Bio-inspired Robotics: Adaptive Object Handling Through Cellular Mechanics and Artificial Intelligence

Syed Hamzah Rizvi
West Lafayette Jr/Sr High School
West Lafayette, Indiana

Unless otherwise noted, all figures were made by the participant

Introduction: Motivation and Problem

Cell Binding and Unbinding

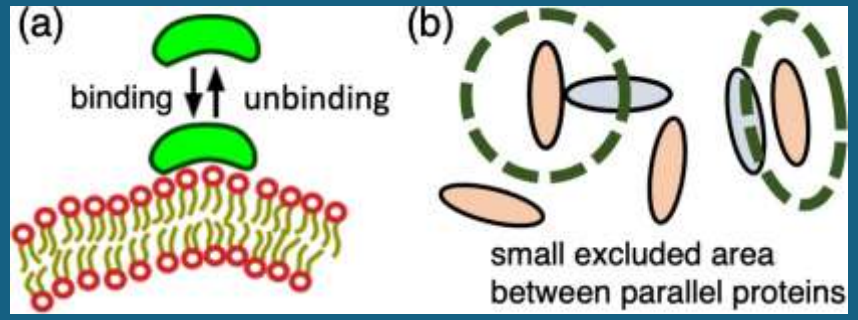

Source: https://visualsonline.cancer.gov/details.cfm?imageid=8295

- Cells regulate force through dynamic binding at adhesion sites
- This allows biological systems to adapt grip strength in real time
- Traditional robotic grippers use fixed stiffness and static control
- As a result, there is poor handling of fragile, irregular, or unknown objects

Why current Robotic Grippers fail

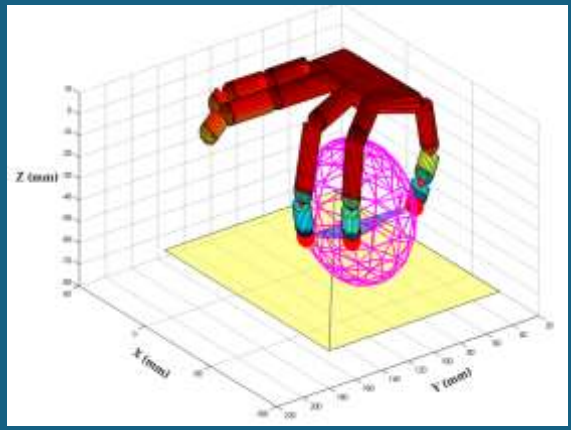

Source: https://www.researchgate.net/figure/Force-closure-grasp-of-the-sphere

- Rigid or PID (Proportional Integral Derivative) grippers apply fixed forces
- Object stiffness is mostly unknown and variable
- This results in slipping of the object (too soft) or damage to the object (too stiff)

Proposed Framework and Methods

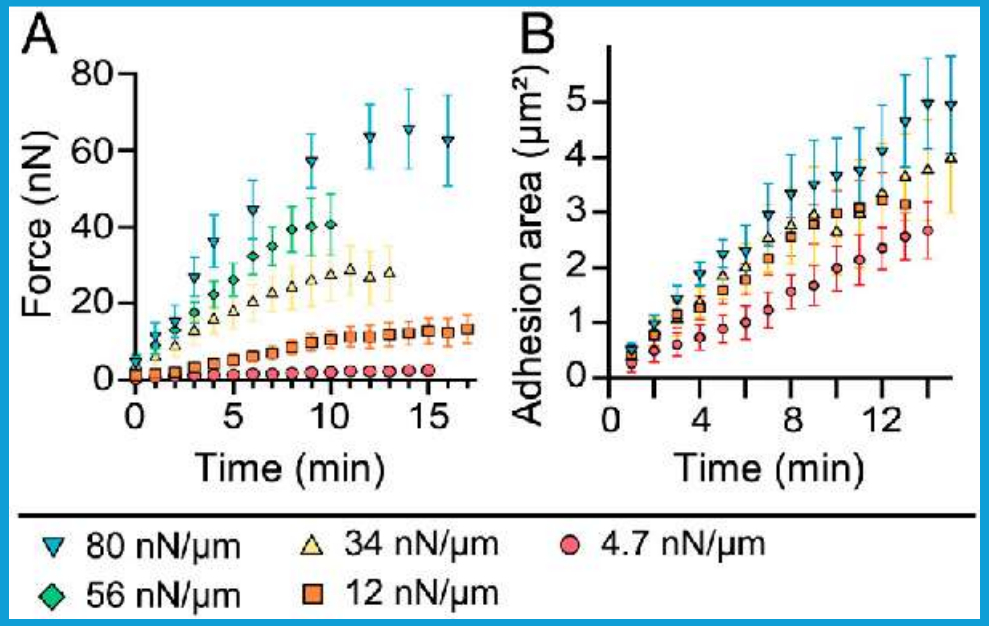

Source: https://www.researchgate.net/figure/Traction-forces-and-focal-adhesion-dynamics

