, including:

1. Threading the catheter into the vein with the full dexterity of one hand (instead of just the dominant hand's index finger) and the full stability of the other hand holding the needle still,
2. Further stabilizing a moving vein, or
3. Palpating throughout the insertion instead of just at the beginning.

Although the practitioner was previously able to accomplish these more advanced forms of help, they required ceasing some of the menial, but important, tasks. As an example, if the practitioner stopped applying skin traction so as to free a hand for catheter threading, the vein would move from its known location.

Freeing the non-dominant hand also allows a person (in theory) to insert an IV on themselves, provided they have the proper training and a suitable vein in the non-dominant hand. Before our device, even a person experienced at inserting IVs in others could not perform a self-insertion because IV insertion is inherently a two-handed procedure as traditionally performed. Although this would be an option of last resort for most patients, there are instances when this might be necessary, including on an active battlefield or on remote scientific outposts (like the International Space Station and Amundsen-Scott South Pole Station).

None of the commercially-available devices discussed in section 5.1.1.1 provide any of the “Helping Hand” functionalities of our device. In fact, the Evena Eyes-On-Glasses were the only product to include any non-visualization, assistive functionalities, and they did not free the non-dominant hand during actual insertion when it is needed most.

5.2 Final Design

Our “eye.V.ease” device consists of two parts: the “LightWand” that transilluminates the patient’s hand and the NIR-goggles that allow the practitioner to see the NIR light. The following sections describe each.

5.2.1 NIR-Goggles

When we were selecting a method of displaying the NIR vein images to the practitioner, our main requirements were that the display method:

1. Have stereo vision to give the depth perception necessary during actual insertion,
2. Be head-mounted so that the display moves naturally with the practitioner, and
3. Provide a real-size (unmagnified) image that was minimally processed before being displayed.

The solution that best fit these criteria was night vision goggles, henceforth referred to as “NVGs”, that are typically worn by hunters. After trying different models, we ultimately selected the Sellmark SM15070 model, which generally retails for \$500 USD. These goggles provide two, independent photocathode tubes, and the stereoscopic effect/depth perception was immediately intuitive. Within 5 minutes of wearing the goggles, we were able to trace the venous pattern on the back of a hand with a marker as easily as tracing on paper with the naked eye. The goggles had many desirable, and some unique, features, including:

1. The tip of each photocathode tube included a M27 x 0.75mm thread to which we attached longpass filters and easy-to-use lens caps via step-up rings.

2. The “generation-1” resolution of 36 lines/mm was more than sufficient to see all of the tiny details on the hand’s surface. Later-generation technology was unnecessary and prohibitively expensive.
3. An auto-shutdown feature turned off the goggles in the event of excess light so as to prevent damage.
4. The goggles included an 805nm NIR emitter and were tuned to see this wavelength, which approximately matched the 808nm laser we wished to use as our transillumination light source. We ultimately covered the emitter to prevent it from overwhelming the transilluminated light.
5. The goggles included a 1:1 magnification (no zoom) and an unusually small minimum-focusing-distance of 300mm, which is shorter than the typical distance between the practitioner’s head and patient’s hand during a normal insertion. This allowed the hand to appear real-size and in focus.
6. Each eyepiece included separate focus and vision-correction rings (± 5 diopters) for users who need glasses and cannot wear contacts. Glasses cannot be worn with the goggles because the NVG eyepieces must rest on the user’s face, contacting the eye socket above and below the eye, for the images in the eyepiece to appear in focus and the stereoscopic effect of the two photocathode tubes to be correct.

The stock head mount for the goggles was awkward and difficult to adjust, so we replaced it with the H5 Head Band from the Oak Tree Group LLC, which allows for fast, independent knob-adjustment of head size in two dimensions, as well as angling the NVGs with respect to the user’s eyes. Whereas switching the NVGs with their original head mount between users typically required 2-3 minutes for resizing, the H5 head band allowed for switching/resizing in as little as 20 seconds.

We found that wearing the goggles so close to the eyes for more than a few minutes resulted in fogging of the eyepieces that presented as a loss of focus. We experimented with various de-fogging solutions and found that “Nikon 8073 Fog Eliminator”, an antifogging-agent-impregnated cloth that is rubbed on the eyepiece, mitigated the fogging significantly.

Recording images through the goggles was difficult. Our best solution was to mount the goggles and a Canon T4I DSLR camera to a rigid board with the camera directly in front of one eyepiece. Although we obtained images that showed the veins fairly well by manually tuning the focus, exposure, ISO, and F-Stop, the veins were much clearer and higher-contrast when viewed through the goggles in person.

We tested our device on twelve adults (mean age = 35, range = [28, 60]), two children (ages 5 and 7), and one infant (7 months old). We could see the veins very clearly in all of the adults, as seen in Figure 5.3, clearly but less distinctly in the children, as seen in Figures 5.5C-G, and not at all in the infant. Figure 5.3C shows the veins of an obese woman (BMI = 39) who had noticeably fatter hands than the other adult test subjects, and there was no difference in visualization quality. We hypothesize that our device could not show the infant’s veins because the combination of their small diameter and deeper placement in the hand allowed for light-scattering to mask them, a phenomenon discussed in Chapter 3.

5.2.1.1 Lens Caps, Longpass-Filtering, and Compatibility with Room Lighting

We encountered two main problems with using the NVGs in a room with bright fluorescent lighting. The first problem was the sheer amount of light, which would blind the goggles and cause them to auto-shutdown. The second problem was that the broad spectrum of wavelengths in the fluorescent lighting overwhelmed our




Figure 5.3: eye.V.ease, Practitioner's View Through the NIR-Goggles. While these pictures show the veins clearly, the in-person view is much superior. The black spot that appears in all Subfigures was due to accidental, direct laser viewing that damaged the goggles during initial development (before adding the safety interlocks). Although not shown here, tests on the foot veins of an adult, Asian female and adult, Caucasian male showed that the device also works well for imaging feet. A-B: 29-year-old, Caucasian male. C: 62-year-old, obese (BMI = 39), Caucasian female. D-F: 28-year-old, Caucasian male (the author). G-I: 32-year-old, Asian male.

specific transillumination wavelength (low signal-to-noise ratio). To solve this, we placed a 720nm longpass filter (Zykor 37mm IR72) in front of each photocathode tube. This filter attenuated the overall amount of light so that the NVGs did not auto-shutdown and boosted the signal-to-noise ratio so that we could see the veins clearly even in a room with bright fluorescent lighting.

It should be noted that the filtered NVGs still could not operate around incandescent, halogen, or natural-sun lighting due to their large infrared component that masked our device's transilluminated light. However, as the majority of lights in corporate settings are fluorescent, the only real limitation is not being able to use the device outside or near a sunny window. In the future, we will replace the longpass filter with an 808nm bandpass filter to increase immunity to non-fluorescent lighting and allow us to use a less powerful laser.

The NVGs are so sensitive to light that keeping their lens caps open around bright, non-fluorescent lights or outside during daylight can permanently damage the photocathode tubes, despite their auto-shutoff feature. For this reason, it is important to have reliable, easy-to-use lens caps. The stock lens caps were unreliable and difficult to insert, so we replaced them with a pair of Leupold 28mm Alumina Flipback Objective Lens Covers (Leupold PN: 58755, distributed by MidwayUSA as PN: 519953). These all-metal lens covers close magnetically, are sealed against moisture, and are spring-loaded such that they flip-open when the user presses on a tab, making them the easiest, most-secure lens caps we could find. We 3D-printed a step-up ring to attach the custom, non-standard thread on the lens covers (1.256inch pitch diameter x 48 TPI) to the standard thread on the longpass filter (M37 x 0.75mm). No vignetting was experienced even after stacking a step-up ring, longpass filter, second step-up ring, and lens cover on the front of each photocathode tube.

5.2.2 LightWand

5.2.2.1 Light Source

5.2.2.1.1 Transillumination vs Reflectance-Based Imaging Configurations

As with developing the NIR imaging system for HaemoBot, we had to choose between a transillumination or reflectance-based goggles/light-source configuration. As we discussed in section 5.1.1, we chose transillumination primarily to avoid the false-vein-detection inherent to reflectance-based systems, as shown in Figures 5.2B-C. However, there were many, additional design reasons for selecting transillumination, including that it allowed for simpler optics that did not require signal-processing.

The decision to use transillumination drove key aspects of the overall design, ergonomics, and usage of the LightWand. Transillumination is much more effective when the patient's hand is in contact with the light source because more light is forced into the hand instead of escaping in the gap between the light source and the hand (decreasing the useful signal's power and increasing the background noise's power). As such, we required the patient to apply active, light pressure (sensed via a set of mechanical interlocks) with his or her hand on the laser/device tip for the device to be powered and enabled. Beyond improving the imaging results, this requirement for active pressure 1. provided a safety check that ensured the laser would only turn on with the hand present and 2. increased the friction between the device and the table-top for added stability. We designed the shape of the device to handle this active pressure, with the device's long, wide base supporting the patient's forearm and hand.

5.2.2.1.2 Laser Choice

We used a laser instead of an LED for space and cost reasons. As we had already developed the infrared vision system for HaemoBot, we were familiar with the available LED components; the space required for the LED, lens, heat sink, and drive circuitry made using an LED prohibitively space-consuming and expensive. As for inexpensive lasers at the required power, the main frequency options were 780nm, 808nm, and 850nm. We ultimately selected an 808nm, 150mW, and TTL-controlled laser (Aixiz PN: AH808-150T12) as it was the closest in wavelength to the 805nm operating-wavelength of the NVGs, provided enough power even for the thickest of hands, and allowed for PWM control of the laser brightness through its TTL-control circuitry.

5.2.2.1.3 Eye and Skin Safety

Although some research suggests that chronic skin exposure to high-power infrared light, specifically wavelengths in the IR-A category (760-1400nm), could cause irritation, inflammation, and accelerated aging of the skin [15] [30], short and infrequent exposure appears to have minimal risk other than from simple thermal heating. That is, as long as the laser does not become so hot that it burns the skin, the skin risk is minimal.

The only real safety concern for our assistive device was eye safety from exposure to the powerful, Class 3B infrared laser. We found empirically that we needed on the order of a 100mW laser to provide a clear image of the veins on the thickest portion of an adult hand in a room with fluorescent lighting, and such a laser could easily damage one's vision in a matter of seconds. One of the most dangerous aspects of infrared light is that the human eye is unable to see it even though it is easily damaged by it. There are three different mechanisms of viewing the NIR light, each of which is discussed below:

1. Exposure to light transmitted through the hand is not a risk as we measured the optical power of such light from our device to be on the order of 0.1mW, with standard laser pointers operating up to 5mW.
2. Light could reflect off of the palm and into one's eye, but we measured the optical power for this scenario at only 0.5-1.5mW, still under the limit for a standard laser pointer.
3. Direct-line-of-sight is the only real safety concern. The practitioner is not at risk from this exposure because the NIR-goggles block the light. However, we included both passive and active safety mechanisms to reduce this risk for the patient and others in the room, as discussed in the following sections.

