Detecting Epileptic Seizures from Electroencephalography

Johns Hopkins Malone Center *for* Engineering *in* Healthcare

Johns Hopkins Whiting School *of* Engineering

Duha Awad [1], Jeff Craley [2], Emily L. Johnson, M.D. [3], Archana Venkataraman, Ph.D. [2]

[1] Department of Electrical and Computer Engineering, University of Maryland, College Park, USA
[2] Department of Electrical and Computer Engineering, Malone Center, Johns Hopkins University, Baltimore, USA
[3] Department of Neurology, Epilepsy Center, Johns Hopkins Hospital, Baltimore, USA

Introduction

Motivation

- Epilepsy affects **1%** of the world population.
- Epilepsy seizure detector would save doctors & specialists time.

Background

- EEG is a non-invasive method that measures the electrical activity of the brain.
- Epilepsy is a neurological disorder caused by abnormal nerve cell activity in the brain that causes episodes known as **seizures**.

Hypothesis

- Using ensemble of features will increase performance in seizure detection.

Method

Raw EEG Data

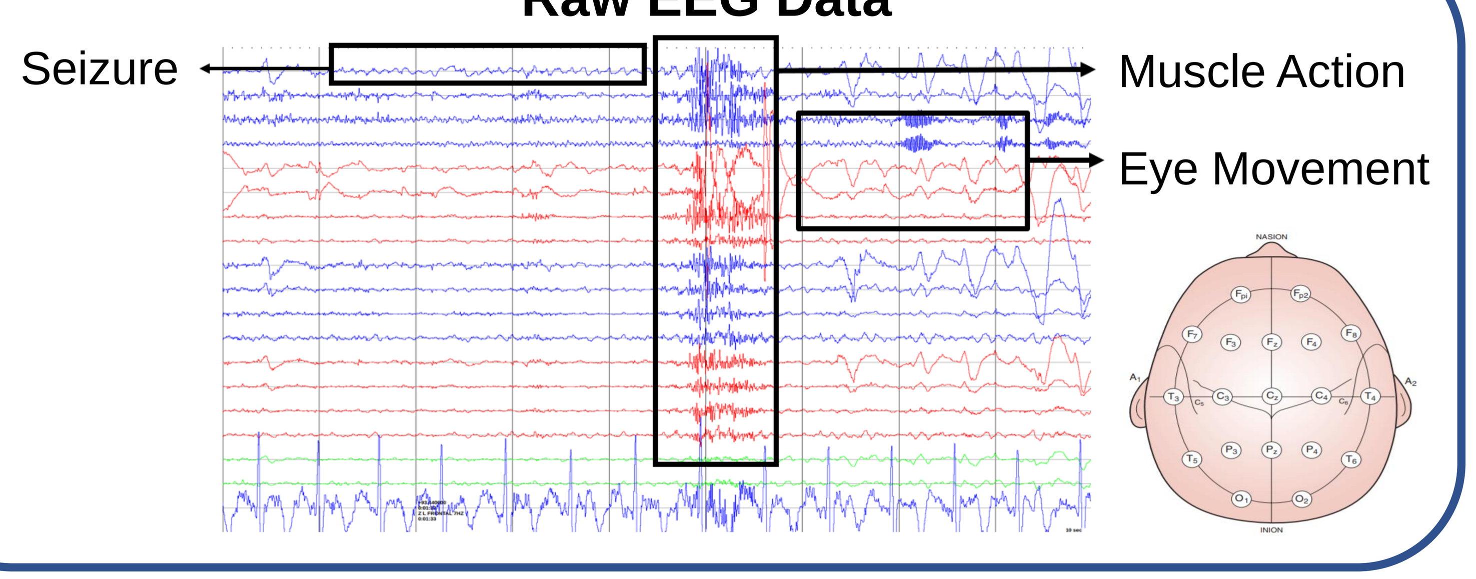

Signal Pre-Processing

Low Pass Filter Cutoff =30 Hz → High Pass Filter Cutoff =1.5 Hz → Windowing

- Low Pass Filter: Removes muscle artifacts
- High Pass Filter: Removes DC trends & slow artifacts
- Windowing: 2 seconds window

Feature Extraction

Time Domain	Power and Zero Crossings
Cross Channels	Cross Correlation, Pearson r Correlation, Mean Phase Coherence
Spectral	Fast Fourier Transform and Band Pass

Classification

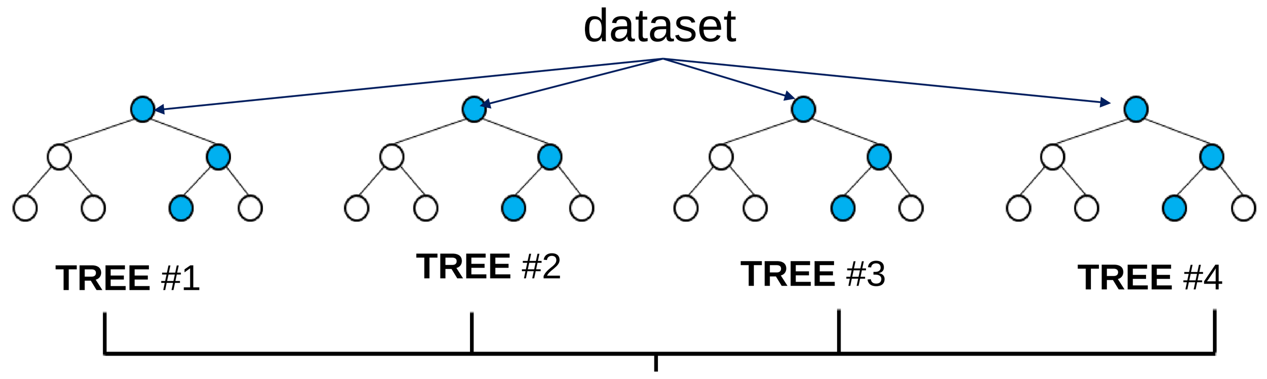

Features

Time Domain:

Power (P)

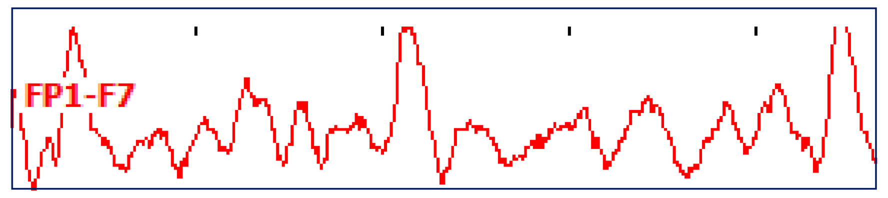

$$\sum_{i=1}^{N} X^2[i]$$

Zero Crossings (ZC)

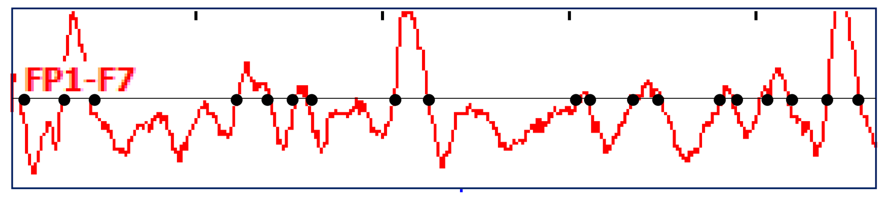

Count places the signal crosses the x-axis

Cross Channels:

Cross Coherence (CC), Pearson r (PR), Mean Phase Coherence (MPC)

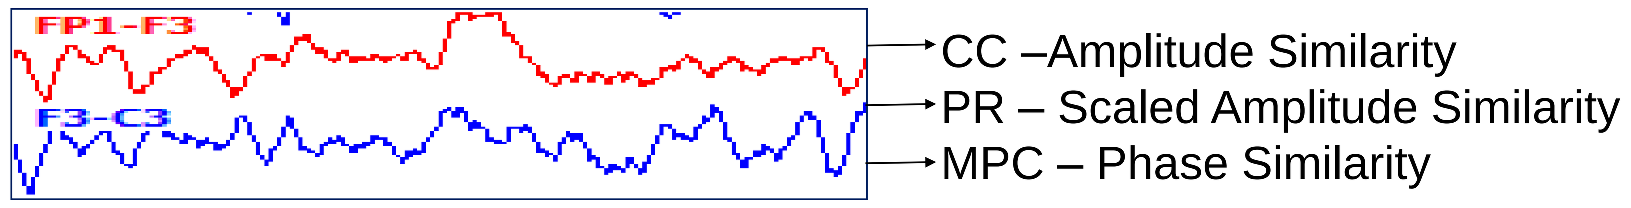

Spectral:

Fast Fourier Transform (FFT)

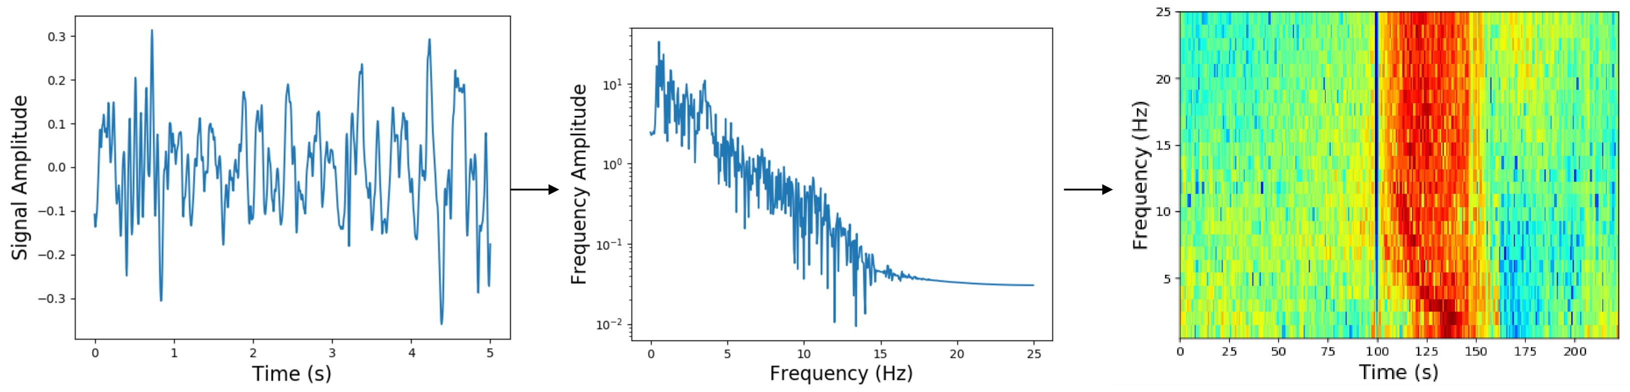

Band Pass Filter (BP)

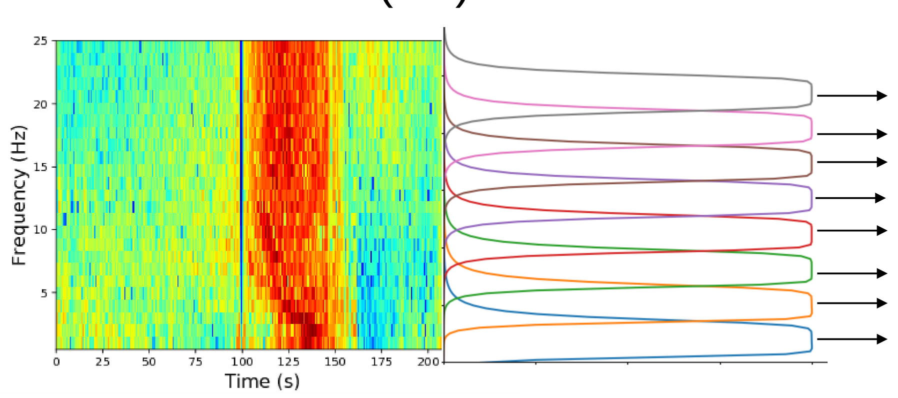

Eight bandpass filters for computing power across different frequency bands

Evaluation

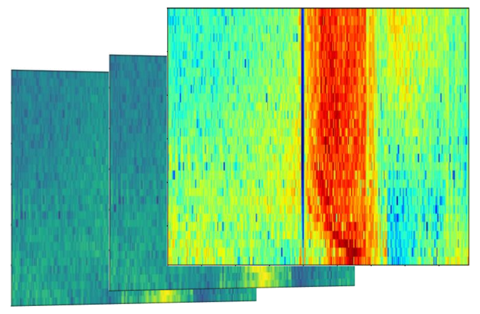

Sensitivity training/test evaluation | Specificity training/test evaluation

- Sensitivity: 1. Non-Seizure 2. All seizure recordings but 1 → y_{train}; Test seizure recording → y_{test}
- Specificity: 1. Seizure 2. All non-seizure recordings but 1 → y_{train}; Test non-seizure recording → y_{test}

Sensitivity: True seizures detection

Specificity: False seizure detection

Results

Data Acquisition: Boston Children's Hospital.
Data: 20 epileptic patients; 154 seizure recordings and 495 non-seizure recordings. Each recording is ~ an hour long.

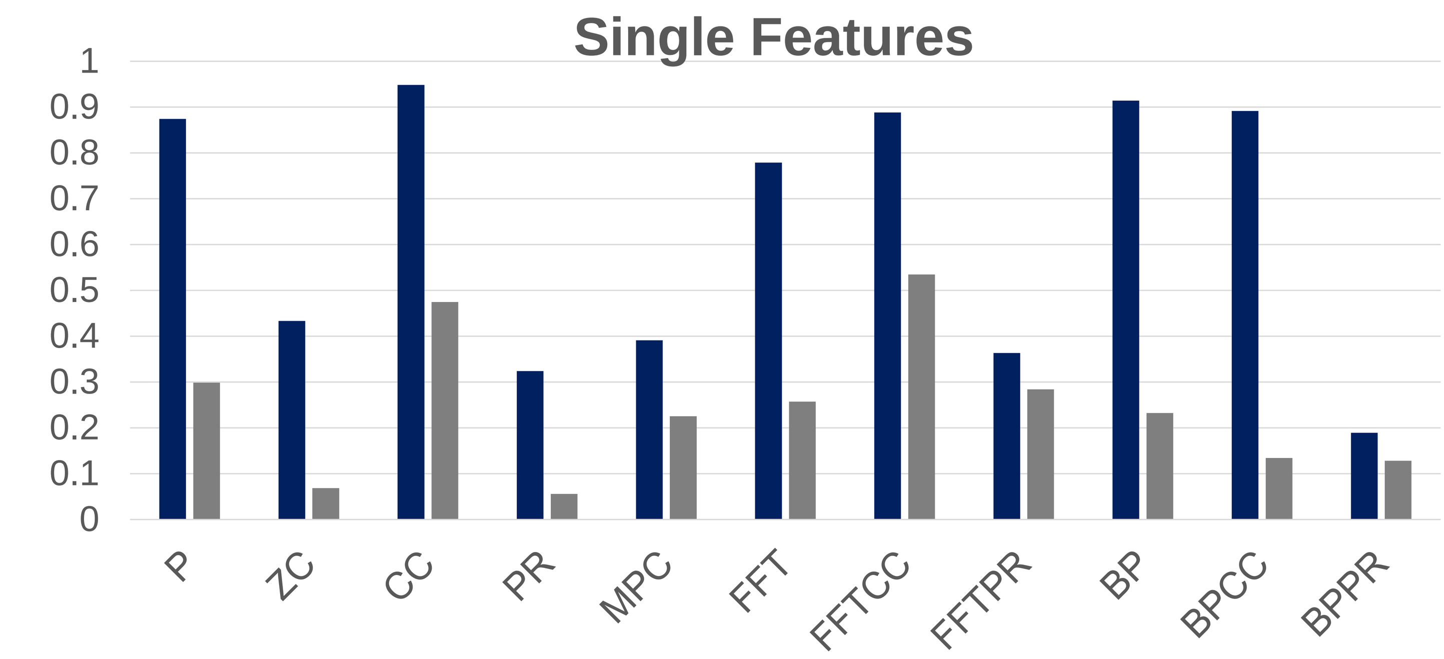

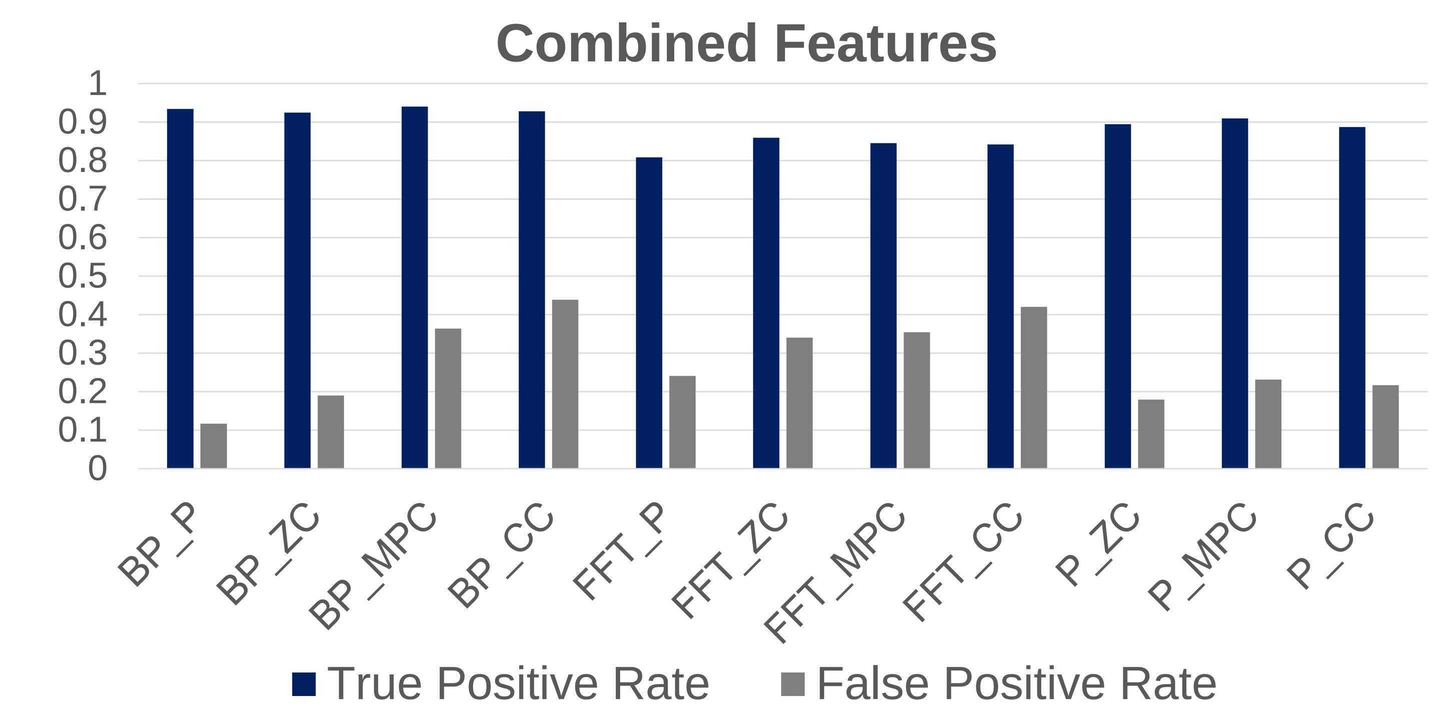

Conclusion

- We can extract features from EEG signals for automated seizure detection.
- Using a combination of features enhanced true positive rate and reduced the false positive rate.
- Spectral features combined with power or zero-crossings performed the best.

Further Research

- Noninvasive Seizure Localization for Focal Epilepsy.
- Seizure detector could be further used to an automated anti-seizure therapy.

Acknowledgments

This work was funded by JHMI Synergy Award. A special thanks to Jeff Craley, Niharika Shimona D'Souza, Ravi Shankar, Dr. Archana Venkataraman, Dana Walter-Shock, and the CSMR–REU family.

References

1. Shoeb, A.H., Guttag, J.V.: Application of machine learning to epileptic seizure detection. In: International Conference on Machine Learning. (2010) 975–982.
2. Orosco, L., Garces, A., & Laciar, E.: Review: A Survey of Performance and Techniques for Automatic Epilepsy Detection. (2013).

Induced 1-Norm Subspace Clustering via Accelerated ADMM

Evan Frenklak[1,2], Chong You[2], Guilherme França[2], Benjamin B. Haro[2], René Vidal[2]

[1]Electrical Engineering and Computer Science Department, University of California Berkeley

[2]Mathematical Institute for Data Science, Johns Hopkins University

Vision Lab @ JHU http://www.vision.jhu.edu

Motivation

- High dimensional datasets usually contain multiple low-dimensional subspaces.

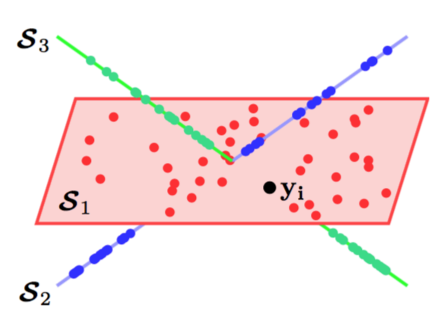

- Subspace clustering is the problem of learning a union of subspaces from unlabeled data, with applications in motion segmentation and facial recognition

Sparse Subspace Clustering

- Each data point can be expressed as a linear combination of the entire dataset excluding itself. That is:

$$Y = [y_1 \dots y_N]$$

$$y_j = Yc_j \qquad c_j(j) = 0$$

- SSC finds the sparsest such representation by solving:

$$\min_C \|C\|_1 + \frac{\lambda}{2}\|Y - YC\|_F^2$$

- If the sparsity of c_j is upper bounded by the dimension of its respective subspace, SSC tends to perform more accurately

Challenge

- The L1 matrix norm is equal to the sum of the matrix's columns' 1-norms:

$$\|C\|_1 = \sum_{j=1}^{N} \|c_j\|_1$$

- L1 regularization allows for greater sparsity in some columns over others, so solutions lack consistent column-wise sparse regularization in the SSC optimization problem

$$\| \quad \|_1 = \left(\blacksquare + \blacksquare + \cdots + \blacksquare \right)$$

Induced 1-Norm Subspace Clustering

SSC-1,1 Optimization Problem: $\min_C \|C\|_{1,1} + \frac{\lambda}{2}\|Y - YC\|_F^2$

- The induced matrix 1-norm is defined as the largest 1-norm of all columns:

$$\|C\|_{1,1} = \max_j \|c_j\|_1$$

- Thus, the induced 1-norm acts as a threshold, upper bounding each column to produce similar regularization across all columns

$$\| \quad \|_{1,1} = \left(\blacksquare + \blacksquare + \cdots + \blacksquare \right)$$

Prox. Operator	Method	Complexity
L1 Matrix Norm	Closed Form	$O(N^2)$
Induced 1-Norm	Iterative	$O(N^2 \log N)$

Table 1: The induced 1-norm's proximal operator is more computationally intensive. Accelerated optimization methods can help reduce total iterations required to solve induced 1-norm SSC.

Lemma 1: Minimum Regularization Parameter for Induced 1-norm SSC

The solution of the SSC-1,1 problem is nontrivial, $C \neq 0$, provided:

$$\lambda \geq \frac{1}{\mu_{1,1}(Y)} \qquad \text{where} \qquad \mu_{1,1}(Y) = \sum_{j=1}^{N} \|x_j\|_\infty \qquad X = Y^TY - diag(Y^TY)$$

Accelerated ADMM

Initialize $A_0, C_0, \hat{C}_0 \Delta_0, \hat{\Delta}_0, \hat{k} = 0$

for $k = 0, 1, 2, \dots$ **do**:

$$A_{k+1} = (\lambda Y^T Y + \rho I)^{-1}(\lambda Y^T Y + \rho \hat{C}_k - \hat{\Delta}_k)$$

$$M_{k+1} = A_{k+1} + \rho^{-1}\hat{\Delta}_k$$

$$C_{k+1} = \text{prox}_{\rho^{-1}\|\cdot\|_{1,1}}(M_{k+1} - diag(M_{k+1}))$$

$$\Delta_{k+1} = \Delta_k + \rho(A_{k+1} - C_{k+1})$$

$$\gamma_{k+1} = \hat{k} / (\hat{k} + 3)$$

$$\hat{\Delta}_{k+1} = \Delta_{k+1} + \gamma_{k+1}(\Delta_{k+1} - \Delta_k)$$

$$\hat{C}_{k+1} = C_{k+1} + \gamma_{k+1}(C_{k+1} - C_k)$$

if $\|A_{k+1} - A_k\|_2 < \|A_k - A_{k-1}\|_2$ **do**:

$\quad \hat{k} = 1$

else **do**:

$\quad \hat{k} = \hat{k} + 1$

Accelerated variables interpolate based on difference from previous iteration

Based on a differential equation model, the restart scheme allows acceleration to regain momentum during the algorithm's convergence

Experiments

Synthetic Data

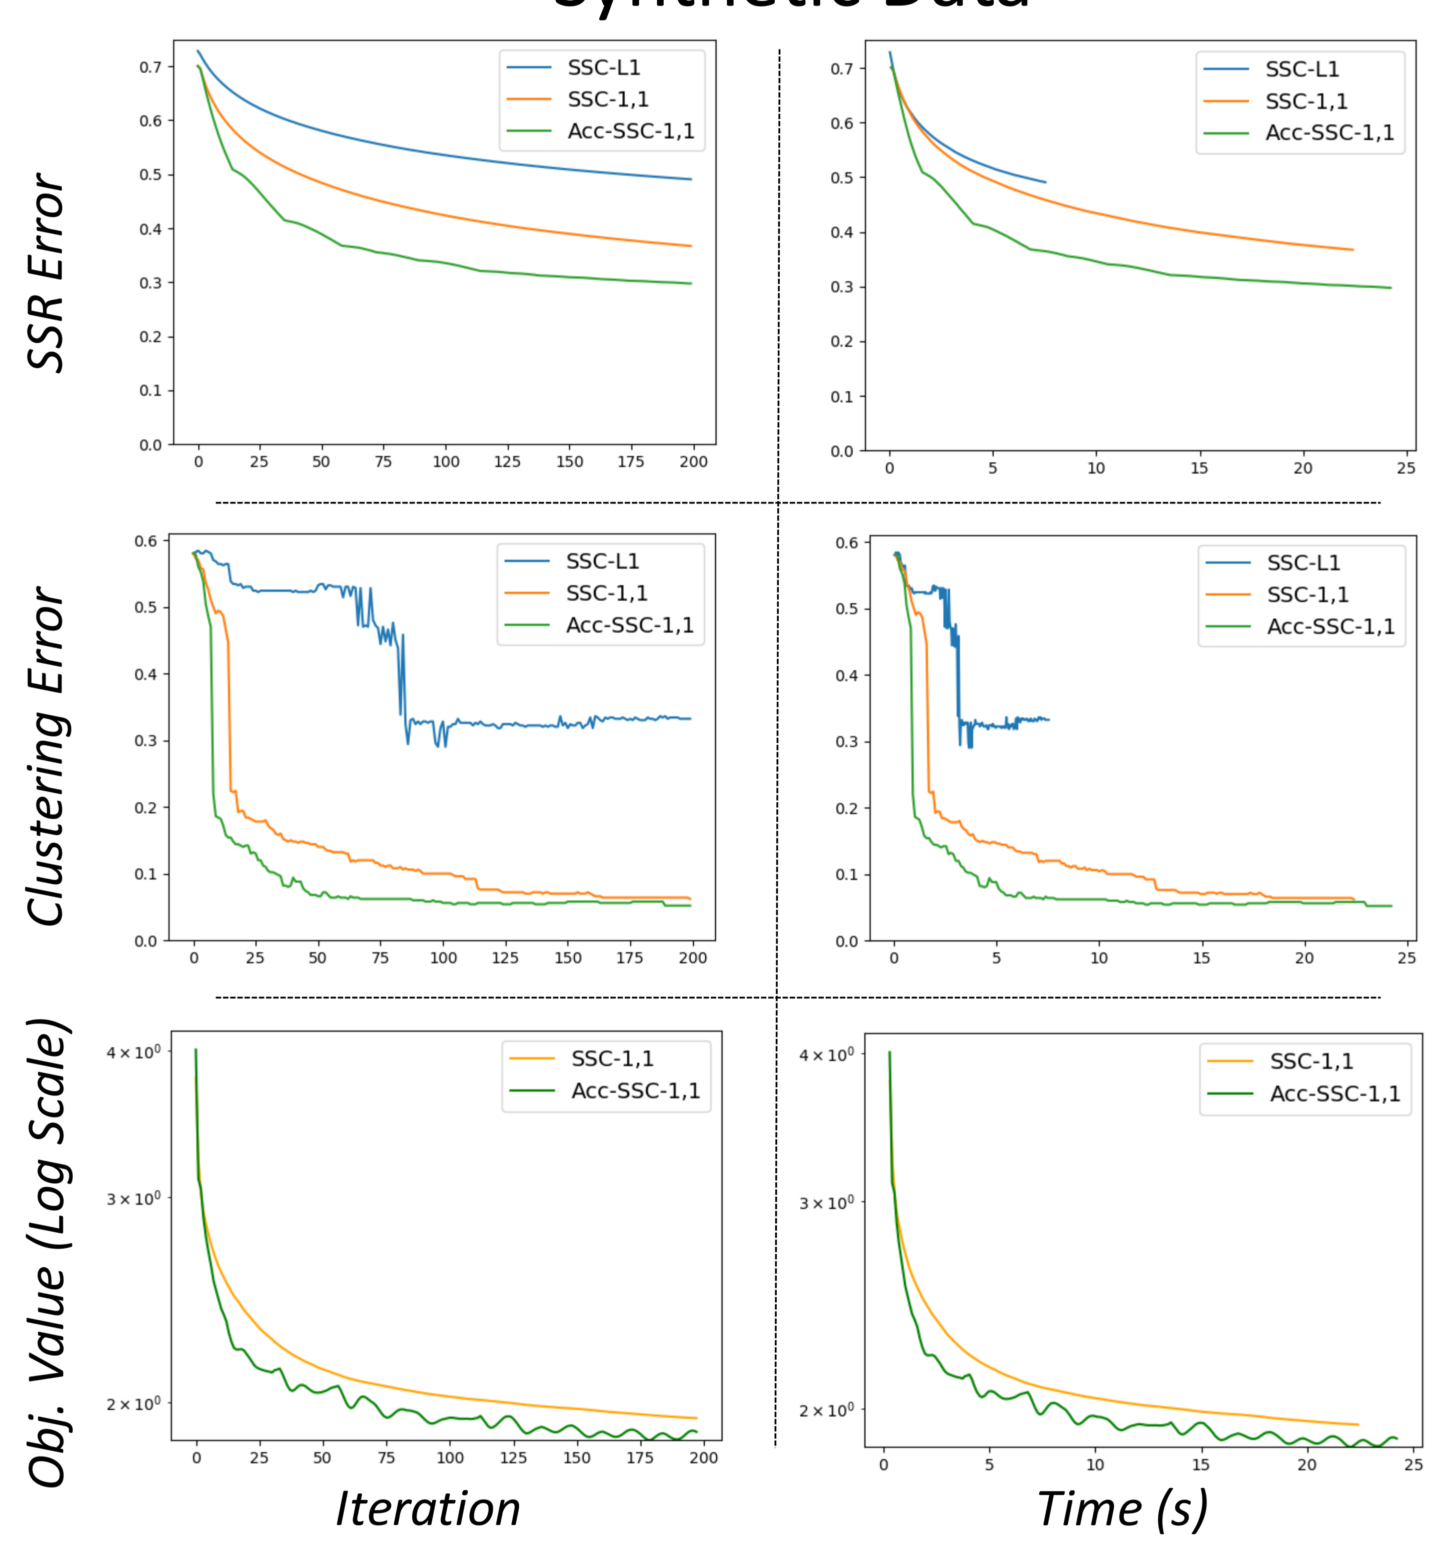

Figures 1- 6: 5 subspaces each of dimension 6 were randomly generated with ambient dimension 9. On a dataset containing 100 points sampled from each subspace (N = 500 total), convergence plots demonstrate reduced error and increased time per iteration for SSC-1,1 compared to SSC-L1. Acceleration of SSC-1,1 led to faster convergence with a similar time per iteration. $\rho_{L1} = 1000,\ \lambda_{L1} = 1000 / \mu_{L1}(Y),\ \rho_{1,1} = \lambda_{1,1} = 1000 / \mu_{1,1}(Y),$

