

Use the Force, Luke!

Learning to Predict Physical Forces by Simulating Effects

Kiana Ehsani^{*2}, Shubham Tulsiani¹, Saurabh Gupta³, Ali Farhadi², Abhinav Gupta^{1,4}

¹ FAIR, ² University of Washington, ³ UIUC, ⁴ Carnegie Mellon University

<https://ehsanik.github.io/forcecvpr2020>

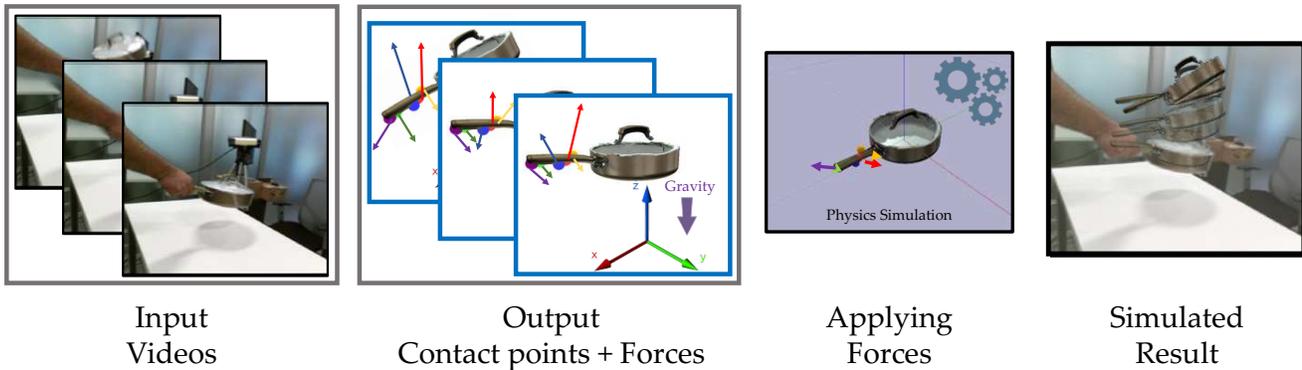


Figure 1: We study the problem of physical understanding of human-object interactions by inferring object motion, points of contacts and forces from RGB videos. We use physics simulation to see how these interactions change the dynamics of the object and try to imitate the motion observed in the video.

Abstract

When we humans look at a video of human-object interaction, we can not only infer what is happening but we can even extract actionable information and imitate those interactions. On the other hand, current recognition or geometric approaches lack the physicality of action representation. In this paper, we take a step towards more physical understanding of actions. We address the problem of inferring contact points and the physical forces from videos of humans interacting with objects. One of the main challenges in tackling this problem is obtaining ground-truth labels for forces. We sidestep this problem by instead using a physics simulator for supervision. Specifically, we use a simulator to predict effects, and enforce that estimated forces must lead to same effect as depicted in the video. Our quantitative and qualitative results show that (a) we can predict meaningful forces from videos whose effects lead to accurate imitation of the motions observed, (b) by jointly optimizing for contact point and force prediction, we can improve the performance on both tasks in comparison to independent training, and (c) we can learn a representation from this model that generalizes to novel objects using few shot examples.

1. Introduction

What does it mean to understand a video of human-object interaction such as the one shown in Figure 1? One popular answer would be to recognize the nouns (objects) and verbs (actions) – e.g., in this case lifting a pot. But such an understanding is quite limited in nature. For example, simply recognizing ‘lifting’ does not tell one anything about how the pot was grasped, or how high it was lifted. To address these shortcomings, there has been a recent push towards a deeper geometric understanding of videos. From estimating contact points on the object [2] to estimating human and object poses [21], these approaches tend to estimate the visible geometric structure. While the high-level semantic labeling (‘lifting’) or the geometric inferences (human and object pose estimation), both provide an answer to *what* happened in the video, it lacks the true physical substance for actionable understanding. For example, just knowing how the pot is moved is not sufficient for the robot to imitate – it needs to also understand *how* the act was accomplished.

In order to obtain a more actionable understanding, we argue that one must account for the physical nature of the task. The answer then, to the question of how the act was done, is rather straightforward from a physical perspective – the object was in contact with the human hand on two sides, and a combination of inward and upward forces applied at these contact points allowed it to be lifted up against gravity. This

* Work done during an internship at FAIR.

understanding of physical *forces* is not only directly useful for an active agent but also completely represents the interaction from the object’s perspective as only external forces cause its motion. In this work, we take a step towards developing such an understanding and present a system that can infer the contact points and forces applied to a known object from an interaction video.