5.2.2.1.3.1 Passive Safety Through Beam Divergence

Because lasers have extremely low divergence, they are able to inflict eye damage up close or across a room. An inexpensive, passive way of mitigating such a long-reaching risk is to add a line-generating lens after the focusing lens on the laser. The line-generating lens spreads the light into a line such that the optical power drops greatly with distance from the lens. Because the hand veins are so close to the lens, they experience very little power loss, whereas the patient's eyes are sufficiently far away that the laser's power at that distance has dropped to nearly-safe levels. As shown in Figure 5.4A, we empirically determined the optical-power-vs-distance curves for our laser with line-generating lenses at different divergence angles using the ThorLabs PM130D Digital Power and Energy Meter. Of the line-angles tested, the 28-degree line-generator was the only lens that retained sufficient power at the distance of the veins (between 15mm and 35mm, depending on the thickness of the hand) while still dropping to safer levels at greater distances.


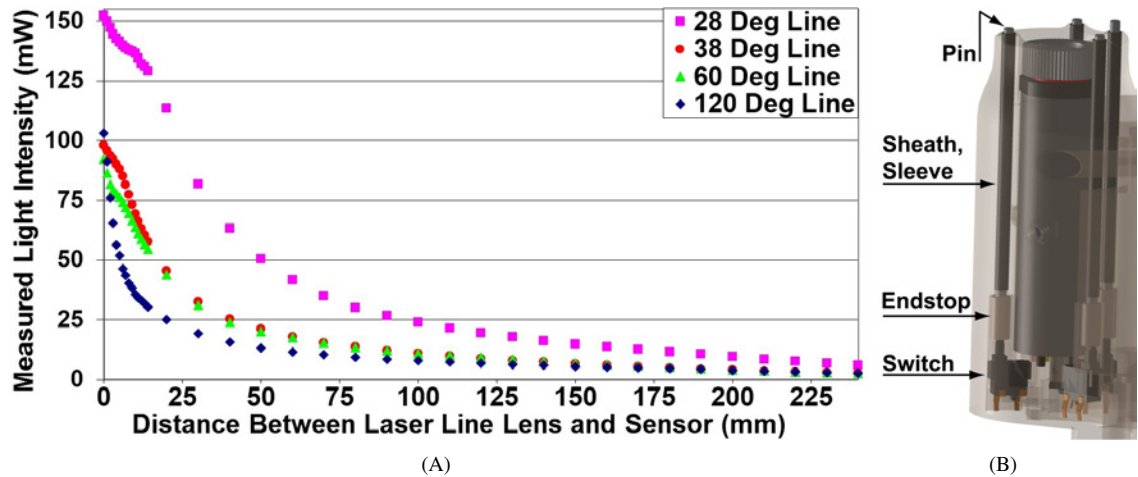


Figure 5.4: eye.Vease, Design for Laser Safety. A: Optical power vs distance for an 808nm, 150mw, TTL-controlled laser with line-generator lenses at various divergence angles. We achieve passive safety at distances near the patient's eyes through beam divergence. B-D: As an active safety precaution, the patient must depress four spring-loaded, mechanical safety interlocks to power the device and enable the laser, ensuring that the hand is completely covering the laser.

5.2.2.1.3.2 Active Safety Through Mechanical Interlocks

As shown in Figures 5.4B-D, we used a set of three mechanical contact switches as safety interlocks in an evenly-spaced ring around the laser's tip to ensure that the laser would only turn on if all three interlocks, and hence the laser, were in firm contact with the hand. Each interlock consisted of a recessed, spring-loaded temporary switch that preloaded a dowel-pin in a long, thin-walled stainless steel tube/sheath. The frictional loss between the dowel-pin and its sheath was very low, and we experienced no sticking of the switches. The geometric spacing, length, and spring-force of the switches meant the the laser would turn off as soon it approached the edge of the hand in any direction or if the hand was tilted such that only 1 or 2 switches were in contact (thereby allowing light to escape near the 3rd switch not in contact). A fourth safety interlock was wired to a power transistor that connected the batteries to the rest of the circuitry such that the batteries were not powering anything but the transistor's gate unless the fourth interlock was depressed. That is, unless

the laser was in firm contact with the hand, only a few nanoamperes of current were being drawn, thereby extending battery life considerably. We used mechanical instead of capacitive switches because our testing revealed that the capacitive sensors were prone to false-positives from touching near, but not on, the capacitive switch's main pad for our specific materials and spacing of the interlocks. Further, we found that overriding all four mechanical interlocks to expose the laser light in an unsafe manner was extremely difficult (requiring both hands), much more so than for capacitive switches that could be triggered by nearby contact.

5.2.2.1.4 Controlling the Laser's Brightness Via Slider-Pot-Set PWM

We needed a convenient method of modulating the laser's brightness to allow a single device model to work on all patients, from children with very thin hands that required little light to shine through them to obese adults with very thick hands that required the laser's full power. Further, variations in thickness within a single hand required changing the laser's brightness, with less light needed near the finger webbing and thinner edges of the hand and more light needed for the thickest part of the palm. Figures 5.5C-G show the need for brightness control in pediatric patients. To achieve this brightness control, we used a "Baby Orangutan B-328" microprocessor/motor-controller board (Pololu PN: 1220) to drive the laser's TTL-control-line via the board's motor-driving Pulse Width Modulation (PWM) capabilities. Compared to optomechanical methods of controlling the brightness (such as using variable cross-polarization), this PWM-control allowed us to command a true-zero ("off") value, to experience no efficiency-loss due to control, and to have an intuitive, linear, and mechanical control-input-vs-brightness-output.

We designed a slide potentiometer ("slider pot") system that set the PWM duty cycle, and hence the laser's brightness, based on the user's thumb movements. As shown in Figures 5.5A-B, a single thumb-adjustment-bar spanned the entire width of the device such that it could be moved from either side for left and right-handed users. The thumb-adjustment-bar was bolted to a pair of linear slide potentiometers that provided better support for the moment loading and redundant analog voltage signals for an electrical safety-check (by requiring both values to match closely) compared to using only a single slider pot. The thumb-adjustment-bar was inset within a wide depression in the VeroWhitePlus main structure to reduce the risk of being damaged in a fall or being bumped accidentally by the practitioner. We determined the desired stroke (20mm) and placement of the sliders through ergonomic testing of five adults with hand sizes ranging from small to large. After testing many different slider pots, we selected the ALPS RS20H111C009 "ultrathin" slider pot as it offered many desirable mounting, alignment, and performance characteristics.

5.2.2.2 Geometry and Ergonomic Design

The main structure of the LightWand was 46mm wide, 56.4mm tall, and 206mm long (from the back end to the laser's central axis). The long, skinny aspect ratio allowed the device to support the hand and most of an adult's forearm while keeping the arm low to the table so that the device would not tip over. The wide, flat base further prevented tipping and provided a large contact-area (and hence traction) between the silicone cover and a supporting table-top. While we had originally wanted for the practitioner's fingers to wrap around the entire device for a secure grip, fitting all components into the wand required a larger cross-sectional area than could be fit in the hand. However, as a compromise we included symmetric, integral ridges in the silicone



Figure 5.5: eye.Vease, Laser Brightness Control. We included a method of controlling the laser's brightness because different patients and regions within the same hand require different amounts of light to show the veins properly. A-B: Brightness control slider/thumb-adjustment-bar. C-F: Images recorded by holding a smartphone's camera to the NIR-goggles' eyepiece. The in-person view is much superior. C-D: 5-year-old, Caucasian male. E-G: 7-year-old, Caucasian female. D, G: Excessive laser brightness causes saturation/pixel-blooming in the NIR-goggles, thereby hiding the veins. C, E, and F: A lower laser brightness allows us to see the veins.

cover, as shown in Figure 5.7A, that provided a secure hold for the practitioner's fingers. The silicone ridges and thumb-adjustment-bar for controlling the laser's brightness were placed symmetrically for ambidextrous operation (left or right handed), as well as holding the device with the palm above or below the device's bottom surface. Figure 5.6A shows an example of the "below grip", and Figure 5.7A shows an example of the "above grip". We positioned the laser at one end of the device so that the patient's fingers would bend down over the end in a natural insertion pose that promotes skin traction. As shown in Figures 5.4B-C and 5.7C-D, the tip of the laser protruded somewhat from the main, top surface that supported the forearm so that the tip could fit into anatomical spaces other than the hand that were tighter or curved, including the ankle and arch of the foot. This protrusion, which included the laser and safety interlocks, had a diameter of 20.4mm and rose to a height of 11.6mm above the main, top surface.

5.2.2.3 3D-Printing the Main Structure

Our main goals for the LightWand's structure were to minimize the number of parts, need for machining, weight, and fabrication cost. As such, the main mechanical structure consisted of a 3D-printed-plastic central body, a 3D-printed-plastic bottom optics cover, a 3D-printed-plastic tape dispenser core, and a top cover laser-cut from 410 stainless steel (chosen because it is magnetic). There were also a few small, internal parts that were laser-cut from stainless steel. Our desire to minimize the part count meant that 1. those parts remaining in the design had extremely complex geometries that could not be machined, and 2. we needed to rely heavily on the use of light, reversible press-fits between components and the main body of the device, such as rectangular LED profiles in rectangular holes. We used a Stratasys Objet24 printer with VeroWhitePlus resin to provide the required precision and material strength/ductility at a reasonable cost.

5.2.2.4 Silicone Sleeve

We encased the device in a 1-piece, nearly all-encompassing silicone sleeve to gain the following benefits:

1. The sleeve served as a protective case in the event of droppage or collision.
2. The sleeve provided a non-porous, durable, and removable surface that decreased the adhesion of bacteria and was easy to sanitize without damaging it.
3. The sleeve provided greater traction with flat mounting surfaces and the practitioner's hand. During table-top usage, contact with the flat bottom surface ensured that the device would not slip. Two symmetric, integral ridges molded into either side of the sleeve, as shown in Figure 5.7A, provided a positive lock with the practitioner's fingers.
4. The sleeve provided an integral plug/flap, as shown in Figures 5.8E-F, that fit into the opening of the battery compartment to seal it against the elements.
5. The sleeve trapped the alcohol bottle in its compartment, as shown in Figure 5.8C, because the hole in the sleeve where the bottle inserted was 1mm smaller than the diameter of the bottle.

As an alternative to using silicone, given the cost and complexity of the molding process, we tried 3D-printing the sleeve out of a flexible material, using Objet's TangoBlackPlus980 material because it was the best option. Unfortunately, the test cover was quite inflexible, ripped easily, had a pronounced texture that trapped contaminants and made cleaning difficult, and degraded in strength considerably after a few weeks.



(A)



(B)

Figure 5.6: eye.V.ease in Use on Ankles and Feet. A: The practitioner is holding the device using a “below grip” while finding veins in the ankle. B: The device is clamped to the bed’s footboard and angled into the bottom of the foot.