Hopkins 155 Motion Data

	SSC-L1	SSC-1,1
Hopkins155 Motion Segmentation, 2 Cluster Sequences		
Mean SSR Error:	4.24%	4.25%
Mean Clustering Error:	1.56%	2.96%
Median Clustering Error:	0.00%	0.00%
Hopkins155 Motion Segmentation, 3 Cluster Sequences		
Mean SSR Error:	7.25%	5.42%
Mean Clustering Error:	4.32%	6.82%
Median Clustering Error:	0.80%	0.56%
Hopkins155 Motion Segmentation, All Sequences		
Mean SSR Error:	4.92%	4.51%
Mean Clustering Error:	2.19%	3.83%
Median Clustering Error:	0.00%	0.00%

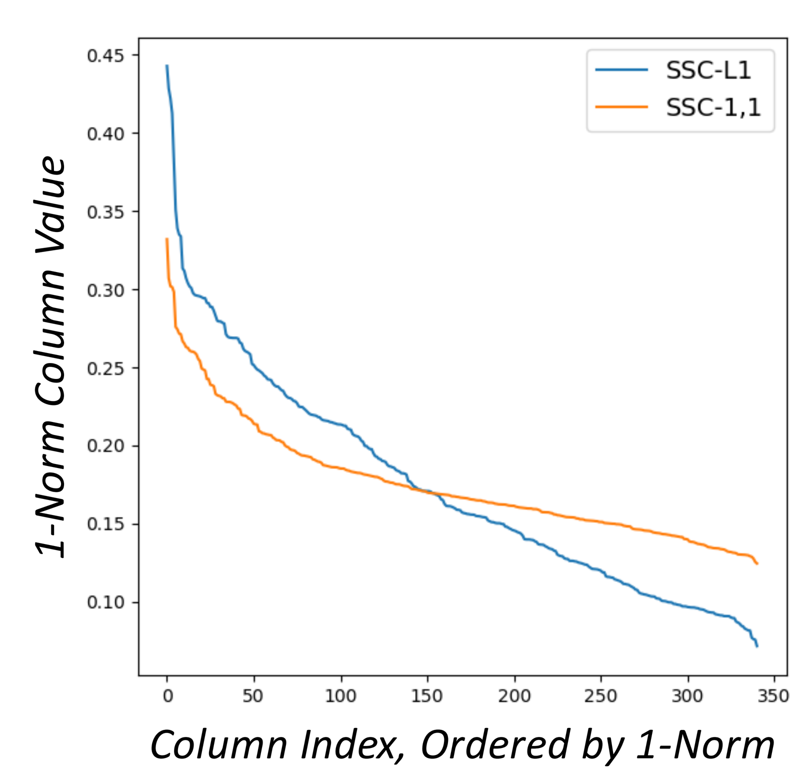

Table 2 (Above): Comparative performance on the Hopkins155 Motion Segmentation dataset (no representation matrix threshold). Induced 1-Norm Regularization reduced or matched Mean Subspace Recovery Error and Median Clustering Error while increasing Mean Clustering Error.

Figure 7 (Left): Descending values of column 1-Norms from final representation matrix for Sequence 85. Induced 1-Norm Regularization produced a more even distribution across columns.

Conclusions

Conclusions

- Our work provides theoretical backing for further research into induced 1-norm subspace clustering
- Induced 1-norm regularization leads to reduced SSR error and reduced median clustering error in some datasets
- Accelerated ADMM improves on ADMM for convergence rate in this problem

Future Work

- Experiment on more data sets to further explore the accuracy improvement potential of the induced 1-norm
- Different optimization approaches to improve convergence in the induced 1-norm SSC problem

Acknowledgements

- NSF grant number 1618637
- Work performed by EF as a summer (2018) student with the Vision Lab at JHU.

References

1) E. Elhamifar and R. Vidal, "Sparse subspace clustering," in 2009 IEEE Conference on Computer Vision and Pattern Recognition. June, 2009.
2) T. Goldstein, et al., "Fast Alternating Direction Optimization Methods," SIAM Journal on Imaging Sciences 7 no. 3, (2014).
3) G. França and D. P. Robinson and R. Vidal, "ADMM and Accelerated ADMM as Continuous Dynamical Systems," ICML (2018)
4) W. Su, S. Boyd, and E. J. Candès, "A Differential Equation for Modeling Nesterov's Accelerated Gradient Method: Theory and Insights," Journal of Machine Learning Research 17 no. 153, (2016).

Cerebellum Parcellation and Data Analysis of the Baltimore Longitudinal Study of Aging

Josue Rodriguez-Nieves [1], Shuo Han [2], Jerry L. Prince Ph.D. [3]
Department of Electrical and Computer Engineering, Inter American University of Puerto Rico, Bayamon, PR, USA
Department of Biomedical Engineering, The Johns Hopkins University School of Medicine, Baltimore, MD, USA
Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, USA

Background

The cerebellum is one of the structures in the central nervous system. It is located at the back of the brain, making up ~10% of the volume in the brain, and contains more than 50% of the brain's neurons. Neuroimaging and clinical studies indicate a contribution of the cerebellum in neural processes beyond the motor domains and also has shown contribution of the posterior lobe in the processing of cognitive functions [1]. Other studies demonstrate consistent activation in the cerebellum posterior lobe for cognitive tasks such as: language, spatial and working memory, while data from clinical populations provide evidence of cognitive deficits after posterior lobe damage [2]. For this reason, detailed mapping of the cerebellum lobules is necessary to describe and establish the functional significance of lobules implicated in cognitive functions of normal subjects [3].

Anterior Lobe   Posterior Lobe   Flocculonodular Lobe

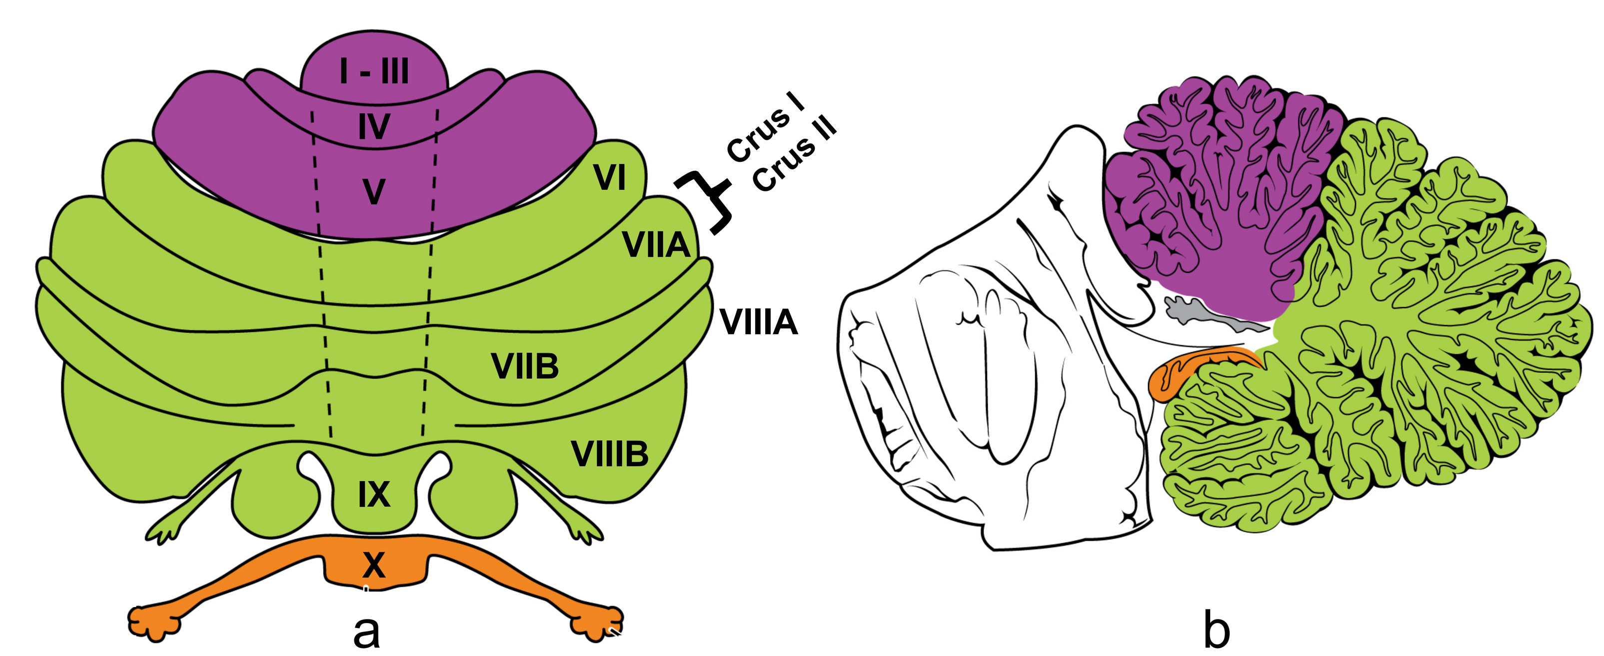

Figure 1. For better understanding of the variations in the cerebellar regions, anatomical and computational studies divide the cerebellum into (a) lobules referred as Roman number from I to X and (b) groups of lobules referred as lobe.

Objectives

- Analyze the correlations between the volumes of the cerebellum regions and the cognitive function test scores in normal subjects.
- Analyze the change of the volumes in the cerebellar regions and cognitive test scores during the aging process in normal subjects.

Hypothesis

- The volumes of different cerebellum regions are correlated with the cognitive test scores and decrease during the aging process in normal subjects.

Methodology

Sample and Image Processing

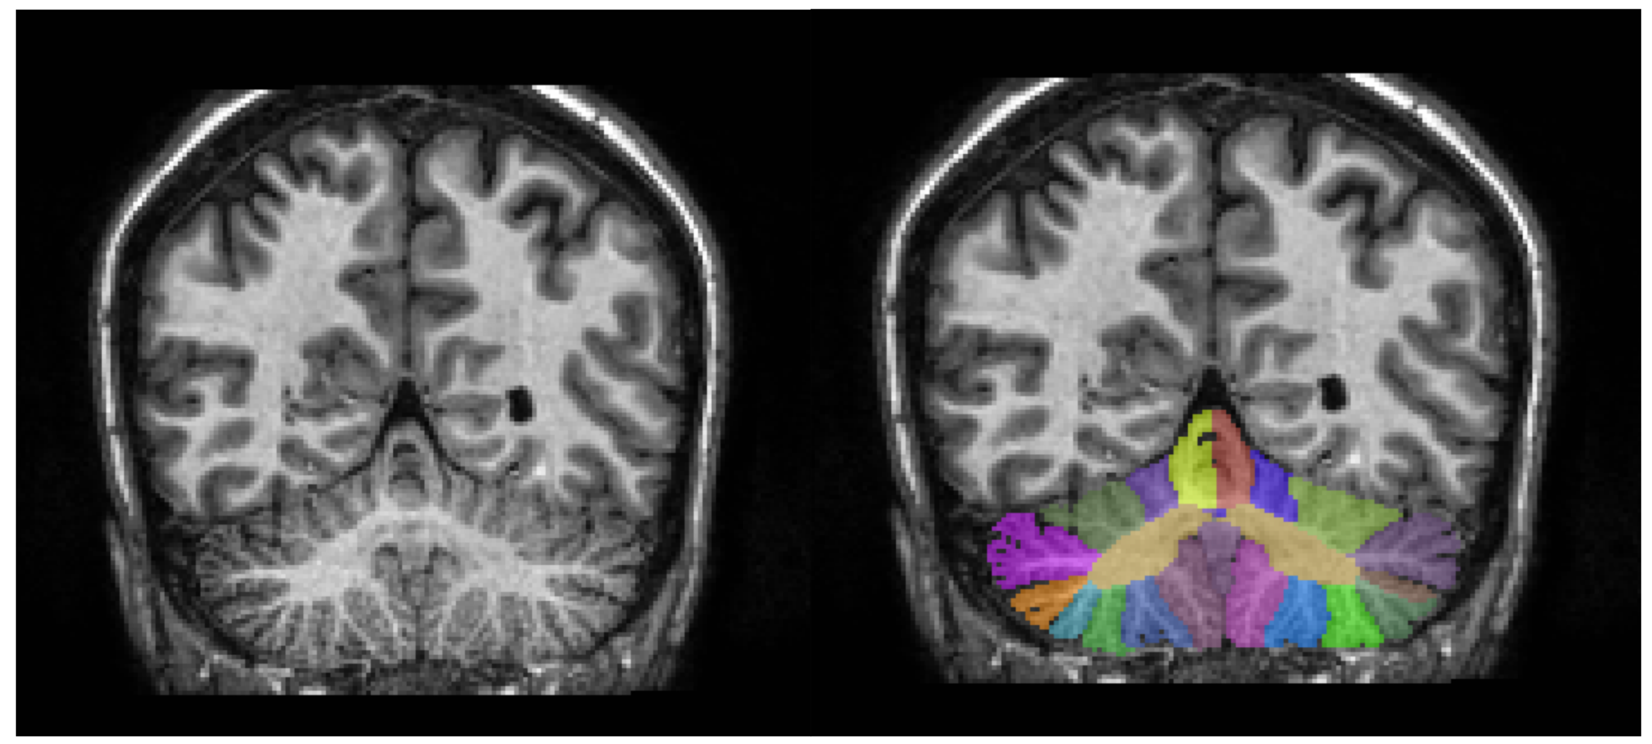

Figure 2. Magnetic Resonance (MR) image was acquired on Phillips 3T Scanner (left) and analyzed using our current deep network algorithm to divide the cerebellum into 28 regions (right).

Statistical Analysis

Partial Correlation

$$\rho_{xy|z} = \frac{r_{xy} - r_{xz}r_{yz}}{\sqrt{1-r_{xz}^2}\sqrt{1-r_{yz}^2}}$$

$\rho = partial\ correlation\ coefficient$
$x = independent\ variable$
$y = dependent\ variable$
$z = control\ variable$
$r = pearson\ correlation\ coefficient$

We use partial correlation to measure the strength and direction of linear relationship between cognitive test scores (y) and volumes in the cerebellar regions (x) controlling the effect of sex, age and intercranial volume (ICV) variables (z).

Linear Mixed Model

$$Y_i = X_i\beta_i + Z_i u_i + \varepsilon_i$$

$Y_i = vector\ of\ continuos\ response$
$X_i = matrix\ of\ fixed\ effects$
$\beta_i = vector\ of\ regression\ coefficients$
$Z_i = matrix\ of\ random\ effects$
$u_i = vector\ of\ random\ effects$
$\varepsilon_i = vector\ of\ the\ residuals$

We use the linear mixed effect model to describe the change of the volume in the cerebellar regions and the cognitive test scores against the aging process in a normal subject.

References

[1] Stoodley, C. J., & Schmahmann, J. D. (2009). Functional topography in the human cerebellum: a meta-analysis of neuroimaging studies. *Neuroimage*, *44*(2), 489-501.

[2] Stoodley, C. J., & Schmahmann, J. D. (2010). Evidence for topographic organization in the cerebellum of motor control versus cognitive and affective processing. *cortex*, *46*(7), 831-844.

[3] Kansal, K., Yang, Z., Fishman, A. M., Sair, H. I., Ying, S. H., Jedynak, B. M., ... & Onyike, C. U. (2016). Structural cerebellar correlates of cognitive and motor dysfunctions in cerebellar degeneration. *Brain*, *140*(3), 707-720.

Results

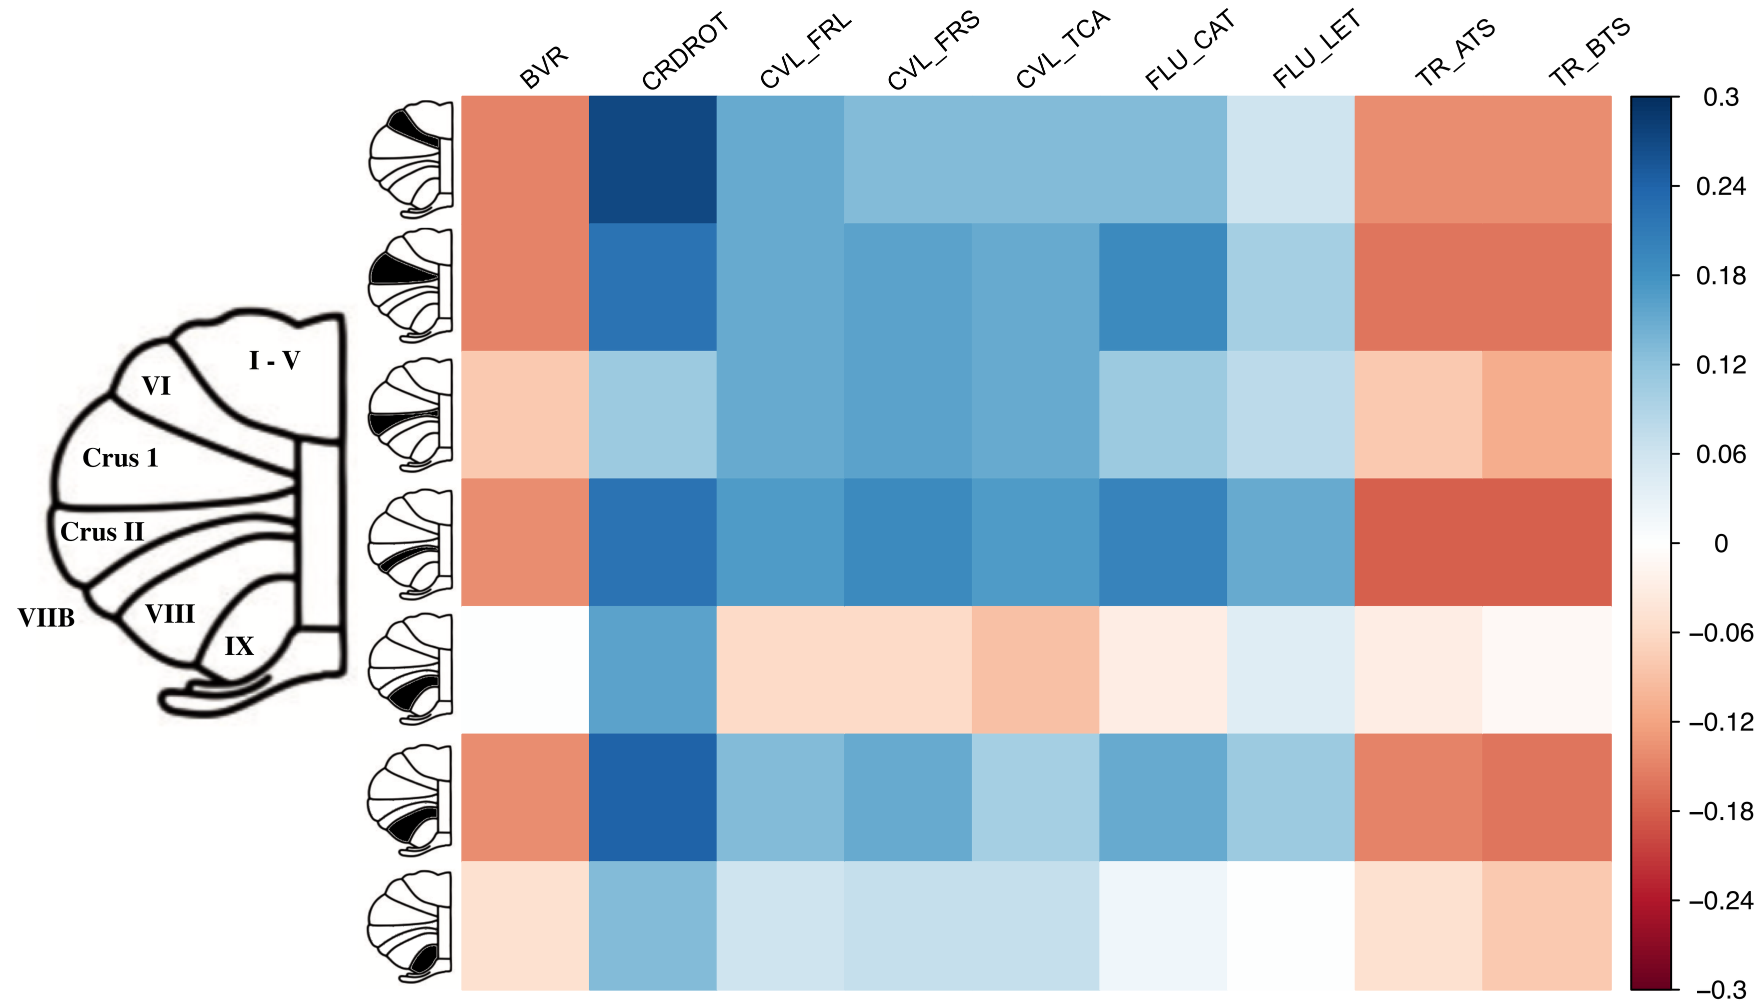

Figure 3. Correlation coefficient (ρ) between regions in the posterior lobe and the cognitive test scores. The rows means different regions of the cerebellar posterior lobe, and the columns means the different cognitive test applied in the BLSA. (BVR: Benton visual retention; CVL: California verbal learning; CRDROT: Card rotation; TR_ATS: Trails making A; TR_BTS: Trails making B; FLU_CAT: Category fluency; FLU_LET: Letter fluency)

Figure 4. Linear mixed model results. (a) shows how the volume of the lobule VI decreases against the age and (b) present how the card rotation cognitive test score decreases against the age.

Discussion and Conclusions

- Parcellation errors in our algorithm could affect the correlation coefficients.
- Volumes of the regions in the posterior lobe are correlated with the cognitive test scores in normal subjects.
- The volumes of the cerebellum regions and the cognitive tests scores decrease in the aging process.

Acknowledgments

Special thanks to the National Science Foundation (NSF) and the Laboratory Computational Sensing and Robotics (LCSR) for supporting this work and the Baltimore Longitudinal Study of Aging, National Institute of Health.

Robot-Assisted Retinal Microsurgery with Real-Time Sclera Force Predictions

Taaha Kamal, Changyan He, Iulian Iordachita

Johns Hopkins Whiting School of Engineering
UT Dallas, The University of Texas at Dallas
Laboratory for Computational Sensing + Robotics
AMIRo Research Lab

Introduction

- Performing surgery on the human eye is an intricate procedure. Unexpected force during surgery could damage the sclera (white part of the eye).
- Robot-assistance can reduce the risk of error but cannot protect against unexpected movement by a user
- Predicting sclera forces in advance can prevent upcoming unsafe manipulations

Purpose

- The project's aim was to keep sclera forces under 140 mN in order to improve the safety of retinal microsurgery
- A neural network was trained to predict sclera forces 25 (ms) in the future and provide active robot feedback if an excessive force is predicted

Methods

- Force sensors on the end of the surgical tool allow us to analyze the force on the eye (Fig 1). The tool's sclera force measurement error is 2.3 mN.
- The data we retrieve from these sensors is used to train a neural network to infer ensuing forces (Fig 2)

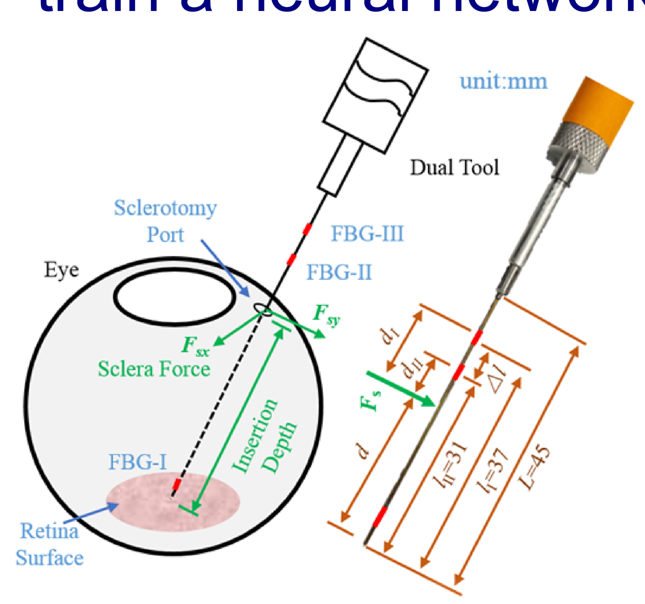

Fig 1: Force sensing tool

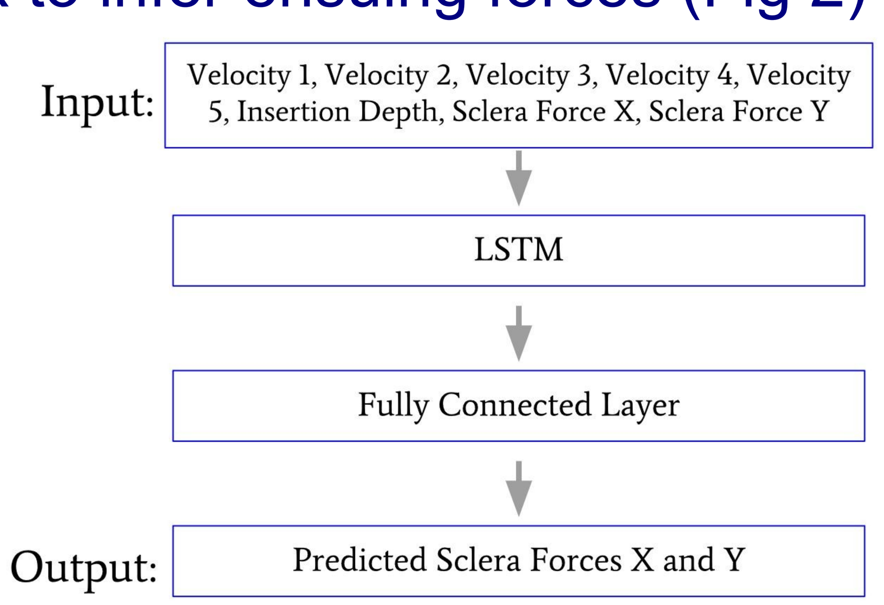

Fig 2: Neural network architecture

- If the predicted sclera forces exceed 140 mN, the velocity is altered to scale down the force to prevent an excessive force on the eye (Fig 3)

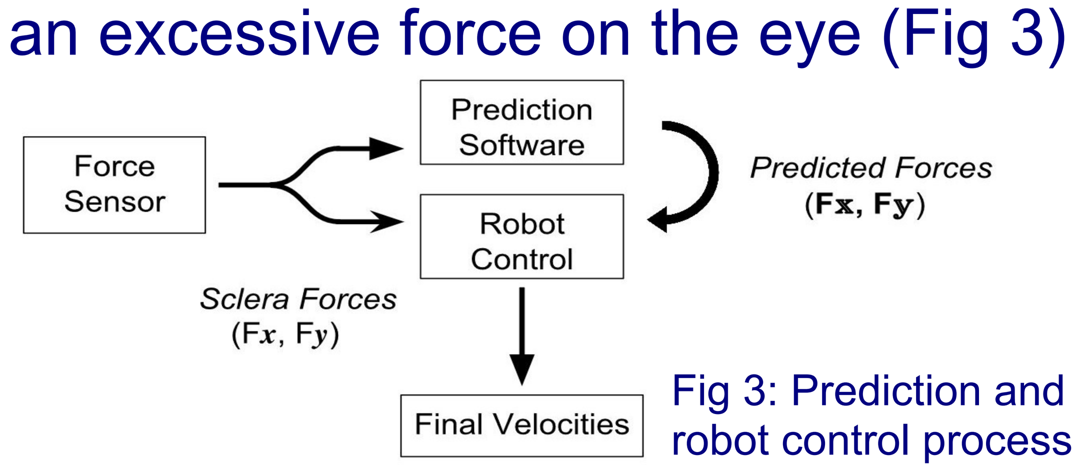

Fig 3: Prediction and robot control process

Methods

Experiment: The force sensing tool is inserted through the scleratomy port (Fig 4) and follows the vessels with and without active control. Sclera forces are then analyzed to determine how reliably the robot's active feedback performed.

