

ABSTRACT

A long-standing question in robot hand design is how accurate tactile sensing must be. This paper uses simulated tactile signals and the reinforcement learning (RL) framework to study the sensing needs in grasping systems. Our first experiment investigates the need for rich tactile sensing in the rewards of RL-based grasp refinement algorithms for multi-fingered robotic hands. We systematically integrate different levels of tactile data into the rewards using analytic grasp stability metrics. We find that combining information on contact positions, normals, and forces in the reward yields the highest average success rates of 95.4% for cuboids, 93.1% for cylinders, and 62.3% for spheres across wrist position errors between 0 and 7 centimeters and rotational errors between 0 and 14 degrees. This contact-based reward outperforms a non-tactile binary-reward baseline by 42.9%. Our follow-up experiment shows that when training with tactile-enabled rewards, the use of tactile information in the control policy's state vector is drastically reducible at only a slight performance decrease of at most 6.6% for no tactile sensing in the state. Since policies do not require access to the reward signal at test time, our work implies that models trained on tactile-enabled hands are deployable to robotic hands with a smaller sensor suite, potentially reducing cost dramatically.

1. INTRODUCTION

In this paper, we use accurate tactile signals from simulation and the reinforcement learning framework to explore the tactile sensing needs in robotic systems. We propose a unified framework to systematically incorporate different levels of tactile information from robotic hands into a reward signal via analytic grasp stability metrics. Furthermore, in Fig. 1 we hypothesize that policies trained with grasp stability metrics on a robotic hand H_f with a *full* tactile sensor suite are deployable to structurally similar but more affordable hands H_r with *reduced* tactile sensing at a small performance decrease.