1. Cells regulate force through stochastic bond dynamics
2. This enables stability across varying stiffness and load

Bell Model for Force-Dependent Bond Dynamics

↓

Monte-Carlo Process

↓

Modelling Gripper-Object Interaction

↓

Feature Extraction

↓

Unsupervised Learning, Clustering, and Statistical Analysis

YCB Benchmarks – Objects and Model Set

Point cloud and textural data overlays on two YCB objects: mustard bottle and power drill.

Source: https://www.ycbbenchmarks.com/object-models/

- For reproducibility purposes, the YCB Dataset was utilized and used throughout the project, paired with Python

Python Libraries Used

- NumPy: numerical arrays, force calculations, stochastic state updates, and implementing clutch and spring dynamics
- NumPy (random module): Monte Carlo replications, sampling uniform random numbers, probabilistic binding and unbinding events
- scikit-learn: K-means clustering, grouping mechanical response features
- pandas: Loading and organizing YCB metadata, managing simulation outputs. Structuring feature tables
- matplotlib: force vs time plots, variance comparisons, clustering visualizations

Bell Model for Force-Dependent Bond Dynamics (Method 1)

- To simulate the clutch-like behavior, each element was modeled as either bound (attached) or unbound (not attached).
- Two equations can be derived from the Bell Model for force-dependent bond dynamics
- In this model, applied forces add mechanical energy to a bond and lower the energy barrier required to break, which increases the unbinding rate exponentially.

Bell Model and Bond Dynamic Equations

$$k_{off}(F) = k_{off,0}\, e^{F/F_b}$$

$$k_{on}(F) = k_{on,0}\, e^{-F/F_b}$$

where:

- F = Current force
- F_b = Characteristic bond force
- $k_{off,0}, k_{on,0}$ = base rates at zero force

Monte-Carlo Process (Method 2)

- Each clutch element is a two-state system: bound or unbound. This type of randomness can be modeled mathematically using a Poisson process, which describes the probability of events occurring in each time interval.
- The transition probabilities follow the definition that a rate multiplied by a small-time interval gives the probability of an event occurring during that time interval. (Figure A)
- This can be simulated to determine whether a clutch breaks or binds. A Monte Carlo step is used, where a uniform random number r ∈ [0, 1] is generated. If a clutch is bound and $r < P_{unbind}$, it unbinds; if it is unbound and $r < P_{bind}$, it binds.

Monte-Carlo Process Extension

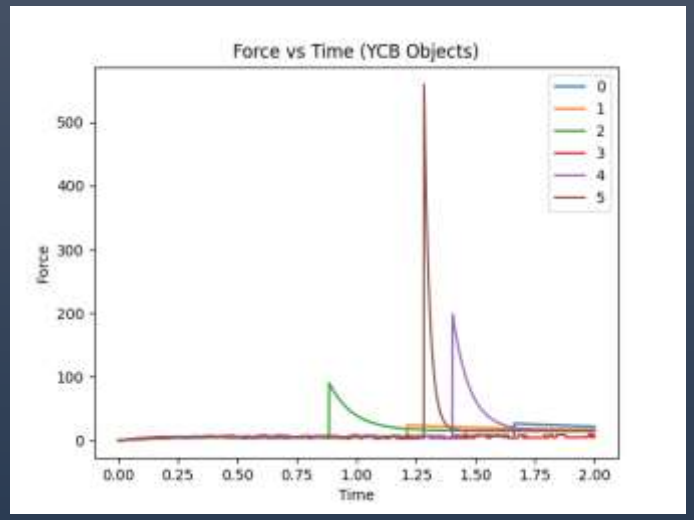

Figure A: Time-dependent force profiles for YCB objects

Modelling the Gripper-Object Interaction as Springs in Series (Method 3)

- To simulate how force is transferred and transmitted from the robotic gripper to the object through molecular clutch elements, the system can be modeled using a simple spring in series system
- Springs are used as they use Hookean mechanics, which can provide a simple yet effective way to model how much force results from a given force and displacement.
- Series springs are preferred over parallel because force travels from the gripper through the clutch system into the object, meaning one must deform after the other under the same applied force