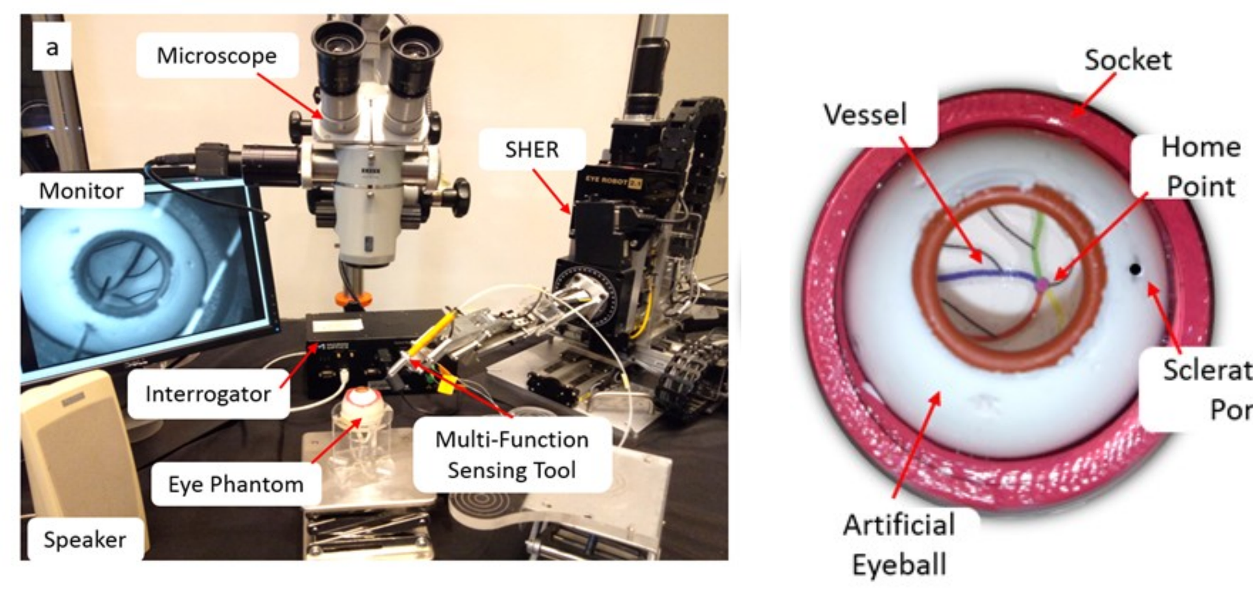

Fig 4: Experimental setup and eye phantom

Results

- Predicting forces 25 milliseconds in the future with collected data resulted in a root mean squared error of: Fx: 6.81 mN, Fy: 6.61 mN
- And max force was: Fx: 120.30 mN, Fy: 198.44 mN

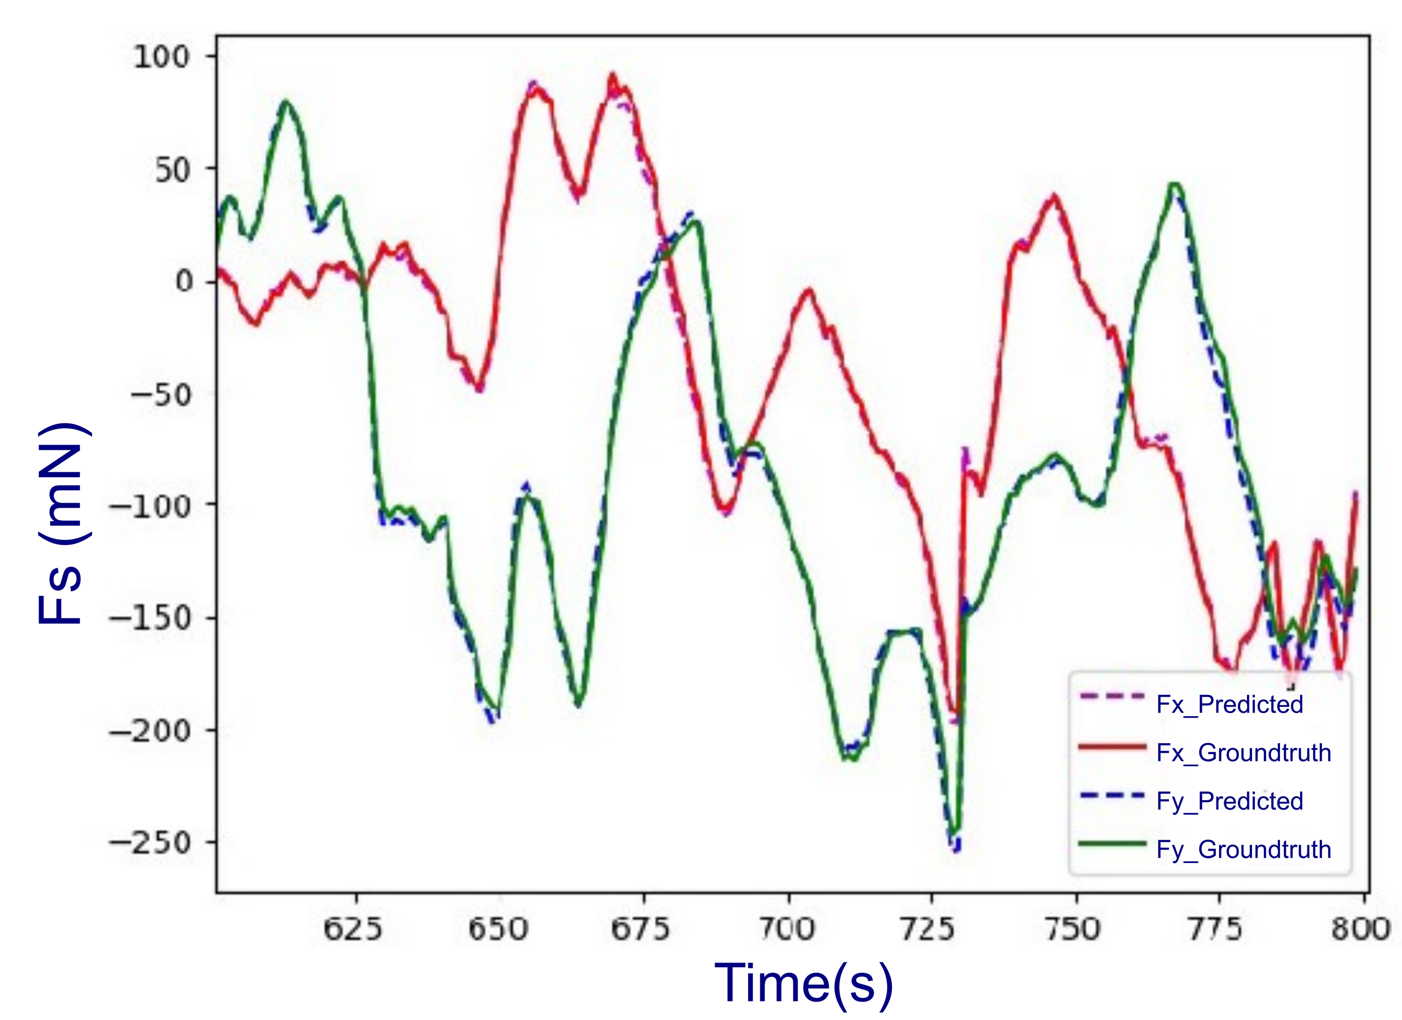

Fig 5: Comparison of groundtruth and predicted sclera forces (Fx and Fy) with active control. Predicting sclera forces 25 milliseconds in the future.

- When tested in real-time, the root mean squared error was: Fx: 14.20 mN, Fy: 16.13 mN
- And max force was: Fx: 221.62 mN , Fy: 150.11 mN

Results

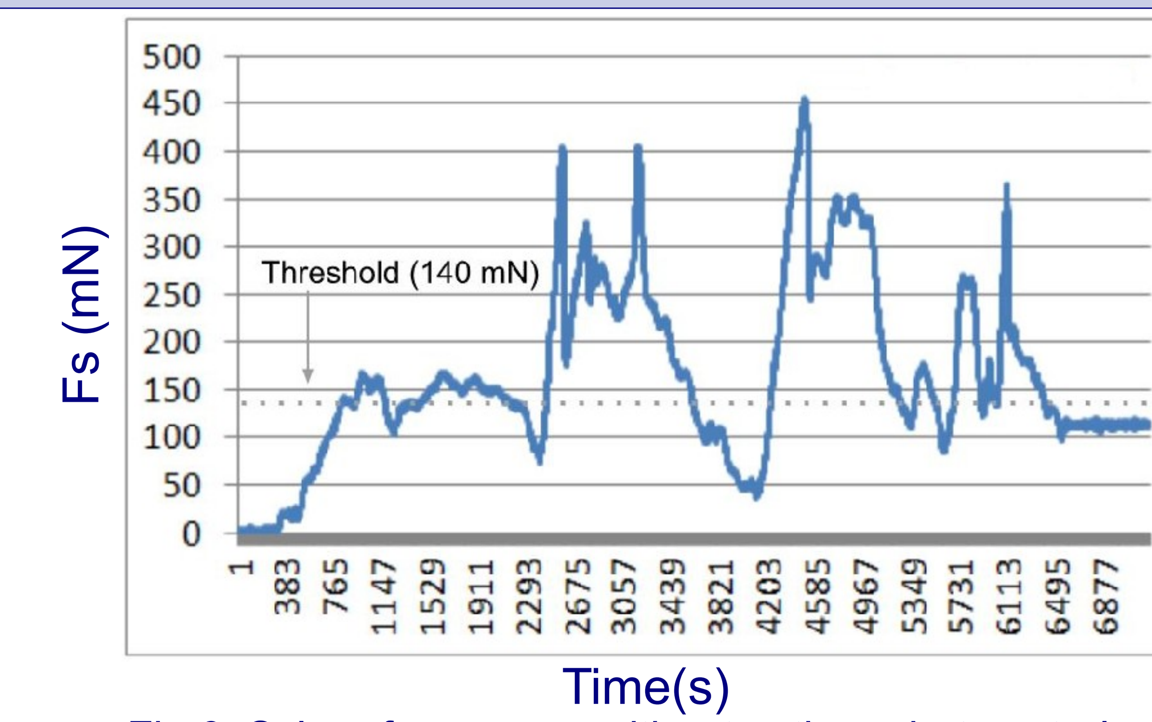

Fig 6: Sclera force norm without active robot control

- Unsafe error percentage: 53.45%

$$\text{Unsafe error percentage} = \frac{(\text{Number of Forces} > 140\text{ mN})}{(\text{Total Number of Forces})}$$

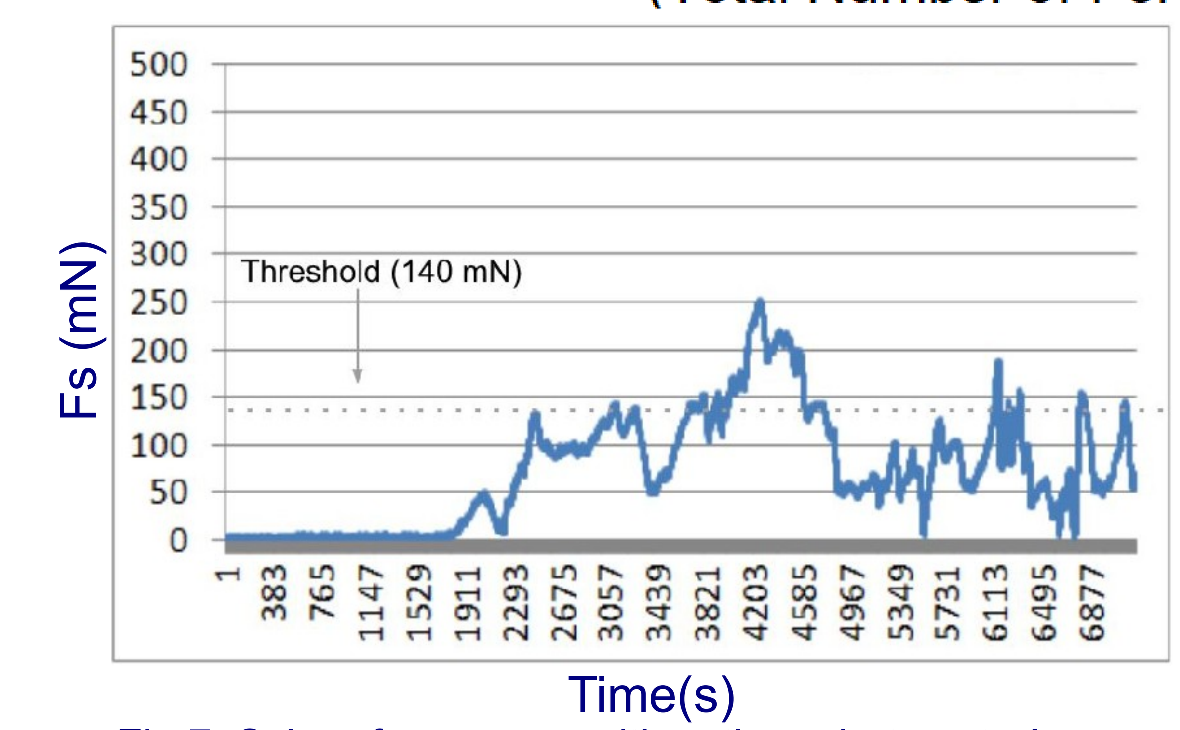

Fig 7: Sclera force norm with active robot control

- Unsafe error percentage: 12.63%
- Active robot control did counter excessive forces, however, the threshold was still passed.
- Increasing prediction time further into the future can increase force insight and decrease error

Conclusions

- The ability to predict future excessive forces and counter them decreases sclera damage
- Robot-assistance with active control enhances safety during retinal microsurgery

Future Directions

- Multi-user study to see how active control is adapted by groups of different levels of surgical expertise

Acknowledgements

U.S. National Institutes of Health under grant number 1R01EB023943-01

Data Acquisition System for Physics-based Robot Obstacle Traversals

Timothy Greco, Yuanfeng (Rick) Han, Yucheng Kang, Yulong Wang, Noah Cowan, Chen Li
Department of Mechanical Engineering, Laboratory for Computational Sensing & Robotics
Johns Hopkins University

Motivation

Many applications require robots to traverse complex 3-D terrain where cluttered obstacles are unavoidable.

Need for Physics-based Control & Planning

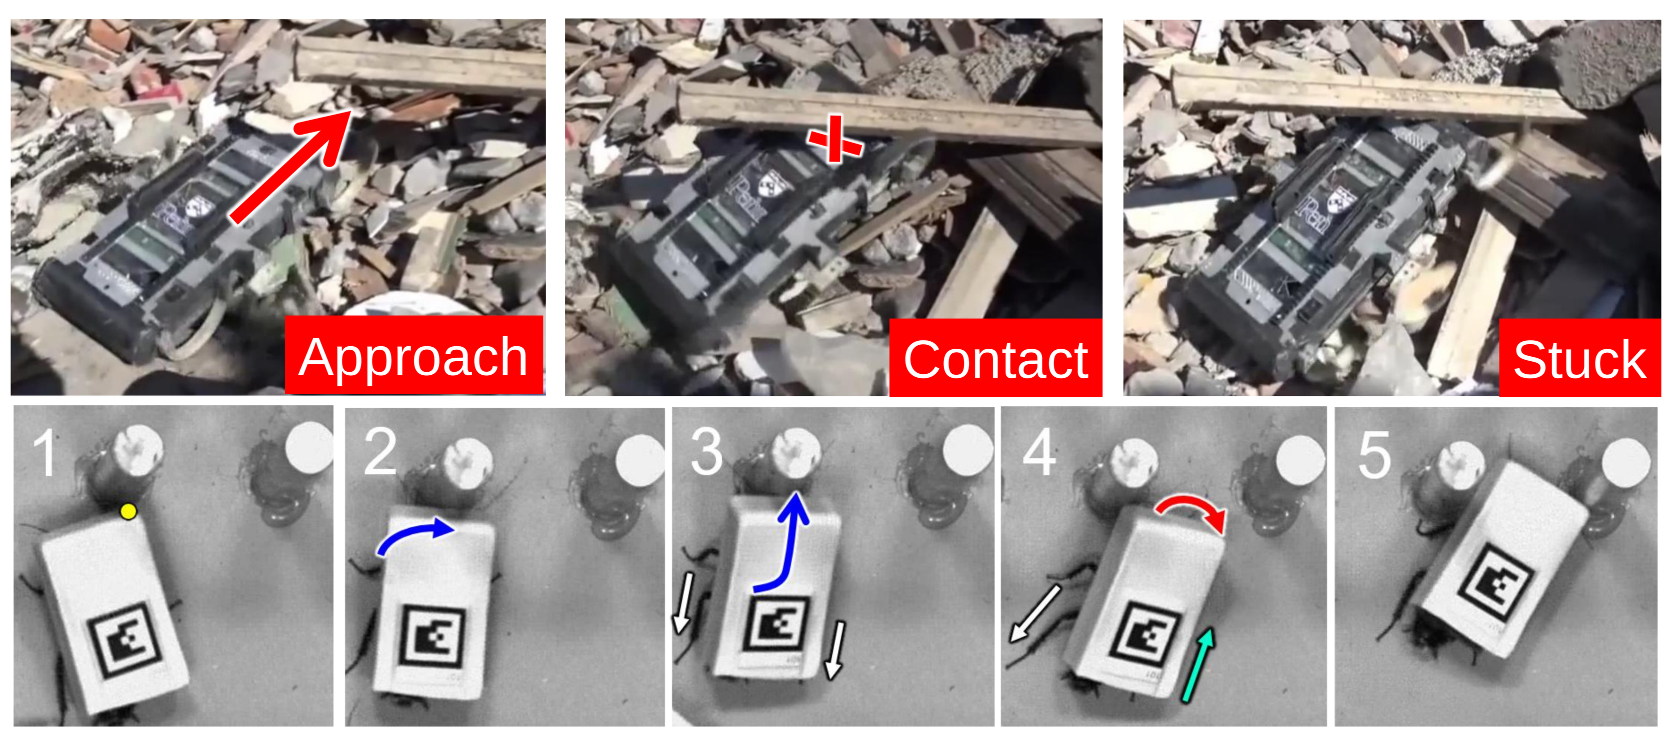

Unlike most robots, animals can actively adjust their motion using mechanical sensing of the physical environment.

Recent Work in Traversal via Physical Contact

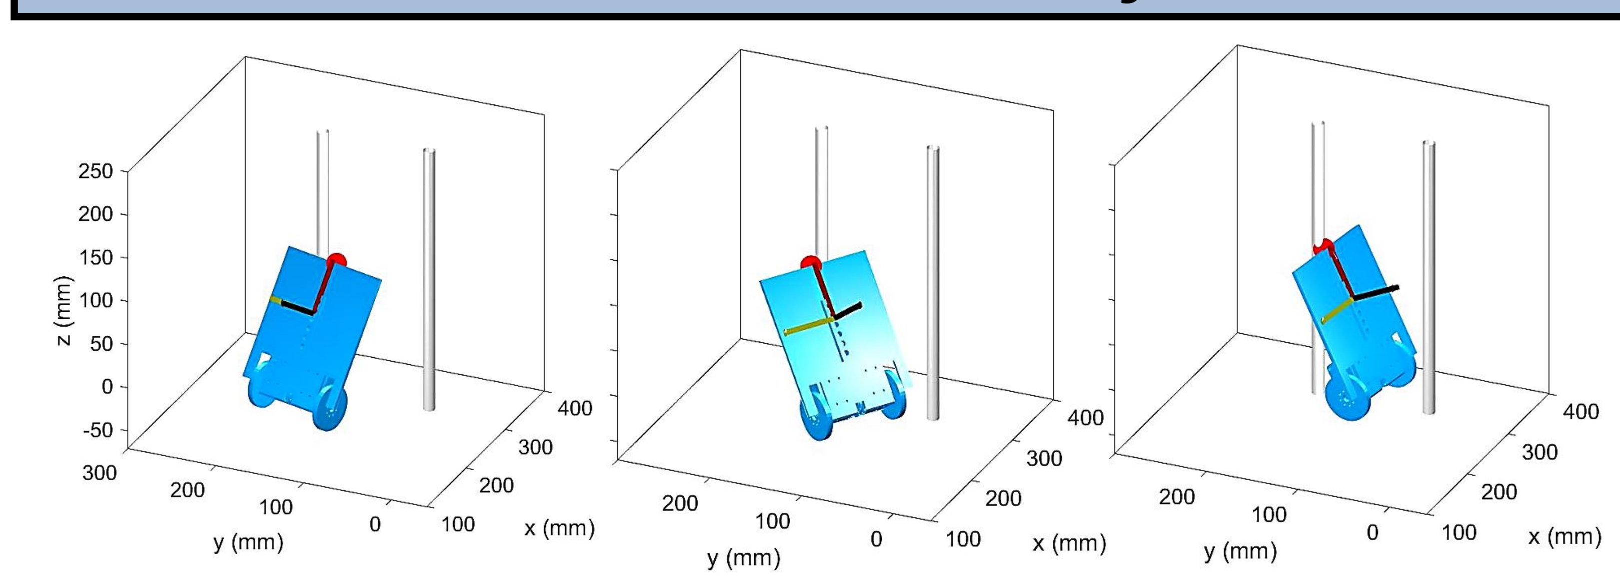

We showed that a wheeled robot can traverse cluttered pillars by properly using physical interaction with obstacles.

Hypothesis

A locomotor (animal/robot) can use internal forces due to physical interaction with the terrain to estimate its terrain-interaction state for control/planning to traverse obstacles.

Experimental Setup

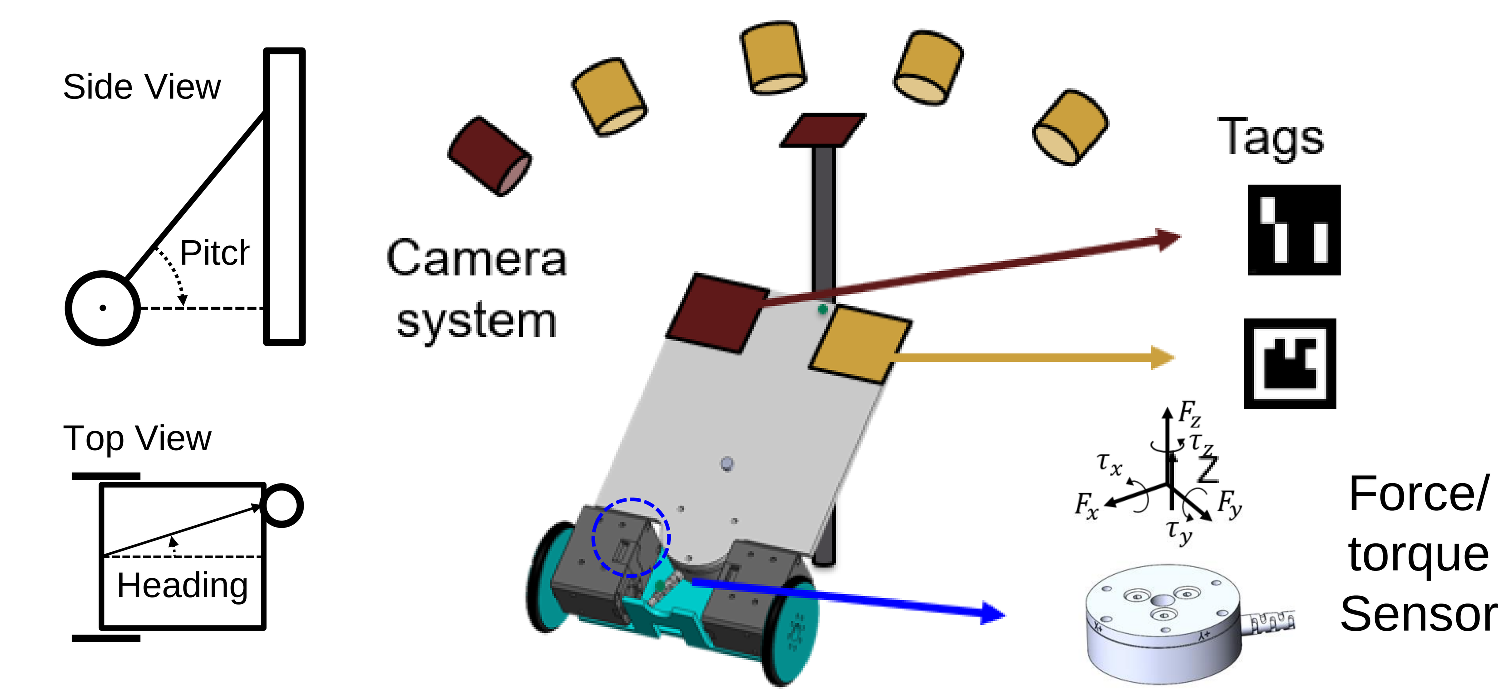

Autonomous Data Acquisition System

An automated experimental system was developed (by Tim Greco) to move the robot and measure its internal force and robot-obstacle interaction state (position and orientation of robot relative to obstacle) over the state space possible (–65º < pitch < –35º, –15º < heading < 15º).

