Robust Grasping by Integrating Machine Learning with Physical Models

Zixi Liu[1], Alexandre Bayle[1], Robert D. Howe[1,2], Lucas Janson[1]
[1] Harvard University, [2] RightHand Robotics, Inc.
http://biorobotics.harvard.edu/research.html

Abstract

Contact sensing is essential for reliable robotic grasping in unstructured environments, but existing methods have not been effective, and requirements for effective sensors are unknown. This project aims to establish the foundation for effective grasp stability prediction and control by developing new ways to integrate machine learning with physical sensor models. **Physical sensor models** will be characterized in grasping experiments and validated against **independent lab gold standard** measurements. Physical models based on mechanical principles (grasp analysis) will be augmented using **parametric and nonparametric machine learning methods**, allowing interpretability and generalizability. Analysis of these models will guide the creation of a new sensor suite that, together with the carefully-crafted models, will form the basis for reliable robotic grasping systems.

Highly Instrumented Robotic Hand

Grasp Parameters	In-Hand	Lab Gold-Standard
Contact Location	Optical tracker (Atracsys Fusion Track 500, resolution: 90 μm RMS) + geometry of fingertip & object	Joint angles (0.0219°/LSB) + kinematic models (±5 mm)
Surface Normal	Optical Tracker + geometry of fingertip & object	Contact location (fingertip frame) + fingertip geometry + kinematics
Contact Force/Torque	Force Torque sensor (ATI Nano17, resolution: 1/160N, 1/32Nmm)	(same as in-hand)

Improved Hybrid Grasp Parameter Estimation

- Pure end-to-end black-box machine learning and pure physics-based methodologies in the literature have unsatisfactory performance when it comes to high-accuracy grasp stability prediction
- Relaxing physics-based approach and improving grasp parameter estimation by allowing noise-induced error terms, which can depend on forces/torques, on the physical measurements and performing gradient descent (blue curve below)
- Further improvement using an ML-based estimation where we use our optimized kinematics grasp parameters in addition to the joint angles and forces/torques as inputs to an ML algorithm (green and red curves use those in ridge regression, neural network respectively)
- Gold-standard grasp parameters, necessary for this ML-based improvement, are obtained with the optical tracker. In the next figure, we observe a direct improvement of the grasp parameter estimation where the ML algorithm was trained on multiple objects and evaluated on a cube.