Gripper-Object Interaction Equations

The equations for effective stiffness can be derived using the series spring equations:

$$x_1 = \frac{F}{k_{spring}} \quad \text{and} \quad x_2 = \frac{F}{k_{object}}$$

Total displacement can also be modelled by:

$$x_{total} = x_1 + x_2 = F\left(\frac{1}{k_{spring}} + \frac{1}{k_{object}}\right)$$

A single equivalent spring would satisfy:

$$x_{total} = \frac{F}{k_{eff}}$$

Setting the two expressions equal and solving gives:

$$\frac{F}{k_{eff}} = F\left(\frac{1}{k_{spring}} + \frac{1}{k_{object}}\right)$$

In this specific model, one spring represents the effective stiffness of the working and engaged clutches, which is labeled as k_{object} . The second spring represents the 6 mechanical stiffness of the object, which is k_{spring} . The resultant equation from the derivation will be modelled once effective stiffness is achieved:

$$F = k_{eff}x$$

Feature Extraction (Method 4)

- From each gripper-object simulation, a set of quantitative mechanical features were extracted to categorize grip behavior and stability.
- Steady-state grip force and estimated stiffness were computed from the final portion of each simulation to quantify average force levels and grip stability
- Time to slip was defined as the first abrupt drop in force, indicating grip failure, while maximum supported force was recorded as the peak force
- Binding fraction was evaluated as the time-averaged proportion of clutch elements in the bound state, linking bond dynamics to force transmission

Steady-State Force (4.1)

- Steady-State force can be defined as the average of the force samples over the final portion of the simulation
- This can be possible by using the standard definition of a time average
- This captures the typical force level after the system has settled. If $F(t_j)$ is the force at the sample time t_j , then the steady-state force is:

$$F_{\text{avg}} = \frac{1}{n+1} \sum_{j=0}^{n} F(t_j)$$

Estimated Stiffness k_{est} (4.2)

- The estimated stiffness can be computed from the slope of the force-displacement curve.
- The two points in this case would be(x_1, F_1) and (x_2, F_2),and on the curve, the effective stiffness would be:

$$k_{\text{est}} = \frac{\Delta F}{\Delta x} = \frac{F_2 - F_1}{x_2 - x_1}$$

Time to Slip t_{slip} (4.3)

- Time to slip is defined as the earliest moment at which the force profile shows an unexpected, sudden drop, indicating that the object has begun to lose grip.
- This is implemented by scanning for the first index when the force decreases:

$$t_{\text{slip}} = \min\{t_j \mid F(t_{j+1}) < F(t_j)\}$$

This feature identifies how long the system can sustain load before slipping occurs, which is important for separating stable and unstable grip conditions.

Binding Fraction (4.4)

- A clutch can only be in one of two states at any given time interval, can either be bound (1) or unbound (0).
- For each clutch *i*, we define a variable:

$$b_i(t) = \begin{cases} 1 & \text{if clutch } i \text{ is bound at time } t \\ 0 & \text{otherwise} \end{cases}$$

- This will be named as an indicator variable, which marks when the clutch is attached. If there are N clutches total, then summing all their states at time t tells us how many clutches are bound.
- The model counts the number of 1s (bound clutches). Unbound clutches contribute to zero in this case. Substitute the calculated count:

$$\sum_{i=1}^{N} b_i(t)$$

- The binding fraction measures the cluster of clutches that are attached at a given time. Each clutch is represented by a binary variable $b_i(t)$, equal to 1 if the clutch is bound and 0 if the clutch is unbound:

$$B(t) = \frac{1}{N} \sum_{i=1}^{N} b_i(t)$$

- Summing these values counts the number of attached clutches, and dividing by the total number of clutches in total gives the fraction

Maximum Supported Force F_{max} (4.5)

- The maximum supported force represents the highest force value during a given simulation before a slip or fall occurs.
- Since the simulation outputs discrete force samples at times $t_0, t_1, \ldots . t_n$ the maximum supported force is simply the largest among all of them.
- Formally, if $F(t_j)$ is the force at time t_j , then:

$$F_{\max} = \max_{0 \le j \le n} F(t_j)$$

- The maximum supported force was defined as the highest force observed during the simulation. If $F(t_j)$ represents the force at t_j, then the maximum force is:

$$F_{\max} = \max_{j} F(t_j)$$

- This directly measures how much load the system can sustain before a slip or fall occurs, making it useful for comparing objects with varying mechanical strengths.