Figure 5.7: eye.V.ease, Design Details, Picture Set 1. A: Silicone sleeve with symmetric, integral ridges/finger-holds for ambidextrous grips (left, right, above, and below). The quick-attachment plate that connects the device to its clamp can be seen at the bottom of the device. B: Silicone cap that protects the laser and safety interlocks when the device is not in use. C: Traction/restraint bar stowed. D: Traction/restraint bar extended for use. E: The tourniquet attaches magnetically to the top, stainless steel plate and can be moved closer or further from the laser to accommodate the different arm lengths of children and adults. F: The tourniquet can be applied and released with one hand.

5.2.2.5 Silicone Cap

We fabricated a silicone cap, as shown in Figure 5.7B, that fit over the laser and safety interlocks to protect them from mechanical damage, keep out dirt, and prevent the device from being powered accidentally while being carried or stored. The free end of the cap contained a small neodymium magnet that was attracted to the device's top stainless steel plate and kept the cap from falling off accidentally.

5.2.2.6 Device Mounting

We intended for our device to be supported (held) mechanically in several different modes:

1. Initial assessment of the potential insertion sites could be conducted with the practitioner holding the device by hand, as shown in Figure 5.6A. We included a wrist strap at the end of the device for preventing droppage by the practitioner.
2. Actual insertion would be conducted with the device sitting on a flat surface, such as a table, or clamped rigidly to a table or part of the hospital bed to provide the required stability. More specifically,
 - 2.1. The device could sit flat on an over-the-hospital-bed table, typically used to serve a meal in bed, with the patient's hand placed on the table. The device's wide, flat, and silicone-covered bottom surface would provide stable, high-traction contact with the table to prevent it from being swiped off accidentally. The tables we tested were adequately stable for IV insertion despite their wheels.
 - 2.2. The device could also be clamped to a rigid mounting surface for stability and angled arbitrarily to obtain proper contact with the insertion site while keeping the patient comfortable. We recessed a quick-release tripod-plate (Strobframe PN: 300-QRC) entirely within the bottom of the device and connected this plate to a highly-dexterous clamp (Pedco Ultraclamp 1.5inch model). The clamp is compact, lightweight, stiff, and has several ball joints that allow for the device to be angled in any configuration needed. The quick-release plate allows for the LightWand to be separated from its clamp effortlessly, whether switching to handheld usage temporarily or storing it compactly. It also allows for mounting on any other clamp, tripod, or support arm with a standard 1/4 – 20 thread. We tested this clamp on a variety of hospital furniture and were able to find sufficiently stable clamping sites for hand/arm insertions on the hospital bed side railing (as shown in Figure 5.1), over-the-hospital-bed table, and wheelchair arm/tubing. The hospital bed footboard provides a clamping surface for foot and ankle insertions, as shown in Figure 5.6B.

5.2.2.7 Traction and Restraint Bar

Applying skin traction is necessary to minimize skin and vein movement during insertion. A common approach is to bend the fingers downward, which naturally tightens the dorsal skin, and then pull the dorsal skin further towards the finger tips using the practitioner's thumb. Our device uses a "traction bar" and medical tape to maintain such traction without the practitioner's hand actively applying it during the insertion.

As shown in Figures 5.7B-C, the traction bar consists of a delrin end-piece with two parallel shafts that slide into the main body of the LightWand, and the bar can stow compactly against the end of the device or extend outward by up to 40mm. The silicone cap is attached to the end-piece so that it does not become separated from the LightWand and lost, and it folds downward and away from the fingers during usage.

When the patient's fingers are bent downward and curled around the traction bar, which is 13mm below the laser-palm interface, the natural traction pose is achieved, as shown at the bottom of Figure 5.1. This curling action further stabilizes the hand and also helps to maintain the required pressure between the hand and the laser for the safety interlocks to be engaged. To apply further traction, a piece of medical tape is placed with one end near the desired insertion site, pulled towards the finger tips to stretch the skin, and then laid down along the fingers with the other end near the finger tips. A piece of 1/2-inch 3M Transpore tape can maintain considerable traction throughout the entire insertion if applied firmly.

A further usage of the traction bar is as a finger/hand restraint for noncompliant infants and children. Applying a few wraps of Transpore tape around the fingers and traction bar traps the fingers and greatly stabilizes the hand against struggling. Such restraint is common in pediatric units, wherein a foam-covered "armboard" is often taped in place across the antecubital fossa to keep the infant's or child's arm straight.

5.2.2.8 Tourniquet

We included a tourniquet as an integral part of the device, and specifically this model (Prestige Medical PN: 4350-roy), because it:

1. Removes the need for the practitioner to obtain a tourniquet prior to insertion,
2. Keeps the device attached to the arm and safe from being dropped or swiped off of a table-top,
3. Removes the need for the practitioner's non-dominant hand, or an additional practitioner, to restrain the arm of a noncompliant infant or child (especially when also restraining the fingers to the traction bar and clamping the LightWand to a secure base).
4. Reduces tourniquet tightening, loosening, and releasing from slow, 2-handed operations to quick, 1-handed operations, as shown in Figure 5.7F, and
5. Uses a latex-free material that avoids allergies and can be sterilized via autoclaving.

The tourniquet is attached magnetically to the top stainless steel plate, and moderate effort can slide it to any distance along the device to fit the particular forearm length of a patient, as shown in Figure 5.7E. Due to the shortness of the device, the tourniquet is applied to the forearm instead of the typical placement on the upper arm, but there is no discernible difference in its ability to engorge the veins (its primary purpose). The tourniquet is positioned close to the laser for infants, at the middle of the device for older children, and at the far, back end of the device for adults. The magnetic attachment is strong enough that the tourniquet cannot be removed without a special tool, preventing accidental detachment and droppage.

5.2.2.9 Tape Dispenser

Many practitioners use tape to secure the catheter to the insertion site. The tape can be used as a temporary restraint until a more secure method, such as a StatLock "stabilization device" can replace the tape, or the tape can be left as the long-term restraint. We included the tape dispenser as an integral part of the device to 1. reduce the number of items that the practitioner had to gather and 2. allow the practitioner to rip off a piece of tape with one hand, as shown in Figure 5.8A. Beyond securing the catheter, the tape is used for the additional, non-standard uses (discussed in Section 5.2.2.7) of 1. applying traction to the dorsal skin, as shown in the bottom portion of Figure 5.1, and 2. restraining the fingers of uncooperative infants/children.




Figure 5.8: eye.Vease, Design Details, Picture Set 2. A-B: Integral tape dispenser for convenience and 1-handed tape ripping. The tape is used for securing the catheter, applying passive traction to the skin, and restraining the fingers of uncooperative infants/children. The tape core attaches to the top stainless steel plate via an embedded neodymium magnet and is positioned by sliding sideways into the device via a rectangular slot. The end of the tape sticks to the serrated tape cutter so that the practitioner can find it easily. C-D: Integral, easily-replaceable alcohol spray bottle for convenience and 1-handed cleaning of the insertion site. The silicone sleeve keeps the bottle in place. E-F: Battery compartment that houses two 3.6V, CR123A rechargeable batteries. The silicone sleeve provides an integral plug/flap that fits into the opening of the battery compartment to seal it against the elements. A sliding stainless steel door with a leaf-spring contact serves as the negative electrical contact and traps the batteries inside of the compartment. To change the batteries, the silicone flap and sliding door must be lifted up and the device tilted vertically so that the batteries fall out.

As shown in Figure 5.8B, our tape dispenser consisted of a removable core, a roll of 1/2-inch-wide 3M Transpore Tape (popular amongst practitioners), and a serrated tape cutter. The core attached to the top stainless steel plate via an embedded neodymium magnet and was positioned by sliding sideways into the device via a rectangular slot. This allowed for a new roll of tape to be installed quickly and without need for adjustments. The end of the tape stuck to the serrated tape cutter so that the practitioner could find it easily.

5.2.2.10 Alcohol Dispenser

We included a small spray bottle of 62% ethyl alcohol in the back end of the device for cleaning/preparing the insertion site, as shown in Figures 5.8C-D. Inclusion of the ready-to-use alcohol bottle removed the need for the practitioner to obtain an alcohol pad prior to insertion or to struggle with opening the pad's packaging with gloved hands. We used an existing hand-sanitizing product that came in pre-filled, disposable bottles so that the entire bottle could be swapped instead of refilling it. We found that the CVS "Hand Sanitizing Spray Pen" (SKU 428505) was the best fit for our device. It can deliver approximately 70 sprays, and replacing a spent bottle with a fresh one can be accomplished with one hand in a matter of seconds.

The chamber in the main housing that holds the spray bottle is closed internally so that leakage from the bottle would not enter the device's main body and short electronics. Similarly, we placed the bottle's nozzle in the back end of the device to keep the nozzle from accidentally dripping onto the laser or safety interlocks. The bottle is far enough below the patient's arm to prevent accidental contact and has such a small aperture that airborne contamination of the alcohol is not a concern. We molded an undersized hole into the silicone sleeve that grabbed the bottle and prevented it from rotating or falling out, and this simple restraint method is what allowed for quick, 1-handed replacement of the bottle, as shown in Figure 5.8C.

5.2.2.11 Battery

Our Baby Orangutan controller board used an onboard 5V regulator, and we powered all electronics from this regulated source, including the laser. In total, all electronics drew a maximum of 0.3A (with the laser at maximum brightness). Therefore, we required a battery with a voltage (or combined voltage from multiple batteries) that could provide over 5V while sourcing at least 0.3A. Further, the aspect ratio of our device lent itself to a long, cylindrical battery that could be stacked in series if multiple batteries were required to achieve the desired voltage. Lastly, we desired a common battery model that could be replaced easily by the end consumer instead of a proprietary or uncommon model that would be difficult to obtain. We preferred the idea of using a removable battery instead of a permanently-installed battery so that the user could change the batteries without needing to wait for the device to charge in case the device was needed and a backup unavailable. In the end, the best fit for our requirements was using two CR123A batteries in series (specifically the Ultrafire 16340 rechargeable CR123A battery, which provided 1000mAh at 3.6V). Such batteries are commonly found in cameras, high-end flashlights, and some handheld medical devices. We used a rechargeable version to reduce the number of batteries that would need to be bought by the consumer.

Our battery compartment was 3D-printed directly into the main body of the LightWand. A stationary, positive electrical contact was located at the far end of the chamber, and a sliding, stainless steel door with a leaf-spring contact served as 1. the negative electrical contact and 2. the mechanical restraint that trapped

the batteries inside the compartment, as shown in Figures 5.8E-F. We designed the geometry of the battery compartment's distal-end/positive-contact to ensure that the batteries could not be installed backwards. To change the batteries, the silicone plug/flap and sliding door were pulled up and the device held vertically for the batteries to fall out (sometimes with the assistance of gentle shaking).