While the goal of being able to infer these forces is desirable, it is unfortunately tedious (if not impossible) to acquire direct supervision for this task. The existing force sensors [22] are not precise enough to provide accurate direction or magnitude measurement, and not compact enough to keep the interaction natural. So, how do we get the supervision? We note that if we can infer physical forces applied to an object, we can also recover a full geometric understanding by simulating the effect of the forces on that object. We build on this insight and present an approach to learn prediction of physical forces, that instead of directly supervising the predicted forces, enforces that their effects match the observations through the interaction video.

To train our system, we collect a dataset of videos recorded from multiple participants grabbing and moving objects. Then we use Mechanical Turk to annotate the keypoints of the objects and contact points on each frame and use this limited information in camera frame to infer the object’s 6DOF pose in world coordinates and the person’s contact points on object mesh. We observe that our approach of learning to predict forces via supervising their effects allows us to learn meaningful estimates and that these can explain the observed interaction in terms of reproducing the observed motion.

Our experiments show that our model learns to infer human contact points on object mesh, and estimate the corresponding forces. We observe that applying these forces on the predicted contact points at each time step in physics simulation we can repeat the behavior depicted in the video. We also show that contact point and force prediction are highly correlated and jointly optimizing improves the performance on both tasks. Finally, we provide interesting evidence that the representation we learn encodes rich geometric and physical understanding that enables us to generalize to interacting with novel objects using only few shot examples.

2. Related Work

Pose estimation. In order to understand physical motions, a network needs to implicitly reason about the object pose. There is a long line of work in this area, with two different approaches: category-based [12, 28, 29] and instance-based [8, 19, 30, 33]. Our work is more aligned with the latter, and we use the YCB object set [3] which provides richly textured objects. While some of our model design decisions is inspired by these works, *e.g.* iterative pose es-

imation [20], our final goal, to infer about physics of the observed motions, is different.

Contact point prediction. Predicting the contact point and hand pose estimation for object manipulation has been studied in the domain of detecting plausible tool grasping [1], human action recognition [10], hand tracking [14], and common grasp pattern recognition [17]. Brahmhatt *et al.* [2] collected a dataset of detailed contact maps using a thermal camera, and introduced a model to predict diverse contact patterns from object shape. While our model also reasons about the contact points, it is only one of the components towards better understanding the physical actions. Moreover, we show that we benefit from force prediction to improve the contact point estimations.

Human Object interaction. Typical approaches for understanding human object interaction use high level semantic labels [11, 13, 34]. Recently there have been some works in understanding the physical aspects of the interaction. Pham *et al.* [26] have used data from force and motion sensors to reason about human object interactions. They also use off-the-shelf tracking, pose estimation, and analytical calculations to infer forces [25]. Hwang *et al.* [18] studied the forces applied to deformable objects and the changes it makes on object shape. Li *et al.* [21] reasoned about human body pose and forces on the joints when the person is interacting with rigid stick-like hand tools. While these methods show encouraging results for this direction, our work concentrates on interaction scenarios with complex object meshes and more diverse contact point patterns.

Predicting physical motions. Recently, learning the physics dynamic has been widely studied by classifying the dynamics of objects in static images [23], applying external forces in synthetic environments [24], predicting post-bounces of a ball [27], simulating billiard games [9], and using generative models to produce plausible human-object interactions [31]. These works are more broadly related to understanding the physical environment, however, their goal to predict how scenes evolve in the future is different from ours. We try to tackle the problem of physically reasoning about the motions observed in the videos.

Recovering physical properties. In recent years there have been efforts in building differentiable physics simulation [6, 16]. Wu *et al.* [32] use physics engines to estimate physical properties of objects from visual inputs. However, in contrast to these approaches aimed at retrieving the properties of the physical world, we assume these are known and examine the problem of interacting with it. Our real-to-sim method is more aligned with the path taken by [4]; that being said, our goal is not bringing the simulation’s and real world’s distributions closer. We rather focus on replicating the observed trajectory in simulation.