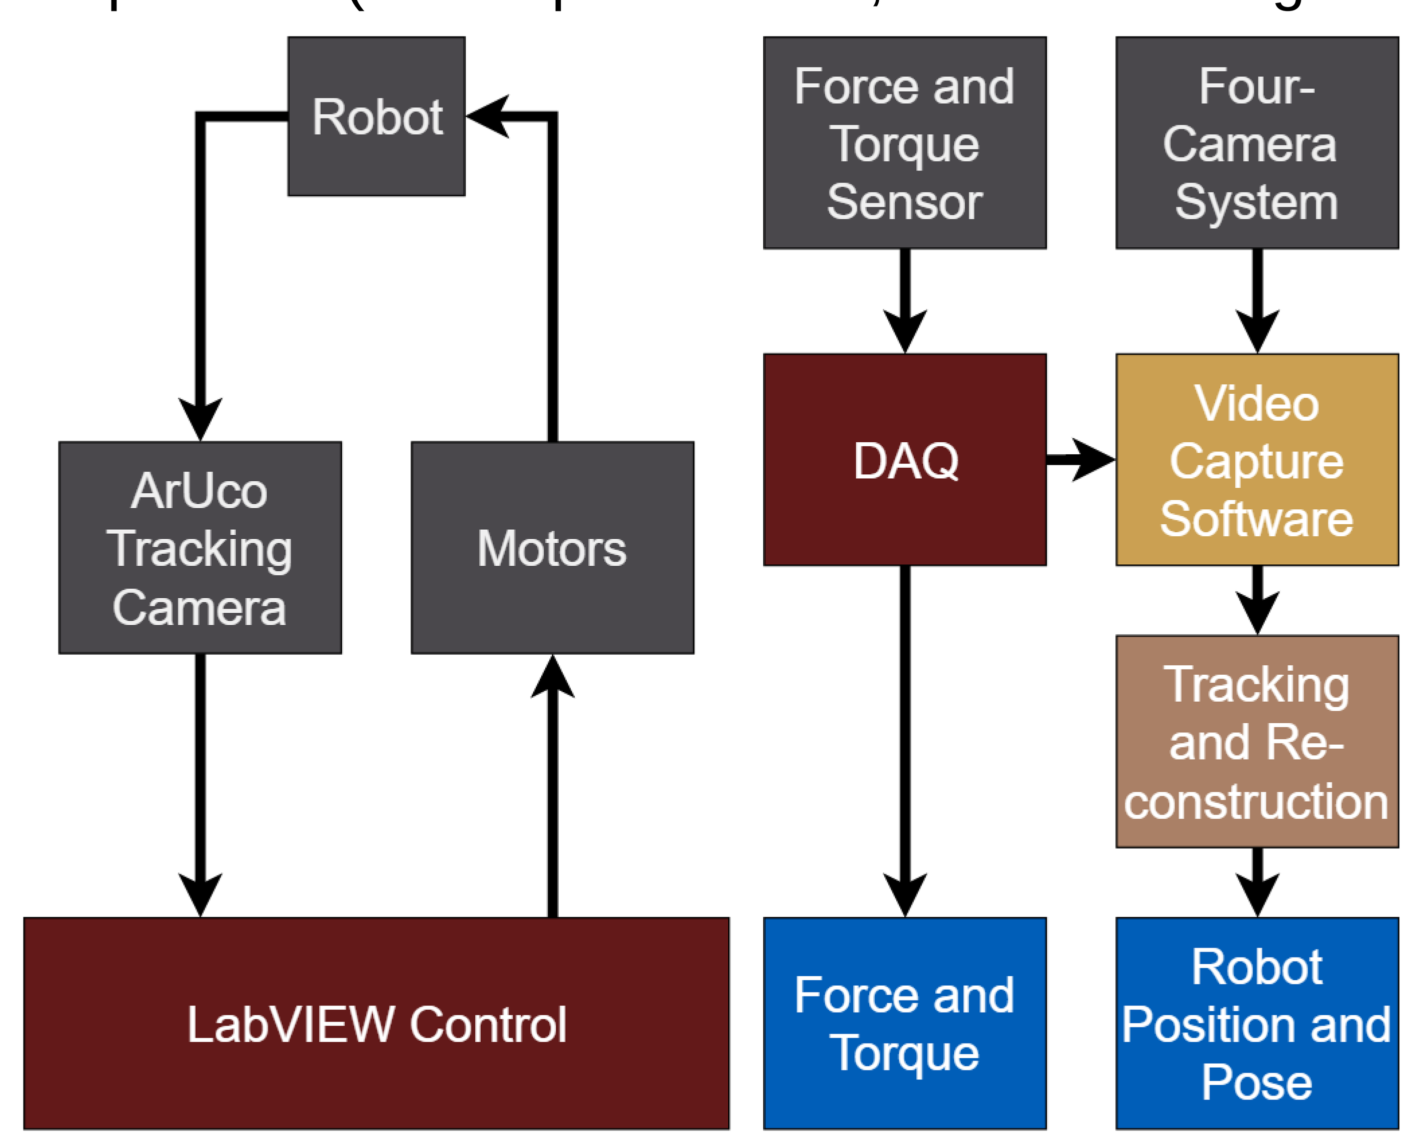

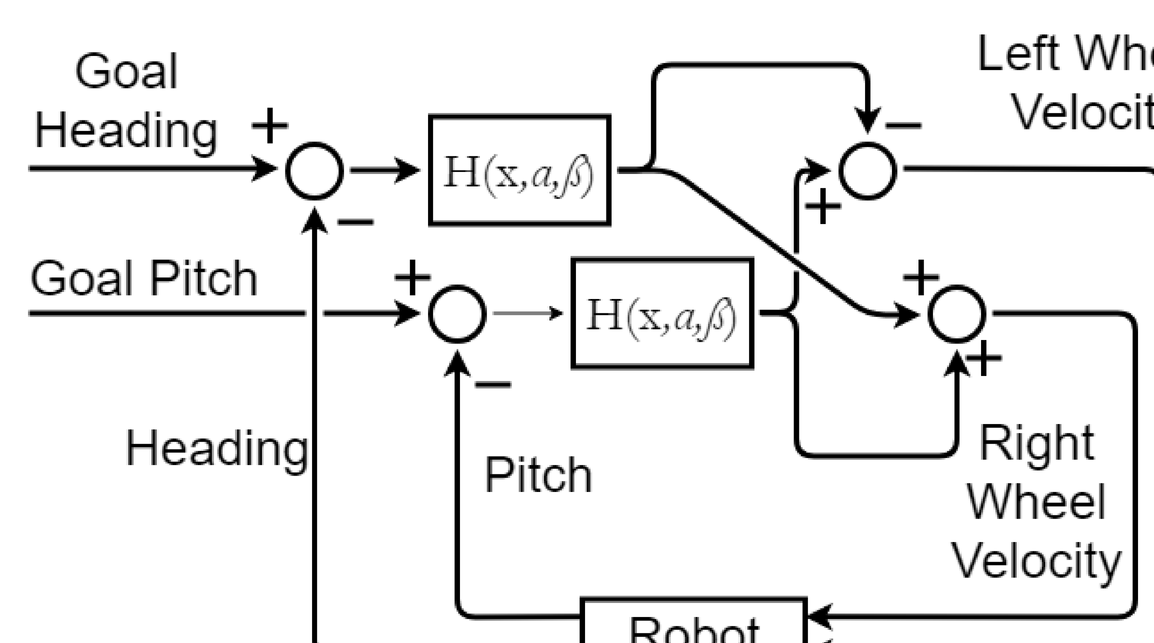

Robot-Obstacle State Estimation Using Forces

$$-F_n \cdot l \cdot \sin\phi + Mg \cdot s \cdot \cos\phi = 0$$
$$F_N \cdot \sin\phi - mg \cdot \cos\phi + F_z = 0$$
$$-F_N \cdot \sin\phi \cdot d + T_x = 0$$
$$\tan\theta = d/l \cdot \cos\phi$$

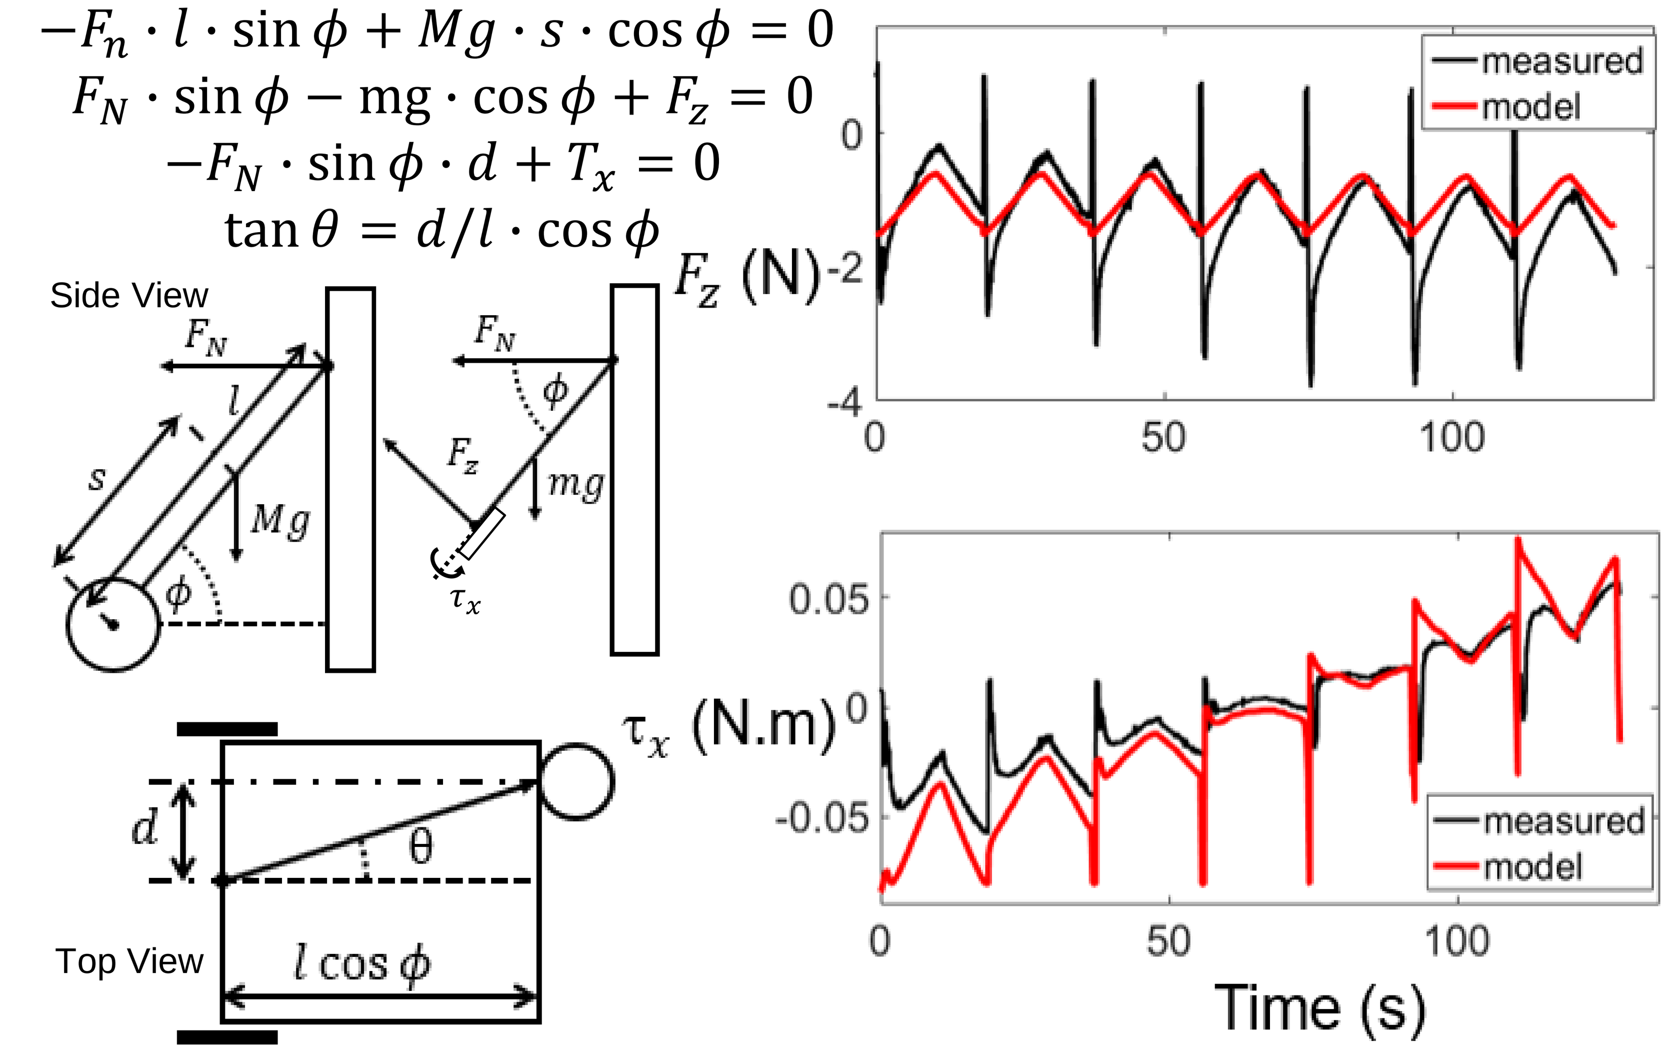

Measured internal forces/torques during traversal of a single obstacle matched well with predictions of a quasi-static force model using robot-obstacle interaction state.

This means that the robot can use internal force to estimate interaction state, confirming our hypothesis.

Future Work: Traversal of Cluttered Obstacles

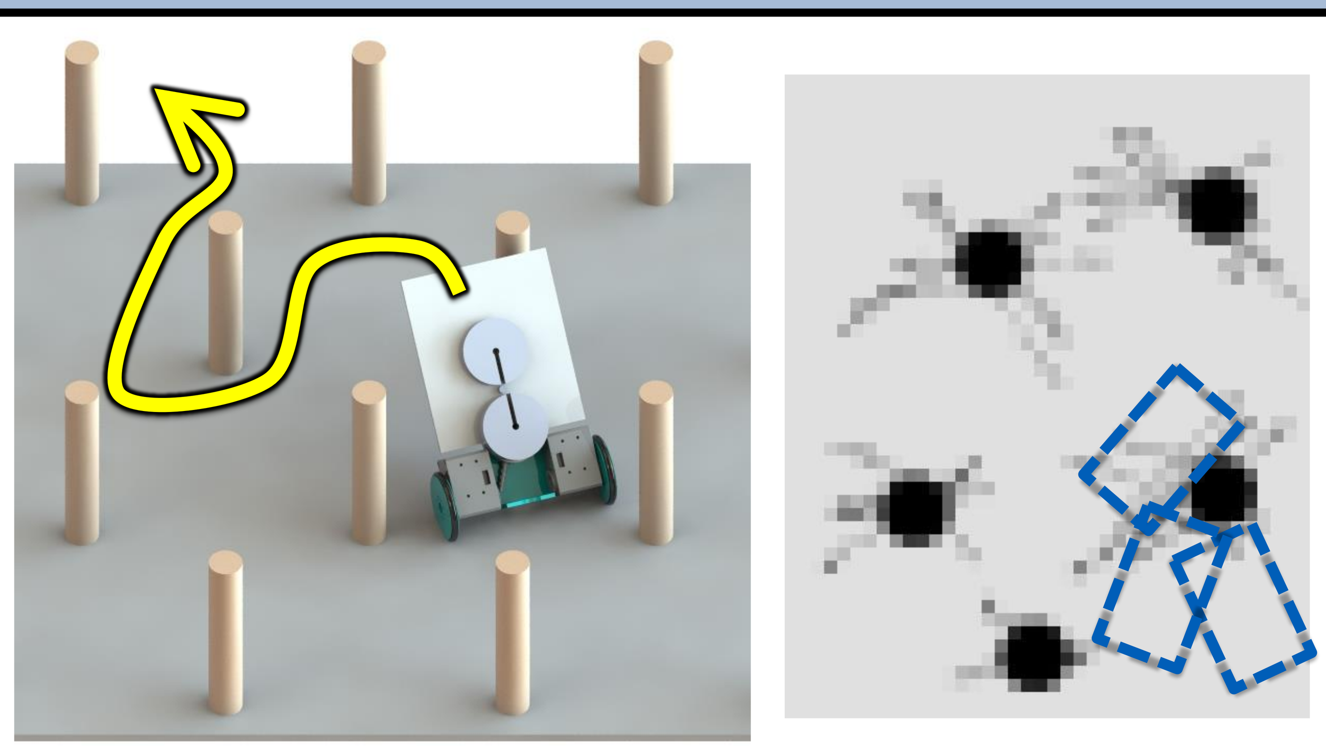

We will develop ability to autonomously traverse cluttered terrain using physics-based control/planning.

Acknowledgements

We thank Marin Kobilarov for discussion, and Yaqing Wang, Qiyuan Fu, and Ratan Othayoth for technical assistance.

This work is funded by NSF CSMR REU (TG) and an Army Research Office Young Investigator Award, a Burroughs Wellcome Fund Career Award at the Scientific Interface, and JHU WSE Startup funds (CL).

Comparing Haptic Modalities in Upper-Limb Prosthetics

JOHNS HOPKINS WHITING SCHOOL of ENGINEERING

Colette McGarvey, Neha Thomas, Jeremy Brown

Johns Hopkins University | Whiting School of Engineering | Baltimore, MD

LABORATORY FOR Computational Sensing + Robotics

Introduction

Amputees often lack haptic feedback in their prostheses, making it both difficult and taxing to perform common activities of daily living. Grasp and lifting an object, for example a cup of water, requires knowledge of grip force and object slip. Since current commercially available myoelectric prostheses lack haptic feedback, amputees are forced to rely almost exclusively on vision, which is known to carry a high cognitive load. Incorporating effective haptic feedback into myoelectric prostheses will thus move one step closer to improving the quality of life for those suffering with limb loss.

Objectives

The purpose of this project is to experimentally compare the utility of two types of haptic feedback, kinesthetic and vibrotactile feedback, in an upper-extremity prosthesis. Our comparison will be based on performance in a stiffness differentiation psychophysics task. We hypothesize that kinesthetic feedback would result in better performance, considering stiffness identification requires knowledge of force. We also hypothesize participants would be most successful in differentiating between hard and soft stiffnesses.

Materials

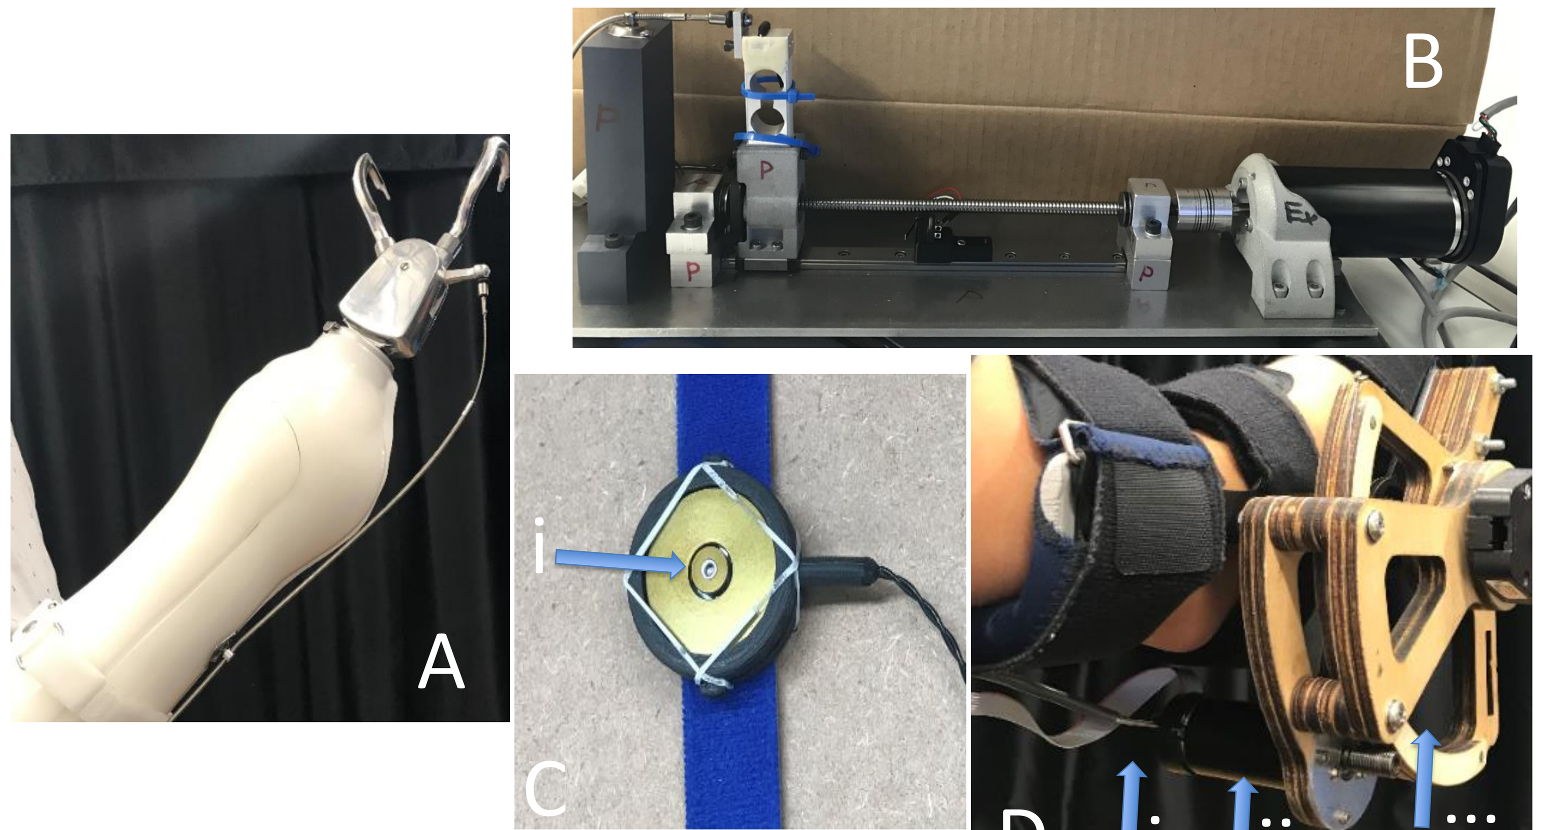

Fig. 1 — A) Myoelectric arm: controlled using EMG signals; works for able-bodied individuals, allowing us to test them instead of amputees (flexing wrist closes prosthetic hand; extending wrist opens prosthetic hand)
B) Prosthesis actuator and load cell
C) Vibrotactile band (vibrotactile feedback): i) C2 tactor
D) Exoskeleton arm (kinesthetic feedback): i) Encoder, **ii)** Maxon motor, **iii)** Capstan drive

- Matlab/Simulink were used to control the myoelectric and exoskeleton arms
- The load cell on the prosthesis drives the signals on the vibrotactile actuator and the exoskeleton torque
- Quanser DAQ board controls signals from load cells and encoders

Methods

Fig. 2: Blocks probed by prosthesis during stiffness identification: Soft (Ecoflex-20), Medium (Ecoflex-30), Hard (Dragon Skin-10)
- No visual feedback from blocks

Fig 3: Experimental setup during stiffness identification: participant wears prosthesis and exoskeleton or vibrotactile band

Protocol:

1. Cross modal matching
- Mapped exoskeleton torque to a vibration amplitude
- Participant wore noise-cancelling headphones playing pink noise to reduce auditory cues

2. EMG calibration
- Electrodes on the wrist flexor and extensor muscles, as well as elbow for grounding

3. Stiffness differentiation task: completed for all three conditions (no feedback, vibrotactile feedback, and kinesthetic feedback); participant was still wearing noise-cancelling headphones
- Participant was presented with 20 pairs of blocks, chosen in a pseudo-random counterbalanced manner from a group of three blocks (soft, medium, hard)
- Participant was instructed to probe each block once and verbalize which block they thought was stiffer

Results

Note: this experiment is still in the pilot stages. For this reason, an example of results for one subject is presented.

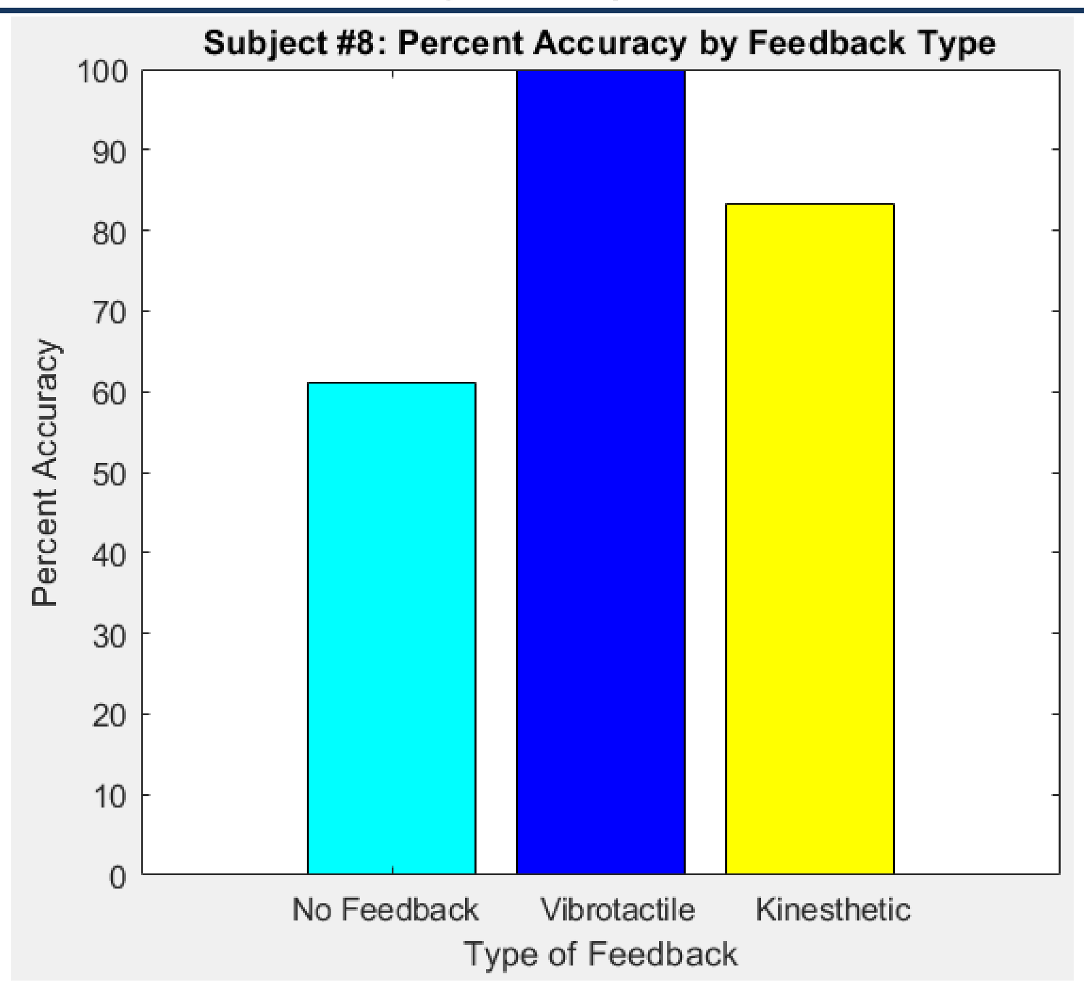

Fig. 4 — Percent Accuracy by Feedback Type

As seen above, the vibrotactile feedback was more effective than kinesthetic in communicating the stiffness of the blocks to the participant. Both feedbacks were superior to no feedback.

Results

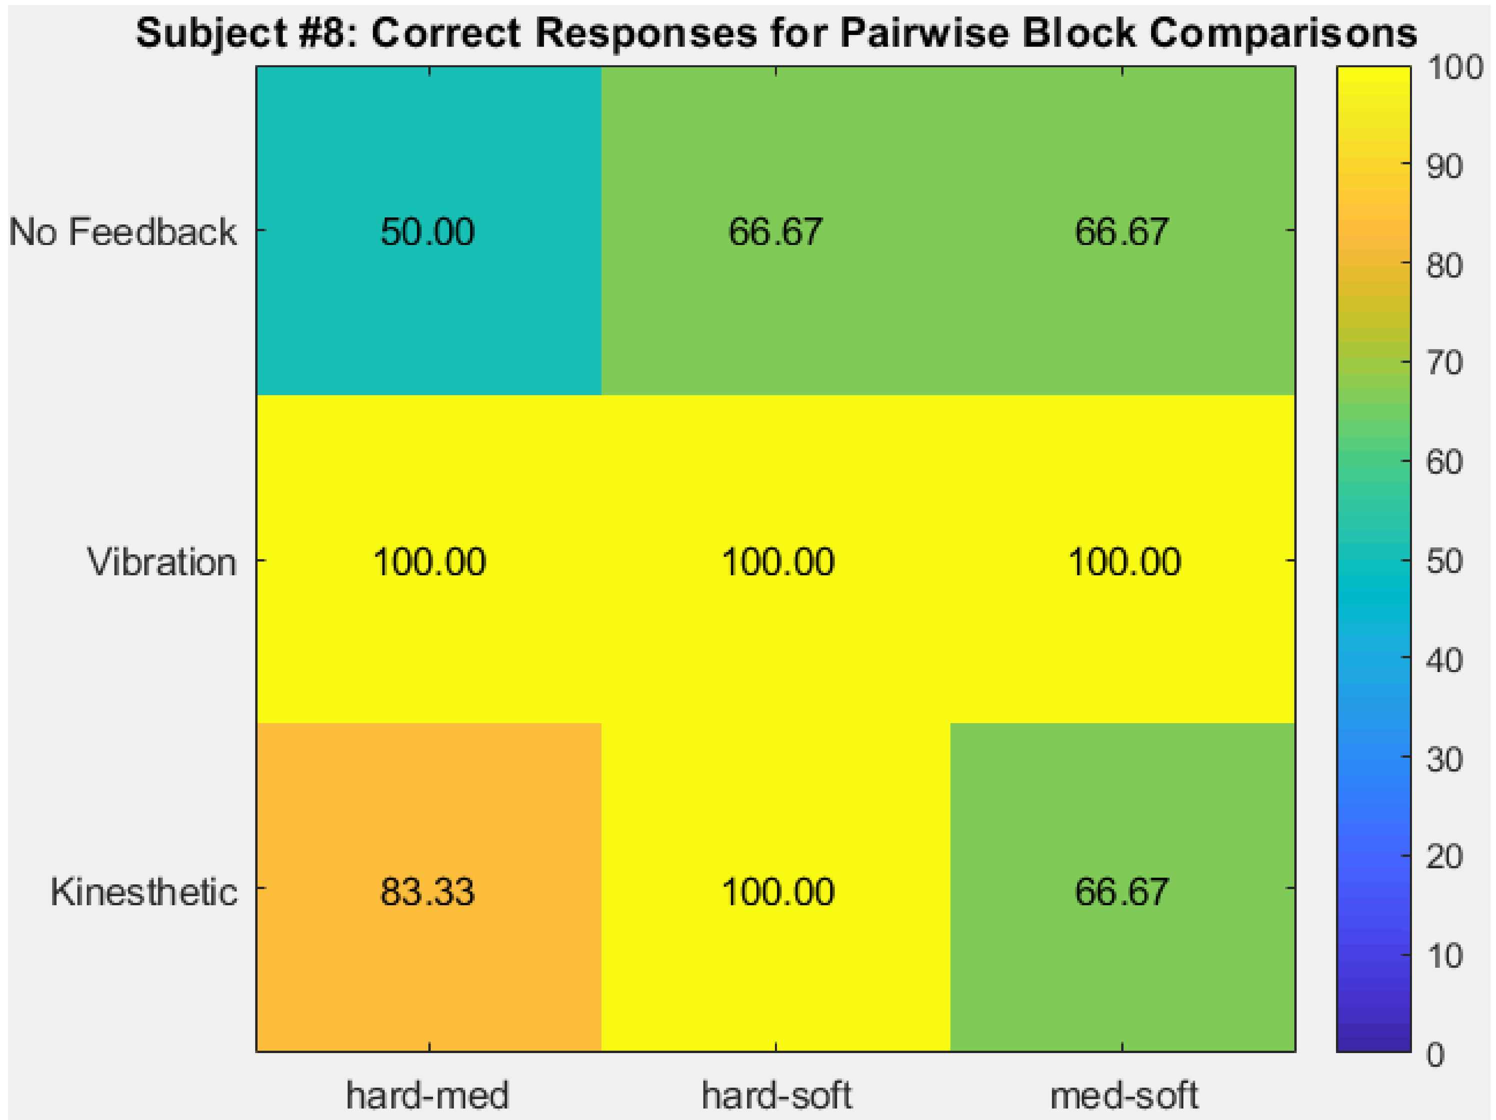

Fig. 5 — Correct Responses for Pairwise Block Comparisons

As expected, the participant had the most success in differentiating between the hard and soft blocks. The yellow regions in the above figure show 100% accuracy in the subject's responses for both the vibrotactile and kinesthetic feedback. It would not be expected for the subject to receive 100% in the no feedback condition.