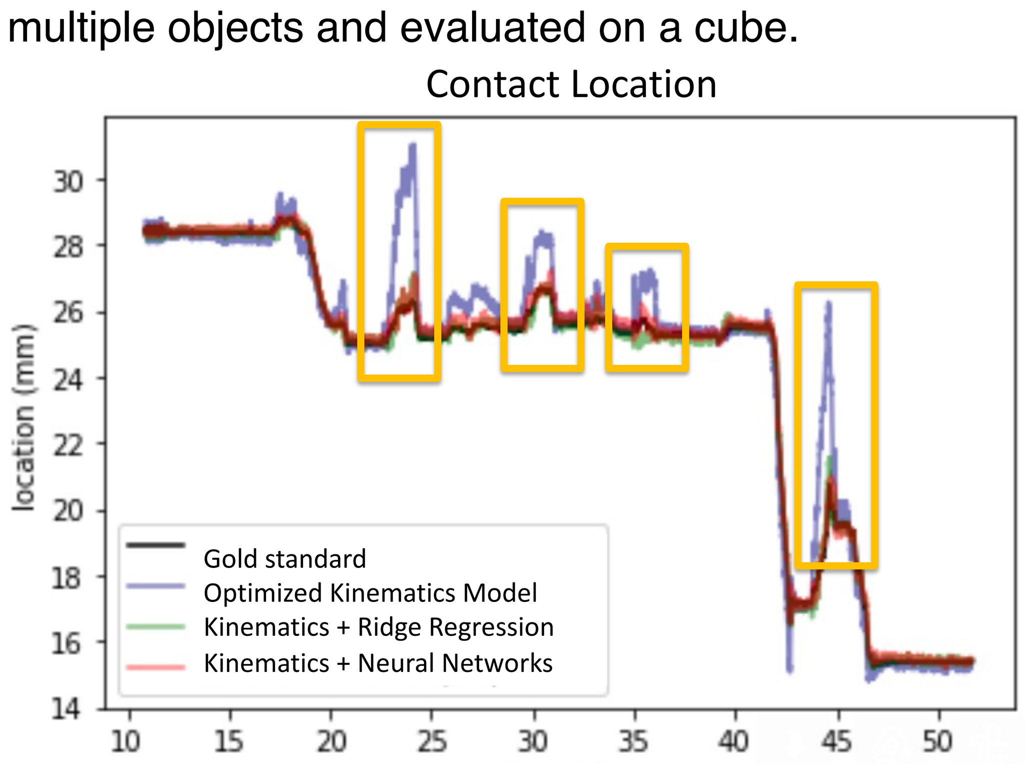

Highlighted in yellow are instances when ML significantly improves upon the optimized kinematics model. In the above experiment, an object is grasped in the robot hand and its stability is disturbed via external force, causing the fingers to briefly slip (See Figure (B)). The spikes in the blue curve are caused by external disturbance.

Improving Physics and ML Grasping

- Improving grasp parameter estimation is one step towards the actual goal of grasp stability prediction
- We want to compare our hybrid approach to the physics-based one and a pure ML approach, we decided to use SVM, which has proved to be very popular in the existing grasp stability literature
- SVM involves no physics thus acting as a black box, while our hybrid model can still benefit from the physics-based approach for grasp stability prediction through the derivation of the friction convex hull from the contact locations and surface normal