Optimal Grip Force (4.6)

- The optimal grip force can be defined as the smallest safe force that prevents slipping without exceeding the material's crush threshold/limit. This leads to the constraint:

$$F_{\text{slip}} < F_{\text{optimal}} < F_{\text{crush}}$$

- and the chosen operating point:

$$F_{\text{optimal}} = F_{\text{slip}} + \epsilon$$

- where ϵ is a small safety margin (5–10% of the slip force). This rule ensures a stable grip without applying unnecessary or damaging force .

Extension of Unsupervised Machine Learning (1)

Solve for μ:

$$n\mu = \sum_{j=1}^{n} x_j$$

$$\mu = \frac{1}{n} \sum_{j=1}^{n} x_j$$

The full K-Means objective function is:

$$J = \sum_{i=1}^{k} \sum_{j \in C_i} |x_j - \mu_i|^2$$

Unsupervised Learning, Clustering, and Statistical Analysis (Method 5)

- **K-means clustering** was used to group simulated mechanical responses into object groups and categories without the need for labeled data (Figure B)

- The algorithm partitions feature vectors into k clusters by minimizing the sum of squared distances between each given point and its corresponding cluster centroid.

- The objective function for K-means is and its derivation:

$$J_i(\mu) = \sum_{x \in C_i} |x - \mu|^2$$

Extension of Unsupervised Machine Learning (2)

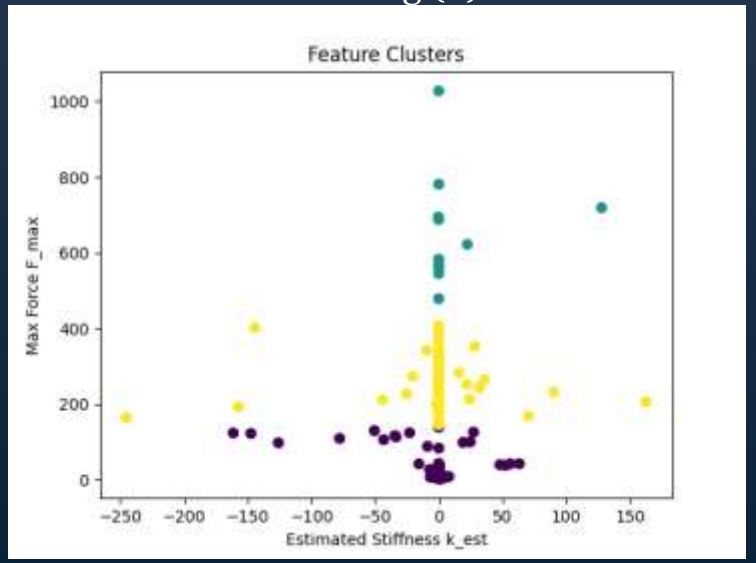

Figure B: Feature space clustering revealing three mechanically distinct object categories

Performance Endpoints (5.1)

- To enable rigorous statistical evaluation of controller performance, multiple performance metrics were considered prior to recording all large-scale simulation experiments.
- Due to binary task completion being infrequent under challenging simulated environments and failing to capture graded differences in grip stability, time to slip was selected as the primary endpoint for statistical inference.
- Time to slip was defined as the cumulative duration for which the object remained stably grasped before the onset of slip, detected as a sudden force drop, and was additionally capped as a time of T seconds for trials in which no slips were observed.

Statistical Analysis (5.2)

- Simulation-generated performance data do not generally satisfy parametric assumptions such as **normality** or **homoscedasticity**; all statistical analyses were conducted using nonparametric inference and uncertainty qualifications.
- The difference in mean performance between each of the controllers was reevaluated across several random permutations to generate an empirical null distribution .
- Two-sided p-values were computed as the proportion of permutations producing a difference as large as the observed value.