As discussed in Section 5.2.2.1.3.2, the patient had to depress four mechanical interlocks simultaneously to power the device and enable laser usage. As this required the patient to apply pressure actively, the device was only drawing power during active insertion and shut off automatically afterwards, thereby extending the battery life considerably. These batteries provided roughly 5 hours of nearly-continuous device usage before needing to be recharged/replaced. In clinical usage, our device's batteries would likely last for 2-3 days because there would be significant breaks between insertions on different patients.

5.2.3 Cost Estimate

The total raw materials for our device cost \$1,339 USD. The individual component costs were:

- \$500 for the NVGs,
- \$164 for the 3D-printing resin,
- \$103 for the laser-cutting stainless steel,
- \$100 for the lens caps,
- \$97 for the specialty screws,
- \$60 for the laser,
- \$50 for the quick-release system,
- \$46 for the head band,
- \$26 for the longpass filters,
- \$23 for the Pedco Ultraclamp, and
- below \$20 each for the remaining 24 line-items.

This estimate does not include usage costs for the laser-cutter and 3D-printer, and such costs could be significant in commercial fabrication.

5.3 Discussion and Thesis Contributions

In this chapter, we developed a simple assistive device for IV insertion, as shown in Figure 5.1, that embodies three main, novel contributions:

1. Provides stereo, NIR vein-visualization that is more anatomically accurate than what is provided by commercially-available devices (no false-vein-detections, as seen in Figures 5.2B-C).
2. Makes IV insertion easier by reducing active insertion from a 2-handed to a 1-handed procedure.
3. Allows theoretically for self-insertion by an individual (dominant hand inserting on the other hand).

Our wide-ranging testing and analysis of commercially-available vein-finding devices is the first of its kind. Although there have been clinical studies of several, individual devices, we are the first researchers to compare and contrast the majority of available products, especially from a technical perspective. Future work includes clinical validation of our device's benefits for both patients and practitioners.

Chapter 6: Development of Anatomically Correct, Hyper-Realistic Hand Vein Phantoms

6.1 Motivation and Overview

To test the operation of HaemoBot, we needed to insert into an inanimate or exanimate hand vein, or “insertion target”. There were two levels of robot testing required: testing as the robot was being designed, fabricated, and coded and testing of the final robot in its completed state. Testing during initial development was prone to mistakes as the code was still being developed and debugged, and these failed tests frequently damaged the insertion target. Thus we required an inexpensive, disposable insertion target that was only realistic enough to provide for simple testing. In its completed state, HaemoBot performed reliably and safely enough to reduce the risk to more expensive, realistic insertion targets that were harder to replace but allowed for more advanced, complete testing.

After surveying many commercial, off-the-shelf hand/vein phantoms and untraditional, but vein-like, products, we found that Twister/“balloon-animal” balloons (specifically the 160Q-size) were ideal for initial code development as they were very inexpensive, easy-to-replace, and fairly similar in size and force characteristics to single, straight sections of vein. However, none of the phantoms we tested were adequate for the more advanced testing on a completed robot that required bifurcations, anatomically-accurate sizes, and hand macrostructure. Our solution was to use the NIR vein images and point-clouds generated by our vision system in Chapter 3 to serve as a model for creating our own, custom phantom that was a near-exact replica of the author’s own hand and hand veins, as shown in Figure 6.1. Creating this phantom required us to develop novel fabrication techniques, such as tracing the 2D, NIR vein patterns in water-soluble wax as a mold-core for dipping-latex, and expand upon known ones, as seen in our use of “external”-rotocasting. The following sections describe the main options for insertion targets, the motivation behind our final selections, and a discussion of the fabrication process and performance of our custom phantom.

6.2 Potential Insertion Targets

6.2.1 Live Human Hand

A live human hand would have provided the most realistic testing possible. However, such testing would be dangerous and administratively-prohibited by FDA and IRB guidelines until HaemoBot’s safe and reliable performance were established using progressively-realistic, inanimate or exanimate insertion targets (from latex tubes to cadaver human/animal veins to live animal veins). Reaching the benchmarks for live-human-testing to be allowed would be a multi-year process.



(A)



(B)

Figure 6.1: Our Custom Phantom: Latex Bifurcation on Silicone Hand. We used life-casting to fabricate the hand as an exact copy of the author's anatomy (down to the finger-prints). We used latex-dipping and "external"-rotocasting of soluble-wax cores that we manually spliced/traced against 2D, NIR vein images from the vision system to create latex venous bifurcations that matched the author's veins almost exactly.

6.2.2 Deceased Human Hand (Cadaver)

A human cadaver would be the most realistic insertion target after live human testing. However, even though cadaver testing is exempt from FDA or IRB approval, it is still time-consuming and expensive. We established a certified cadaver storage/testing room and a partnership with the willed-body program at University of California at San Francisco (UCSF) to provide cadaver arms but ultimately could not proceed with the experiments due to insufficient time and funds, with specimens costing \$600 for an arm up to the shoulder.

It should be noted that the majority of the cadaver specimens were from elderly individuals and that obtaining specimens from pediatric, young-adult, or even middle-aged patients is extremely rare. Testing on such a narrow demographic would be limited in its usefulness as the physical properties of the hands and veins of elderly patients are very different from those of younger patients. Further, while a human cadaver specimen has the correct anatomy of a live human hand, it does not arrive in a state that is ready for insertion. Specimens must be obtained as soon as possible after the rigor mortis has ceased and without first undergoing embalming or freezing, which change the tissue's mechanical properties significantly. The veins have residual, somewhat-coagulated blood and are open where the arm was severed. Before the specimen can be used, the venous system must be flushed clean and the vein ends ligatured shut. The veins then have to be refilled and properly pressurized with fresh animal or artificial blood that has infrared-absorption properties similar to live human blood to work with our imaging system.

6.2.3 Live Animal Vein

Testing on a live animal vein could serve as a reasonably accurate test of some aspects of our system but would be difficult administratively and technically. Such testing requires certification from Stanford's Administrative Panel on Laboratory Animal Care (APLAC), approved facilities for the animals, and major funding of the process. Further, as HaemoBot was designed specifically for human anatomy, retrofitting it for usage with a live animal would be extremely difficult. The first problem is that an animal with similarly-sized veins to those of an adult human would be too large to fit in the central insertion space of HaemoBot. Further, we would need to find a vein that could be seen with our NIR transillumination imaging system. Goats were suggested to us by other researchers familiar with live animal testing because they have large, easily-accessible veins in long, skinny legs that could potentially be transilluminated. However, a goat's size could not be accommodated in the robot, and it appears unlikely that the goat's leg would contain a venous bifurcation for our vein detection algorithm to work properly.

6.2.4 Deceased Animal Vein

Testing on a deceased animal vein is an inexpensive, reasonably straight-forward idea, but it would be a significant time investment. It requires essentially the same approval from Stanford's Environmental Health and Safety (EH&S) office as is required for human cadaver work. However, because the macro-anatomy of the vein and surrounding tissues is so different in animals from that of a human hand, this is not a particularly realistic option compared to testing on a human cadaver. If we could find an animal that had a single, straight section of vein with comparable size (outer diameter and wall thickness) to a human hand vein, we could



Figure 6.2: Commercial Off-The-Shelf IV Phantoms. A: Twister (“balloon-animal”) 160Q-size balloon. B: Latex tubing. C-D: Simulaids IV Hand Simulator 147R next to our custom phantom.

test the mechanics of the vein insertion without the vision system or the behaviors that arise from the vein residing in the anatomical context of a human hand (such as rolling beneath the skin). To test the vision system, we would need to locate a comparably-sized venous bifurcation, and it is unclear whether this would be a possibility. Even if feasible, it would take a large amount of time to dissect an animal, remove the vein, and set it up for insertion (ligaturing, restraining, aligning, and pressurizing it).

6.2.5 Artificial Human Hand/Vein (Phantom)

Using a non-biological, artificial human hand/vein, or phantom, is the least-expensive, cleanest, fastest, and administratively-easiest option of all the insertion targets. There are three main options for commercial off-the-shelf (standard) vein phantoms, and we will now discuss the benefits and detractions of each below.

1. We found that twister/“balloon-animal” balloons, particularly the Qualtex-brand 160Q-size, were very inexpensive, easy-to-use vein phantoms for initial development and debugging. They provided a similar diameter, wall-thickness, and force “pop” profile to real veins. Figure 6.2A shows a successful automatic insertion into a twister balloon. If just the first inch of the balloon was inflated to a larger diameter and the remaining length left at the near-original/uninflated diameter, the fully-inflated section

would serve as an air reservoir for the uninflated section so that minor air-leakage during insertion did not cause the vein to lose patency. The balloons could be filled with either air or artificial blood and still make use of this reservoir approach for maintaining patency. Filling the balloon with artificial blood yielded a more realistic stiffness/force profile and allowed for blood-flash detection. However, for tests that required less realism and no blood-flash, air was cleaner and allowed for faster preparation. The balloons frequently roll sideways under the pressure of the IV needle, similarly to real veins, and this can either be allowed for advanced testing or restrained for simpler testing by embedding the balloon in a shallow groove. In our tests, we placed the balloons halfway in a soft foam groove. The main drawback of using “balloon-animal” balloons, assuming their realism is appropriate for a particular test, is that a single IV stick prevents them from being reinflated without leaking, making them single-use.

2. A more robust, albeit more difficult-to-use, alternative to the “balloon-animal” balloons is latex tubing (McMaster-Carr PN: 5234K24, 4.76mm OD x 3.18mm ID x 0.79mm Wall Thickness). Figure 6.2B shows a successful automatic insertion into such tubing. This tubing has several attractive properties, including that:
 - 2.1. the tubing dimensions are very similar to those of the author’s wrist-veins,
 - 2.2. the tubing can be used for many IV insertions, even when filled with artificial blood, because they are fairly self-sealing from needle sticks,
 - 2.3. they do not require pressurized air or fluid to remain patent, and
 - 2.4. their semi-transparency allows for the user to see the IV needle inside of the vein, which can provide useful insights into the vein/needle interaction.

However, the tubing is more difficult to use than the twister balloons because 1. it does not easily allow for passive pressurization through an integral reservoir of fluid (requiring an external method of pressurization), and 2. its ends are not as easily tied/restrained.

3. Many companies sell hand/arm IV and phlebotomy training-phantoms with replaceable skin, latex-tubing for veins, and some manner of filling the veins with artificial blood (colorized water). Our initial proof-of-concept movie on the first HaemoBot prototype used such a phantom (Simulaids IV Hand Simulator 147R, \$254.50). Figures 6.2C-D show the Simulaids phantom next to our custom phantom for comparison. However, this phantom suffered from many problems, including that it was:
 - 3.1. unrealistically-large (25% larger, overall than the author’s large hands),
 - 3.2. had skin that was overly thick/tough (1.1-1.4mm thick MB1500 vinyl),
 - 3.3. had veins that were overly large/thick (6-7mm OD x 3mm ID x 1.5-2mm Wall), and
 - 3.4. contained no real bifurcations (the hand approximated a bifurcation by placing two, separate tubes next to each other, but the veins were not connected by material that could be inserted into.)