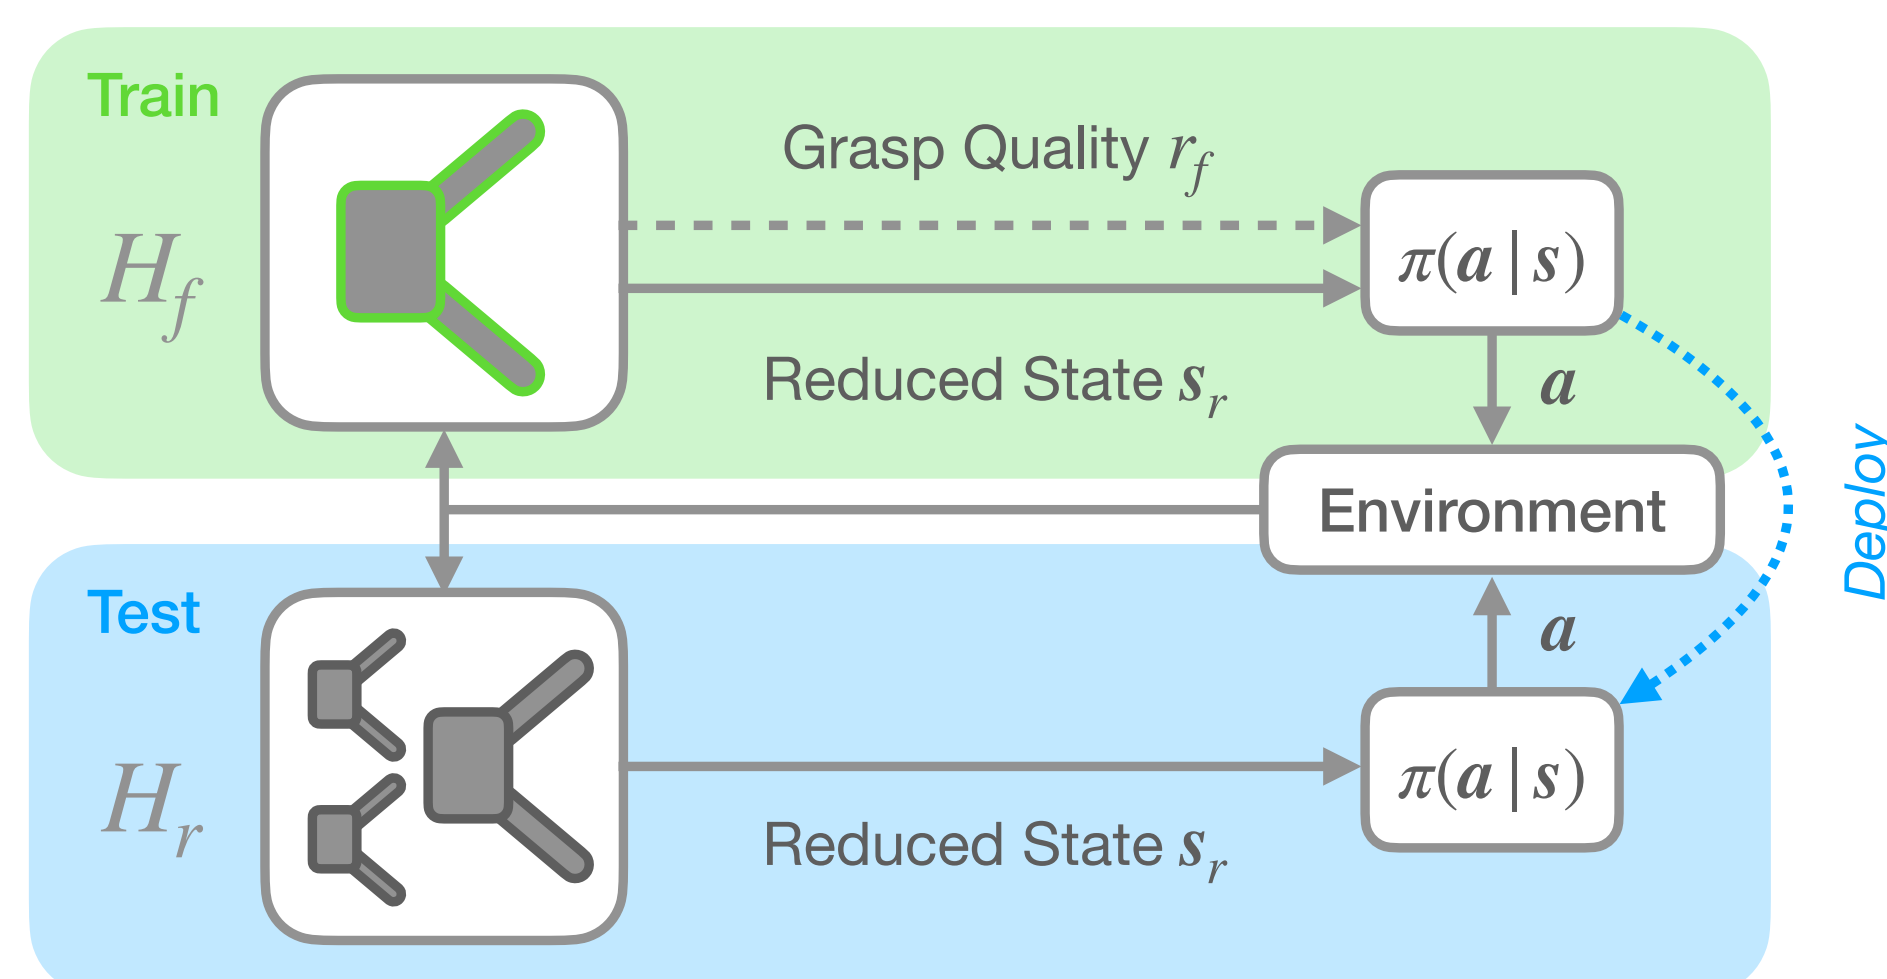


Figure 1: The hypothesized workflow for training and deploying RL-controlled grasping systems.

2. GRASP STABILITY METRICS

Mirtich and Canny [1] define two quality metrics ϵ_f and ϵ_τ that measure a grasp's ability to resist unit forces and torques, respectively. We define additional metrics. As shown in Fig. 2, δ_{cur} measures *current* grasp stability as the average magnitude of the safety margins $\bar{f}_{i,cur}$ to the friction cone. Moreover, δ_{task} measures expected grasp stability during *task execution* by estimating the task contact forces and their expected margins to the friction cone. While ϵ_f and ϵ_τ are a function of contact positions and normals, δ_{cur} and δ_{task} also take the actual contact forces that the contacts currently apply to the object into account.




Figure 2: Grasp with current contact forces $f_{i,cur}$ and tangential force margins $\bar{f}_{i,cur}$ to the friction cones.

3. TACTILE SENSING AND REWARD

In our first experiment, we estimate the relevance of contact position, normal, and force sensing for the **reward signal** in robotic grasp refinement. Fig. 3 shows an overview of one grasping episode. In stage (A), we initialize the world. Thereby, we randomly generate a new object, wrist error tuple (O, E). We vary the object O category (cuboid, cylinder, sphere), mass, and dimensions.