- To calculate uncertainty in estimated performance improvements, bootstrap confidence intervals were computed.
- Trials were resampled with replacement within each controller group, and the difference was recalculated across bootstrap replicates. The 95% confidence interval was defined by the 2.5th and the 97.5th percentiles of the bootstrap distribution.
- For continuous-valued secondary endpoints, standardized effect sizes (Cohen's d) were further reported where they were appropriate. **All statistical analyses were computed in Python using fixed random seeds to ensure reproducibility.**

Statistical Analysis Extension

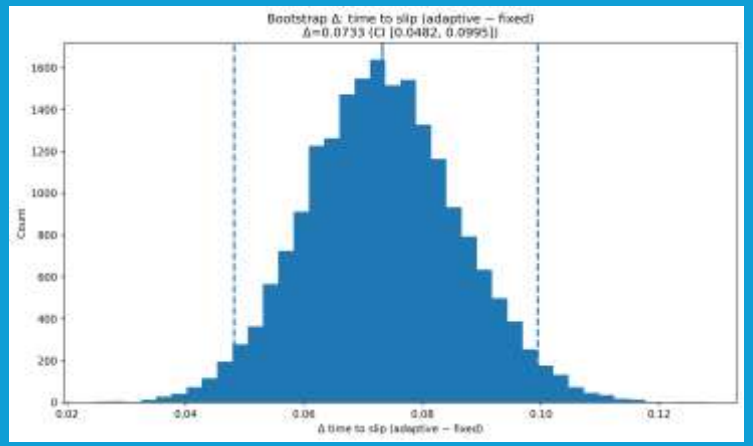

Figure C: One of the supplementary figures that shows the confidence intervals between adaptive and fixed controllers

Results

- To ensure that the simulation parameters reflect realistic object mechanics, material categories (texture, softness, etc.), and frictional properties were input using metadata from reputable datasets such as the Yale-CMU-Berkley (YCB) object dataset
- Objects in the YCB dataset were then categorized into three groups based on manufacturer-listed stiffness properties: soft (foams, plastics), medium (composites), and stiff (metals).
- These simulations were then mapped to the simulation parameters k_{object} and F_{crush} to ensure physical realism in the generated training data.

Fixed vs PID (YCB Object Dataset) Figure D

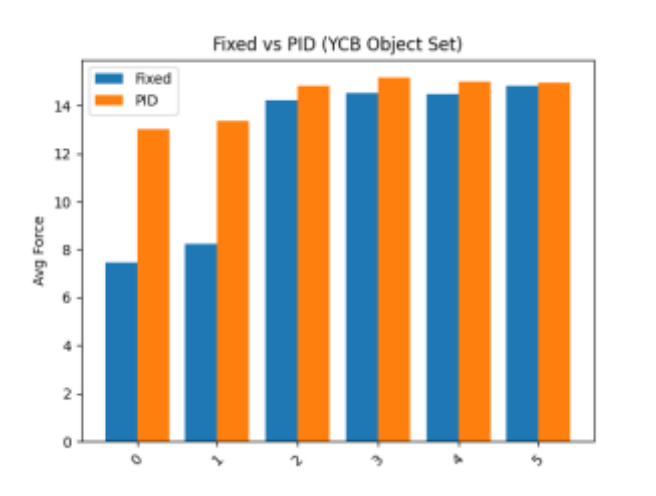

Figure D: Comparison of Fixed and PID controller performance on YCB objects

Fixed vs PID (YCB Object Dataset) Results A

- This bar chart (Figure D) shows a comparison between the models, the fixed model, and the PID model. The PID consistently produces higher average forces.
- Traditional controllers apply excessive or insufficient force due to the lack of information about the stiffness of the materials.
- PID over-grips across all YCB objects, showing poor adaptability

Adaptive Controller Performance (Results B)

- The histogram (Figure E) highlights the distribution of average forces applied by the adaptive controller when gripping various YCB objects.
- The adaptive controller assigns a unique grip force to each identified cluster by the K-Means, and the distribution forms distinct peaks representing learned optimal force requirements.
- The low-force peak corresponds to soft objects such as foam and rubber, the mid-force peak represents medium objects like plastic and wood, and the high-force peak indicates stiff and dense objects, including ceramic and metal. (Continued on Next Slide)

Fixed vs PID (YCB Object Dataset) Figure E

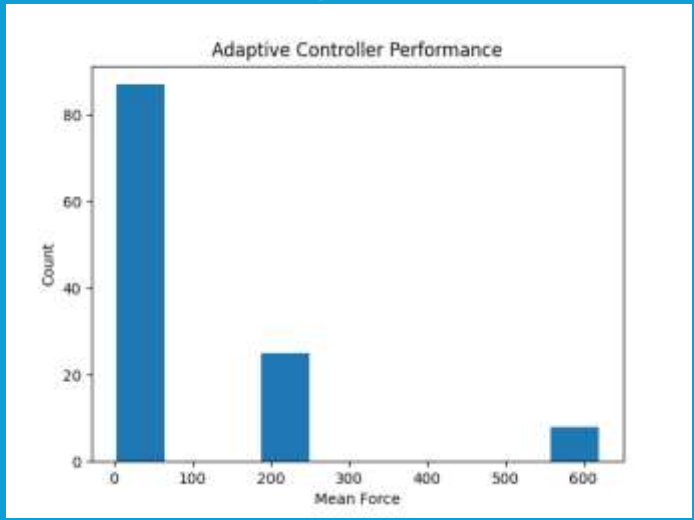

Figure E: Distribution of forces applied by the adaptive controller across YCB objects