We attempted to insert the same 20-gauge IV needle that we used in the clinical experiment of Chapter 2 into the Simulaids phantom, but the needle was bent and the catheter-tip damaged beyond use on every attempt. To penetrate the phantom without damaging the needle/catheter, we had to switch to a larger-bore, stronger 12-gauge needle which is rarely used clinically. We conjecture that this phantom’s skin

and veins were made thick and tough to facilitate manufacturing and increase its insertion-life over a thinner, more delicate version. For instance, a thinner skin would be more likely to rip when installing it on the hand and would not survive as many punctures. Also, thinner latex tubing would not be as self-sealing as thicker tubing after a puncture so that the artificial blood would leak.

A common detraction to all of these options is that they lacked venous bifurcations, which our vein-detection algorithm requires to operate and our needle-insertion mechanical tests require for better realism. Overall, we found these phantoms essentially unuseable for realistic testing with our robot.

There are companies, such as Blue Phantom, that will custom-build a hand/vein phantom to the customer's material and geometric specifications, providing a more life-like and realistic phantom than the cheaper, standardized models. We submitted a quote for a transparent hand with an embedded, nonreplaceable venous bifurcation and no skin, and the company quoted a lead time of 3 weeks and a cost of \$4,000 USD. Ultimately, the cost was too high for our budget given that the phantom likely would be functional for only a small number of insertions, had no skin, and was only representative of a single hand/vein anatomy.

Having found the commercially-available, non-biological phantoms to be inadequate for advanced testing, we concluded that we must fabricate our own phantom.

6.3 Fabrication of Our Custom Phantom

Our hand phantom consists of three main elements: a hand structure, veins, and a skin (not installed in the prototype shown in Figure 6.1). While many of the fabrication processes we used are known in other industries, particularly practical special effects, we have contributed new processes, expanded on some of the known ones, and combined them all uniquely towards this purpose of making realistic artificial hands and veins that simulate IV insertion correctly.

Due to time and financial constraints, we were unable to complete a useable skin for the first phantom prototype shown in Figure 6.1. However, we developed a fabrication-process for creating a single-piece, anatomically-detailed, and custom-molded skin complete with real human hairs punched into its thickness that we hope to refine as part of our future work.

6.3.1 Hand Fabrication

Fabricating the main hand structure is the easiest and most-common process, using standard life-casting and moldmaking techniques. We were able to fabricate the hand shown in Figure 6.1 in a matter of 45 minutes of active labor and 1 hour of waiting.

We began by having a volunteer insert his or her shaved/mold-released arm up to the shoulder in a large container of Smooth-On AlgaSafe alginate. After 15 minutes, the hand was removed from the "set" AlgaSafe, and the casting material was poured into the negative copy of the hand/arm. If a fast-curing silicone, exact replica of the hand is being fabricated without any additional features, such as routing-channels for tubing, the completed hand can be removed after 1 hour. However, if a slower-curing silicone or additional mechanical features are required, then HydroCal white gypsum must be poured into the AlgaSafe to produce an intermediate, positive copy. After setting for 30 minutes, the finished HydroCal copy can be removed by



Figure 6.3: Silicone Venous Network Patterned On Author’s Anatomy. Although we were able to fabricate the venous network in silicone, the force “pop” was not close enough to that of a real vein.

tearing the soft AlgaSafe from it.

The HydroCal hand can be used directly to make a multi-part, silicone female-mold or can be modified first to add geometric features to the hand (whether building positive features via a sulphur-free clay or cutting negative features via a diamond spherical burr on a high-speed rotary tool). One such additional feature could be cutting a shallow groove that embeds the vein partially in the hand structure to ensure proper placement and to help reduce rolling of the vein to a manageable, more-realistic level.

We chose platinum-based silicone as the hand material because of its durability, minimal-shrinkage, chemical inertness, ability to capture minute anatomical details (including fingerprints, pores, and hair follicles), and life-like mechanical properties. Adding Smooth-On “Slacker” (a silicone-modifier) in a 2:1 mass-ratio to the silicone increased the realism of the hand’s feel. We also colorized the silicone lightly with a pre-made human skin tone. Finding that the fast-curing silicones cured too quickly to allow for proper degassing, we obtained significantly better results using the slow-curing or regular-curing silicones, degassing them at our leisure, and then heat-curing them to completion in 45 minutes.

6.3.2 Vein Fabrication

Our requirements for the artificial vein were that it:

1. be a venous bifurcation that closely-matched the author’s actual anatomy,
2. be seamless and smooth internally and externally with the same material used for the entire vein, and
3. have the approximately-correct force “pop” when inserted into by an IV needle.

The second requirement precluded us from splicing together segments of pre-made tubing with an adhesive because the internal and external edges would have been sharp, and the epoxy used to join the segments would have had different material-properties from the tubing itself. We considered molding the vein in a traditional 3-part mold (2 outer halves and an inner core), but even small amounts of deflection in the parts made obtaining consistent wall-thickness throughout the vein difficult. Using precision-machined parts would have solved this problem, but the machining would have been prohibitively expensive or impossible given the complicated geometries of the veins. As such, we decided that dipping/coating a removeable core (the inner




Figure 6.4: Force vs Time Profile For Automatic Insertion Into Our Custom, Latex Vein Phantom.

space of the vein) into a material that could then be cured in place around the core was the only solution. Once the vein material was cured, the core would be removed, leaving a hollow, seamless venous bifurcation.

6.3.2.1 Vein Material

In practical special effects, silicone and latex are the chief materials for making artificial body parts, and the standard processes that we implemented, modified, and used as inspiration for our new processes were all developed for those materials. Of the available subtypes of latex, dipping-latex, a less-commonly used, liquid alternative to foam-latex, was the most appropriate latex-subtype for making veins. The main material choice was ultimately between platinum-based silicone and dipping-latex, henceforth referred to simply as latex.

Silicone has many attractive mechanical, chemical, optical, and ease-of-fabrication properties, and we were able to make silicone venous bifurcations patterned on the author’s anatomy with good results, as shown in Figure 6.3. However, we found that the insertion force “pop” when inserting IV needles into silicone did not feel realistic, regardless of the thickness or durometer of the silicone. Latex veins felt much closer to the insertion force “pop” we desired, which is why we opted in the end to use dipping-latex for the vein material. Figure 6.4 shows the force-vs-time profile for an automatic insertion into our custom vein phantom. The shape of the force profile, the presence of identifiable needle-“pop” and catheter-“pop”, and the magnitude of the catheter-“pop” (1 N for our phantom compared to the clinical average of 1.462 ± 0.497 N), are all quite similar to the force profiles that we measured clinically in Chapter 2. It should be noted that this force profile for our custom phantom is nearly identical to those measured for the Twister/“balloon-animal” balloons and latex tubing, likely because they are all made of similar latex material. An added benefit of using latex instead of silicone is that the latex is far better at “self-sealing” after an IV puncture so that the artificial blood does not leak from the vein. It is possible that there exists a different formulation of silicone that would produce a more accurate “pop” and have better self-sealing properties, in which case we would switch to silicone for the vein material to obtain its many other benefits.

Often used to make surgical gloves, dipping-latex is specifically formulated to cling to a mold-core that is dipped into it and coat that core evenly with a thin layer. To prevent micro-holes from penetrating the entire thickness of the finished latex part, it is best to build up the full wall thickness through many thin layers instead of using fewer, thicker layers. The layer-thickness added per dip relates to the “total latex content” of the dipping-solution and the chemical formulation of the latex coagulant used to prepare the mold-core.

Coagulant is a specially-formulated liquid in which a mold-core is dipped prior to dipping the mold-core in the actual dipping-latex (allowing the coagulant to dry first), and this dried-layer of coagulant improves the ability of the latex to cling to the mold-core and to cling evenly. Dipping-latex is already “cured” chemically and simply needs to dry to enter its final, solid state, meaning that it can only be used to form layers up to a few millimeters thick and cannot be used for molds wherein the geometry would restrict significant air flow (e.g. internal surfaces with small inlets like a soda bottle).

6.3.2.2 Core Material

We needed a core material that could be injected into a core-making mold as a liquid, solidified for use as the core during the dipping process, and re-liquified or dissolved for removal from the completed vein.

We found a water-soluble wax (Freeman Sol-U-Carv Wax, Gesswein PN: 263-1051) that dissolved quickly in water, and adding heat and 10% citric acid (by mass) accelerated the process so that complete wax-removal took only a few minutes. The main downsides of using the water-soluble wax were that:

1. The considerable water content in the dipping-latex reacted with the core immediately and dissolved the core before the latex could dry/solidify. We solved this problem by sealing the soluble-wax-core with three light coats of Krylon Crystal Clear Gloss, which is an acrylic lacquer that provides a thin, water-tight boundary. An added bonus of using the paint sealant is that it smoothed any irregularities in the core’s surface, providing a perfectly smooth vein interior.
2. The low melting point of the wax meant that we could not heat the coagulant solution that promoted latex adhesion to the core, thereby reducing its effectiveness considerably.

A final alternative that we researched but have not yet tried as a core material is low melting-temperature metals, typically some alloy of Bismuth, Lead, Indium, Tin, and Cadmium. These metals can be bought for a specific melting temperature (McMaster-Carr sells them for 117, 134, 136, and 142 degrees Fahrenheit). The advantages of this as a core material over soluble-wax are that 1. no sealing would be required, 2. the core would be strong enough to withstand fine sandblasting to promote latex adhesion, and 3. the materials could allow for greater heating of the coagulant solution and better latex-adhesion.

6.3.2.3 Making the Vein-Core Geometry

We formed the soluble wax into the desired vein-core geometry using two, different processes, each of which had its own advantages:

1. injecting liquid wax into a 3D-printed mold designed from real 3D point-cloud, NIR vein-data and
2. manually bending and splicing straight, cylindrical sections of solid wax to trace the vein pattern over a 2D, real-size, NIR image of the veins.

6.3.2.3.1 Molding the Vein-Core From 3D Point-Cloud, NIR Vein-Data

The most anatomically-correct method of creating the vein-core is to 3D-print the mold based on real 3D point-cloud data obtained through our NIR imaging system. Using the 3D point-cloud data, we can create a 3D spline of the vein path and the outer diameter tapering profile in Solidworks, create a swept/blended extrusion along those paths, split the extrusion into two parts so that we can extract the core without breaking it or the mold, add registration keys for keeping the two halves precisely located, and add an injection site

for the wax. It should be noted that there are many practical, design issues that must be addressed to release the wax core from the mold without bending or breaking it, including providing a smooth surface finish on the mold, cooling one half of the mold much faster than the other half to allow for safe mold-opening, and proper cooling of the mold assembly to promote self-release of the wax-core without making it brittle. Further, careful design of the insulation for the wax-injection system is needed to keep the wax from cooling too rapidly and blocking the narrow mold cavities, syringe nozzle, or tubing. Figure 6.5A shows a wax core made using a 3D-printed mold, but in this case the vein geometry was designed directly in Solidworks as a controlled test instead of taken from a 3D point-cloud of real veins. While we were able to fabricate a proof-of-concept vein-core using this method, we ultimately decided to use a 2D-approximation method that was appreciably simpler and faster.