Stochastic Friction Models for Practical Grasping and Manipulation

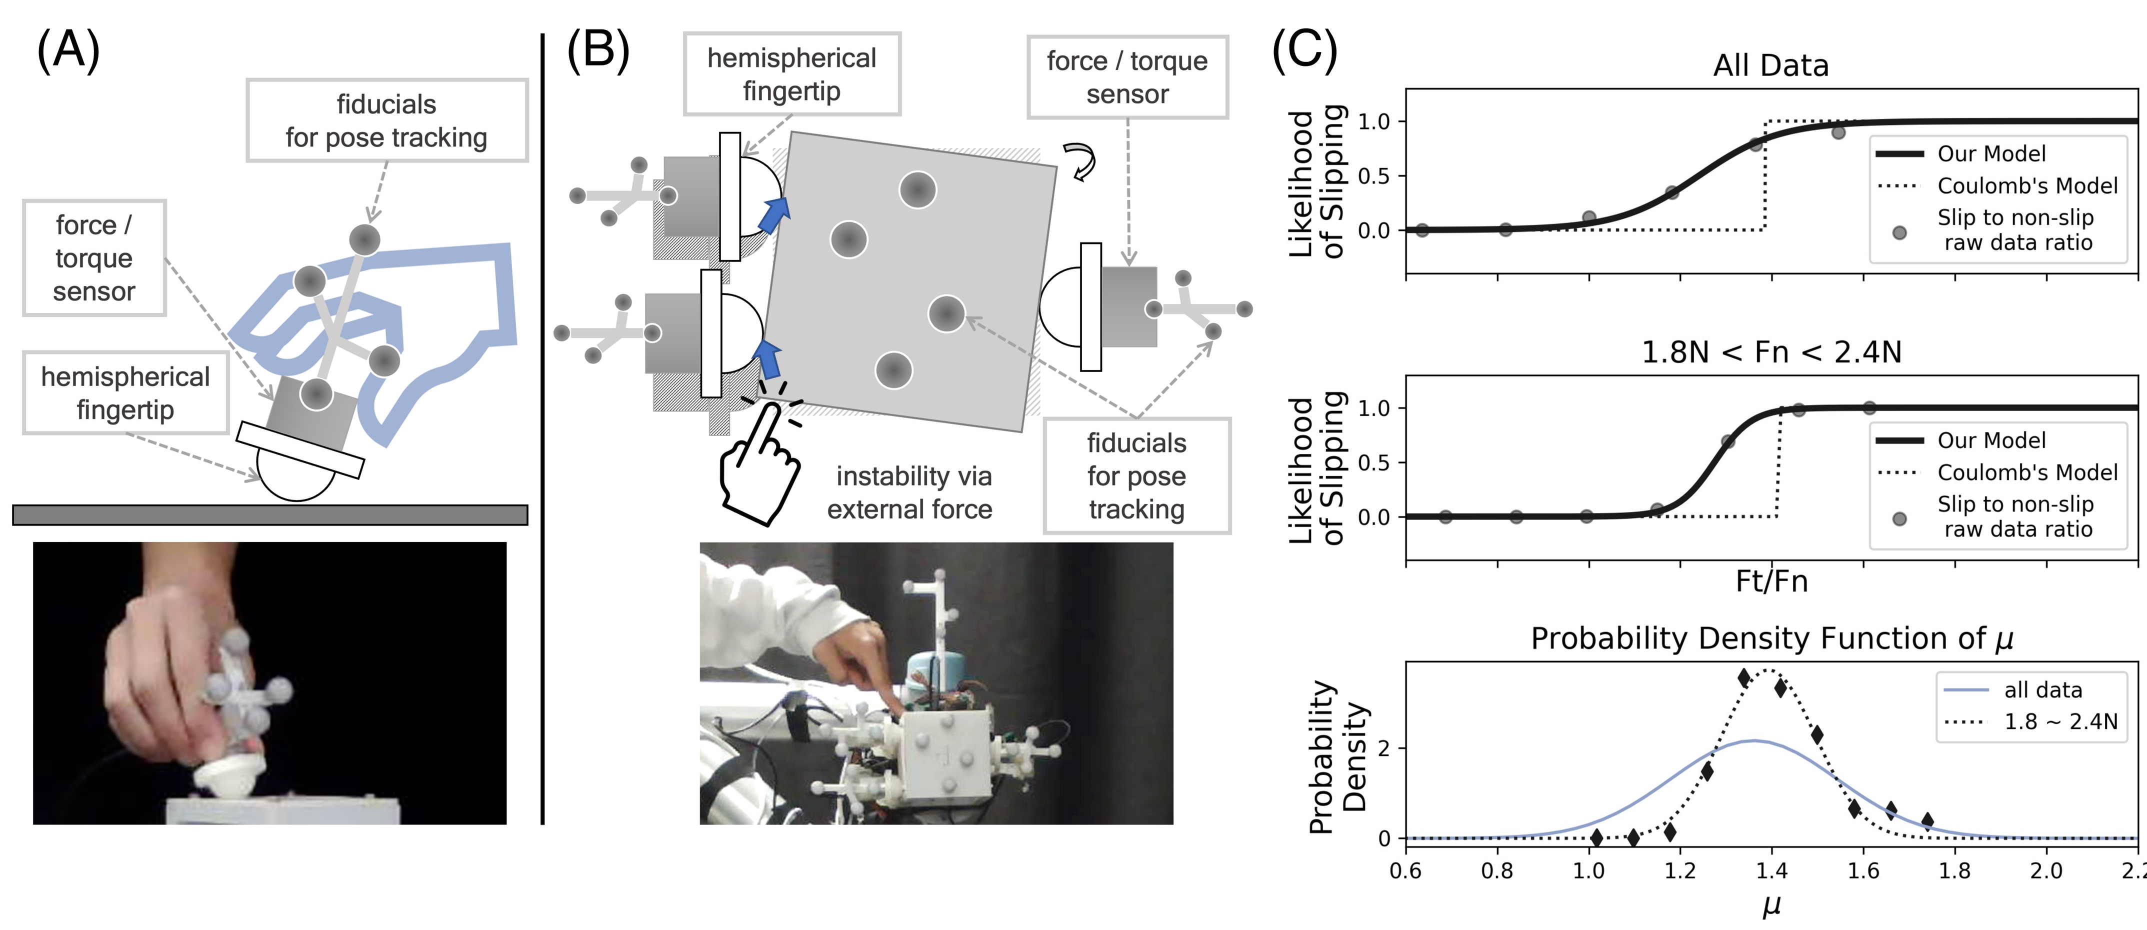

(A) Handheld experiment: a human operator holds a robot finger and slides across surfaces in varying contact force, velocity, and contact location. (B) Robot hand experiment: the robot hand grasps an object and a human operator applies external force to cause instability, leading to slips on one or more fingertip(s) without pushing the object out of the hand. (C) An example of using prior information on Fn to narrow the distribution of μ. The top plot shows the likelihood of slip (raw data for the slip-to-non-slip ratio of points in a certain Ft/Fn range) vs. Ft/Fn for all data for the heavyweight paper. The middle plot shows the ratio for the force range 1.8N < Fn < 2.4N. The bottom plot shows the probability density function of μ in the above plots.

The coefficient of friction is a key component of reliable grasping and manipulation. It is typically estimated in robotics applications using Coulomb's law of friction as a constant coefficient of friction from the literature, even though actual friction behavior is **variable** and depends on many factors. Here we conducted sliding experiments with robot fingers and a hand, and show that rubber friction varies strongly with normal force Fn and contact velocity v, and includes a significant stochastic component. **We present a framework for modeling the coefficient of friction μ as a distribution rather than a single constant** and show how this distribution can be narrowed when given a prior on Fn or v. For a given distribution, the likelihood of slipping is a continuous function with respect to the tangential- to-normal force ratio, instead of a step function according to Coulomb's law. By modeling friction as a function of Fn and v, we demonstrate that friction parameters can be estimated using regression models from a single sliding stroke of the fingertip against the object surface and that strokes that span a larger range of Fn-v space provide better friction estimates. These results can be applied to grasp control to enable a quantitative trade-off between the likelihood of slipping vs. grasp force levels and to sliding manipulation planning by clarifying the relationship between desired velocity and anticipated force levels. Application to machine learning has the potential to enhance reinforcement learning and sim-to-real transfer by providing more accurate representations of frictional behavior.