Future Work

Continue fine-tuning the details of the protocol. We are currently investigating whether hand dominance has an effect on our results. We are also modifying our control algorithm to include a bias adjustment for the EMG signals, as they have been observed to drift during the experiment, causing difficulties in controlling the myoelectric arm.

Conclusion

For this pilot participant, we found that vibrotactile feedback provides more utility than kinesthetic feedback in differentiating between objects of various stiffnesses. Both feedback types were superior to no feedback. This is a promising result that hopefully will hold true when conducting actual experimentation. This would further support the presumption that it would be beneficial to incorporate haptic feedback into myoelectric prostheses.

Acknowledgements

I would like to acknowledge my research mentor, Neha Thomas (second-year Ph.D. student at Johns Hopkins University), as well as my lab's Principal Investigator, Dr. Jeremy Brown, for their support, encouragement, and guidance throughout the conduction of my research for this project. This project is funded by an NSF REU supplement to grant# CHS1657245.

Development of the JHU ROV II

Laboratory for Computational Sensing + Robotics

Katherine Mao, Florian Pontani, Aiden Devaney, Charlie Watkins, Eli Pivo, Tyler Paine, Louis Whitcomb

Johns Hopkins Whiting School of Engineering

Background and Goals

JHU ROV (**R**emotely **O**perated underwater **V**ehicle)

- Test bed for experimental research in ROV control and navigation algorithms
- 15 years old

JHU ROV II – Existing Design

- Continuation of previous two year's REU work

Project Goals

- Assemble power housing Mark I
- Finalize CPU housing design
- Design JHU ROV II frame & trim
- Manufacture oil compensator

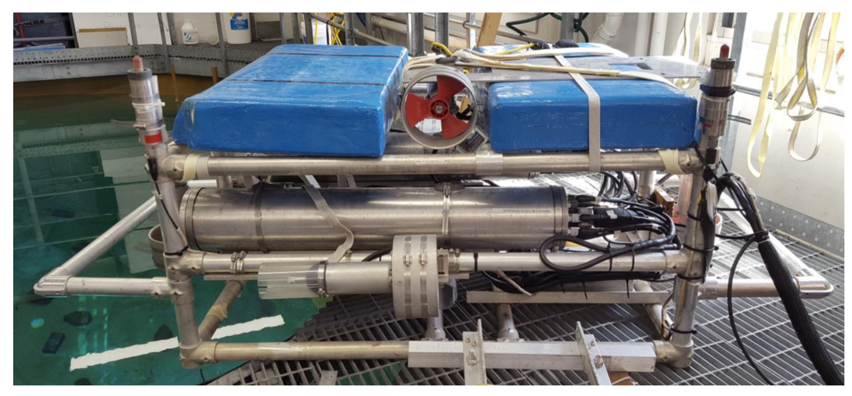

JHU ROV sitting on deck in the Dynamics System and Control Laboratory (DSCL)

Power Housing

Mark I

- Remove battery wheels
- Power JHU ROV II via tether
- Extra connector needed to supply power

Mark II

- Redesign power distribution boards & charging circuitry
- Modify battery wheels
- Power JHU ROV II via tether or batteries

Modified Connector Layout

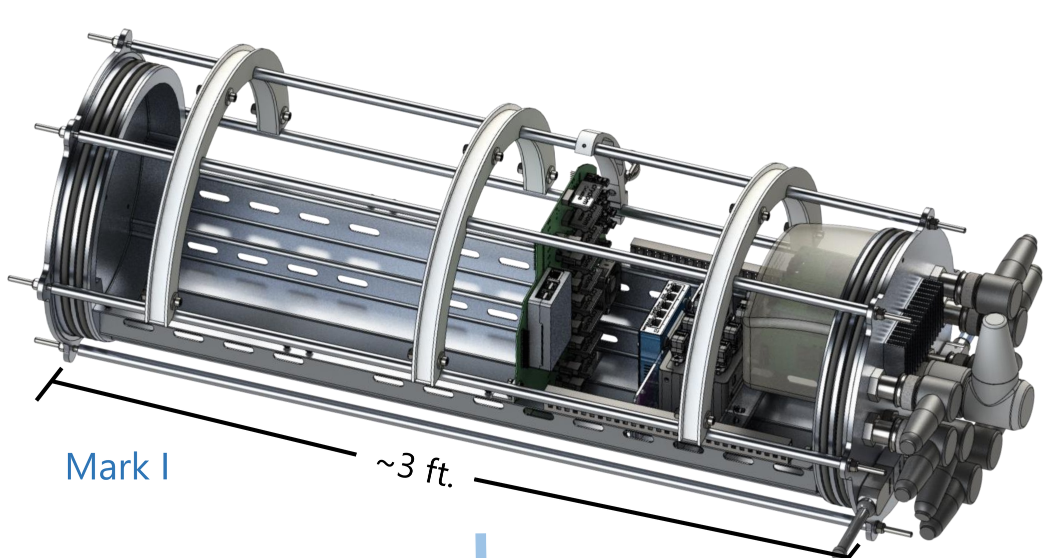

Mark I

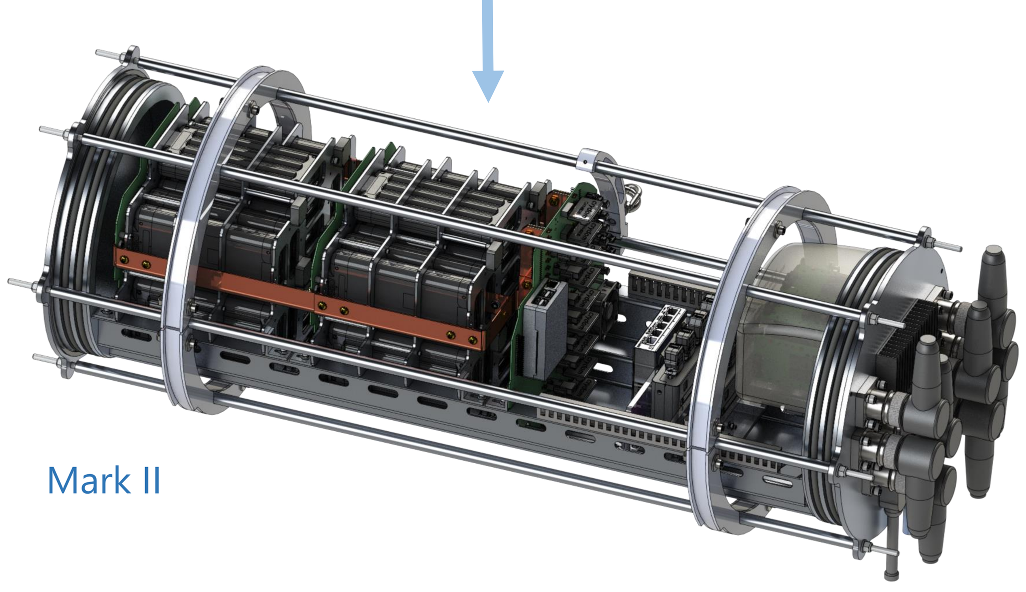

Mark II

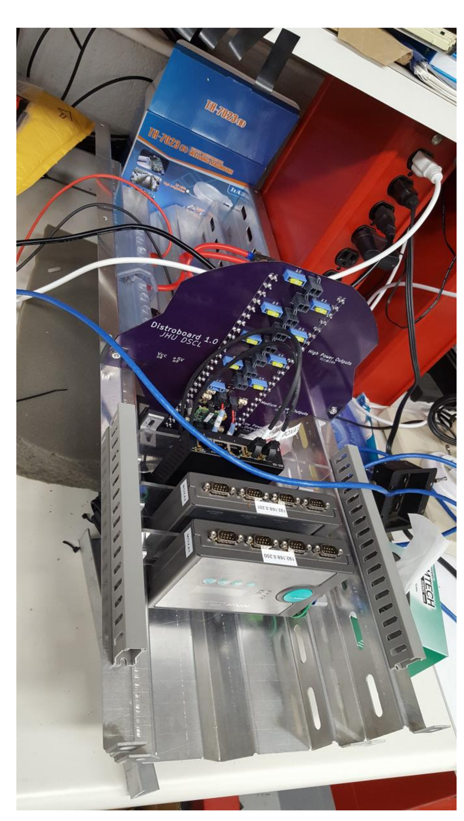

Mark I Circuit Testing

CPU Housing

Updates

- Extended length of housing from 20" to 26"
- Designed mounting for Vicor DC-DC converters
- Added chassis roller wheels
- Modified connector layout

Vicor DC-DC Converter mounting

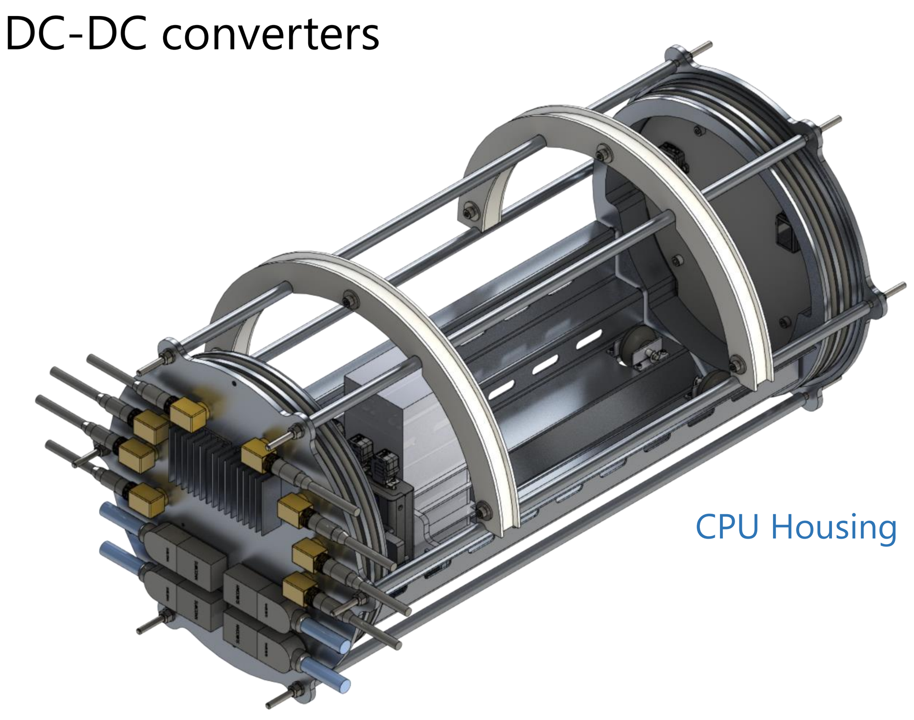

CPU Housing

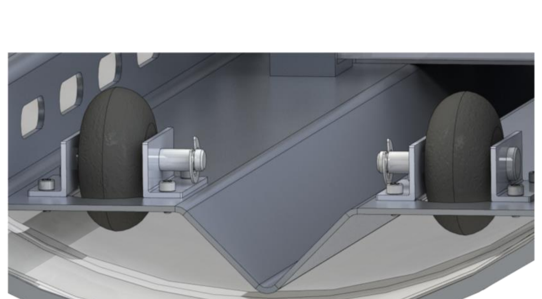

Chassis Roller Wheels

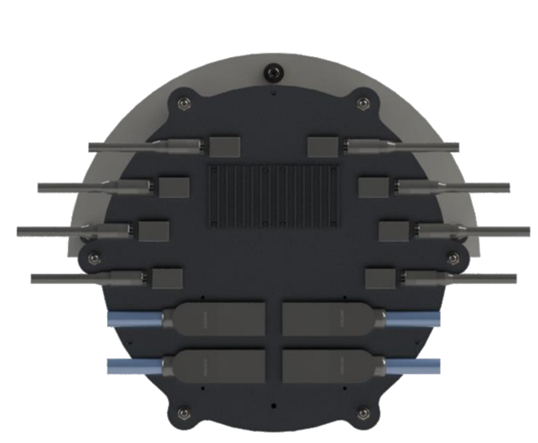

Modified Connector Layout

Frame

Objectives

- Improve frame design for JHU ROV II
- Layout housings & sensors onboard JHU ROV II

Previous State of Design

- JHU ROV I – 52" x 33" x 22"
- Smaller than JHU ROV I – 42" x 35" x 29"
- No room for wiring
- Old thruster model

Features

- 1-1/4 Pipe size speed rail
- Dimensions: 43" x 48" x 33"
- 6 thrusters, 6 degrees of freedom
- Improved design for housing mounting
- Inertial Measurement Units (IMU) and Doppler Velocity Log Sonar (DVL) aligned along frame centerline
- Hard-lift point for transportation
- Bumper guards for thrusters

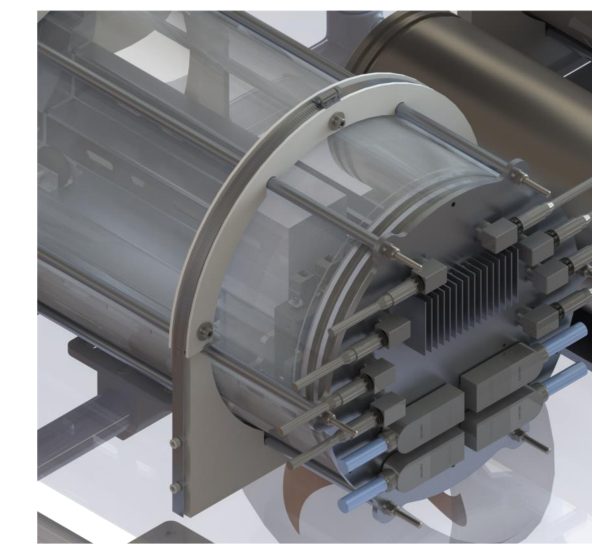

Housing Mounting Holders

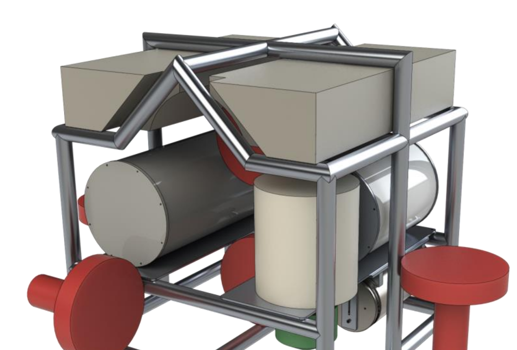

Preliminary JHU ROV II Frame Design

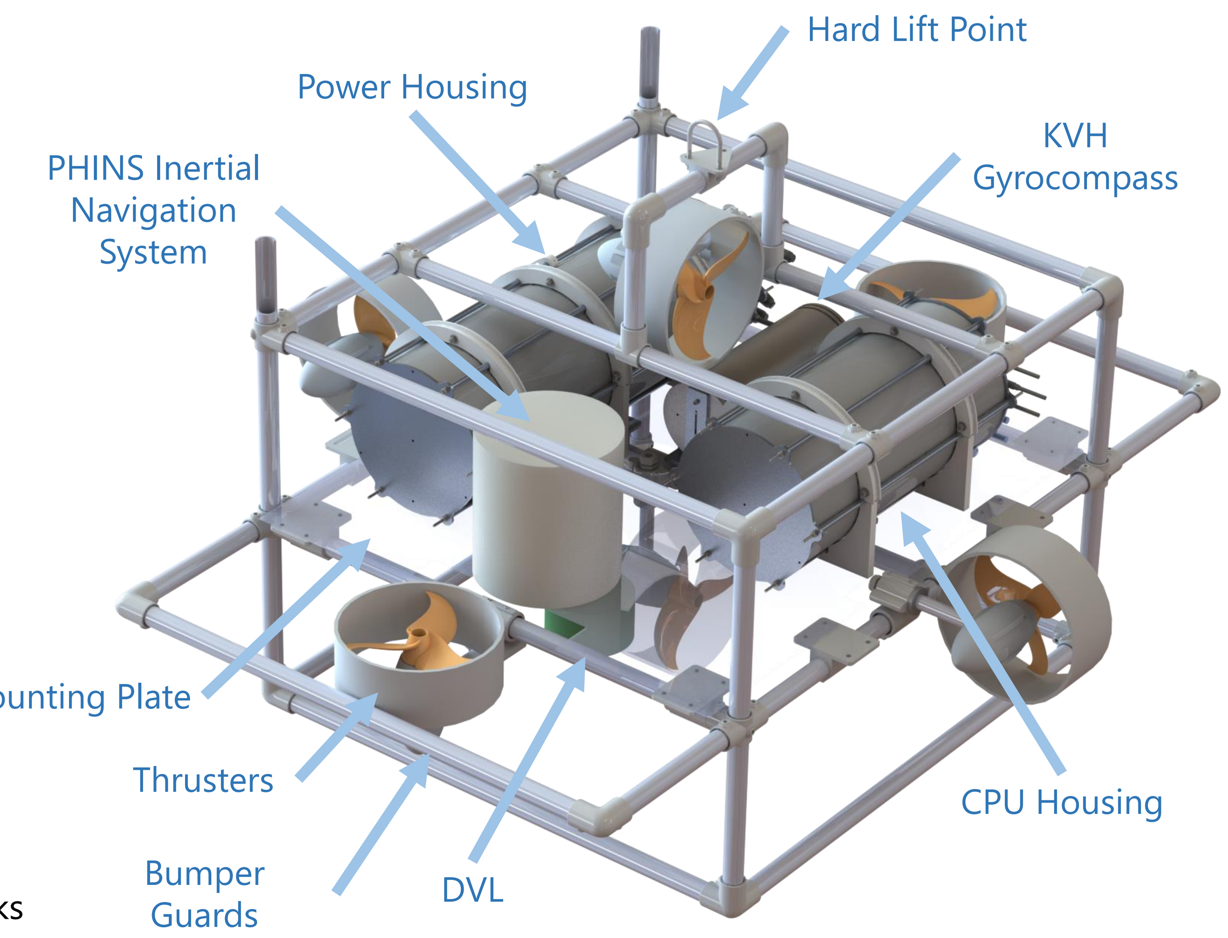

Latest JHU ROV II Frame Design

Not Pictured
- Flotation blocks
- SHARPS system
- Depth Sensor
- Oil Compensator

Oil Compensator

Background and Goals

- Torqeedo Travel 1003 L Thruster
- Original lip seal failed waterproofing test
- Attempted replacement seals failed waterproofing test
- Flood internals with mineral oil

Torqeedo Travel 1003 L Thruster

Design

- Six thrusters linked to central compensator
- Watertight rubber sock holds mineral oil
- Spring (k = 12 lbs/ft) compresses oil
- Quick disconnect connectors from thruster
- Flooded housing

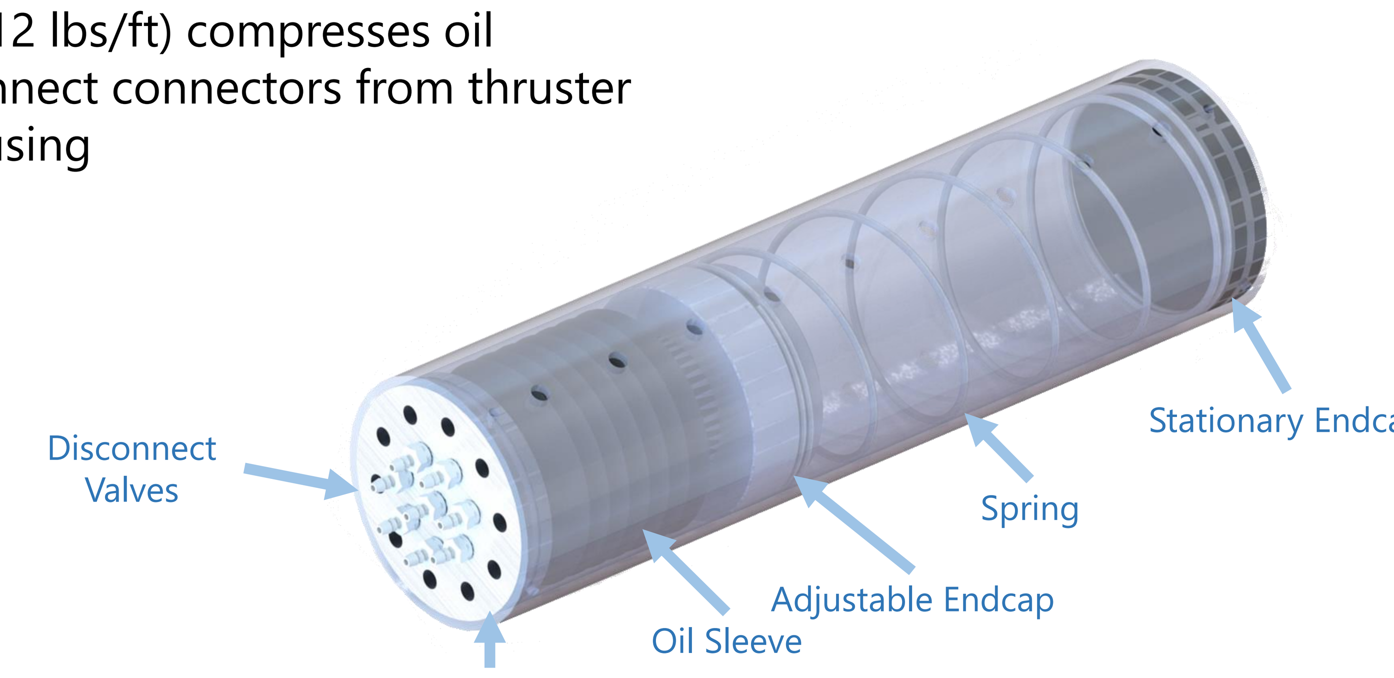

Future Work

- Design & Manufacture Power Housing Mark II
- Assemble & Test CPU Housing
- Manufacture & Assemble Frame
- Test Oil Compensation System

Acknowledgements

CSMR

Thanks to Dana Walter-Shock, the CSMR REU program coordinator, and to Zak Harris, Andrew Spielvogel, Rachel Hegeman, and Andy Cohen

FPGA Firmware Development for the Da Vinci Research Kit

Mikiyas Bokan , Prof. Peter Kazanzides
Laboratory For Computational Sensing and Robotics
Johns Hopkins University

Introduction

Filters play a significant role in digital signal processing systems. Filters improve signal quality by discriminating certain part of a signal and retaining another part of a signal. Therefore, filters are used to remove noise from signal.

Filters can be classified as analog and digital based on implementation.

The main purpose of filters in the Da Vinci Research Kit is to clean potentiometer signals and motor currents.

Currently the DVRK systems in the field are equipped with Analog filters that implement a low pass filter with a cutoff frequency of 60hz.

A low pass filter is a type of filter which eliminates signals whose frequency is above the cutoff frequency while retaining signals whose frequency is below the cutoff frequency.

The signals that are filtered in this project are potentiometer signals. Potentiometers have two major purposes in the DVRK system. These are preloading the encoder values and for performing safety checks.

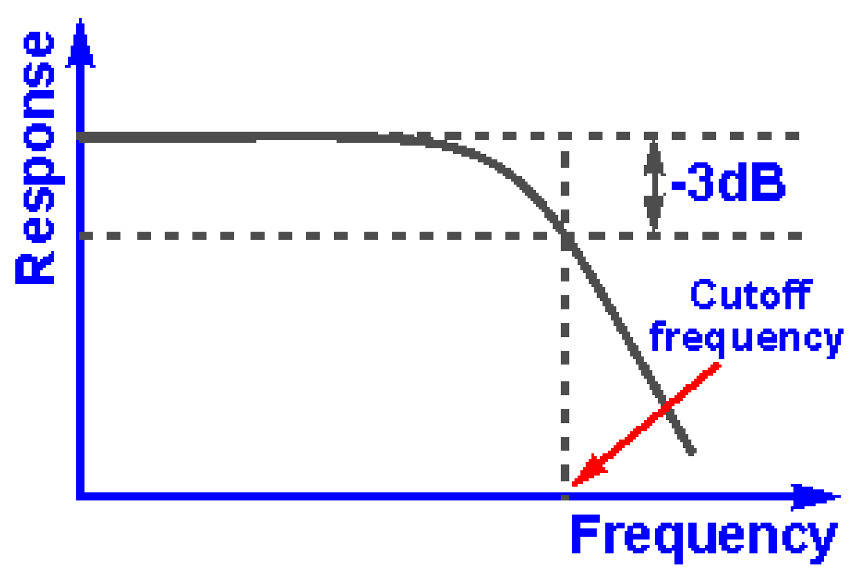

Fig 1. Frequency diagram for a low pass filter

Problem statement

While a 60hz analog filter performs significant amount of filtering, the cut off frequency can't be adjusted via software.

Hence analog filters aren't suitable in situations where we need to change the cut off frequency as desired.

Another issue that arises from analog filters is that they can sometimes be affected by thermal and environmental factors. Moreover, filtering at 60hz is associated with a significant amount of delay during operation. Hence, filtering at a low frequency makes multiplexing between setup joints slow.

Method

In this research project we propose replacing analog filters with digital filters. Digital filters perform filtering by executing a set of mathematical operations on a sampled discrete time signal. One the main advantages of using a digital filter is that the cutoff frequency can adjusted to produce various types of filters. Digital filters can be reprogrammed via software. All of these features makes digital filters more suitable for the DVRK system. For this project digital filters are implemented in Field Programmable Gate Array (FPGA). The particular type of FPGA used is the IEEE-1394 FPGA Controller. FPGAs have a special hardware known as DSP48 slice which is dedicated for digital signal processing. FPGAs expedite computation because of their parallel architecture.

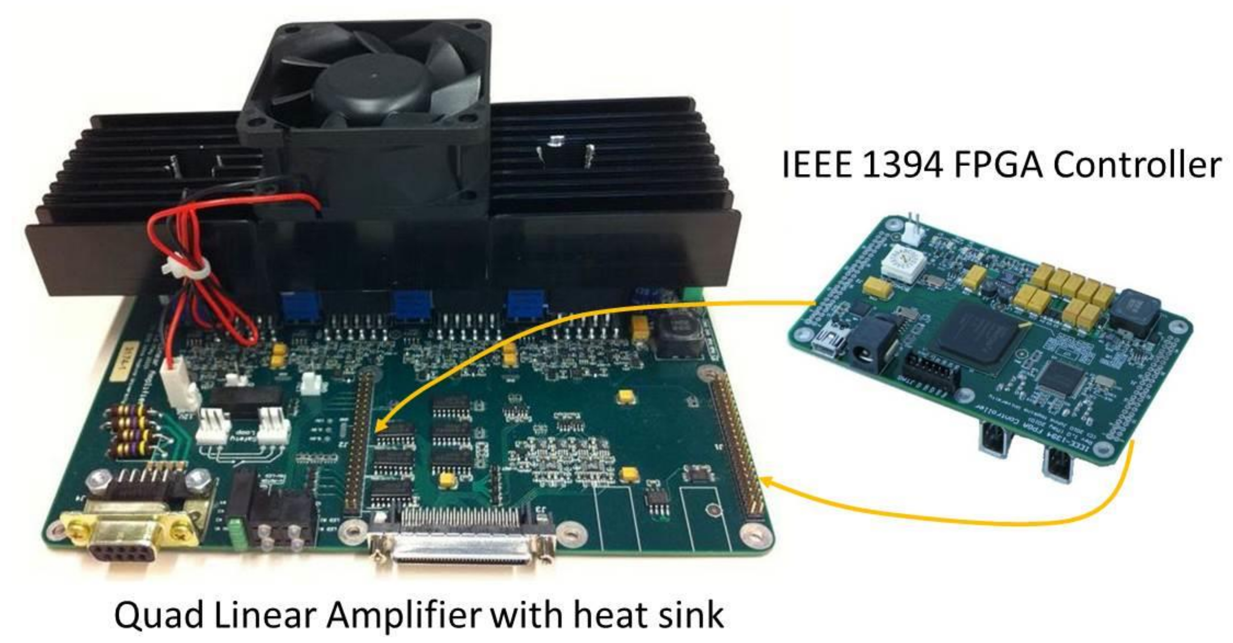

There are two broad categories of digital filters based on implementation. These are Finite Impulse Response (FIR) and Infinite Impulse Response(IIR) filters. For this project we chose FIR filters. FIR filters are filters in which the output is expressed as the weighted sum of past inputs. A hardware description language called Verilog is used to program the FPGA. The Xilinx ISE software package comes with an FIR compiler that allows a user to produce the desired filter.