6.3.2.3.2 Splicing/Tracing the Vein-Core Against a 2D, NIR Vein Image

We also developed a lower-cost, faster method of creating a vein-core that approximates the vein geometry reasonably well. If we treat the vein geometry as 2D instead of 3D, we can use either of the 2D stereo images as the geometry template instead of the 3D point-cloud. Although the accuracy of the vein map suffers a small amount from this approximation, we can fabricate the veins much more quickly and inexpensively.


We first calibrated a normal printer such that we could print images that were true-size. To do this, we placed a calibration grid on the hand that appeared in our infrared images to provide scale. We then printed the image and measured the actual size of the grid on the printed paper, giving us a scaling factor that, once applied to our original image, always resulted in a printed image wherein the calibration grid and veins were correctly-sized. We then layed cylindrical segments of wax on top of the veins on the printed paper and joined their ends together via precisely placed heat (essentially tracing in wax). Printing the scaled, mirrored image onto a plastic transparency and splicing the wax segments on top of the smooth, unprinted, correctly-unmirrored side of the transparency prevented the wax from sticking to the texture of the printed image. Figures 6.5B-C show manual splicing/tracing of the wax core on top of the 2D, NIR vein image. Figure 6.5C shows the actual soluble-wax core that we used in creating the prototype shown in Figure 6.1.

We started the manual splicing process by making long, straight wax cylinders that served as our splicing stock. We found that clear PVC tubing was an excellent mold for these cylindrical sections and could be removed quickly and gently using a wire-insulation-cutter. Because plastic tubing only comes in certain diameters, we selected diameters as close to the vein diameters as possible for each section of the vein.

The main tool used for splicing the wax segments together was a hot-knife, or essentially a soldering iron that heats a replaceable razor-knife/Xacto knife at the tip. The main cutting process of this knife is melting of the wax, not cutting, so it can be used for very delicate work at low forces. Because the leading edge of the knife is a razor, the heat can be placed very precisely for splicing without melting nearby areas. With two segments trimmed at the appropriate mating angle and placed touching each other, the hot-knife was inserted fully between their intersecting faces to melt each segment's face simultaneously. After waiting for a fraction of a second, the hot-knife blade was removed from the wax, and the melted faces of the segments bonded together quickly. After 10 seconds, the two segments were bonded well enough to be treated as a single piece



(A)



(B)

(C)

Figure 6.5: Making the Soluble-Wax Core. A: Via a 3D-printed mold designed from 3D point-cloud, NIR vein-data. B-C: Via manual splicing/tracing of solid wax cylinders against a 2D, NIR vein image printed in true-size. All vein-data was generated by the vision system in Chapter 3.

of wax. However, the surface of the transition between the two segments was often not smooth, gradual, and filleted enough to serve as a good mold-core. To achieve a smooth, rounded transition between the two segments, the hot-knife's side was used to melt the surface of the wax near the transition area and drag wax over the area. The surface tension of the wax helped the dragged wax to pool into a naturally-smooth, gradual transition. Although this process appears complicated, we were able to complete each splice in roughly 20-30 seconds after only 15 minutes of practice. To achieve segments that curve, we used a heat-gun to bend a straight section while the segment was unspliced and far from the rest of the vein system.

Once the wax vein network had been spliced together, we added a laser-cut plywood frame around the network that allowed us to dip the network in latex and attach it magnetically to the rotocaster while only touching the plywood. Attaching the ends of the wax vein network to the plywood frame followed the same splicing procedure and yielded a strong bond due to the plywood's texture.

6.3.2.4 Defeating the Drip: Obtaining Uniform Thickness Though Constant Rotation

As soon as the soluble-wax-core had been dipped into the latex and removed, the latex started to drip off of the core. Local differences in geometry and coagulant-adhesion-quality led to more or less latex in some areas, which ultimately led to a vein with uneven wall thickness. Since dripping relies on gravity to pull the material downward and away from the vein, we prevented its effects by constantly changing the direction of gravity with respect to the vein's orientation. We designed and fabricated a rotocaster, shown in Figure 6.6, that rotated the centrally-located vein in all directions until the latex had fully dried/solidified. We were not removing the dripping tendency but were changing its direction so quickly that the dripping was smaller in amount and more uniform.

Although we built what is typically referred to as a rotocaster, we inverted its normal usage. Rotocasting and rotational molding are standard fabrication processes for making large, hollow structures with thin walls. In both processes a liquid is placed in a hollow mold that is then rotated slowly about two axes until the liquid has fully solidified through self-curing (rotocasting) or cooling (rotational molding) and can support its own weight. The liquid clings to the inner surface of the mold and gradually builds up a thickening layer that will solidify into a hard shell. In our case, we were still rotating a mold with a liquid material until it hardened, but the material was on the mold's external, instead of internal, surface. For this reason, we will refer to our modified version of the process as "external"-rotocasting.

While rotocasters are commercially available, we found none that met our budget, forced-air-cooling, and quick-mold-attachment needs. As such, we designed and fabricated our own rotocaster, using laser-cut plywood (3/4-inch, AC-grade Baltic Birch) as the main structure. The rotocaster consisted of two rotating frames, an "outer" frame that rotated about a fixed-axis in the horizontal plane and an "inner" frame that rotated about a spinning-axis orthogonal to the outer frame's fixed-axis. The laser-cut frames used no adhesives and relied only on screws/Tee-nuts and press-fit, modified finger joints for structural integrity. By combining these outer and inner rotations, we could span all possible gravity vectors for a mold riding at the center of the inner frame.

Each axis was driven by a scooter motor with a timing-belt reduction that was tensioned via a set of




Figure 6.6: Rotocaster, Motion Sequence, Overhead View. The rotocaster spun the latex-dipped vein-pattern about two axes constantly so that the latex could never drip for too long in any given direction, ensuring a uniform wall-thickness for the finished vein.

slidable idlers, as shown in Figure 6.7A. The inner frame was rotated by a flying motor (mounted on the outer frame), which was counterbalanced by a dummy motor on the opposite corner of the outer frame and received power through a pair of custom slip-rings. As shown in Figure 6.7B, each slip-ring consisted of a strip of 0.010 inch spring-steel that was flexed against the surface of the rotating shafts that supported the outer frame, one contact on each shaft. The slip-rings never failed or required cleaning, even after operating for hundreds of hours.

Because liquid latex needs to dry to solidify, we included a set of box fans to blow on the mold from two, orthogonal directions, reducing the drying time from hours to minutes. The inner frame included two sets of aluminum T-slot-framing for mounting molds of varying sizes, from a small wax vein mold to a full-length mold of the arm. Because we needed to dip the wax vein-core in latex and then immediately spin it to keep the latex from shifting, we needed a method of attaching the dipped vein-core to the rotocaster quickly. Our solution was to use a set of opposing magnets on the rotocaster and the wax vein-core's plywood frame that snapped together for a strong connection almost instantly. Figures 6.7C-D show this magnetic attachment, along with examples of silicone and latex veins that had just finished the rotocasting process.

The process we developed for “external”-rotocasting of the latex also worked well for silicone. Since the silicone is more viscous than the dipping latex, we added an additional, first step of spinning both axes quickly so that the centrifugal forces would even out the silicone layer before we started the long, slow rotation phase during which the silicone cured in place.




Figure 6.7: Rotocaster Design Details. A: Scooter motor and timing-belt transmission. B: Custom slip-rings for powering the inner frame motor. C-D: Magnetically attached vein-cores. C: Silicone vein network that has finished the rotocasting process. D: Latex vein network that has finished the rotocasting process.

6.3.2.5 Closing Vein Segment Ends and Filling With Artificial Blood

Once the rotocasting/solidification process was completed, the latex veins were placed in the heated citric acid solution to dissolve the wax over the course of 5 minutes. Final cleaning, including removing the Krylon Crystal Clear inner skin, took an additional 10 minutes, and the veins were dried with an additional 10 minutes in an oven. We used superglue/cyanoacrylate to close the segment ends permanently.

We filled the veins with artificial blood by using two IV needles. A first IV needle was inserted into the bottom of the largest vein branch and its catheter attached to a artificial-blood-filled syringe. To allow for air to escape the venous system as the artificial blood was injected, a second IV needle was inserted into the end of another branch as that segment was filled, with this process repeated for each branch in the network.

6.4 Discussion and Thesis Contributions

In this chapter, we surveyed commercial off-the-shelf IV phantoms and other alternatives, as shown in Figure 6.2, identifying Twister-160Q-size/“balloon-animal” balloons as a good solution for simple robot testing but not finding a viable option for more advanced testing that required real bifurcations, anatomically-accurate sizes, and hand macrostructure. To address this unmet need, we developed methodologies and tools for fabricating a hand/vein phantom, shown in Figure 6.1, that copied real patient anatomy almost exactly, including seamless, properly-sized, latex bifurcations that had realistic force “pop” and a silicone hand that had accurate anatomy from the macrostructure down to the fingerprints.

We developed the novel fabrication technique, shown in Figures 6.5B-C, of splicing/tracing water-soluble wax as a mold-core for dipping-latex based on 2D, NIR vein images of real patients. This allowed for nearly-identical copying of real vein patterns. Our phantom vein had a realistic force profile, as shown in Figure 6.4, with similar shape, magnitude, and needle/catheter force-“pop”s to the force profiles that we measured clinically in Chapter 2.