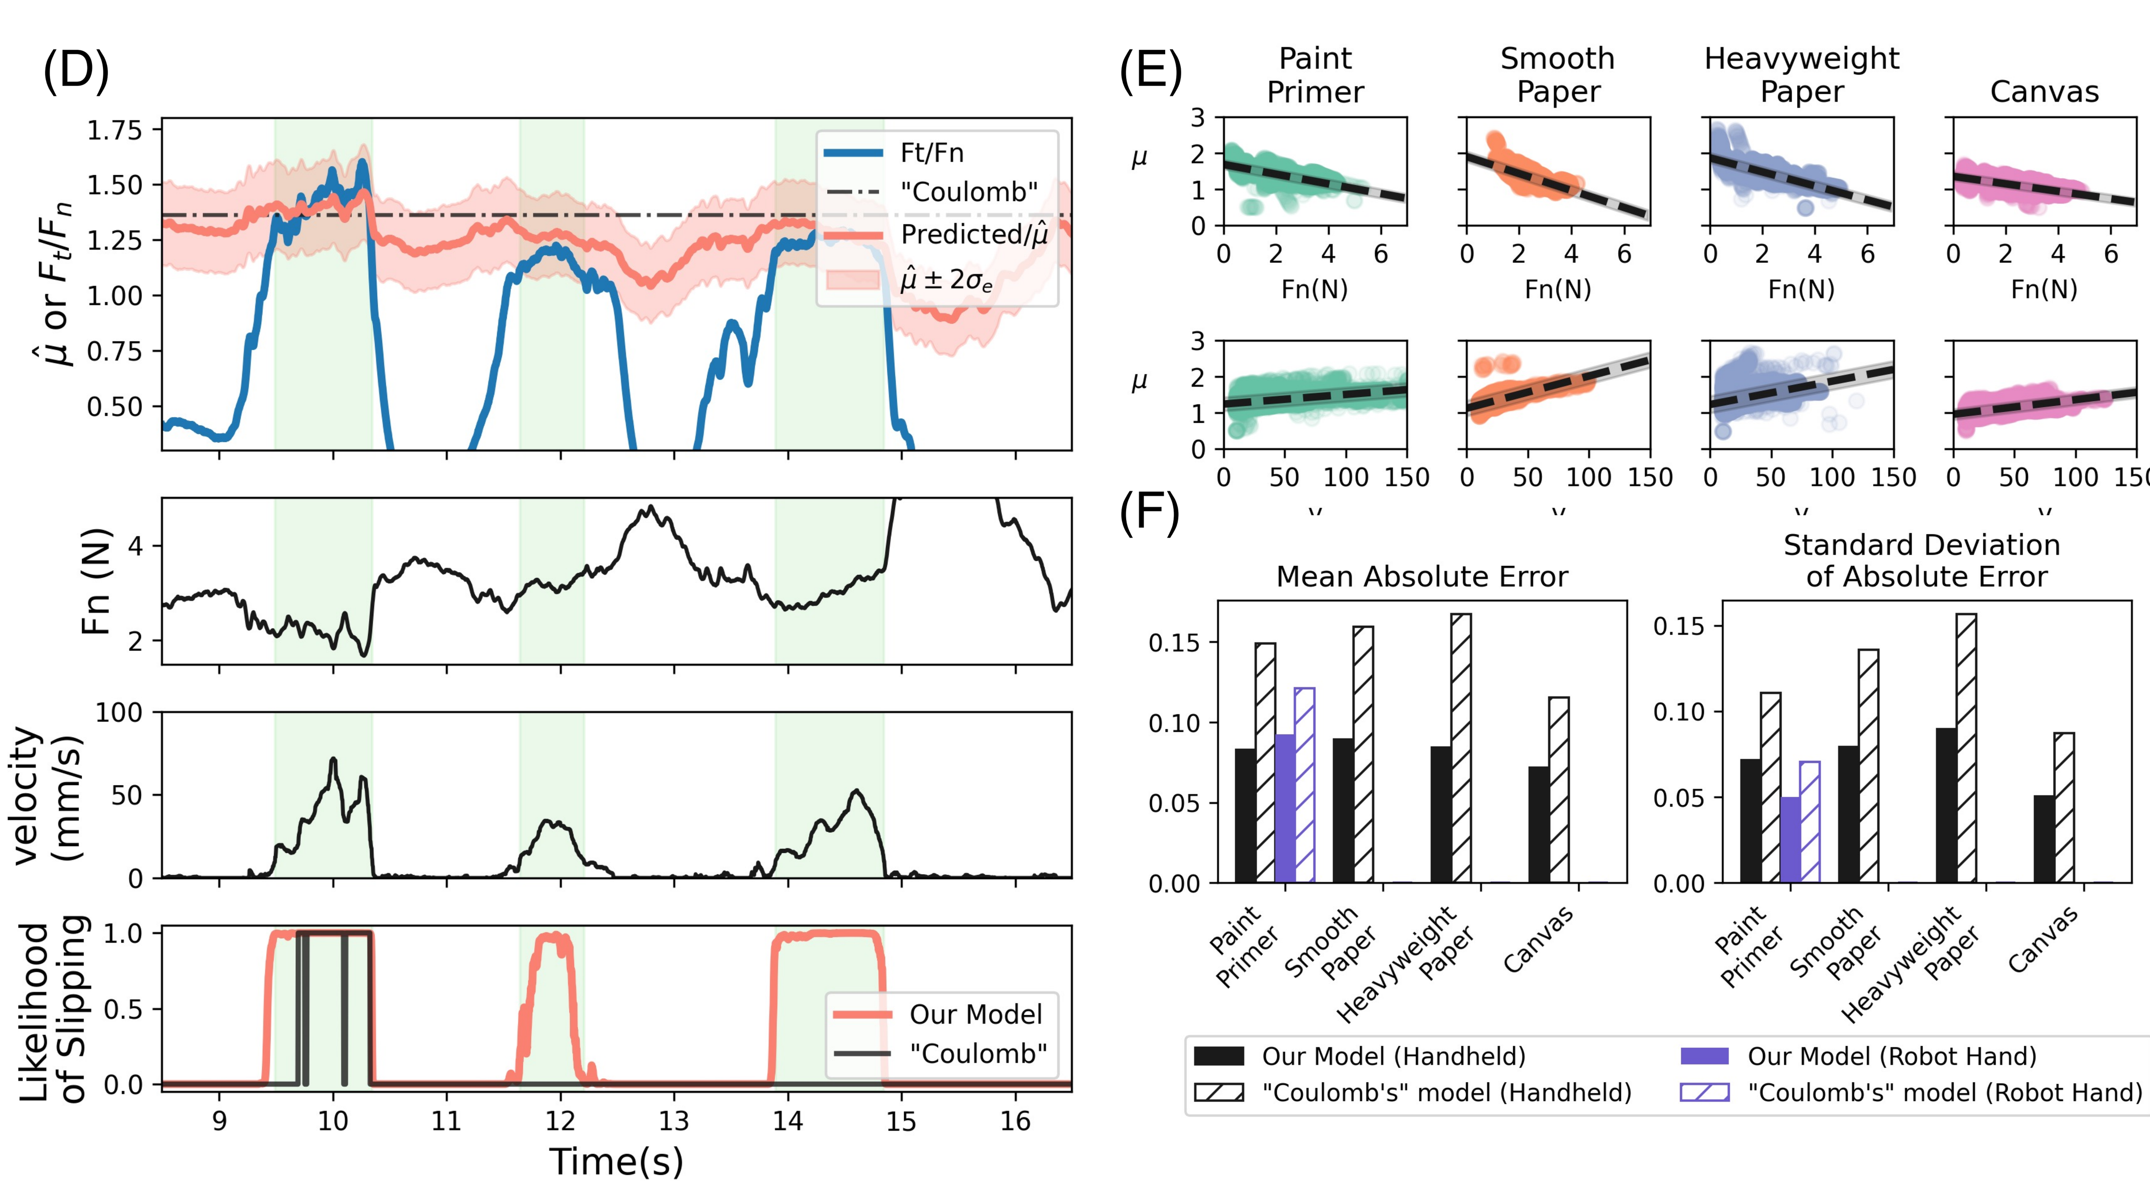

(D) An example of the variation of Ft/Fn during and across slipping episodes and the likelihood of slip from our model vs. Coulomb's model. Green highlights denote slipping intervals. Note that our model predicts $\hat{\mu}$ with low error and reports a continuous likelihood of slipping. (E) The coefficient of friction with respect to the normal force (top row), and contact velocity (bottom row) across four materials. Dashed line shows linear regression to the data. (F) The mean and standard deviation of absolute error of the estimation of the coefficient of friction for our linear regression model vs. Coulomb's model. Our model outperforms Coulomb's model in every case. The results from the robot hand experiment are on par with the handheld experiment.

2023 FRR & NRI Principal Investigators' Meeting
May 2-3, 2023
Award ID#: 1924984