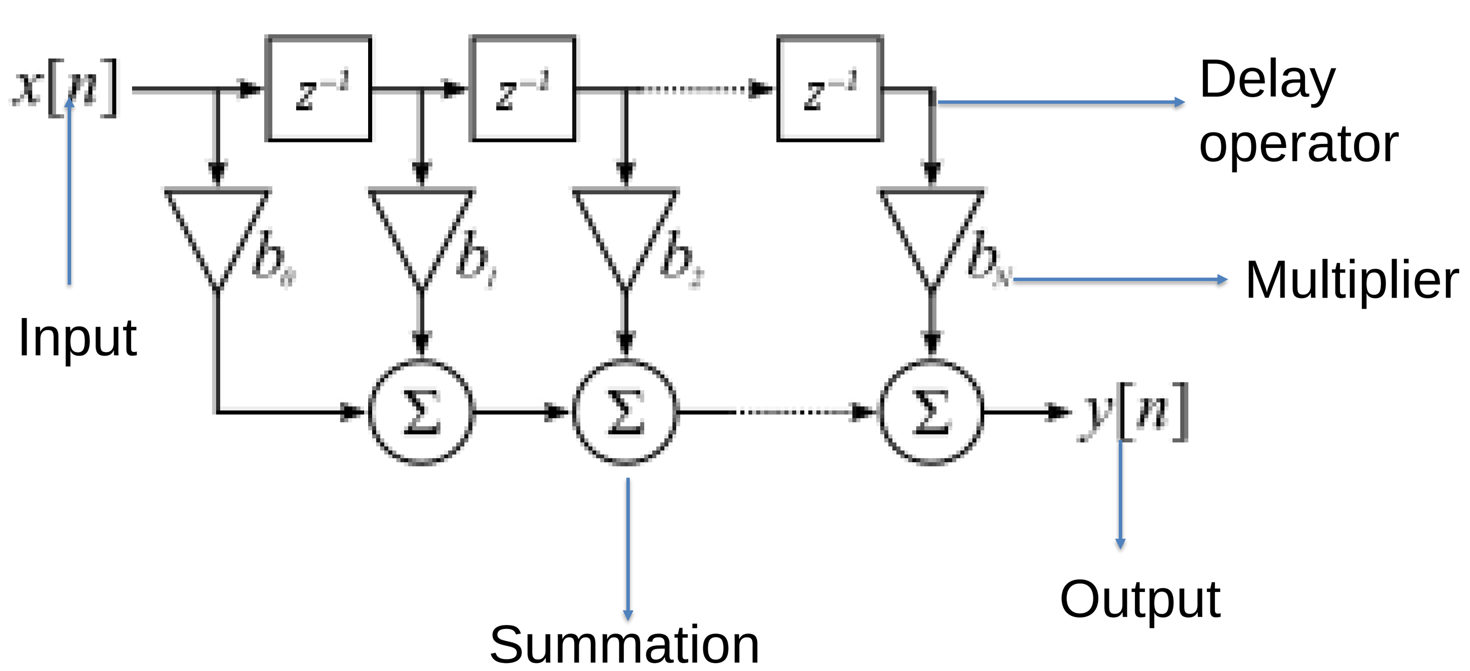

Fig 3. Block Diagram for FIR filters

Results

We have implemented a low pass filter with a cutoff frequency of 5khz. A software package called Chipscope is used to perform low level debugging. The figure below shows a 5khz lowpass filter passing a 1khz input signal. The yellow sine wave is the input 1khz signal while the blue sine wave is the filtered output. The sampling frequency is 0.02MHZ while the delay is 0.5 millisecond.

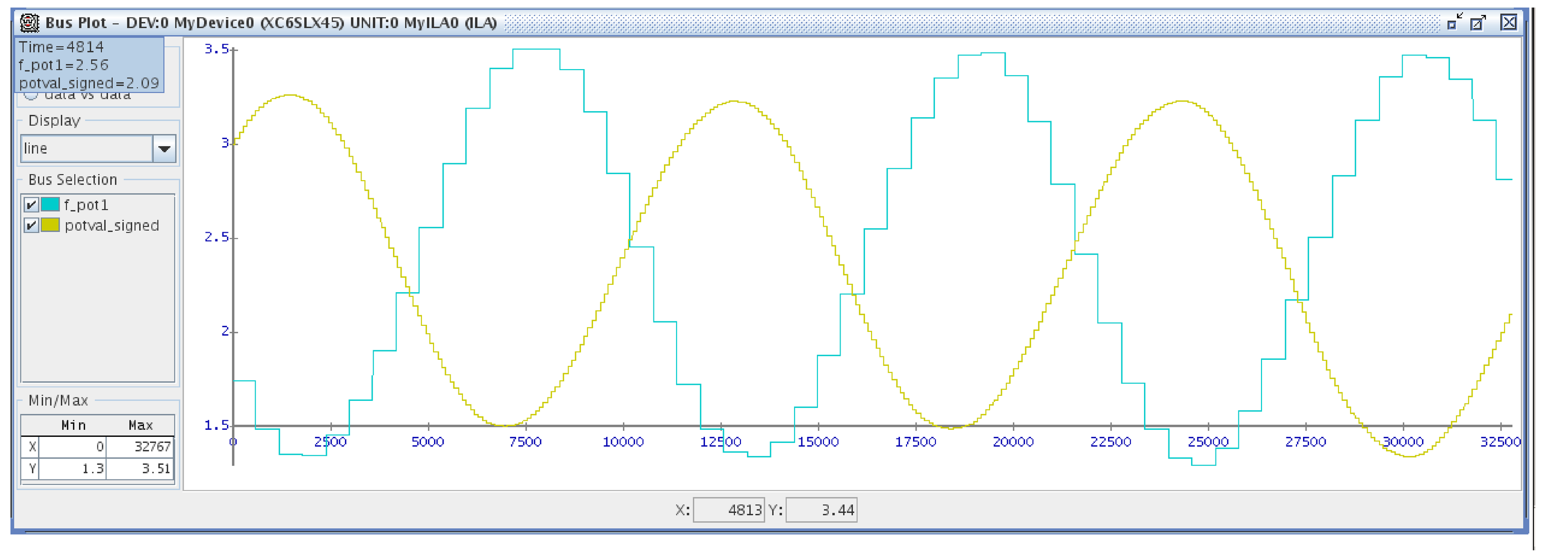

Fig 4. A 5khz lowpass filter passing a 1khz signal

We then applied a 10khz signal through the function generator. As expected the 5khz low pass filter removed this input signal. Just as the previous case the yellow signal is the 10khz input while blue signal is the eliminated signal.

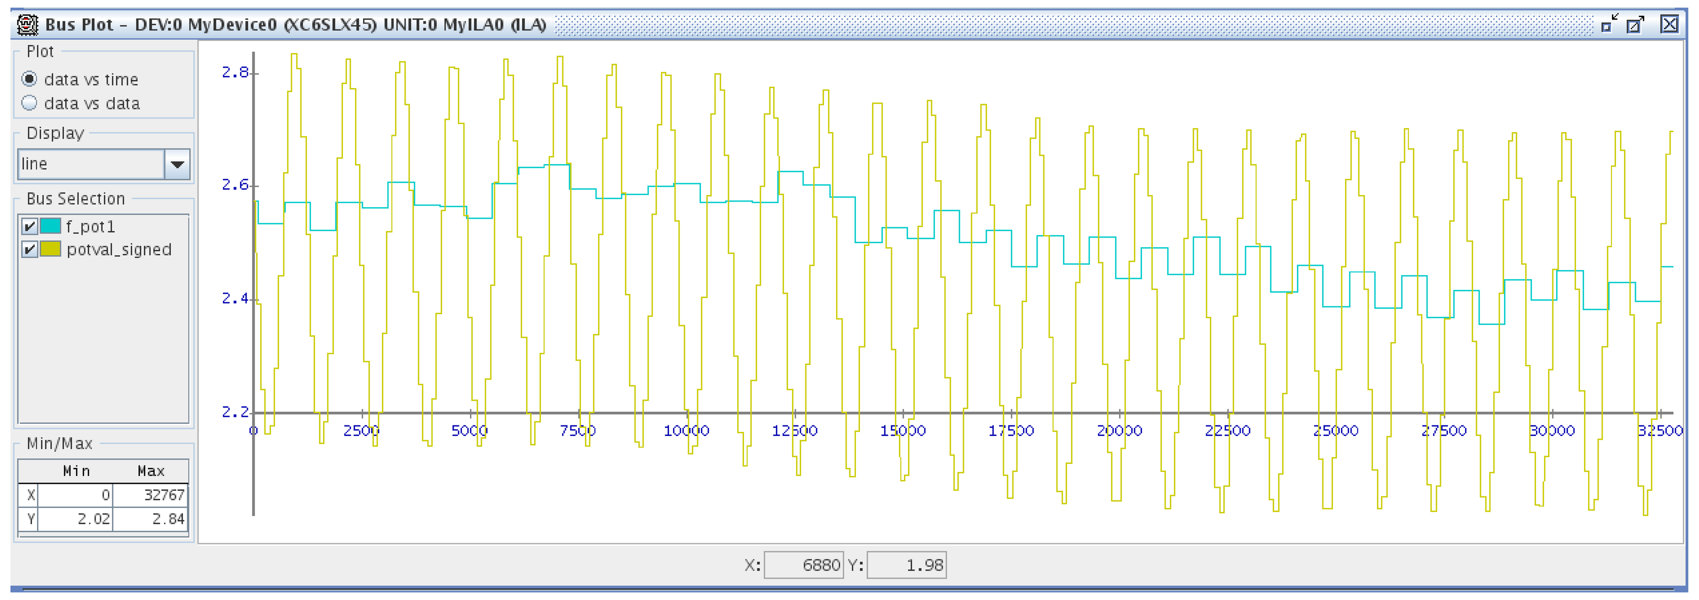

Fig 5. A 5khz signal eliminating a 10khz signal

Conclusion

We have successfully implemented a 5khz low pass filter. Unlike the analog filter this digital filter can be reprogrammed to produce different types of filters. However, more work should be done in eliminating the delay between input and output. Finally this digital filter can used to clean potentiometer and motor current signals.

Acknowledgment

1. Professor Peter Kazanzides
2. Keshuai Xu

Joint Object Detection and 3D Orientation Estimation

Ming Yang Lu, Siddharth Mahendran and René Vidal
Department of Biomedical Engineering, Center for Imaging Science
Johns Hopkins University

Problem Statement

Given a 2D image as input, predict bounding boxes, associated category labels and 3D Orientations of all objects in the image.

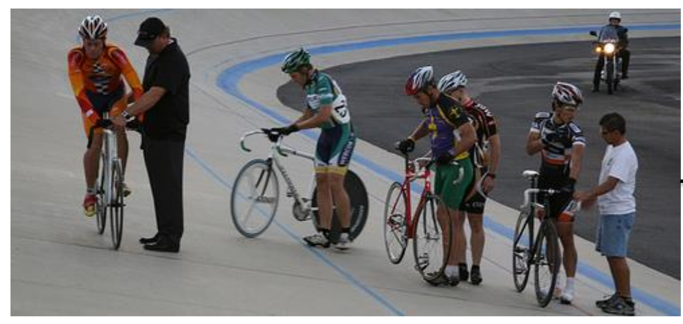

CNN → Bounding Boxes, Class Labels, 3D Orientations

Motivation

Locating objects with bounding boxes are sometimes insufficient.
3D pose information is needed for tasks such as scene reconstruction and autonomous driving.

☑☐ Object location
☑☐ Object Class
☒ Object Spatial Orientation

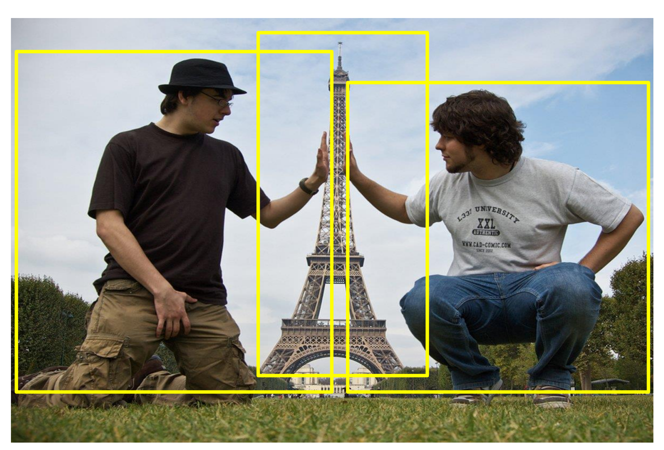

Current algorithms are able to predict 3D pose of an object if given bounding boxes around objects with object class labels:

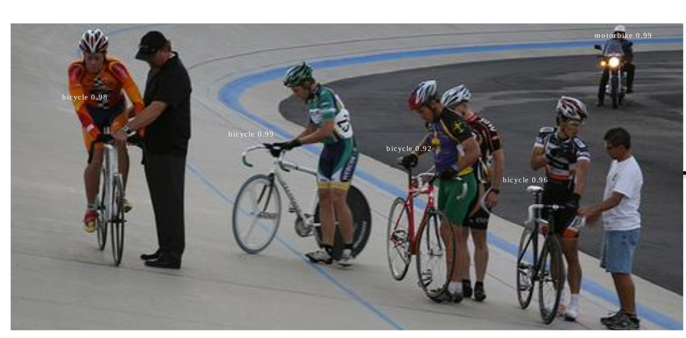

Pose Network → 3D Orientations

Other algorithms are capable of predicting bounding boxes and associated class labels given an image as input:

Detection Network → Bounding boxes, Class Labels

Next Step:
We want to combine the two independent tasks of detection and orientation estimation to perform them jointly with a single integrated pipeline.

References

[1] Ren et al, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks", NIPS 2015.

[2] He et al. "Mask R-CNN", ICCV 2017.

[3] Mahendran et al, "3D Pose Regression using Convolutional Neural Networks", ICCVW 2017.

[4] Y. Xiang, R. Mottaghi, and S. Savarese. Beyond PASCAL: A benchmark for 3d object detection in the wild. WACV, 2014.

Network Architecture

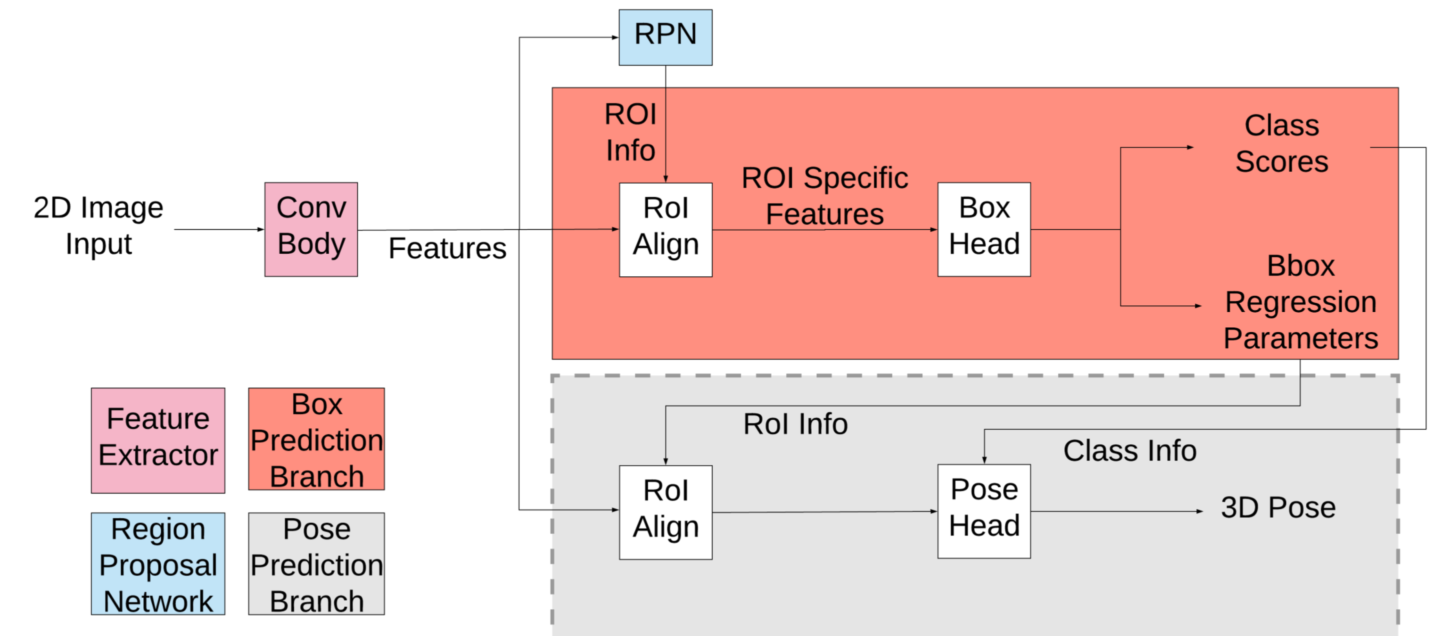

- Modified Mask-RCNN [1] [2] object detection pipeline (shown in Red) for the task of Joint Object Detection and 3D Orientation Estimation. Our contribution shown in gray.
- Detection and pose estimation share feature extractor backbone.
- Detection bounding boxes specify the Region of Interests used for per category pose estimation.
- Detection class predictions are used to select class specific pose predictions.

Loss Functions & Pose Representations

Multi-task Loss

$$L(y,y^*,c,c^*,b,b^*) = L_P(y,y^*) + L_C(c,c^*) + L_B(b,b^*)$$

y, y^*: Orientation c, c^*: Class b, b^*: Bounding Box targets

*Symbols w/ * denote ground truth, w/o * denote network predictions.*

L: Total loss
L_P: Pose specific loss
L_C: Classification softmax loss
L_B: Bounding box smooth L1 regression loss

Pose Representation and Pose Specific Loss [3]

3-dof pose: Azimuth *az*, Elevation *el*, Camera-tilt *ct*

Rotation Matrix: $R = R_Z(ct)R_X(el)R_Z(az)$

Axis Angle Representation: $y = logR$

Axis Angle Representation better captures the geometry of the underlying orientation space!

Euclidean Loss:

$$L_P(y,y^*) = \|y - y^*\|^2$$

Geodesic Loss:

$$L_P(y,y^*) = \cos^{-1}\left[\frac{trace(R^T R^*)-1}{2}\right]$$

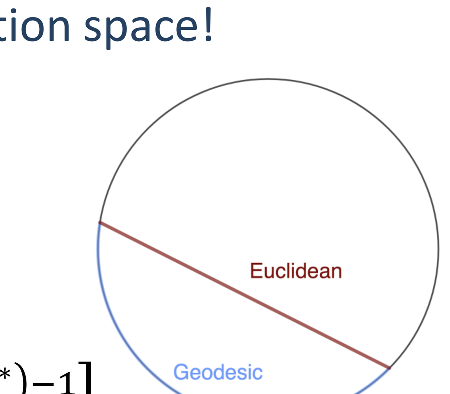

Results

Dataset
- Pascal3D+ [4] consists of ImageNet and Pascal VOC2012 images with 3D pose annotations.
- ImageNet trainval and VOC2012-train images are used for training and VOC2012-val images are used for testing.

Metrics
- mAP: measures detection performance. Higher is better.
- ARP: an extension of mAP for joint object detection and orientation estimation. Higher is better.

Quantitative Results

Step 1: Finetune for Detection Only

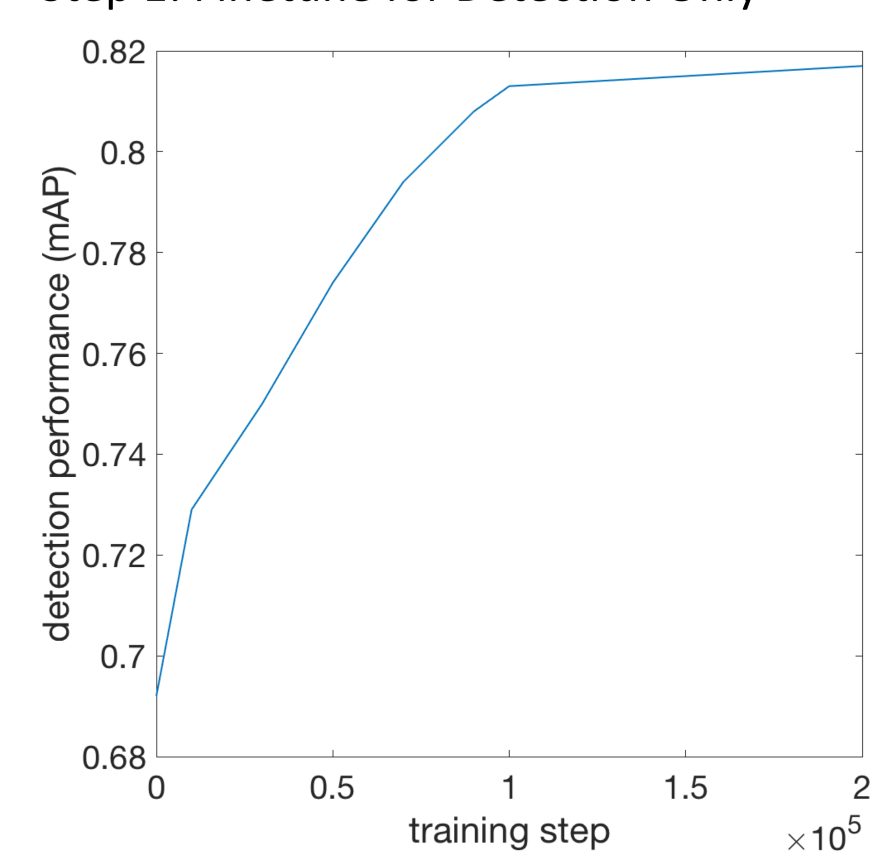

Step 2: Training Pose Network Only using Euclidean Loss

Step	mAP	ARP
0	0.817	0.262
200000	0.817	0.377

Step 3: Training Pose Network Only using Geodesic Loss

Step	mAP	ARP
0	0.817	0.377
200000	0.817	0.422

Step 4: Finetuning End to End using Geodesic Loss

Step	mAP	ARP
0	0.817	0.422
110000	0.818	0.457

Qualitative Results

Inference demonstration of joint detection and pose estimation using the Pascal3d+ dataset

orientation prediction

class and confidence score

bounding box

Acknowledgements

This research was supported by NSF grants 1527340 and 1834427.

Spray-Coating Flexible Optoelectronics for Solar Energy Harvesting and Sensing on Any Surface

Laura Shimabukuro[1], Botong Qiu[2], Lulin Li[2], Yida Lin[2], Susanna M. Thon[2]

[1]Department of Electrical Engineering, Sacramento City College [2]Department of Electrical and Computer Engineering, Johns Hopkins University

Introduction

- [1]Solution processed semiconductors offer the promise of low cost, production-scale optoelectronic devices, such as flexible solar cells and photodiode infrared sensors.
- [1]However, their current uses are limited to lab scale production methods, such as spin-coating and dip-coating.
- Spray-coating has emerged as a more efficient method to produce large-scale flexible devices that can be used on any type of surface.
- **Goal:** In this project we aim to build the first colloidal quantum dot solar cell in which *all cell layers, including electrodes*, are deposited via spray-coating as a proof-of-principle demonstration of a manufacturing method that can be applied to any surface.

Figure:
(a) Schematic of CQD solar cell architecture
(b) Custom-Build Spray-Coating System

(a)

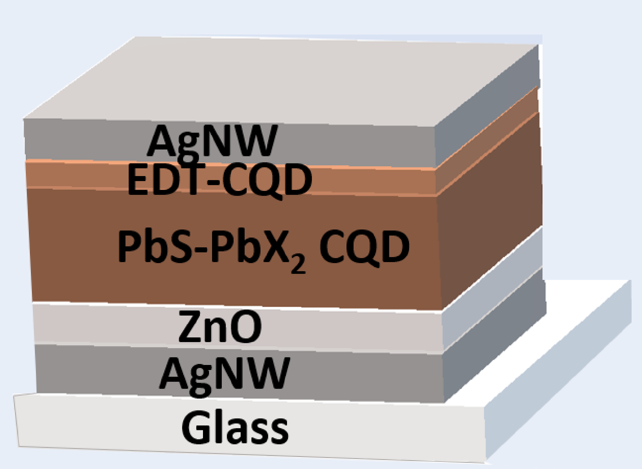

(b)

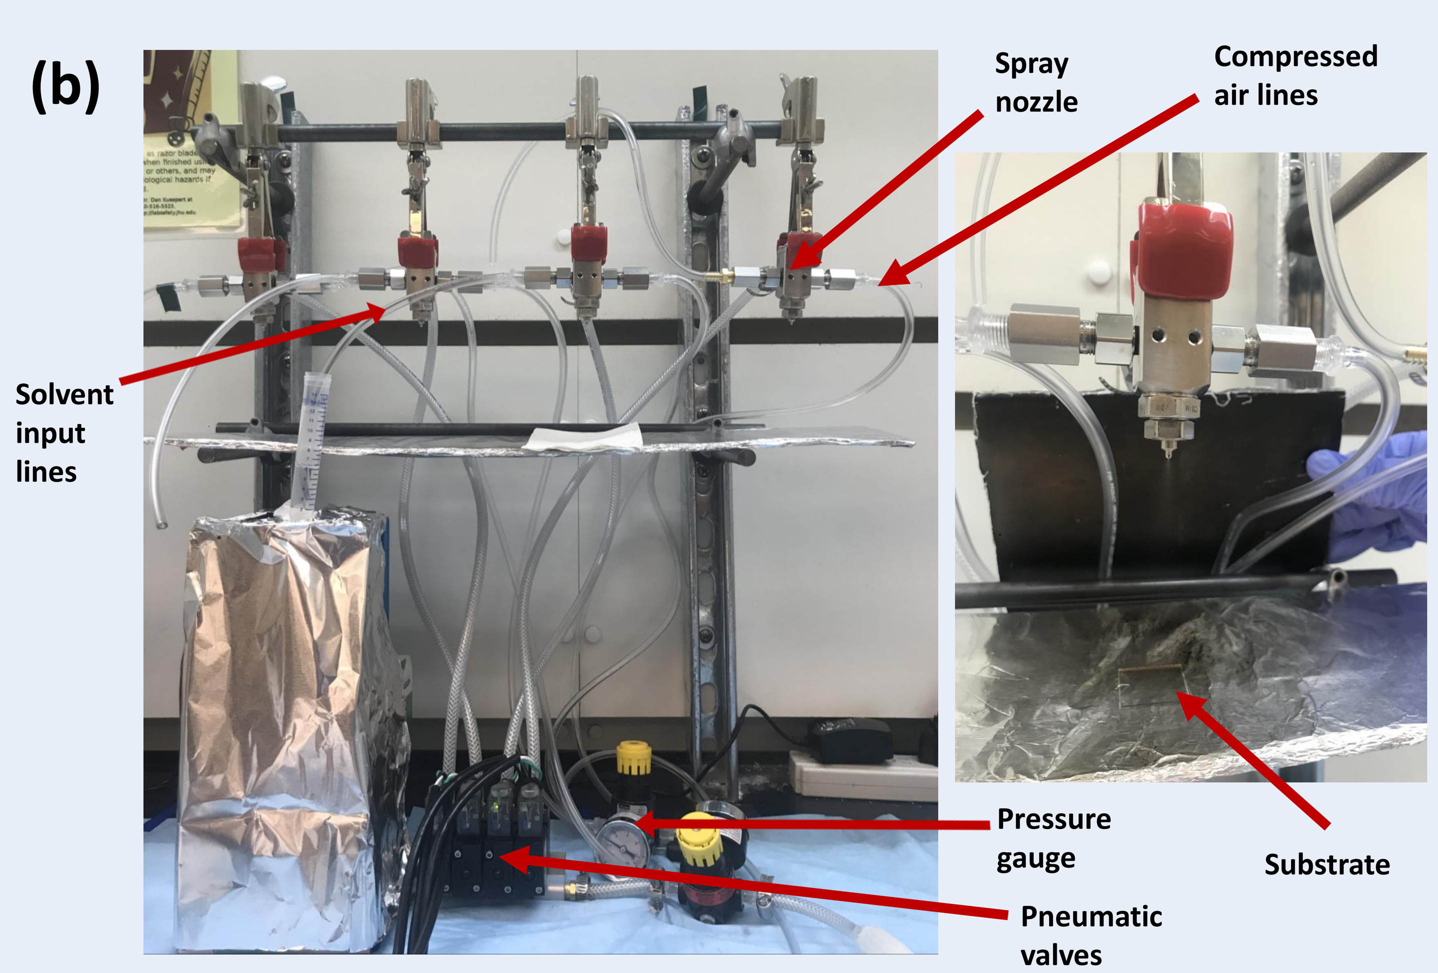

Methods

Electrode Spray Optimization

Silver Nanowires

The three main parameters we tested to produce uniform and optimal density layers were:

- Distance between the nozzle and substrate
- Spray pattern
- Solution concentration

We optimized each parameter to achieve the best tradeoff between transparency and conductivity.