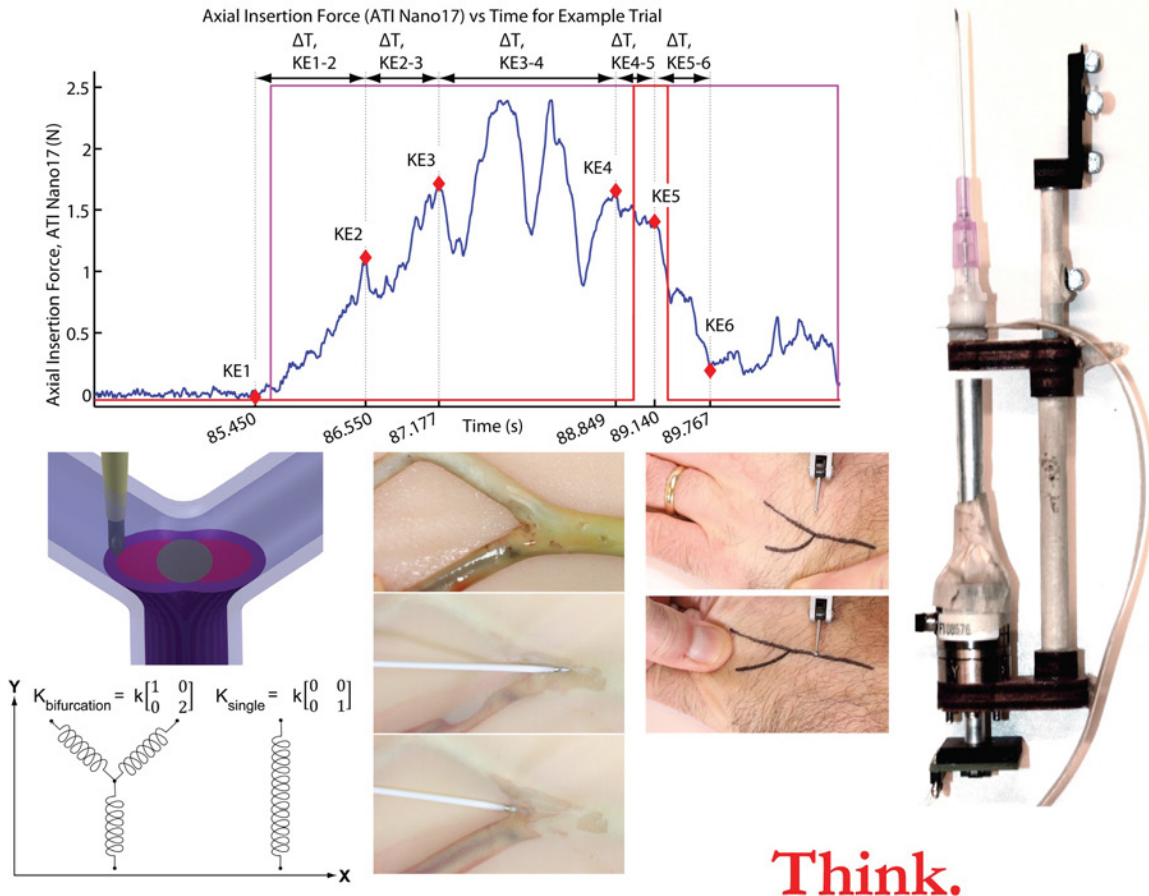
We expanded upon a known fabrication technique in our process of “external”-rotocasting, designing a process and rotocaster that met our unique needs for applying a thin and uniform external layer of dipping-latex to the soluble-wax vein-core, as shown in Figures 6.6- 6.7 This same process and hardware work even better for fabricating silicone veins, as shown in Figure 6.3, and we hope in the future to identify a silicone formulation that yields the proper force-“pop” so that we can switch to silicone as a more durable, transparent, and easier-to-use vein material.

Chapter 7: Conclusion, Contributions, and Future Work

Recall that the question driving this research was whether or not we could raise the IV insertion success rate of a robot beyond that of a human practitioner. To answer that question, we had to ask and answer a simpler one first: can a robot insert an IV autonomously at all? We then deconstructed the insertion tasks of a human practitioner into the “Think. See. Do.” paradigm and translated each element into a set of technical questions about how a robot would accomplish these same tasks. In the course of our research, we expanded our paradigm to include the two additional items: “Do Simpler”, which spoke to our desire to distill the lessons from our robot into a simpler, less -expensive assistive device, and “Do, But Practice First”, which spoke to our need for a realistic vein phantom to facilitate the development and testing of our device.

We successfully answered all of the questions and met all of the objectives in each category of the “Think. See. Do. Do Simpler. Do-But Practice First.” paradigm, as detailed in Figures 7.1- 7.5. We demonstrated that our robot can 1. autonomously locate and calculate a 6D needle pose to insert on a venous bifurcation and 2. autonomously insert an IV into a realistic phantom of a bifurcation by detecting the force “pop”. Due to time constraints, these two demonstrations were performed as separate experiments on different computers instead of as one, contiguous vein-detection-and-insertion procedure as in a real IV insertion. However, the vision system was designed and built to communicate the 6D needle pose to the robot for targeting, and it is a straight-forward matter of code integration and standard camera-to-robot-frame-calibration to connect the vision and insertion systems. As such, we believe that proving them separately still proves that their sum is a valid, affirmative answer to the question of whether a robot can insert an IV by itself.


The only question that remains is whether a robot could be more successful, on average, than a human practitioner at inserting IVs. Answering this question constitutes the bulk of our future work. Performing a controlled, robot-vs-human insertion-experiment on hand phantoms, cadavers, and ultimately live humans is the only way to see if there is a statistical difference in success rates.



We made five main, novel contributions in Chapter 2:

1. Empirically-determined a first-of-its-kind, complete list of human IV insertion steps that provides a framework for the robot to perform the whole process, not just a small portion.
2. Developed a novel methodology for parsing IV insertion data to identify the four penetration-events, or “pops”, from a real IV insertion, proving that a human can parse insertions and providing a framework for robotic parsing.
3. Empirically-measured normative values for force, motion, and intravenous pressure overall and at key moments in real IV insertions on live patients using a uniquely-ergonomic, instrumented IV needle. These values provided 19 “Key Discoveries” that explain properties of insertion more broadly (such as non-axial forces being unnecessary and insertion being slow and interrupted).
4. Observed and analyzed clinical failure modes and recovery techniques, many of them never before documented or explained, that could be used by the robot to prevent and recover from failures.
5. Developed three novel, technical explanations for why bifurcations can be easier insertion targets than single, non-branching sections of vein and two explanations for applying skin traction.

Figure 7.1: Chapter 2 (“Think”) Contribution Summary.

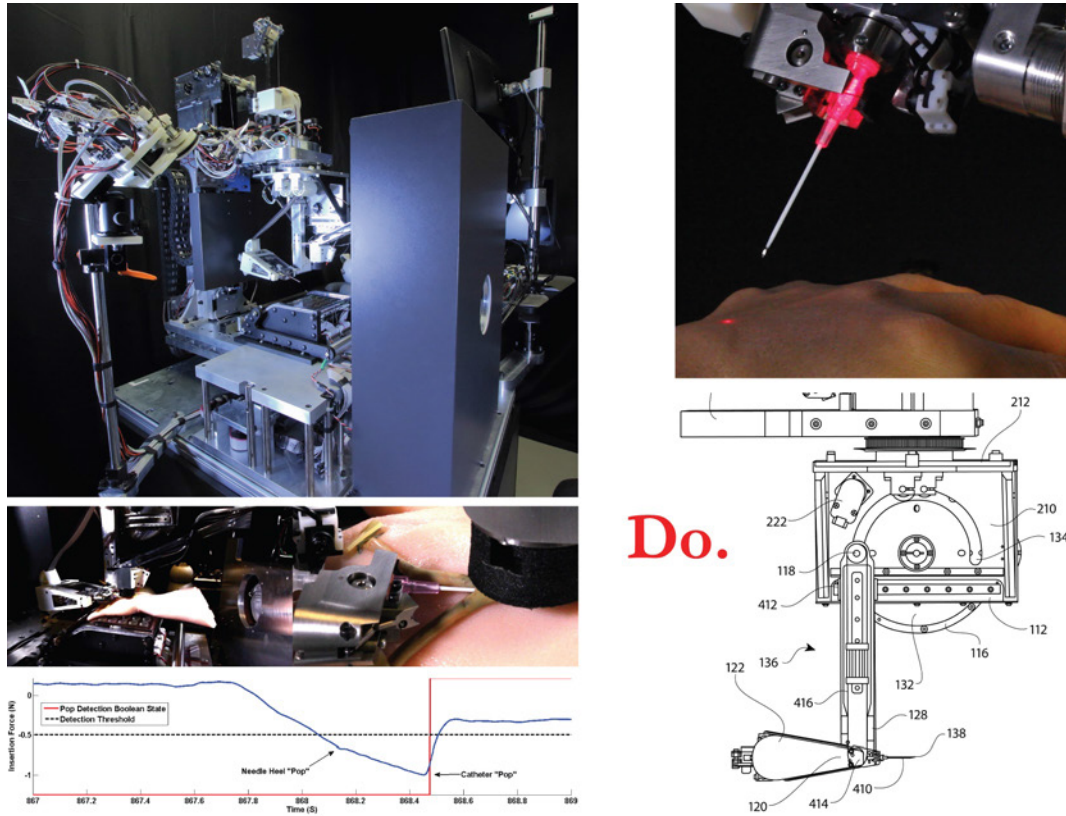


See.

We made three main, novel contributions in Chapter 3:

1. Proved that a robot can locate a vein and compute a 6D needle pose by itself.
2. Developed novel NIR transillumination and stereo camera hardware capable of generating first-of-their-kind, high-resolution, high-contrast vein point-clouds.
3. Developed novel computer-vision algorithms for locating bifurcations in a 2D image and then generating a 6D needle pose in the 3D point-cloud.


Figure 7.2: Chapter 3 (“See”) Contribution Summary.



We made five main, novel contributions in Chapter 4:

1. Proved that a robot can perform all of the many, complicated 1-handed and 2-handed mechanical steps of IV insertion by itself.
2. Proved that the robot can perform these steps based on similar sensory cues (e.g. the force “pop”) that we observed with human practitioners.
3. Proved that we can we do all of this on realistic hand, arm, and vein geometries (not just rubber tubes).
4. Developed a patented method for shining a targeting-laser through an IV needle while maintaining sterility.
5. Developed a patent-pending design for a Remote Center of Motion that has many superior fabrication, installation, and performance properties over existing RCM designs.

Figure 7.3: Chapter 4 (“Do”) Contribution Summary.





Do Simpler.

In Chapter 5, we developed a simple assistive device for IV insertion with three main, novel contributions:

1. Provides more-accurate, stereo vein-visualization than other existing devices.
2. Makes IV insertion easier by reducing it from a 2-handed to 1-handed procedure.
3. Allows theoretically for self-insertion.

Figure 7.4: Chapter 5 (“Do Simpler”) Contribution Summary.



Do-But Practice First.

In Chapter 6, we developed a methodology and tools for fabricating hand and vein phantoms with realistic force “pop” characteristics and patient-specific, realistic anatomy (size, geometry). This non-cadaver phantom provides realistic but inexpensive practice insertions.

Figure 7.5: Chapter 6 (“Do, But Practice First”) Contribution Summary.