3. Approach

Given a video depicting a human interacting with an object, our goal is to infer the physical forces applied over the course of the interaction. This is an extremely challenging task in the most general form. For example, the object geometry may vary wildly and even alter over the interaction (*e.g.* picking a cloth), or forms of contact may be challenging (*e.g.* from elbowing a door to playing a guitar). We therefore restrict the setup to make the task tractable, and assume that the interaction is with a known rigid object (given 3D model), and only involves a single hand (five fingers apply the force). Given such an interaction video, our goal is then to infer the forces applied to it at each time-step along with the corresponding contact points.

Formally, given a sequence of images $\{I_t\}$ depicting an interaction with a known object, and additional annotation for (approximate) initial object pose, we predict the person’s contact points in object coordinates frame $C_t = (c_t^0, \dots, c_t^4)$ (each representing one finger), and the forces applied to each contact point (f_t^0, \dots, f_t^4) . As alluded to earlier, it is not possible to acquire supervision in the form of the ground-truth forces over a set of training interactions. Our key insight is that we can enable learning despite this, by instead acquiring indirect supervisory signal via enforcing that the simulated effect of predicted forces matches the observed motions across the video. We first describe in Section 3.1 a dataset we collect to allow learning using this procedure. We then describe in Section 3.2 how we can extract supervisory signal from this data, and finally present our overall learning procedure in Section 3.3.

3.1. Interaction Dataset

We collect a dataset of object manipulation videos, representing a diverse set of objects, motions, and grasping variations. We leverage the objects from the YCB set [3], as these have the underlying geometry available, and record a set of videos showing participants manipulating 8 objects. To enable learning using these videos, we collect additional annotations in the form of semantic keypoints and pixel locations of contact points to (indirectly) allow recovering the motion of the objects as well as the 3D contact points for each interaction.

We describe the data collection procedure in more detail in the supplementary, but in summary, we obtain: a) annotations for 2D locations for visible keypoints in each frame, b) 6D pose of the object w.r.t. the camera in each frame, though these are noisy due to partial visibility, co-planar keypoints, *etc.*, and c) 3D contact points on the object mesh over each interaction video. There are 174 distinct interaction videos in our dataset (111 train, 31 Test, 32 Validation), 13K frames in total. The 8 objects used are: pitcher, bleach bottle, skillet, drill, hammer, toy airplane, tomato soup can and mustard bottle. We show some examples from

the dataset in Figure 2. We will publicly release the dataset and believe that it will also encourage future research on understanding physical interactions.

3.2. Supervisory Signal via Physical Simulation

Given per-timestep predicted forces f_t and corresponding contact points C_t , we show that these can get supervisory signal by simulating their effects, and comparing the simulated motion against the observed one.

Discrepancy between Simulated and Observed Motions.

A rigid body’s ‘state’ can be succinctly captured by its 6D pose, linear velocity, and angular velocity. Given a current state s_t and the specification of the applied forces, one can compute using a physics simulator \mathcal{P} , the resulting state $s_{t+1} \equiv \mathcal{P}(s_t, f_t, C_t)$. Therefore, given the initial state, and predicted forces and contact points, we can simulate the entire trajectory under these forces and obtain a resulting state at each timestep.

One possible way to measure the discrepancy between this resulting simulated motion and the observed one is to penalize the difference between the corresponding 6D poses. However, our annotated ‘ground-truth’ 6D poses are often not accurate (due to partial visibility *etc.*), and this measure of error is not robust. Instead, we note that we can directly use the annotated 2D keypoint locations to measure the error, by penalizing the re-projection error between the projected keypoints under the simulated pose and the annotated locations of the observed ones. We define a loss function that penalizes this error:

$$\mathcal{L}_{keypoint}(l_t^{kp}, s_t) := \|l_t^{kp} - \pi(R_t, T_t)\|^2 \quad (1)$$

Here, l_t^{kp} is the annotated 2D location for keypoints, π is the projection operator which transforms 3D keypoints on the model to the camera frame under the (simulated) rotation and the translation at $s_t = (R_t, T_t)$.

Differentiable Physics Simulation. To allow learning using the objective above, we require a differentiable physics simulator \mathcal{P} . While typical general-purpose simulators are unfortunately not differentiable, we note that the number of input variables (state s_t , forces f_t , and contact points C_t) in our scenario is low-dimensional. We can therefore use finite difference method to calculate the gradients of the outputs of the simulation with respect to its inputs.