Binding Fraction vs Time (Figure F)

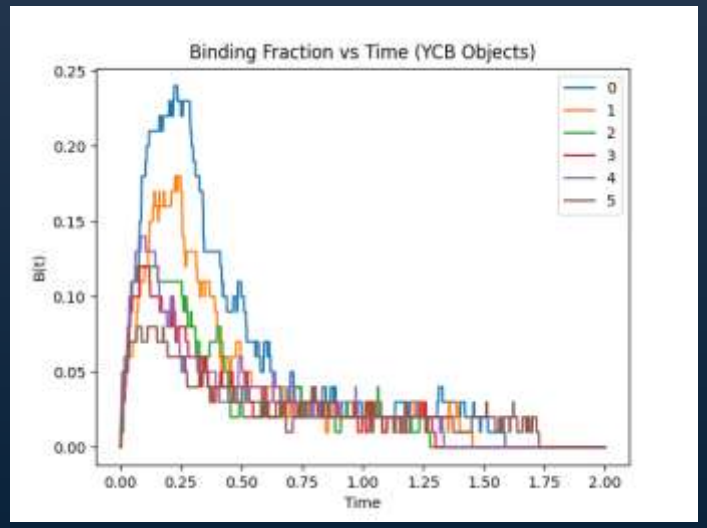

Figure F: Binding fraction evolution over time under fixed-force control

Binding Fraction vs Time (Results C)

- The figure (Figure F) shows the binding fraction B(t) over time for YCB objects under a fixed-force controller. It shows the convolution of the binding fraction, defined as the number of clutch elements attached at any given moment.
- All objects exhibit an initial rise in binding, followed by a sudden peak and a collapse corresponding to slip or rapid unbinding events.
- The behavior matches predictions from the molecular clutch model and confirms that the simulation responds realistically to varying object stiffness.

Force vs Time (YCB Objects) Results D

- The plot (Figure A) displays the time-dependent force applied to each object during a given gripping simulation.
- Each curve corresponds to a different object in the YCB dataset. It varies significantly depending on the object's stiffness.
- Continuation and Figure A on the next slide.

Force vs Time (YCB Objects) Figure A

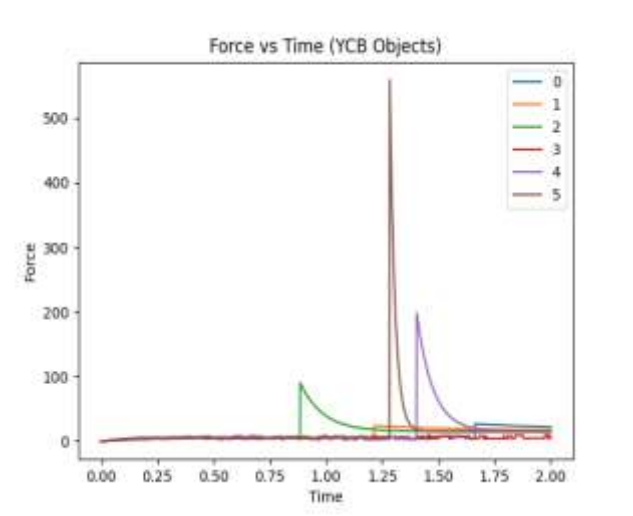

Figure A: Time-dependent force profiles for YCB objects

Feature Clusters (Results E)

- This is the clustering of extracted features using the K-Means algorithm, plotted by estimated stiffness and maximum force F_{max}. (Figure B)
- Each point represents a simulation trial encoded by two meaningful features: estimated stiffness (x-axis) and maximum supported force (y-axis).
- K-Means identifies three major clusters, separating objects into varying mechanical categories.