AgNW Spray Parameters and Optimization Results

Trial	Distance (cm)	Concentration (mg/mL)	Spray Time (min)	Dry Time (min)	# Spray Cycles	Transparency (avg %)	Sheet Resistance (Ω/sq)
1	5	0.5	5	15	1	50 %	18.2
2	8	0.5	6.5	30	13	75 %	21.9-36.5
3	8	0.5	4	15	1	75 %	474.3
4	6	0.5	3	15	1	80 %	693.2
5	8	0.125	10	40	20	95 %	To high to measure

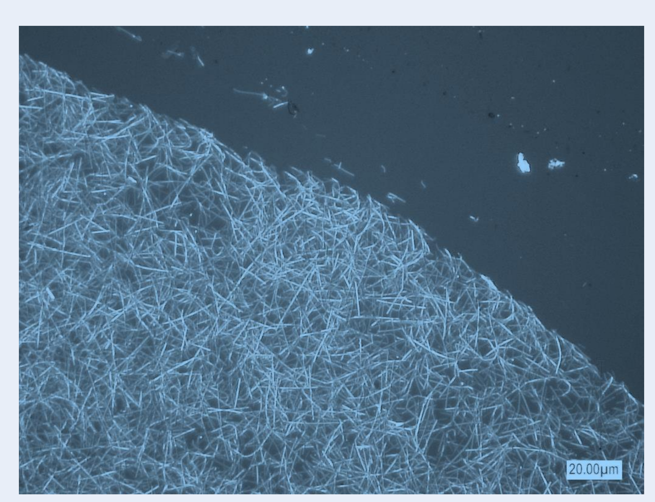

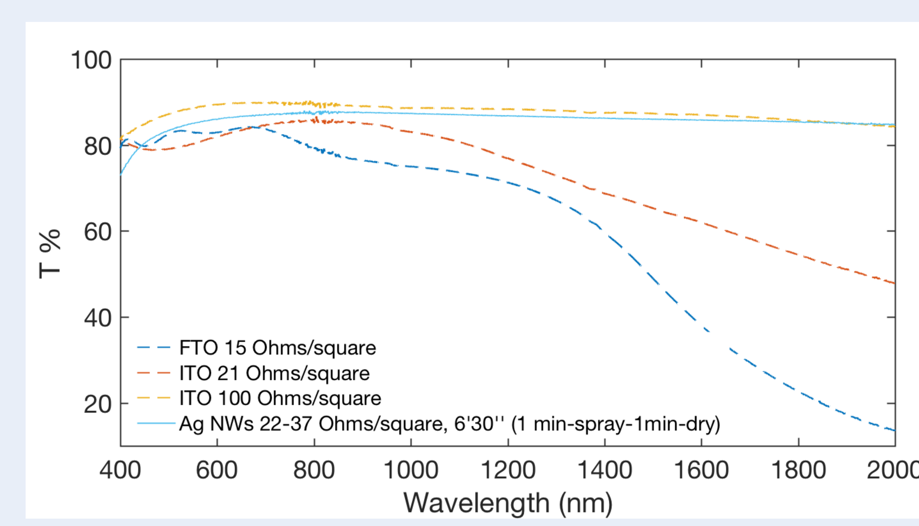

8 cm, 6’30’’ (1min-spray-1min-dry)
0.5 mg/mL AgNW

Zinc Oxide

We tested each parameter to achieve a continuous spray distribution and uniform thickness, similar to the results from spin coating.

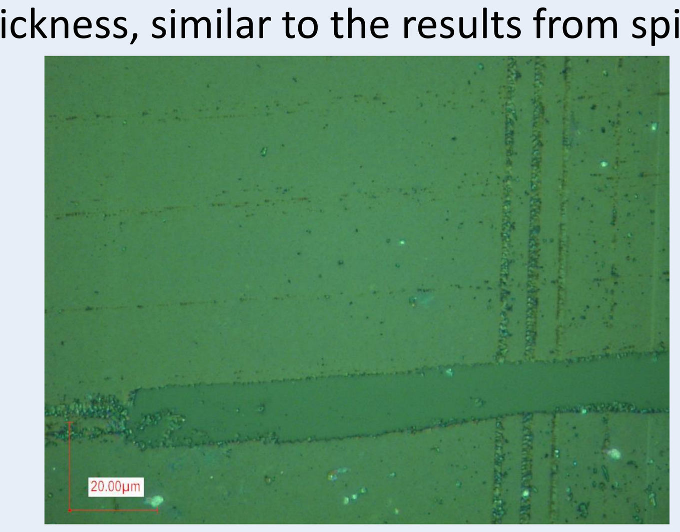

Spin-coated ZnO

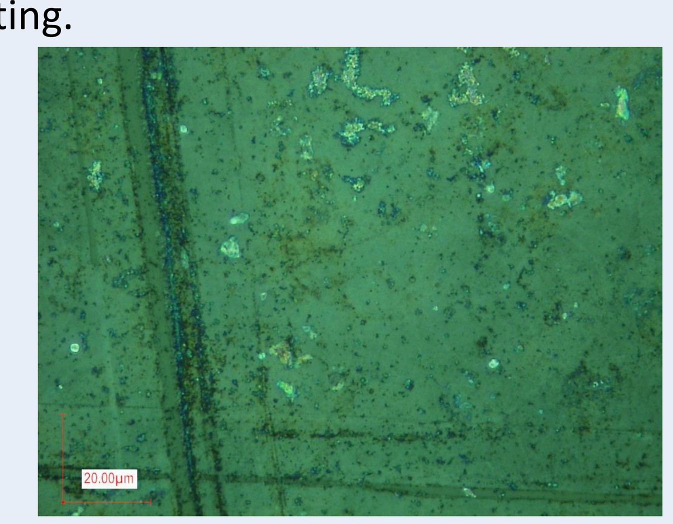

Spray-coated ZnO

Absorber Layer Spray Optimization

Quantum Dots

We began optimizing the spray coating of an untreated quantum dot layer. The tested parameters include:

- Gas Pressure
- Spray pattern
- Solution concentration

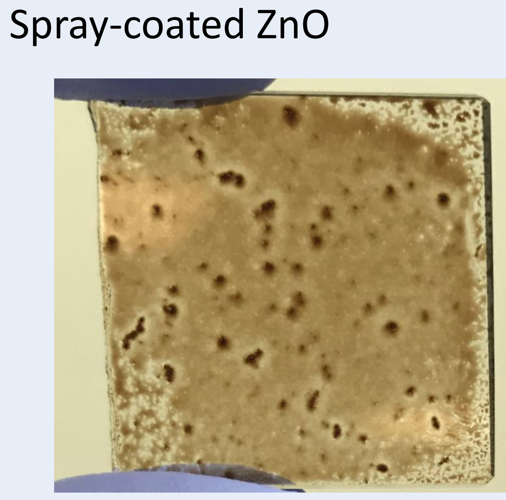

Results

	AgNW	ZnO	Quantum Dots
Best Samples		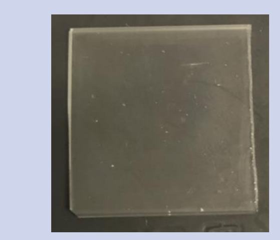	
Measurement Results	• 21.9-36.5 Ω/sq • > 75% avg. transmission	• Thickness: ~200 nm	• Thickness comparable to spin coated films achieved

Conclusion

- Our optimized AgNW sample (1’ spray, 1’ dry) had an average transparency of 75% across the visible region and an optimal sheet resistance of 21.9-36.5 Ω/sq, comparable to the best conventionally-deposited solar cell transparent electrodes.
- We achieved a spray-cast ZnO layer of comparable thickness and optical density to conventionally-deposited materials for spraying conditions of 10’ of continuous spray from a 2mg/mL ZnO solution.
- We achieved our goal of the first demonstration of a completely spray-deposited CQD solar cell. This proof-of-principle demonstration should enable flexible deployment of solar energy harvesting and sensing technologies on a variety of urban infrastructures and surfaces in remote locations.
- Future work includes optimizing the spray parameters for the remaining layers of the solar cell and achieving comparable performance in all-spray and conventional devices.

References

1. I. J. Kramer, J. C. Minor, G. Moreno-Bautista, L. Rollny, P. Kanjanaboos, D. Kopilovic, S. M. Thon, G. H. Carey, K. W. Chou, D. Zhitomirsky, A. Amassian, E. H. Sargent. Efficient Spray Coated Colloidal Quantum Dot Solar Cells. *Advanced Materials* 10.1002/adma.201403281 (2015).
2. K. Zilberberg, F. Gasse, R. Pagui, A. Polywka, A. Behrendt, S. Trost, R. Heiderhoff, P. Görrn, T. Riedl. Highly Robust Indium-Free Transparent Conductive Electrodes Based on Composites of Silver Nanowires and Conductive Metal Oxide. *Advanced Functional Materials* 10.1002/adfm.201303108, 1671-1678 (2014)
3. D. Choi, H. Kang, H. Sung and S. Kim. Annealing-Free, Flexible Silver Nanowire–Polymer Composite Electrodes via a Continuous Two-Step Spray-Coating Method. *Nanoscale* 10.1039/c2nr32221h (2012)
4. M. Yuan, M. Liu, E. H. Sargent. Colloidal Quantum Dot Solids for Solution-Processed Solar Cells. *Nature Energy*, Vol. 1, Article 16016, 10.1038/nenergy.2016.16 (2016)

Optimization Landscape of a Neural Network Model for Activity Recognition

Ksenia Lekomtceva[1,2], Connor Lane[2], Effrosyni Mavroudi[2], René Vidal[2]
[1]Montgomery College, [2]Center for Imaging Science, Johns Hopkins University

Motivation and Problem Statement

- The introduction of deep learning has led to remarkable improvements in performance across many tasks, including activity recognition.

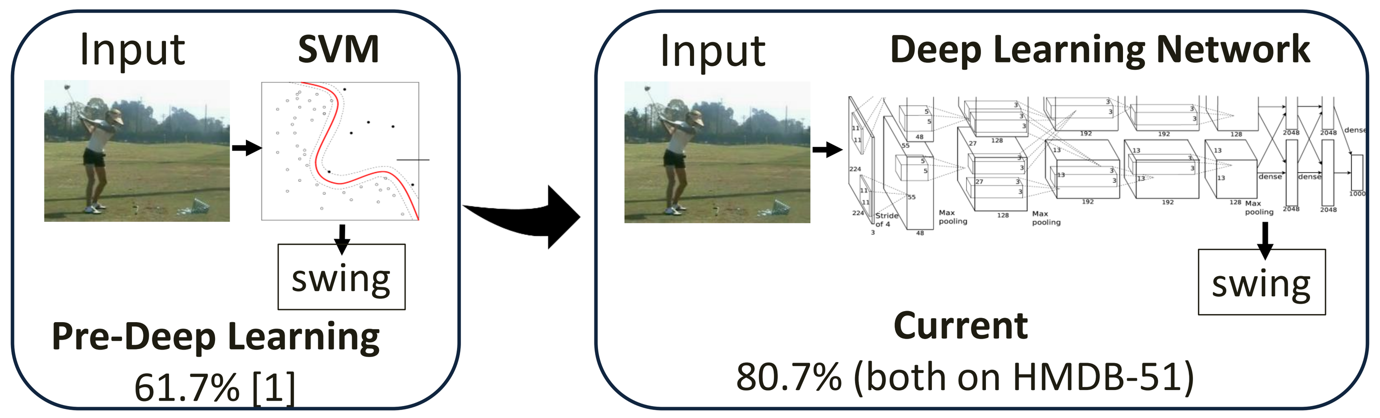

- However, a solid theoretical understanding of deep learning models' worst-case performance is needed to avoid failures in sensitive applications.

Problem Statement: Recent theoretical work provides performance guarantees for some neural networks. Do they perform better?

Model and Optimization Problem

A single hidden layer neural network model with multiple tunable parameters W:

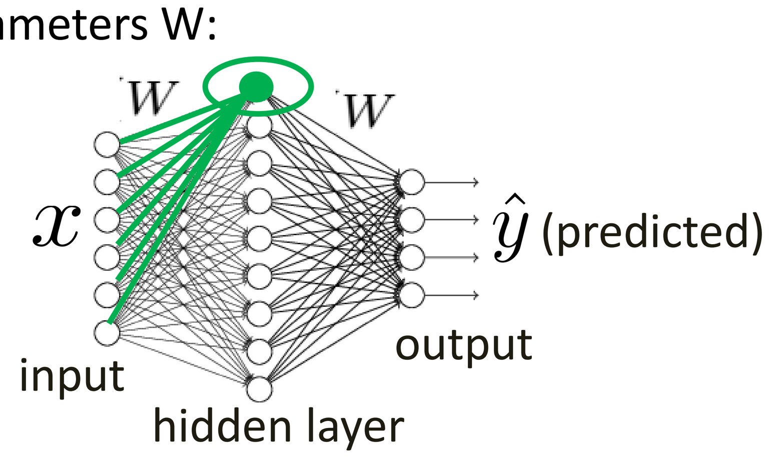

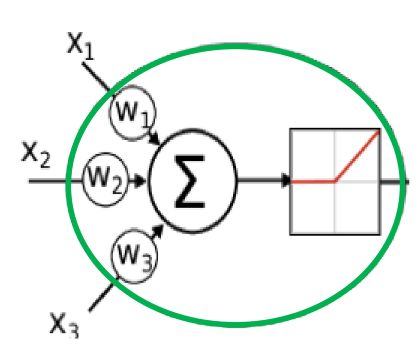

Neural network training solves the optimization problem, given training data:

$$\min_W \sum_{i=1}^N l(\underbrace{y_i}_{\text{(target)}},\ \overbrace{\Phi(W,\ x_i)}^{\hat{y}_i}) \quad \text{(loss)}$$

Optimized by making small repeated updates to the parameters W:

$$W_{k+1} = W_k - \eta \partial W_k$$

This procedure can potentially get stuck in bad local minima due to the non-convex character of the optimization problem (right):

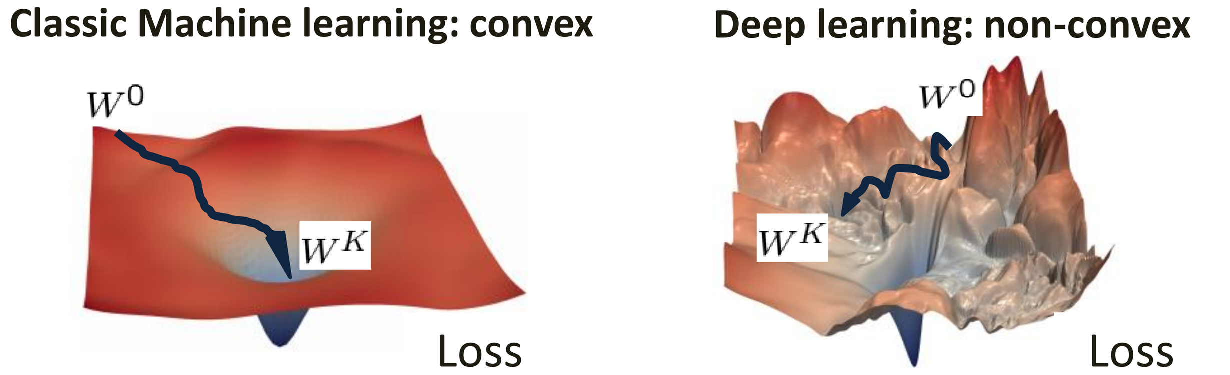

References:
[1] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in CVPR 2017, pp. 4724–4733.
[2] B. Haeffele and R. Vidal, "Global optimality in neural network training," in CVPR 2017, pp. 7331–7339.
[3] E. Mavroudi, L. Tao, and R. Vidal, "Deep moving poselets for video based action recognition," in WACV 2017, pp. 111–120.

Optimization Theory

According to recent theoretical work [2], we can guarantee that bad local minima do not exist if the following three properties are satisfied:

1. Over-parameterization *(many hidden units):*

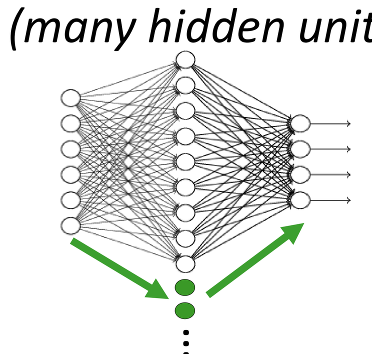

➕ 2. Regularization *(control weight size):*

$$\min_W \sum_{i=1}^N l(y_i,\ \Phi(W,\ x_i)) + \boxed{\lambda\Theta(W)}$$

➕ 3. Positive Homogeneity *(predictable effect of scaling):*

$$\Phi(\alpha W,\ x_i) = \alpha^p \Phi(W,\ x_i)$$

➡ Bad local minima (★) do not exist:

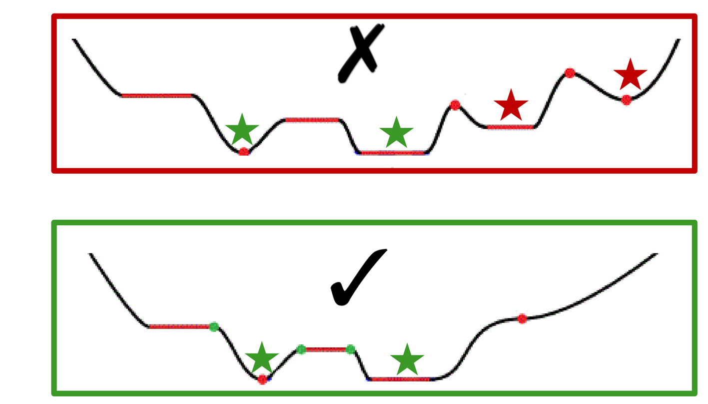

Project Hypotheses

Does satisfying the theoretical properties improve performance in practice?

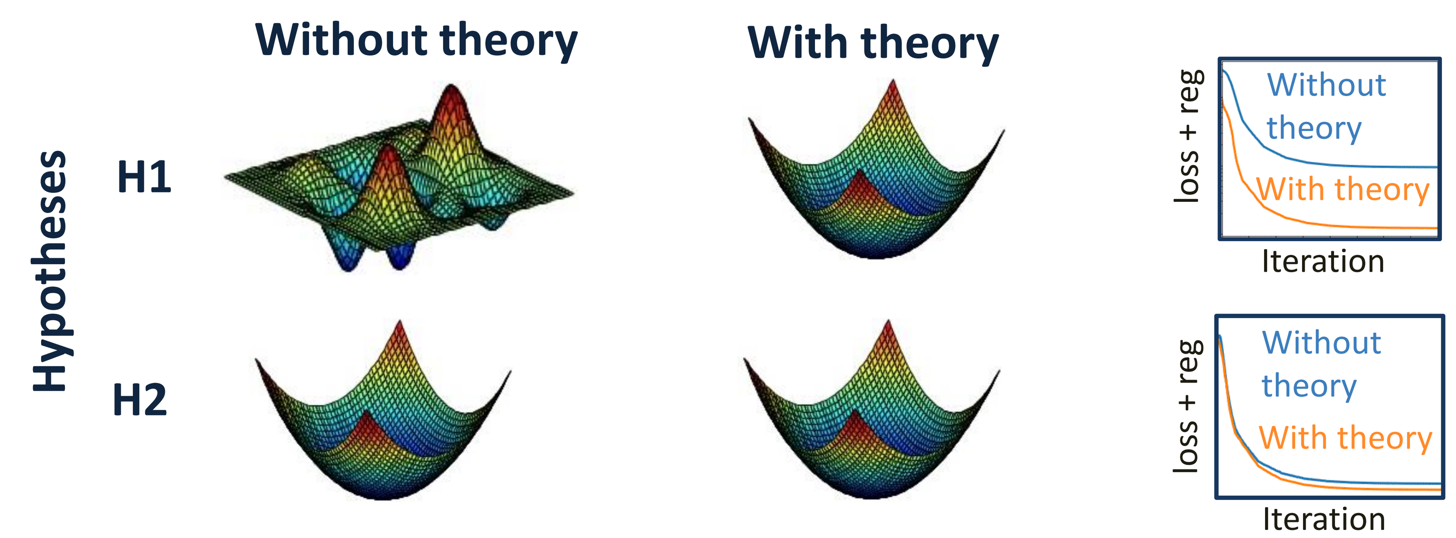

Synthetic data Experiment and Results

First, the hypotheses were tested in a well-controlled "ideal" setting:

Task: mirror behavior of "black-box" network

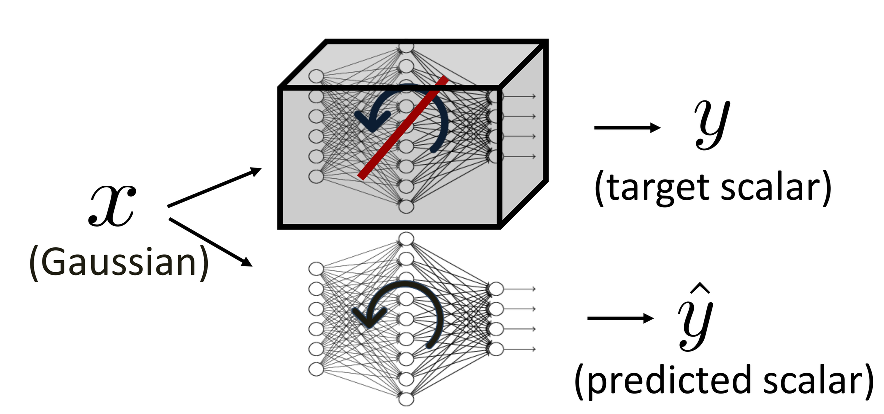

- Unlimited randomly generated data
- Access to "ground truth" parameters W that generated the data

The plots show the (regularized) loss of the model with 60 (ground truth size), 120, 240, and 480 hidden units:

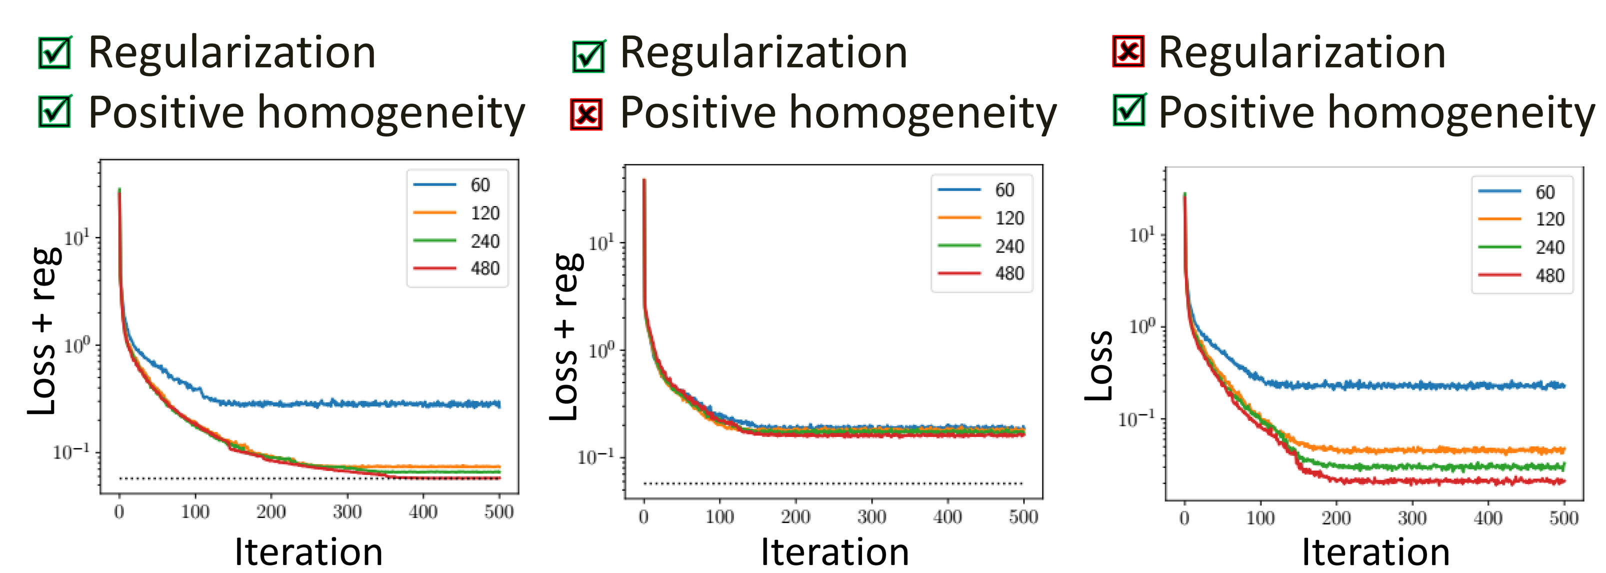

Conclusion: H1 supported

- Over-parameterization with positive homogeneity improves optimization.
- Regularization not necessary → is it possible the theory can be generalized?

Activity Recognition Experiment and Results

Next, the hypotheses were tested in a realistic activity recognition setting using the Deep Moving Poselets [3], where the hidden layer captures discriminative movements of body parts:

Task: classify actions from short video clips

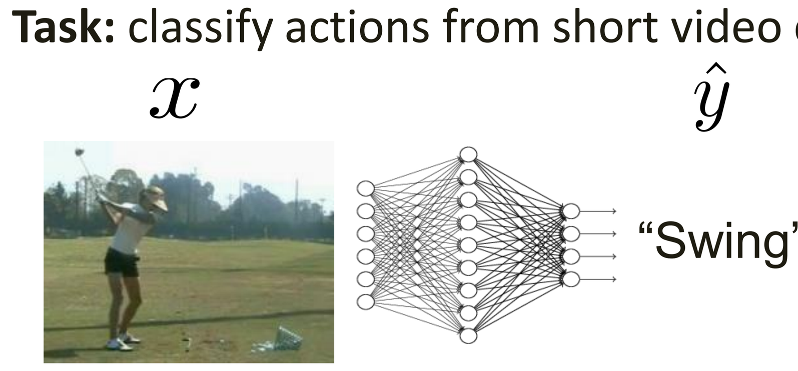

- Only 227 training examples
- No "Ground truth"

The plots show the (regularized) loss for the model with 200, 400, 800, and 1600 hidden units:

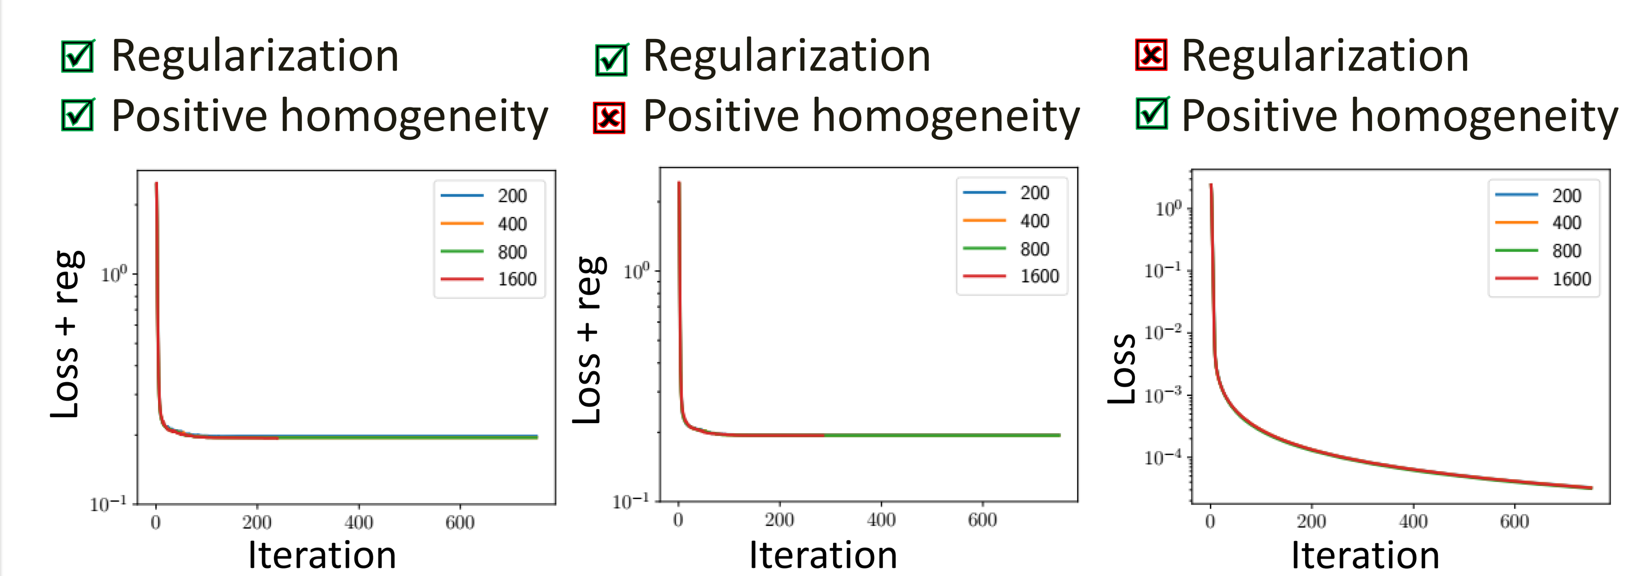

Conclusion: H2 supported

- Over-parameterization, positive homogeneity, regularization do not affect optimization.
- Possible that optimization landscape is too easy to begin with.