In order to calculate the derivative of output with respect to input, we need to calculate the partial derivatives $\frac{\partial S_{t+1}}{S_t}$, $\frac{\partial S_{t+1}}{f_t}$ and $\frac{\partial S_{t+1}}{C}$. We use the approximation

$$\frac{df}{dx} \cong \frac{f(x+h) - f(x-h)}{2h} \quad (2)$$

where h is a small constant. As $s_t \in \mathbb{R}^{13}$, $f_t \in \mathbb{R}^{k \times 3}$, and $C_t \in \mathbb{R}^{k \times 3}$ ($k = 5$ is the number of contact points), we can

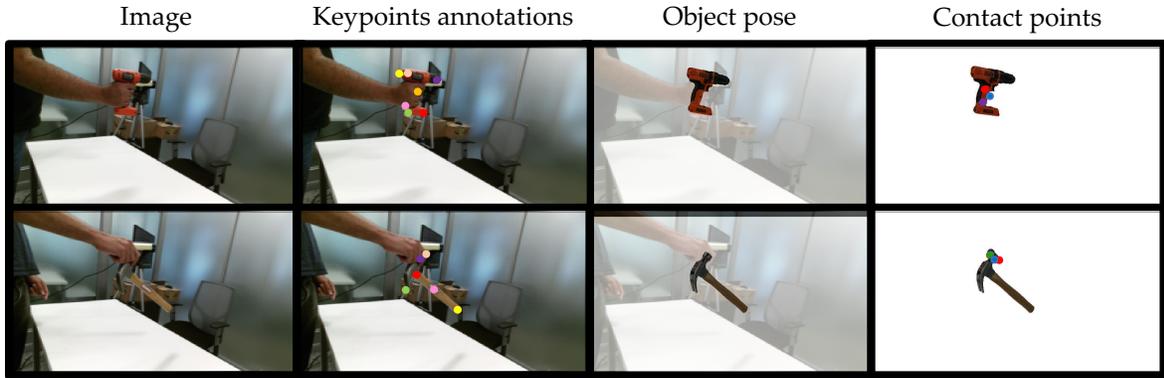


Figure 2: Dataset. Showing two sample frames (drill and hammer) from our dataset. We collect annotation for semantic keypoints of the objects (Column 2) and human contact points. This data helps us to calculate object 6DOF pose (Column 3) and contact points on object mesh (Column 4).

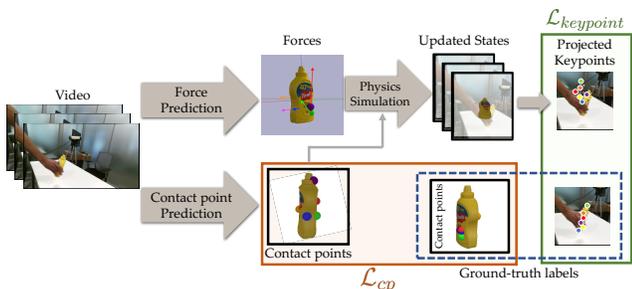


Figure 3: Training schema. Given a video as input, the model predicts the forces and their corresponding contact points. We then apply these forces on the object mesh in physics simulation and jointly optimize for keypoint projection loss $\mathcal{L}_{keypoint}$ and contact point prediction loss \mathcal{L}_{cp} with the aim to imitate the motion observed in the video.

compute the gradients w.r.t the input using $(13+3k+3k)+1$ calls to the simulator \mathcal{P} (the last call is for calculating $f(x)$). We use the PyBullet simulator [5] for our work, and find that each (differentiable) call takes only 0.12 seconds.

3.3. Putting it together: Joint Learning of Forces and Contact Points

Given a video and the initial state of the object, we encode the sequence into a visual embedding, and use this embedding to predict the contact points and their corresponding forces. We then apply these forces in the physics simulation to infer the updated state of the object, and use it in addition to the sequence embedding to iteratively predict the subsequent forces to be applied. This can help the network to adapt to the possible mistakes it might have made, and change the forces in the next steps accordingly (Figure 4). To train our model we have two objectives: (1) to minimize the keypoint re-projection error, that help with reducing the discrepancy between the object trajectory in simulation and the one seen in the video (Equation 1), and (2) to minimize the error in contact point prediction in compare to the

ground-truth (Figure 3). The objective we use for optimizing the contact point estimation is defined as,

$$\mathcal{L}_{cp}(C_t, \hat{C}_t) := \sum_{i=0, \dots, k} \|c_i - \hat{c}_i\|^2, \quad (3)$$

where k is the number of contact points, and C_t and \hat{C}_t are the ground truth and predicted contact points at time t .