Feature Clusters (Figure B)

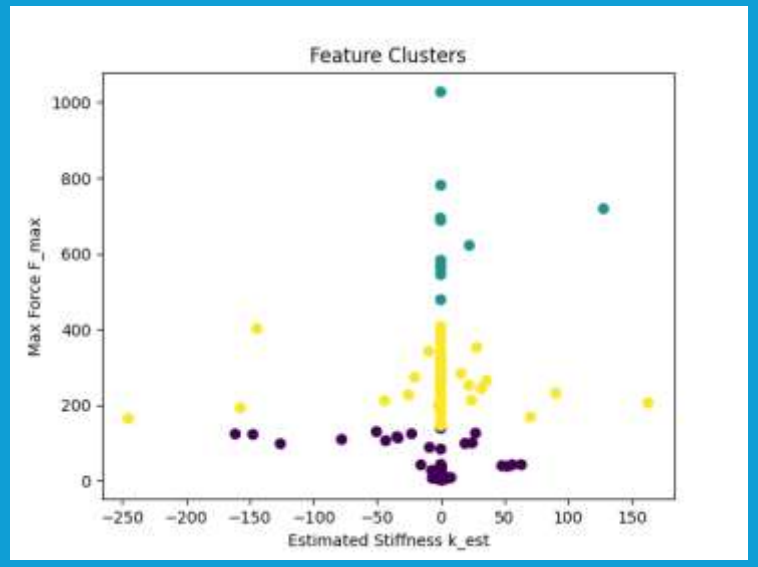

Figure B: Feature space clustering revealing three mechanically distinct object categories

Statistical Validation of Controller Performance

- To quantitatively evaluate whether observed performance differences between controllers were statistically robust, nonparametric methods were implemented to aggregate simulation outcomes.
- Statistical inference focused primarily on time-to-slip, the designated endpoint
- Across repeated Monte Carlo trials conducted under randomized object characteristics and environmental disturbances, **the adaptive controller showed a longer mean time-to-slip in contrast to baseline controllers.** (Continued on Next slide)

- Permutation testing revealed that the seen differences in time-to-slip were unlikely to arise under the null hypothesis of no controller (permutation testing indicated a statistically robust improvement for adaptive versus fixed-force control: two-sided permutation; Figure G)
- The 95% confidence interval for the difference in mean time-to-slip excluded zero, indicating that the observed improvement was stable across resampled experimental conditions rather than isolated trials.
- The differences between the adaptive and PID controller is **smaller** and, in some "regimes," did not reach statistical significance, reflecting the ability of PID controllers to regulate and control force under moderately hard, uncertain conditions
- However, secondary endpoints, including maximum supported force and force variability, revealed consistent reductions in excessive applied force implementation by the adaptive controller.
- Due to the binary nature of the success rate and due to the limited sensitivity under challenging simulated conditions, it was not used as a primary inferential metric.

Two-sided Permutations (Figure G)

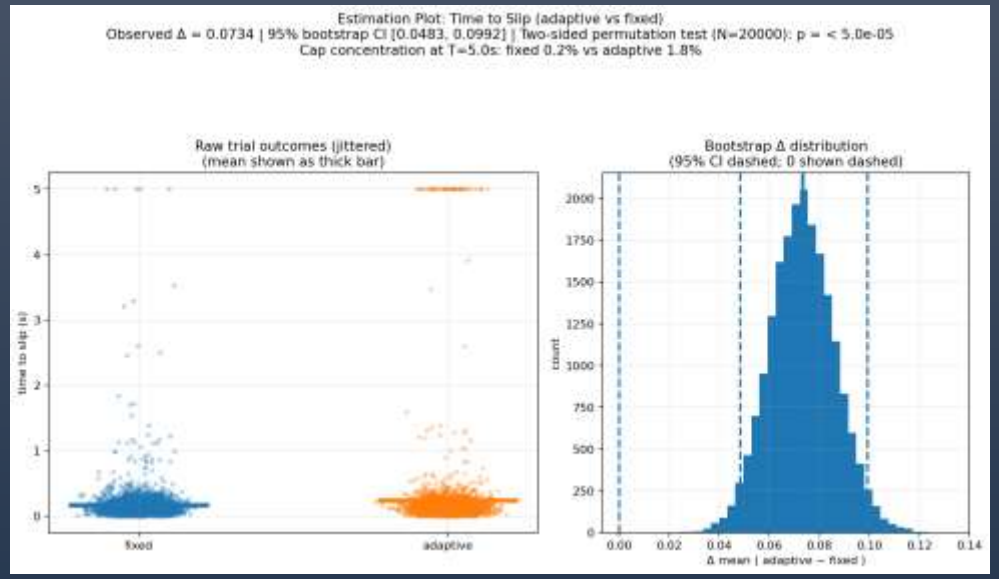

Estimation plot of time-to-slip for the adaptive controller compared with fixed-force control. The adaptive controller increases mean time-to-slip (, 95% bootstrap CI [0.0483, 0.0992] and the difference is unlikely under the permutation null