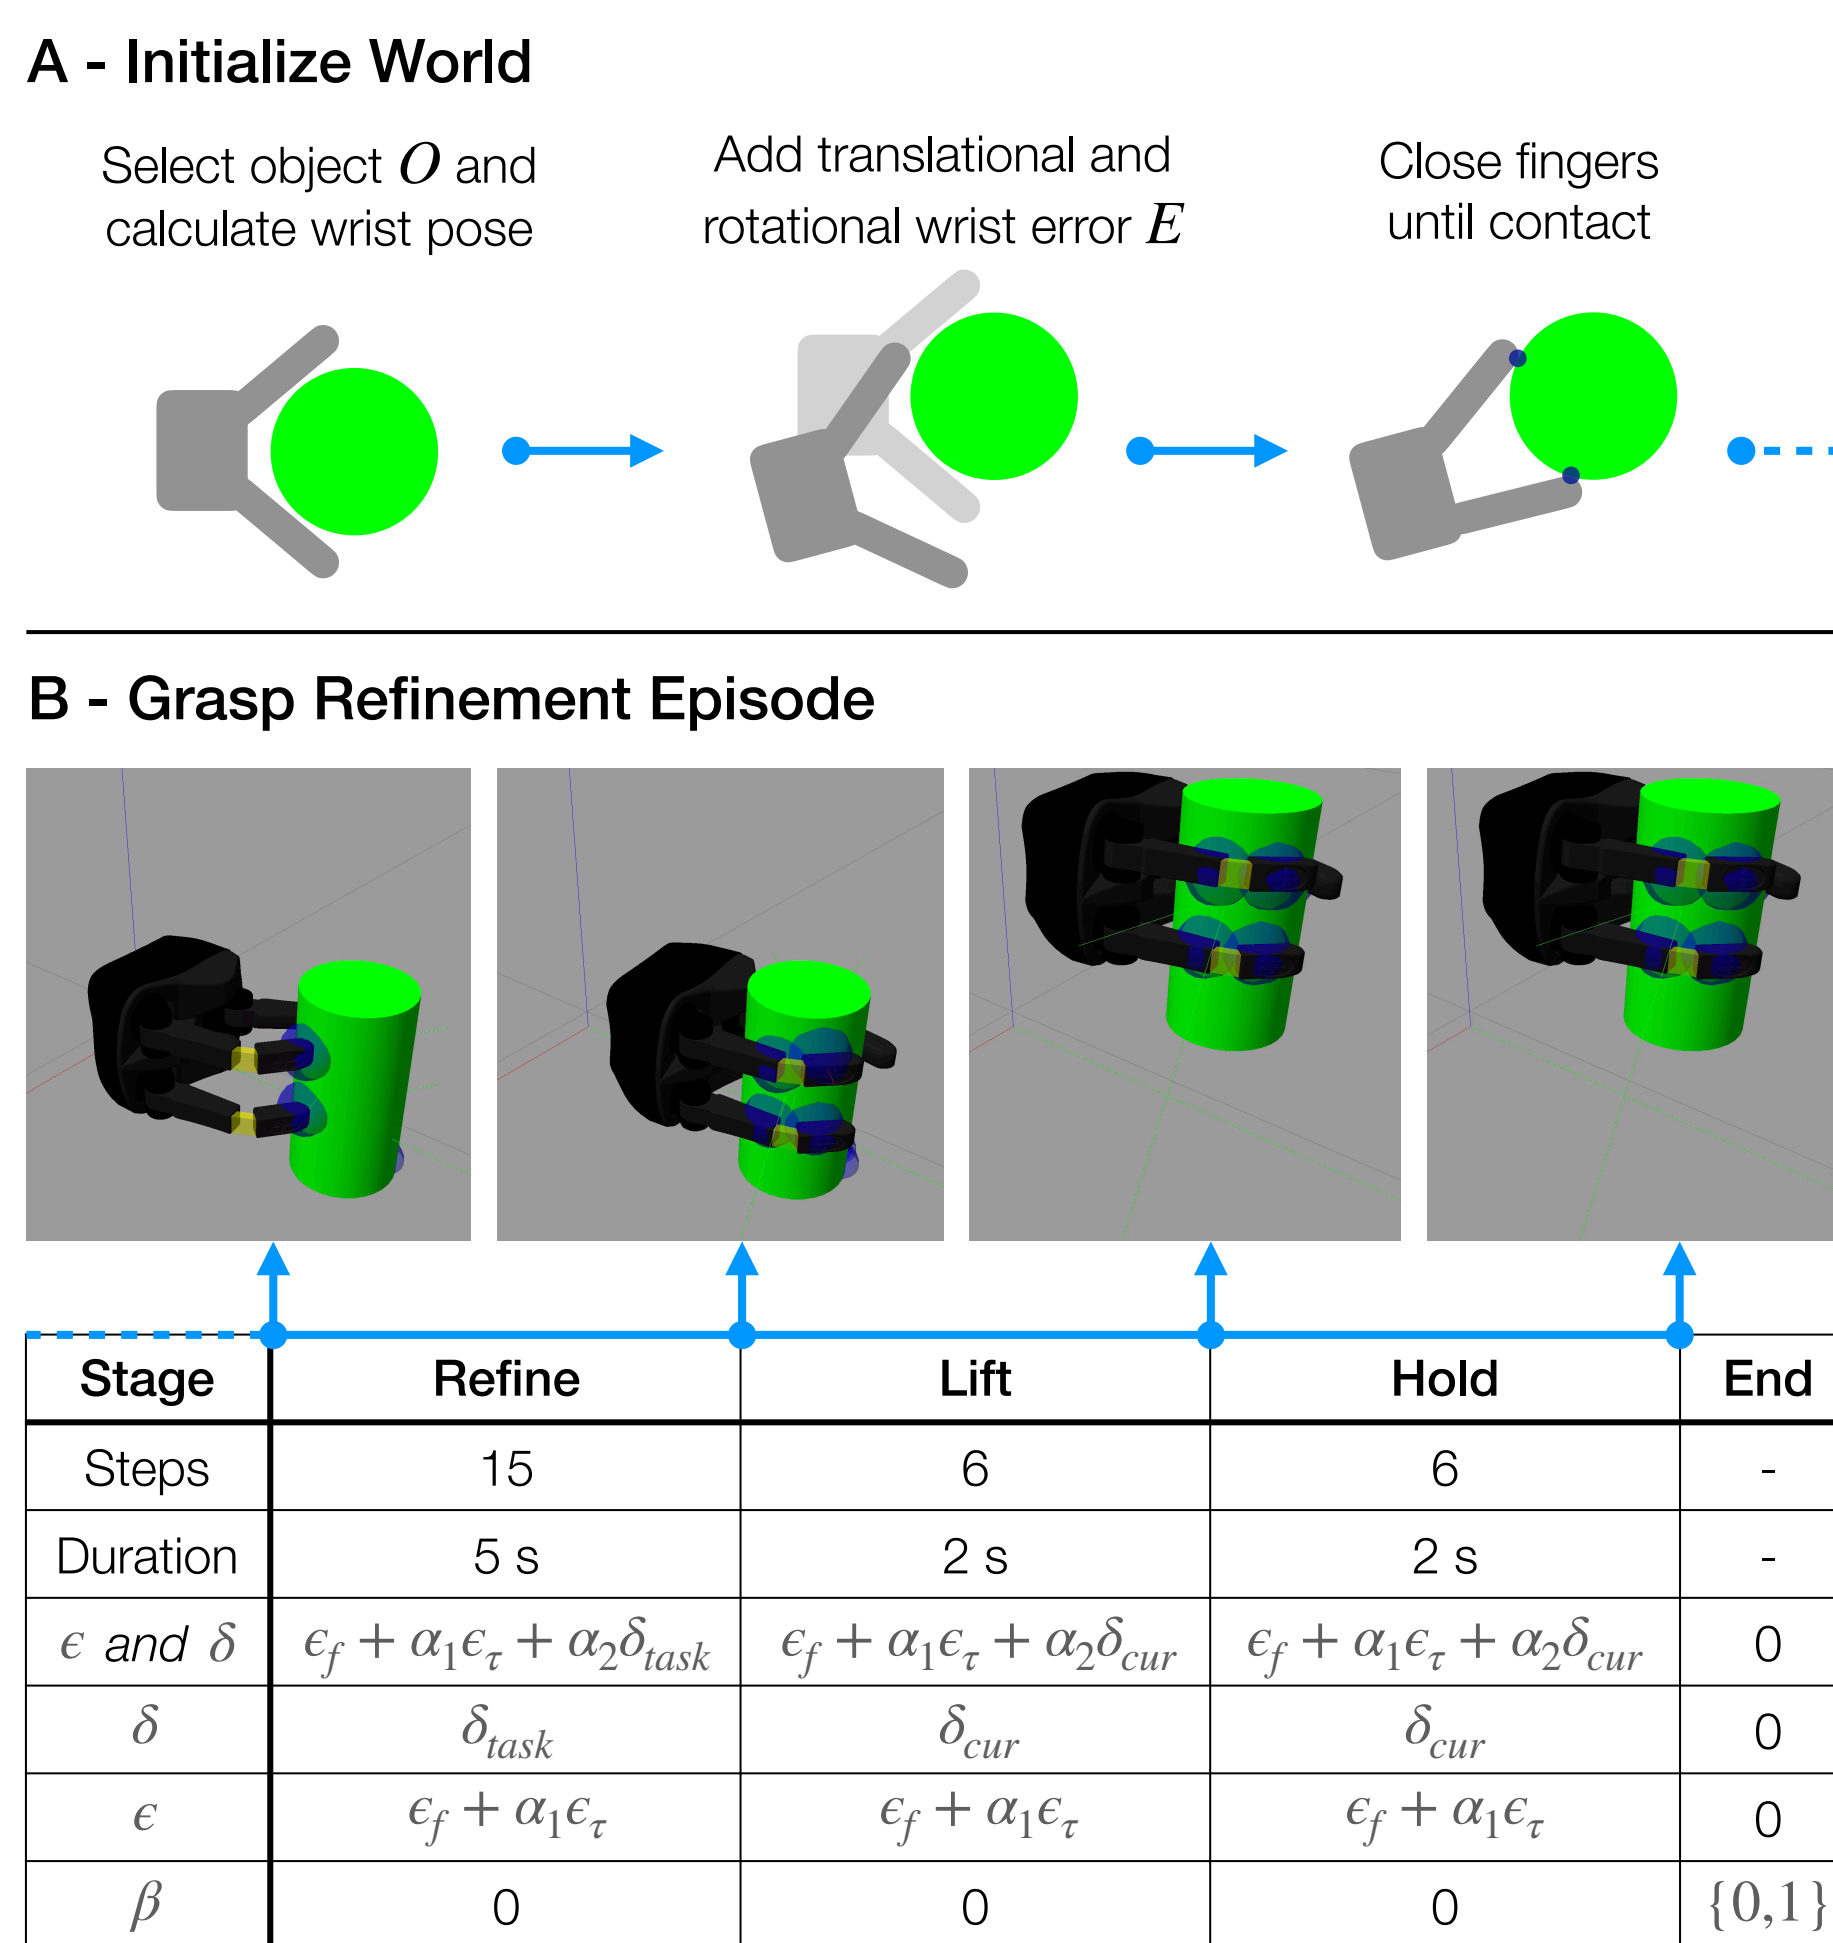


Figure 3: Overview of one algorithm episode. The weighting factors of $\alpha_1 = 5$ and $\alpha_2 = 0.5$ were empirically determined.

We uniformly sample the translational wrist error from $[-5, 5]$ cm and the rotational error from $[-10, 10]$ deg for each axis, respectively. We add the wrist pose error E to an optimal sideways facing grasp pose to simulate calibration errors and close the fingers of the robotic hand in the erroneous wrist pose until the fingers make contact with the object. Consequently, the grasp refinement episode (B) starts. We split the grasp refinement algorithm into four stages (Refine, Lift, Hold, End). As shown in the table, we compare the following reward configurations: (1) both ϵ and δ , (2) only ϵ , (3) only δ and (4) the baseline β . Fig. 3 shows that δ refers to δ_{task} in the *refine* stage to measure expected grasp stability before lifting and δ_{cur} in the *lift* and *hold* stages to measure current stability. Further, ϵ is a weighted combination of ϵ_f and ϵ_τ . While these rewards provide stability feedback after every algorithm step, the baseline β gives a sparse reward after the holding stage, indicating if the object is still in the hand (1) or not (0).