Conclusions and Future Research

- For unlimited synthetic data, we find evidence that satisfying the theoretical conditions helps training avoid bad local minima.
- However, for activity recognition task, satisfying the theoretical conditions has no effect on training performance.
- In future work, we will explore possible explanations for this result, including dataset size, and data "realizability".

Acknowledgements:
This research was supported by the NSF under Grant No. 1004782

In Vivo Photoacoustic Image Guidance Of Abdominal Surgery

PULSE
PHOTOACOUSTIC & ULTRASONIC SYSTEMS ENGINEERING LAB

JOHNS HOPKINS
WHITING SCHOOL of ENGINEERING

Kelley M. Kempski[1], Alycen Wiacek[2], Jasmin E. Palmer[3], Michelle Graham[2], Eduardo González[4], Bria Goodson[5], Derek Allman[2], Huayu Hou[2], Jin He[6,7], Muyinatu A. Lediju Bell[2,4,8]

[1]Biomedical Engineering, University of Delaware, [2]Electrical and Computer Engineering, Johns Hopkins University, [3]Mechanical Engineering, Massachusetts Institute of Technology, [4]Biomedical Engineering, Johns Hopkins University, [5]Biology, Delta State University, [6]Surgery, Johns Hopkins Medicine, [7]Oncology, Johns Hopkins Medicine, [8]Computer Science, Johns Hopkins University

INTRODUCTION

Intraabdominal hemorrhage results in morbidity in 23-46% of liver resections and mortality in 30-58% of pancreaticoduodenectomies [1,2].

Photoacoustic (PA) imaging (pulsed laser light coupled with ultrasound [US]) is one solution for visualizing blood vessels to reduce risk of accidental injury to an artery or vein intraoperatively [3,4].

The objective of this study was to determine the feasibility of visualizing various abdominal blood vessels *in vivo,* including the right gastroepiploic vein (RGEV) and artery (RGEA), superior mesenteric vein (SMV), and hepatic veins (HV), using PA imaging.

METHODS

PA imaging and Doppler were used to identify various blood vessels in the pancreas (Fig. 2A) and liver (Fig. 2B) of two pigs during open surgery.

All PA data was beamformed and contrast was calculated by selecting the region-of-interest (ROI) inside the vessel and a corresponding ROI in the background at the same depth.

Short-lag-spatial-coherence (SLSC) beamforming was applied to determine if signal inside the vessel was due to the blood or artifact from the fiber bundle [5].

An *ex vivo* experiment was performed to help explain *in vivo* pancreas observations. The *ex vivo* phantom consisted of bovine liver tissue that surrounded tubes filled with human blood (Fig. 1B).

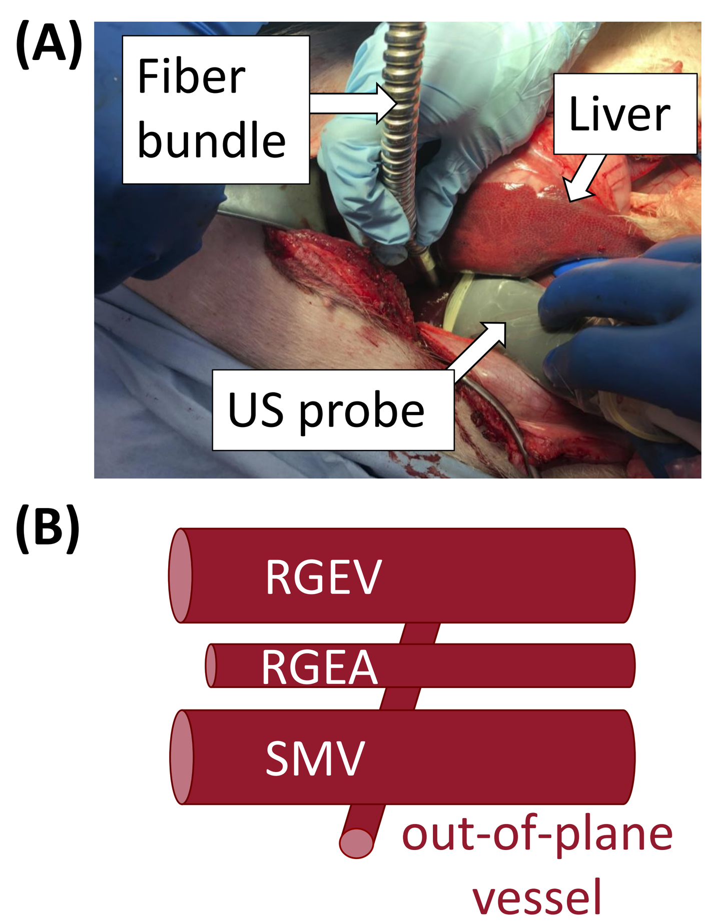

Fig. 1. Set-up for (A) animal and (B) ex-vivo experiments. The fiber bundle delivered the laser light.

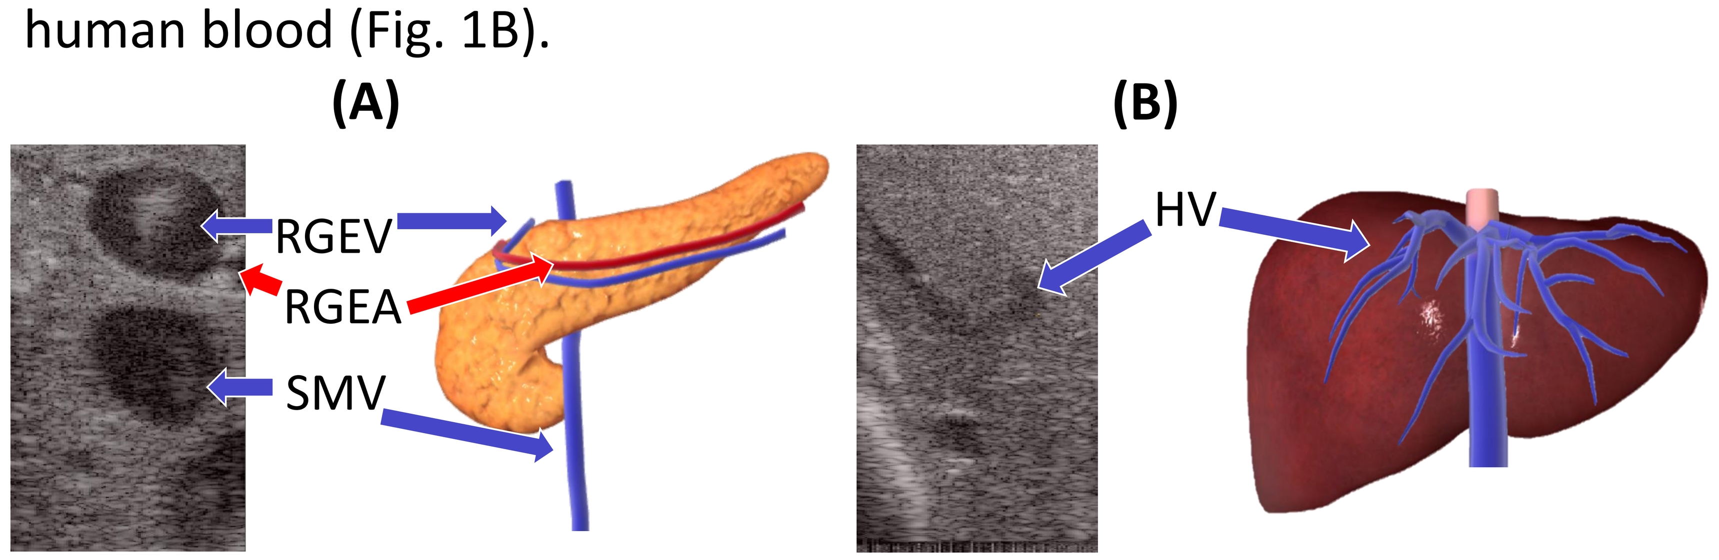

Fig. 2: Identification of blood vessels in (A) pancreas and (B) liver.

IN VIVO PANCREAS RESULTS

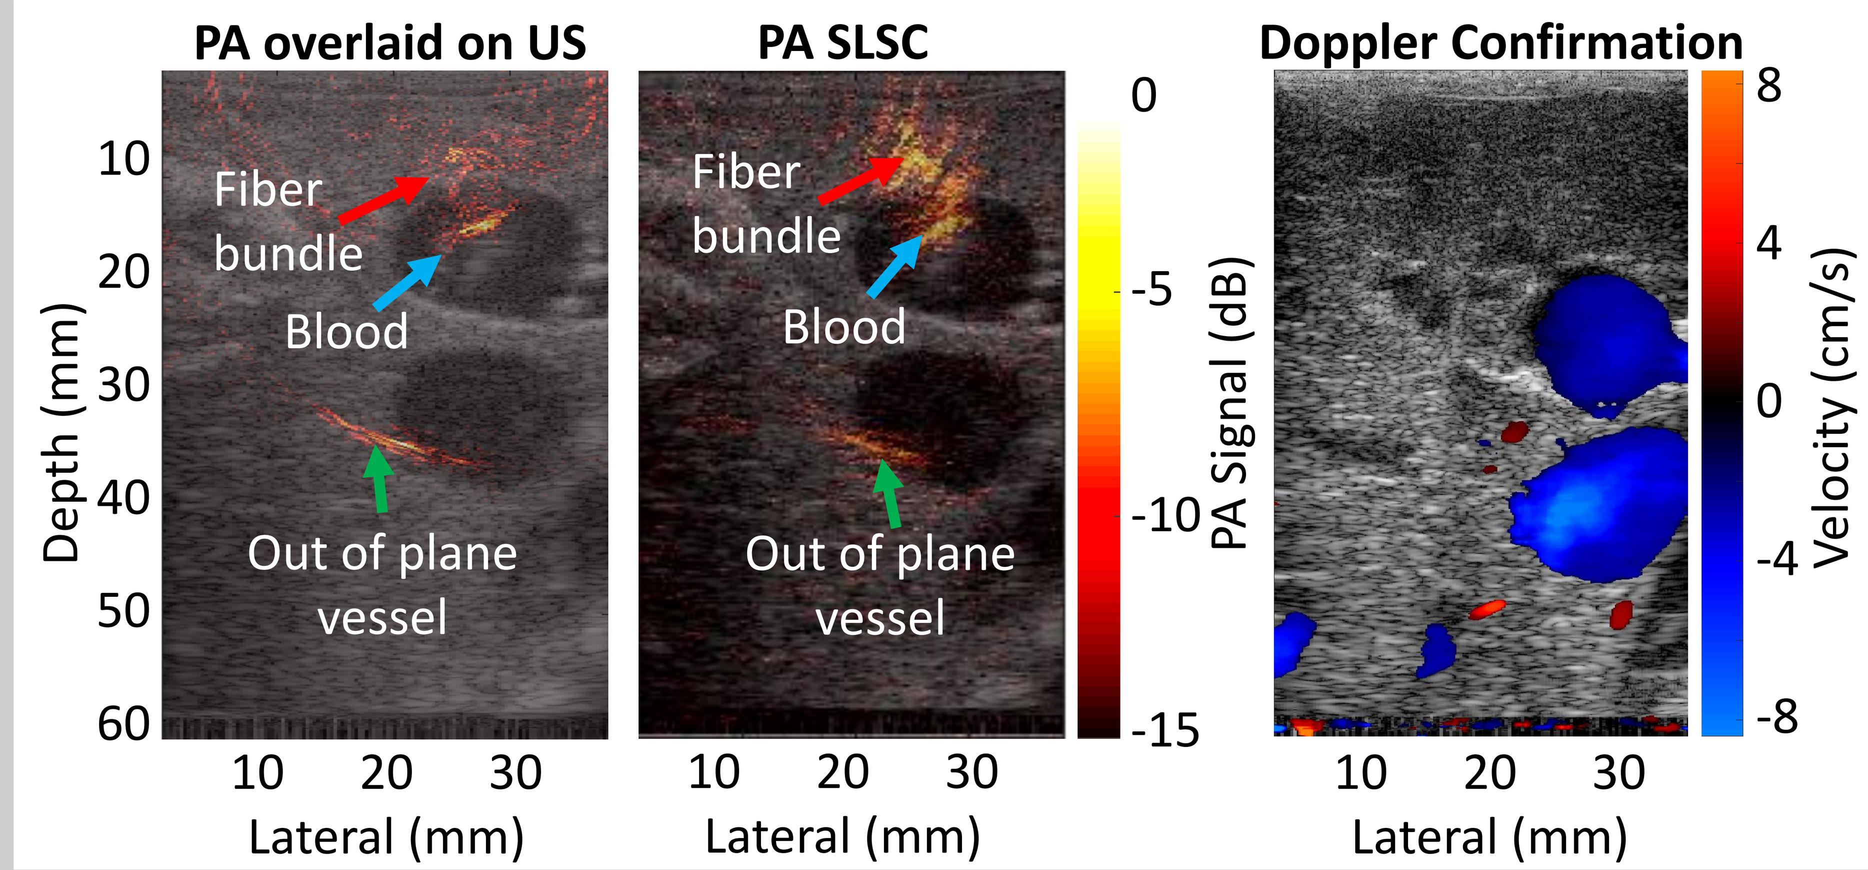

IN VIVO LIVER RESULTS

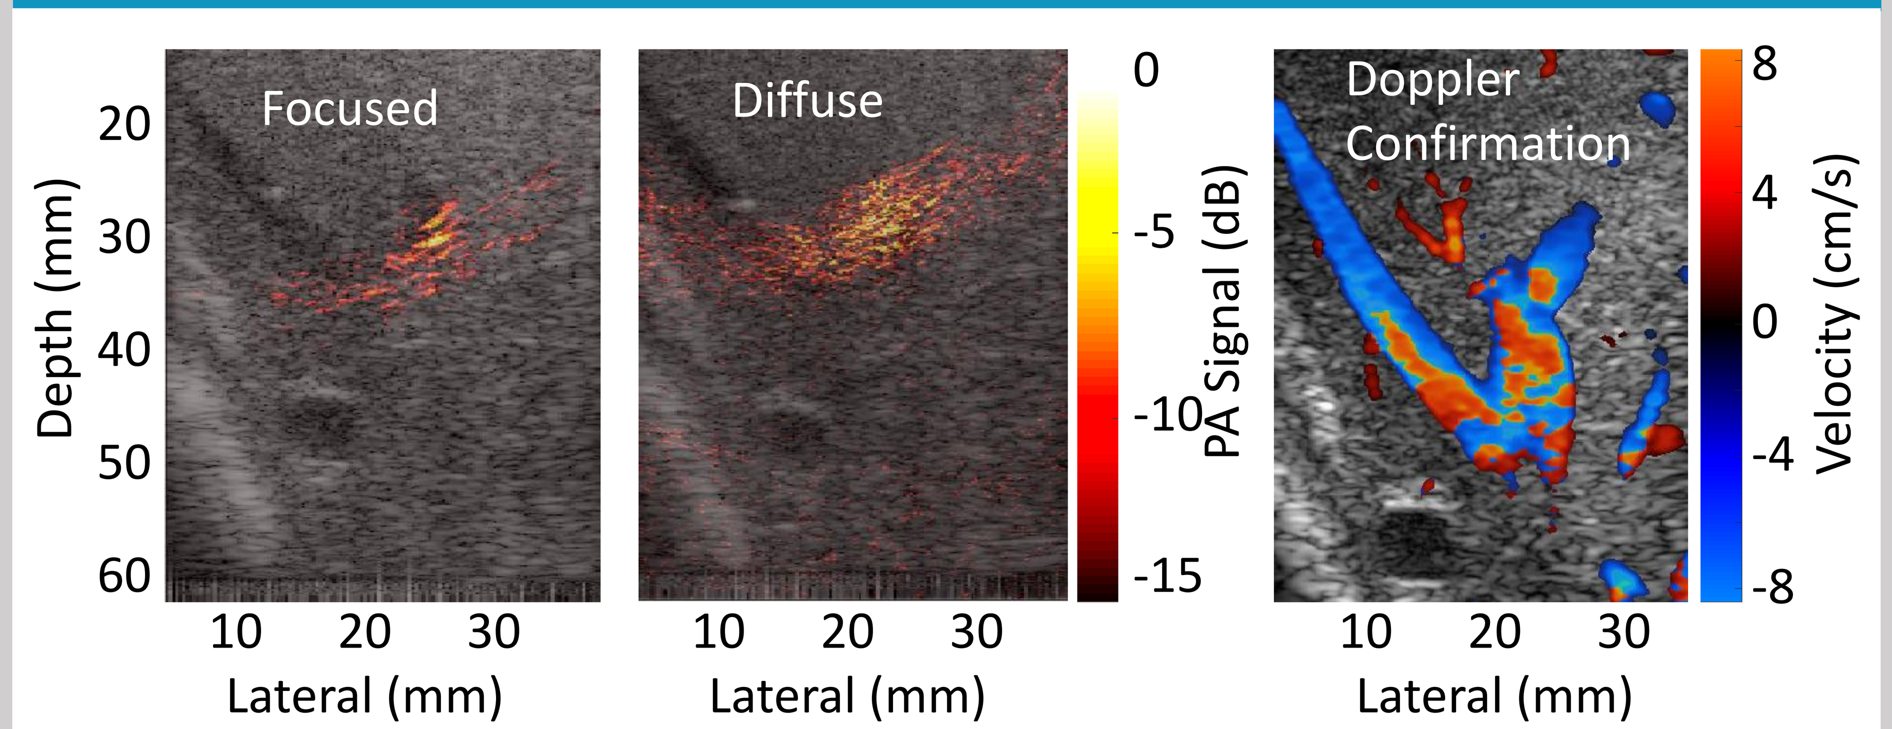

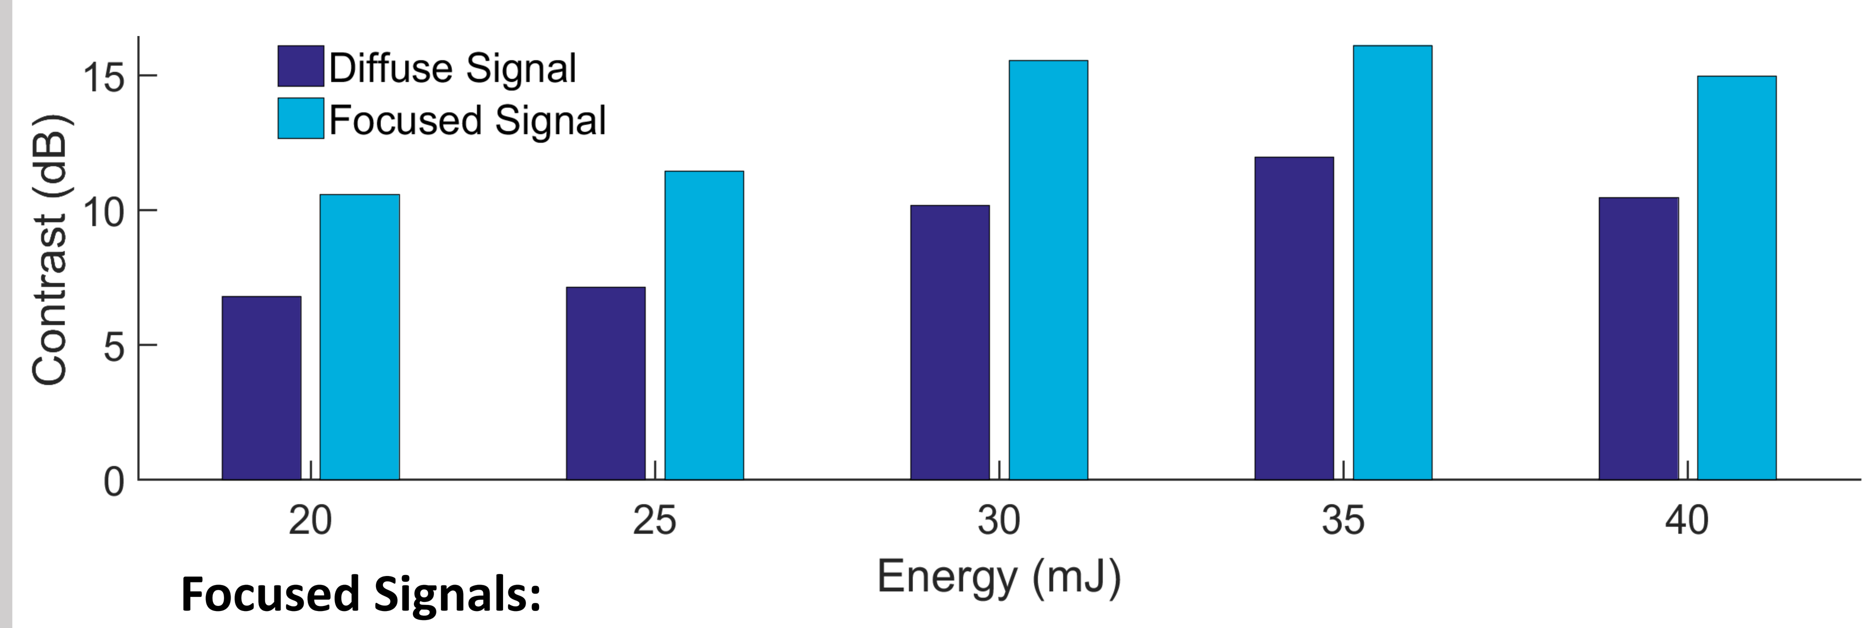

Focused Signals:

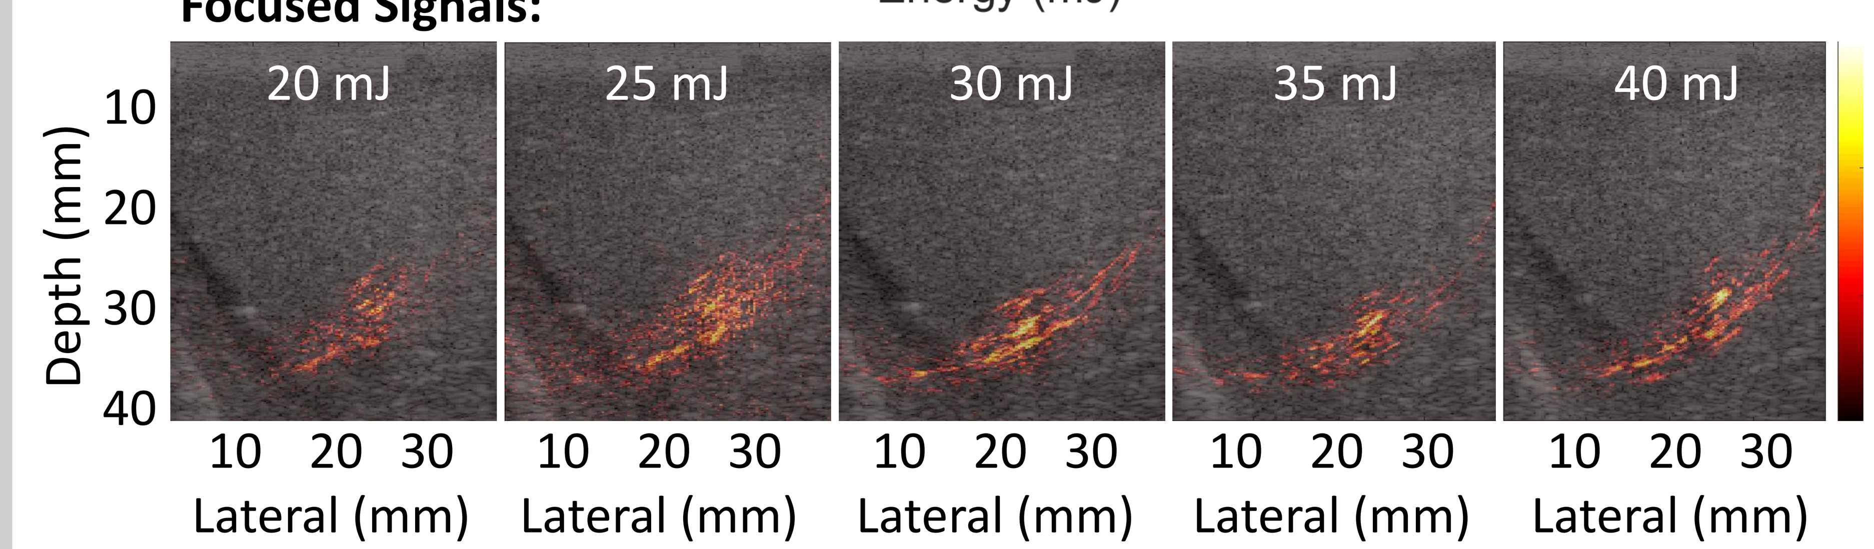

OUT OF PLANE VESSEL CONFIRMATION

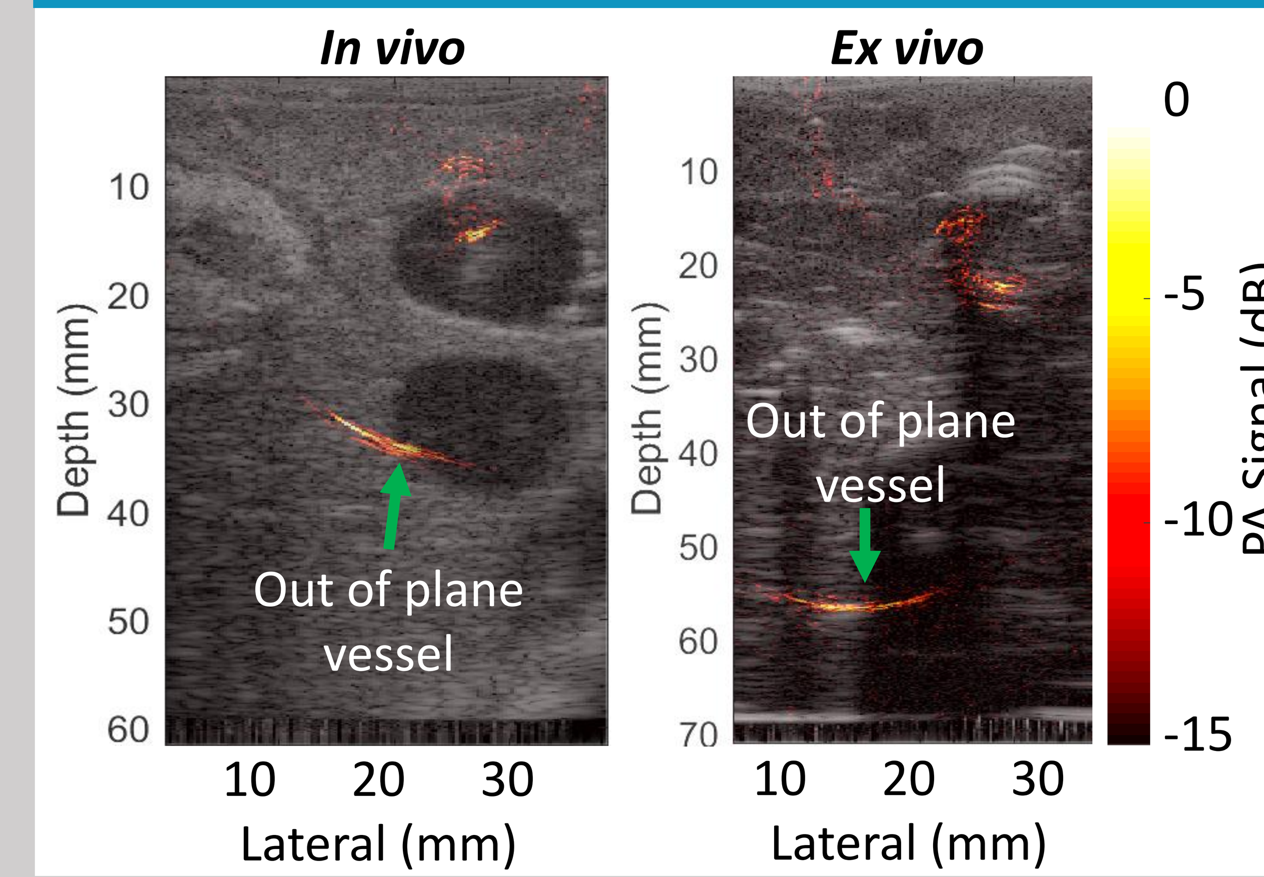

CONCLUSIONS

This study is the first to demonstrate the feasibility of *in vivo* vessel visualization to enable photoacoustic-guided pancreatic and liver surgery.

An energy of 30 mJ was sufficient to identify major blood vessels with 15.5 dB contrast in the liver.

We differentiated blood signals from tissue-related signals based on the diffuse or focused appearance of the photoacoustic signals.

Future directions:

- Use larger light delivery system
- Determine effect of PA imaging on blood loss
- Explore laparoscopic approach

REFERENCES

[1] Romano et al. *HPB Surg.* 2012.
[2] Rumstadt et al. *Annals of surgery.* 1998.
[3] Beard et al. *The Global Organ Shortage.* 2013.
[4] Bouchard et al. *IEEE Trans Ultrason Ferroelectr Freq Control.* 2014.
[5] Bell *Ultrasound Med Biol.* 2014.

ACKNOWLEDGEMENTS

This study was partially funded by the Computational Sensing & Medical Robotics REU program at Johns Hopkins University and NSF CAREER Award #1751522. The authors thank Nicholas Louloudis, Sue Eller, Ivan George, Dr. Lingdi Yin, and Dr. Liang Wang for animal care and surgery support.