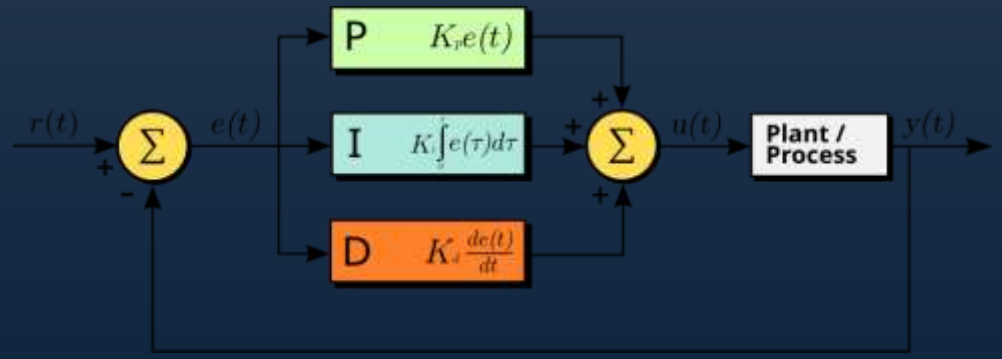

A block diagram of a PID controller in a feedback loop. $r(t)$ is the desired process variable (PV) or setpoint (SP), and $y(t)$ is the measured PV.

Conclusion

- This study demonstrated that force adaptation in robotic gripping can be obtained through a framework inspired by molecular clutch mechanics and assisted by unsupervised learning.
- When evaluated against traditional baselines, the adaptive controller demonstrated improved grip stability, reducing premature slip events and avoiding excessive force application. Importantly, performance enhancements were quantified using time-to-slip as the primary endpoint, a continuous and physically interpretable measure of grip stability.
- These results support the central hypothesis that clutch-based models paired with unsupervised mechanical classification can drastically improve performance without requiring labeled data or predefined material models.
- Nonparametric permutation testing and bootstrap resampling confirmed that the observed gains over fixed force control were statistically robust and not attributable to stochastic simulation noise.
- Improvements over PID control were smaller and, in specific regimes, approached the boundary of statistical significance, reflecting the partial adaptability of PID strategies under moderate to high uncertainty.
- This was further supported by statistical analysis, which confirmed that the adaptive strategy produced significantly less force variance than the fixed force and PID controllers, indicating more efficient and reliable grip behavior

Constraints and Future Goals

- Although the adaptive strategy showed meaningful improvements over already existing traditional controllers, the approach remains constrained by its simplified spring mechanics, **reliance on simulated clutch dynamics**, and the absence of real sensor noise or hardware validation. Such hurdles mean that the results **show theoretical feasibility rather than immediate real-world performance.**
- Keeping this in mind, the framework does provide a strong foundation for future projects, including the implementation of this idea onto robotic grippers, expansion to soft-robotic systems, or application to automated material classification tasks. This project constructs a clear pathway for extending clutch-based control into broader robotics.

References/Citations

- [1] Bell, G. I. (1978). Models for the specific adhesion of cells to cells. Science, 200(4342), 618–627. https://doi.org/10.1126/science.347575
- [2] Calli, B., Singh, A., Bruce, J., Walsman, A., Konolige, K., Srinivasa, S., Abbeel, P., & Dollar, A. M. (2015). Yale-CMU-Berkeley dataset for robotic manipulation research. International Journal of Robotics Research, 36(3), 261–268. https://doi.org/10.1177/0278364917706967
- [3] Chan, C. E., & Odde, D. J. (2008). Traction dynamics of filopodia on compliant substrates. Science, 322(5908), 1687–1691. https://doi.org/10.1126/science.1163595
- [4] Gillespie, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. The Journal of Physical Chemistry, 81(25), 2340–2361. https://doi.org/10.1021/j100540a008
- [5] Halliday, D., Resnick, R., & Walker, J. (2013). Fundamentals of Physics (10th ed.). Wiley.
- [6] MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability (Vol. 1, pp. 281–297). University of California Press.
- [7] ˚Astr¨om, K. J., & H¨agglund, T. (1995). PID Controllers: Theory, Design, and Tuning (2nd ed.). Instrument Society of America.