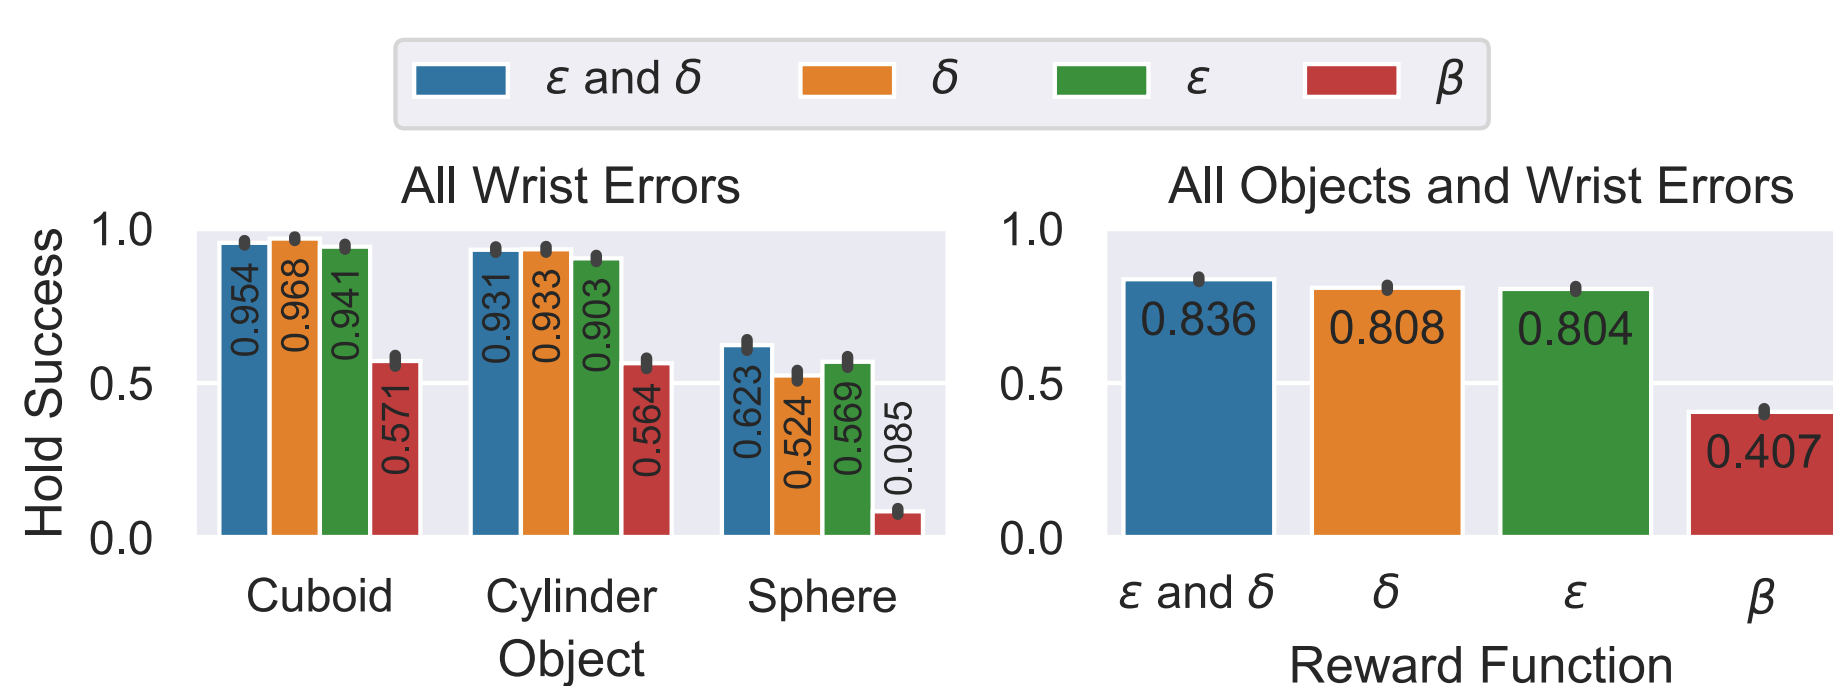


Figure 4: Test results for reward frameworks. We average performance over 40 models trained with different seeds for each framework. The error bars represent ± 2 standard errors.

In Fig. 4, we find that combining the geometric grasp stability metric ϵ with the force-agnostic metric δ yields the highest average success rates. The results demonstrate that information about contact positions and normals encoded in ϵ combines well with the force-based information in the δ reward. This result motivates building physical robotic hands capable of sensing these types of information.

4. TACTILE SENSING AND STATE

Our second experiment gradually decreases tactile resolution in the **state vector** to find realistic training and deployment workflows for grasping algorithms. We compare four contact sensing frameworks. The *full* contact sensing framework receives the same state vector as in the first experiment (full force vector at each link, contact positions, and normals, and joint positions). In the *normal* framework, we only provide the algorithm with the contact normal forces and omit the tangential forces. In the *binary* framework, we only give a binary signal whether a link is in contact (1) or not (0). Finally, we solely provide the joint positions in the *none* framework. The reward function in these experiments is ϵ and δ from Fig. 3. Hence, all contact sensing frameworks receive contact information indirectly via the reward.