Bibliography

- [1] IV-Therapy.net, Dec 2014.
- [2] W. H. Aellig. Venoconstrictor effect of dihydroergotamine in superficial hand veins. *European Journal of Clinical Pharmacology*, 7(2):137–139, 1974.
- [3] R. R. Anderson and J. A. Parrish. The optics of human skin. *Journal of Investigative Dermatology*, 77(1):13–19, 1981.
- [4] D. Becker. Image processing algorithms for retinal montage synthesis, mapping, and real-time location determination. *IEEE Transactions on Biomedical Engineering*, 45(1):105–118, 1998.
- [5] M. B. Boeson et al. Peripheral nerve injury from intravenous cannulation: a case report. *AANA Journal*, 68(1):53–57, Feb 2000.
- [6] R.D. Brewer and J.K. Salisbury. Visual vein-finding for robotic IV insertion. In *2010 IEEE International Conference on Robotics and Automation*, pages 4597–4602, May 2010.
- [7] R.D. Brewer and J.K. Salisbury. Remote center of motion mechanism and method of use, February 14 2013. US Patent Application 13/572,452.
- [8] R.D. Brewer and J.K. Salisbury. Assistive method and visual-aid device for vascular needle insertion, May 6 2014. US Patent 8,715,233.
- [9] A. Can. Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. *IEEE Transactions on Information Technology in Biomedicine*, 3(2):125–138, 1999.
- [10] C. Chang and S. Chatterjee. Quantization error analysis in stereo vision. In *1992 Conference Record of The Twenty-Sixth Asilomar Conference on Signals, Systems and Computers*, pages 1037–1041, Oct 1992.
- [11] L. L. Chapman et al. Veinviewer-assisted intravenous catheter placement in a pediatric emergency department. *Academic Emergency Medicine*, 18(9):966–971, 2011.
- [12] A. Chen et al. Portable robot for autonomous venipuncture using 3D near infrared image guidance. *Technology*, 1(01):72–87, 2013.
- [13] W. Cheong et al. A review of the optical properties of biological tissues. *IEEE Journal of Quantum Electronics*, 26(12):2166–2185, 1990.
- [14] F. B. Chiao et al. Vein visualization: patient characteristic factors and efficacy of a new infrared vein finder technology. *British Journal of Anaesthesia*, 2013.
- [15] S. Cho et al. Effects of infrared radiation and heat on human skin aging in vivo. In *Journal of Investigative Dermatology Symposium Proceedings*, volume 14, pages 15–19. Nature Publishing Group, 2009.
- [16] T. de Boer et al. Laboratory study on needle-tissue interaction: towards the development of an instrument for automatic venipuncture. *Journal of Mechanics in Medicine and Biology*, 07(03):325–335, 2007.
- [17] A. Delalleau et al. A nonlinear elastic behavior to identify the mechanical parameters of human skin in vivo. *Skin Research and Technology*, 14(2):152–164, May 2008.

- [18] D. Delpy. Quantification in tissue near-infrared spectroscopy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 352:649–659, 1997.
- [19] E. D. Devlo-Favre et al. Implementation of near-infrared technology (Accuvein AV-400) to facilitate successful PIV cannulation. *Journal of Infusion Nursing*, 37(4):250–258, July/August 2014.
- [20] A.P. Dhawan. Apparatus and method for skin lesion examination, September 15 1992. US Patent 5,146,923.
- [21] M. R. Dunn and S. Conrad. Utilization of infrared trans-illumination as an aid for peripheral arterial access. *Academic Emergency Medicine*, 5(Supplement 1):34–35, 2005.
- [22] E. Eriksson et al. Penetration forces in cannulation of the dorsal veins of the hand: I. A comparison between polyurethane (Insyte) and polytetrafluoroethylene (Venflon) cannulae. Results of a study in volunteers compared with those from an in vitro study. *Acta Anaesthesiologica Scandinavica*, 35(4):306–314, May 1991.
- [23] L. Gagnon. Procedure to detect anatomical structures in optical fundus images. In *Medical Imaging 2001*, pages 1218–1225. International Society for Optics and Photonics, 2001.
- [24] R. Goldman et al. Three dimensional imaging of veins, July 19 2011. US Patent 7,983,738.
- [25] M.L. Harris and F.J. Ball. Vein imaging systems and methods, February 13 2014. US Patent Application 13/802,423.
- [26] M.L. Harris et al. Vein imaging systems and methods, February 6 2014. US Patent Application 13/802,604.
- [27] R.J. Harris et al. Systems and methods for autonomous intravenous needle insertion, July 26 2012. US Patent Application 13/335,700.
- [28] W. G. Haynes et al. Direct and sympathetically mediated venoconstriction in essential hypertension. Enhanced responses to endothelin-1. *Journal of Clinical Investigation*, 94(4):1359–1364, Oct 1994.
- [29] H. A. Hess. A biomedical device to improve pediatric vascular access success. *Pediatric Nursing*, 36(5):259–263, 2010.
- [30] A. M. Holzer et al. The other end of the rainbow: infrared and skin. *Journal of Investigative Dermatology*, 130(6):1496–1499, Jun 2010.
- [31] S. L. Jacques. Optical properties of biological tissues: a review. *Physics in Medicine and Biology*, 58(11):R37–R61, 2013.
- [32] A. Jensen et al. Improving the quality of near-infrared imaging of in vivo blood vessels using image fusion methods. In *Proceedings of the IASTED International Conference*, volume 661, page 081, 2009.
- [33] J.F. Jensen. Remote center positioning device with flexible drive, October 6 1998. US Patent 5,817,084.
- [34] B. Khanlari et al. Local effect of transdermal isosorbide dinitrate ointment on hand vein diameter. *European Journal of Clinical Pharmacology*, 57(10):701–704, Dec 2001.
- [35] A. Kienle et al. Why do veins appear blue? A new look at an old question. *Applied Optics*, 35(7):1151, 1996.

- [36] M. J. Kim et al. Efficacy of VeinViewer in pediatric peripheral intravenous access: a randomized controlled trial. *European Journal of Pediatrics*, 171(7):1121–1125, Jul 2012.
- [37] B. Kochner. Course tracking and contour extraction of retinal vessels from color fundus photographs: Most efficient use of steerable filters for model based image analysis. *SPIE Conference on Image Processing*, 3338:755–761, 1998.
- [38] R. Lininger. Pediatric peripheral iv insertion success rates. *Pediatric Nursing*, 29(5):351–355, 2003.
- [39] L.J. Love et al. Force-based needle insertion for medical applications. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2592–2597, Oct 2009.
- [40] S. Masoorli et al. Danger points. How to prevent nerve injuries from venipuncture. *Nursing*, 28(9):34–39, Sep 1998.
- [41] I. Meglinski. Quantitative assessment of skin layers absorption and skin reflectance spectra simulation in the visible and near-infrared spectral regions. *Physiological Measurement*, 23:741–753, 2002.
- [42] N.A. Mullani. Transilluminator light shield, March 22 2007. US Patent Application US 11/227,195.
- [43] N.A. Mullani. Transillumination having orange color light, January 25 2011. US Patent 7,874,698.
- [44] C. Nachev et al. Simplified method for measuring compliance of superficial veins. *Cardiovascular Research*, 5(1):147–156, 1971.
- [45] E.R. Nosari. Method of locating vein, December 21 1976. US Patent 3,998,210.
- [46] H. Ogden-Grable. Phlebotomy puncture juncture: Preventing phlebotomy errors-potential for harming your patients. *Laboratory Medicine*, 36(7):430–433, 2005.
- [47] D. Okuno et al. Development of an automatic blood sampling system: control of the puncturing needle by measuring forces. In *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 4, pages 1811–1812 vol.4, Oct 1998.
- [48] V. Paquit et al. Near-infrared imaging and structured light ranging for automatic catheter insertion. *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, SPIE 6141:19, 2006.
- [49] V. Paquit et al. Combining near-infrared illuminants to optimize venous imaging. In *Medical Imaging*, pages 65090H–65090H. International Society for Optics and Photonics, 2007.
- [50] V. Paquit et al. 3D multispectral light propagation model for subcutaneous veins imaging. In *Medical Imaging*, pages 69130D–69130D. International Society for Optics and Photonics, 2008.
- [51] V. Paquit et al. 3D and multispectral imaging for subcutaneous veins detection. *Optics express*, 17(14):11360–11365, 2009.
- [52] J. Payne. Alternative techniques for venipuncture. *The American Journal of Nursing*, 72(4):702–703, 1972.
- [53] T. S. Perry. Profile: Veebot. *Spectrum, IEEE*, 50(8), 2013.
- [54] K. A. Peterson et al. Does the use of an assistive device by nurses impact peripheral intravenous catheter insertion success in children? *Journal of Pediatric Nursing*, 27(2):134 – 143, 2012.
- [55] J. Proehl. *Adult emergency nursing procedures*, page 197. Jones and Bartlett Publishers, 1993.

- [56] Lenhardt R. Local warming and insertion of peripheral venous cannulas: single blinded prospective randomised controlled trial and single blinded randomised crossover trial. *British Medical Journal*, 325(7):409–412, 2002.
- [57] H. Saito et al. Detection of needle puncture to blood vessel by using electric conductivity of blood for automatic blood sampling. *Sensors and Actuators A: Physical*, 125(2):446 – 450, 2006.
- [58] H. Saito and T. Togawa. Detection of needle puncture to blood vessel using puncture force measurement. *Medical and Biological Engineering and Computing*, 43(2):240–244, Mar 2005.
- [59] T. R. Savarimuthu and A. Sorenson. Detection of vessels in human forearms using 2D matched filtering. In *Signal and Image Processing 2008*, pages 294–299, 2008.
- [60] A. Sorenson et al. Tactile sensing methods for automated blood samples on humans. In *Proceedings of the 11th International Conference on Climbing and Walking Robots. Special Session on Manipulation and Grasping*, pages 875–889, 2008.
- [61] P. Szmuk et al. The VeinViewer vascular imaging system worsens first-attempt cannulation rate for experienced nurses in infants and children with anticipated difficult intravenous access. *Anesthesia and Analgesia*, 116(5):1087–1092, May 2013.
- [62] C. Tsai. Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images. *IEEE Transactions on Information Technology in Biomedicine*, 8(2):122–130, 2004.
- [63] C. Tsai et al. Near-infrared absorption property of biological soft tissue constituents. *Journal of Medical and Biological Engineering*, 21(1):7–14, 2001.
- [64] R. L. Wesly et al. Static linear and nonlinear elastic properties of normal and arterialized venous tissue in dog and man. *Circulation Research*, 37(4):509–520, Oct 1975.
- [65] P.G. Yock and A.R. Selfridge. Apparatus for use in cannulation of blood vessels, December 19 1989. US Patent 4,887,606.
- [66] G. Zaccanti et al. Optical properties of biological tissues. In *Photonics West'95*, pages 513–521. International Society for Optics and Photonics, 1995.
- [67] S. X. Zhang and H. M. Schmidt. Clinical anatomy of the subcutaneous veins in the dorsum of the hand. *Annals of Anatomy*, 175(4):381–384, Aug 1993.
- [68] A. Zivanovic and B. L. Davies. A robotic system for blood sampling. *IEEE Transactions on Information Technology in Biomedicine*, 4(1):8–14, 2000.
- [69] A. Zivanovic and B. L. Davies. The development of a haptic robot to take blood samples from the forearm. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2001*, volume 2208 of *Lecture Notes in Computer Science*, pages 614–620. Springer Berlin Heidelberg, 2001.