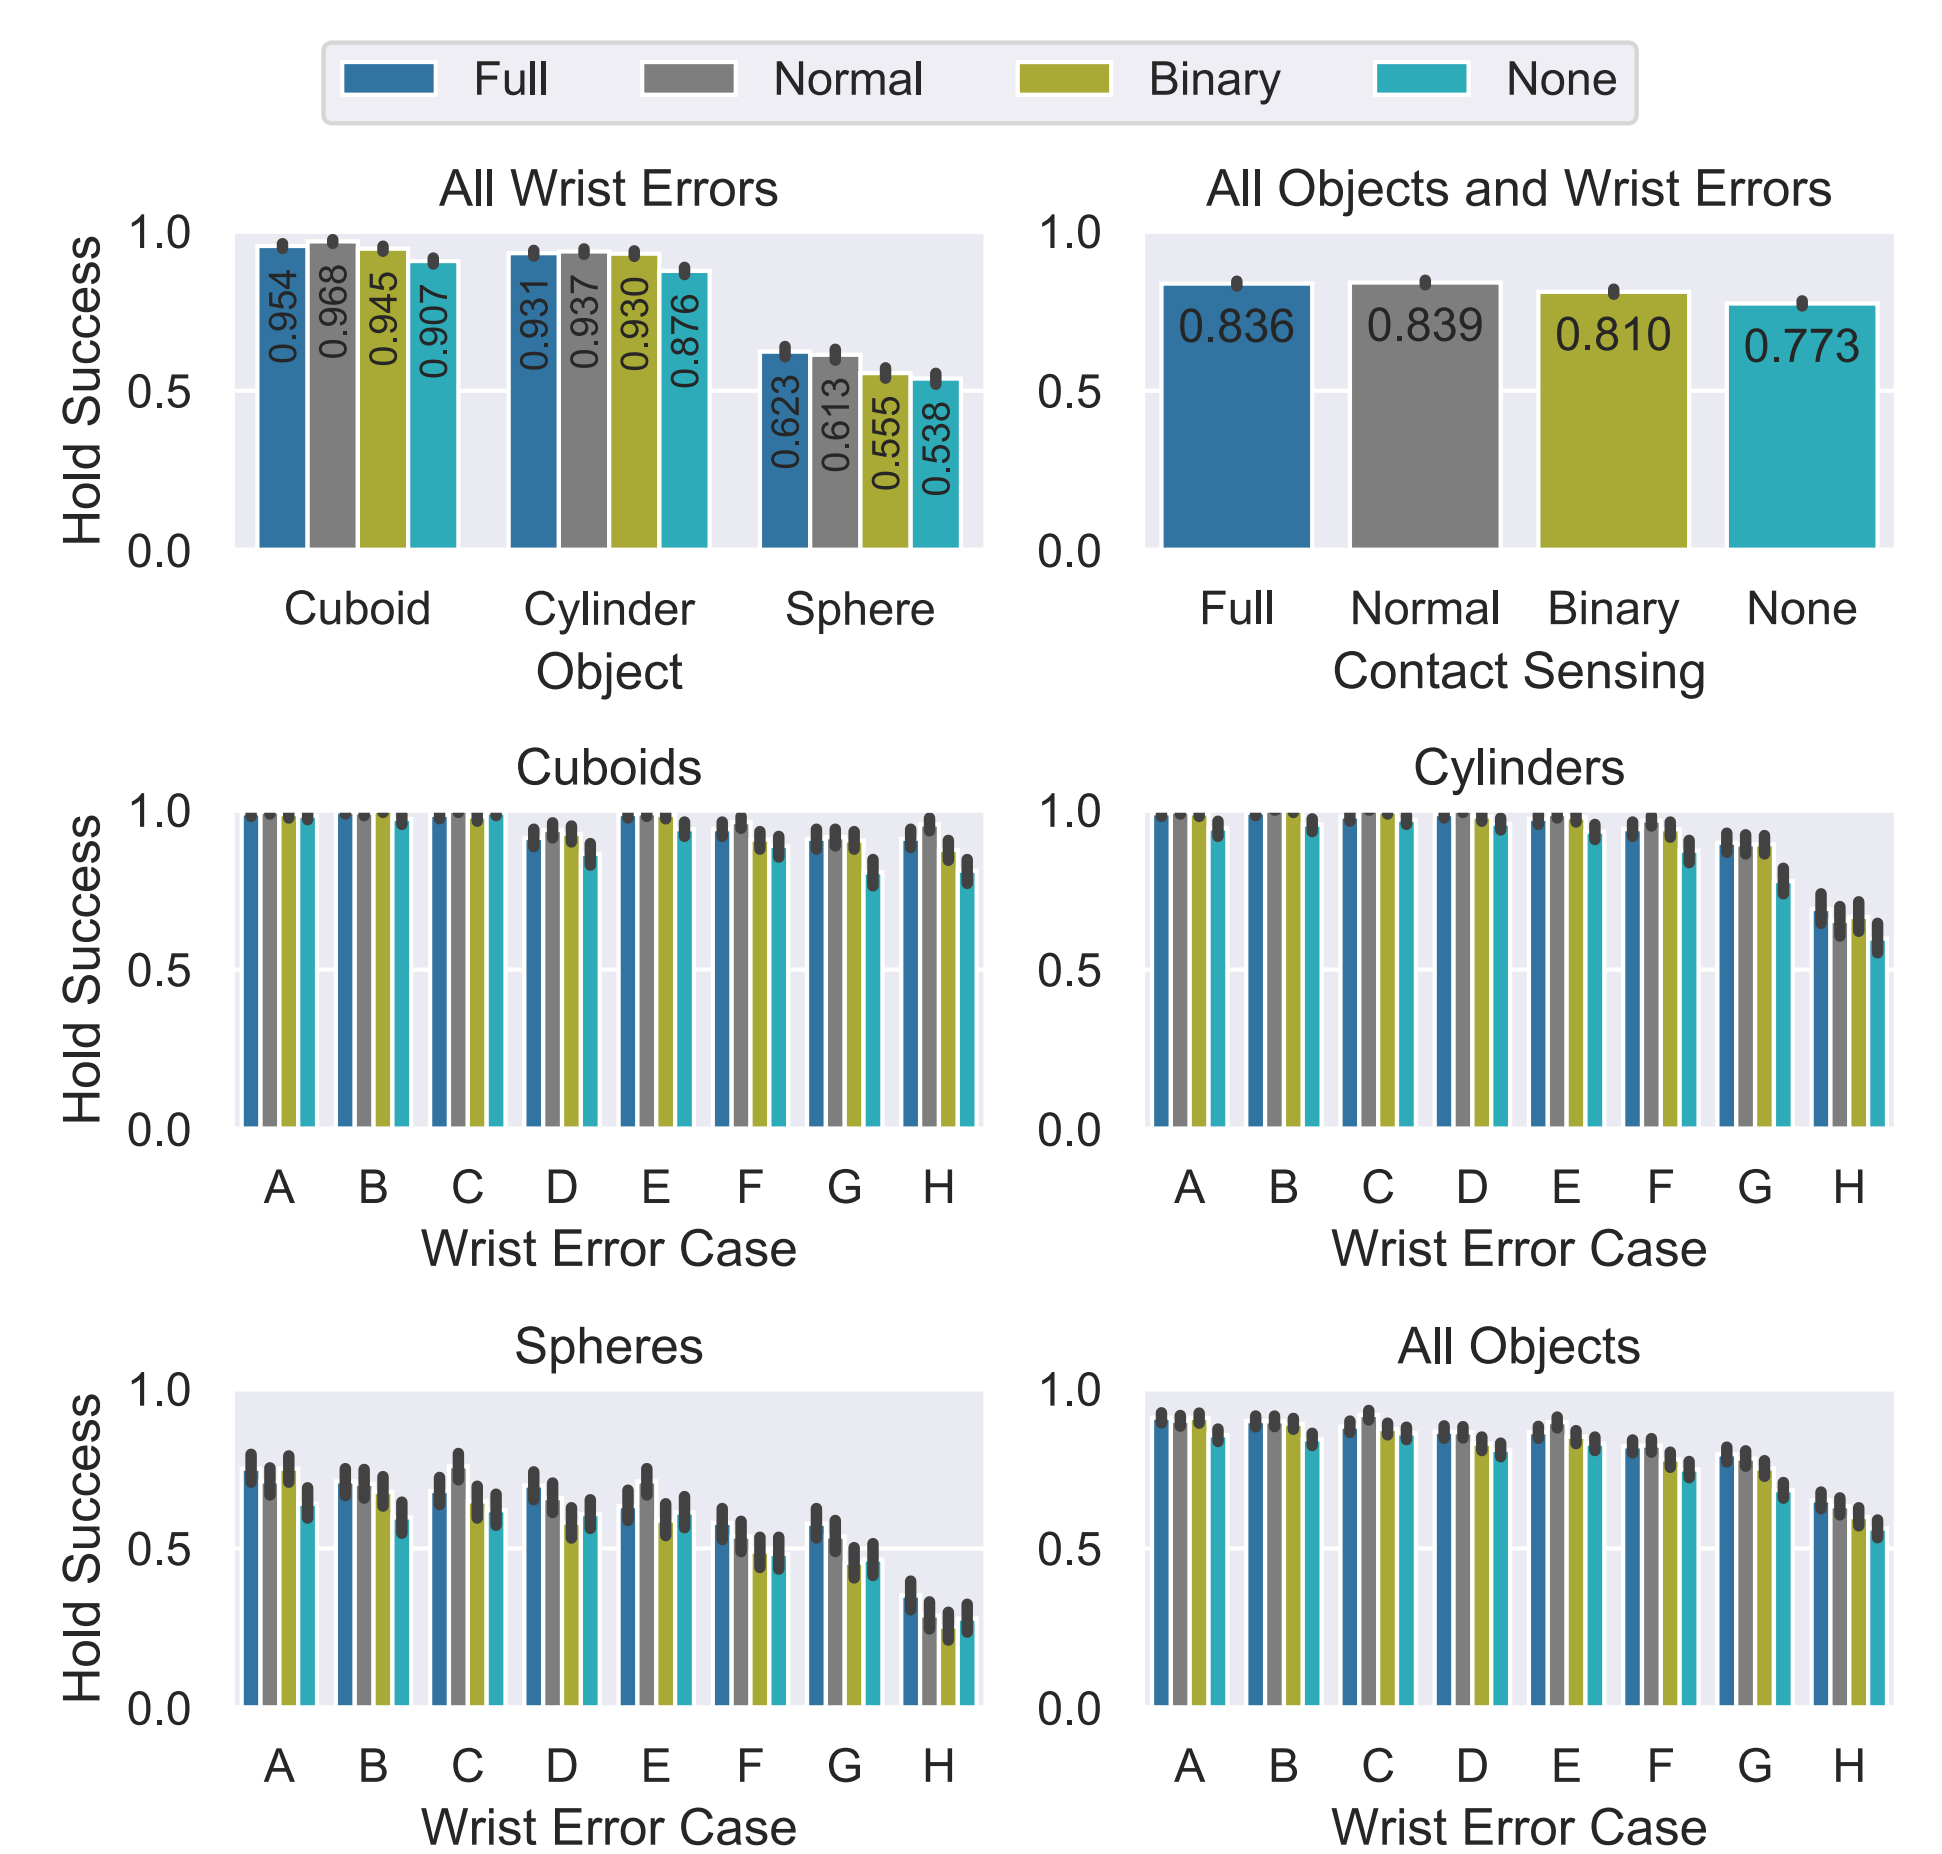


Figure 5: Test results for contact sensing frameworks. Wrist error case A means no translational and rotational wrist pose offset. Case H means 7 cm translational and 14 deg rotational L2 offset. Performance decreases for larger wrist errors.

Fig. 5 shows that the frameworks which receive contact feedback (*full*, *normal*, *binary*) outperform the *none* framework by 6.3%, 6.6% and 3.7%, respectively. The improvements, however, are small and suggest that an affordable *binary* contact sensor suite, or even no contact sensing at all, may be suitable if a small decrease in performance is tolerable. This result supports our hypothesis that RL grasping algorithms are deployable to hands with reduced contact sensor resolution at little performance decrease when incorporating rich tactile feedback at train time.

ACKNOWLEDGEMENTS, REFERENCES

This material is based upon work supported by the US National Science Foundation under Grant No. IIS-1924984 and by the German Academic Exchange Service. An extended paper including the material in this abstract has been submitted for publication.

[1] B. Mirtich and J. Canny, "Easily computable optimum grasps in 2-d and 3-d," in *IEEE ICRA 1994*.