We note that the contact point decoder gets supervisory signals both from contact point loss and the keypoint loss. We believe this constrains contact point prediction to generate physically plausible motions as seen in videos. In experiments, we show this joint loss leads to improvement even in contact point prediction.

Training details. The backbone for obtaining primary image features is ResNet18 [15] pre-trained on ImageNet [7]. We take the features (which are of size $512 \times 7 \times 7$) before the average pooling. For all the experiments we use batch size 64, Adam optimizer with learning rate of 0.001, videos of length 10 frames, with frequency of 30 fps, $h_f, h_s = 0.01$ and $h_c = 0.05$ for approximating gradients (refer to Equation 2) for force, state and contact point respectively.

We calculate the direction of the gravity in world coordinates and use it in physics simulation to ensure realistic behavior. To train our model we first train each branch in isolation using \mathcal{L}_{cp} for contact point prediction and $\mathcal{L}_{keypoint}$ for force prediction modules, then we jointly optimize for both objectives and train end to end.

4. Experiments

The area of physical understanding of human-object interaction is largely unexplored, and there are no established benchmark datasets or evaluation metrics. In this work, we use our novel dataset to provide empirical evaluations, and in particular to show that: a) our results are qualitatively meaningful; and (b) individual components and loss terms are quantitatively meaningful (ablations). We will

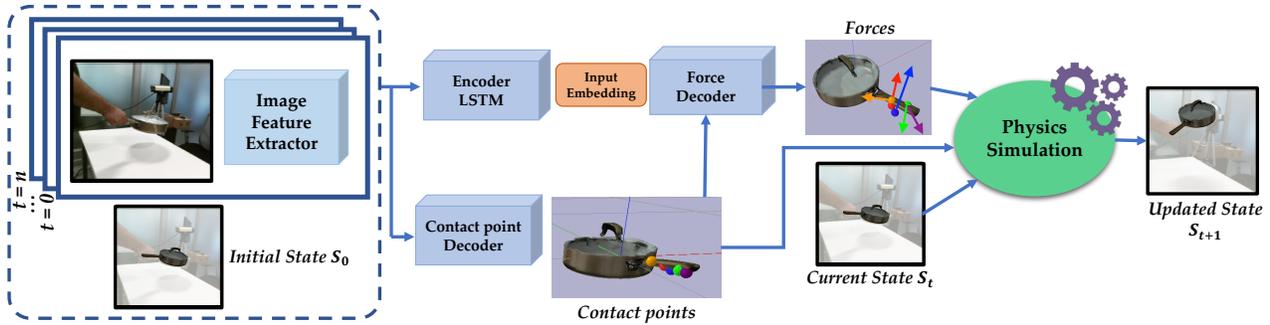


Figure 4: Model overview. Given a video of a person moving an object, along with the initial pose of the object, our network predicts the human contact points and the forces applied at those for each time step. The implied effects of these forces can then be recovered by applying them in a physics simulation. Using the gradients through this simulated interaction, our model learns how to optimize for its two objectives: minimizing the error in projection of the object to camera frame, and predicting the accurate contact points.

also demonstrate that a physical understanding of human-object interactions leads to improvement even in individual components such as contact point estimation and estimating object poses; and (c) finally we will demonstrate the generalization power by showing that our network learns a rich representation that can use few examples to generalize to manipulating novel objects.

Evaluation Metric: The goal of our work is to obtain physical understanding of human-object interactions. But getting ground truth forces is hard, which makes it impossible to quantitatively measure the performance by just force values. So instead of measuring forces, we evaluate if our predicted forces lead to similar motions as depicted in the videos. Therefore, for evaluating our performance, we use the CP and KP error (Equations 1 and 3) as evaluation metrics. The former measures the $L1$ distance between the predictions and ground truth for contact point (in object coordinates), and the latter measures error in keypoint projection (in image frame). Original image size is 1920×1080 and the keypoint projection error is reported in pixels in the original image dimension. Rotation and translation error are the angular difference (quaternion distance) in rotation and $L2$ distance (meters) in translation respectively.

4.1. Qualitative Evaluation

We first show qualitative results of our full model: joint contact point and force prediction. The qualitative results for force prediction and contact point predictions are shown in Figure 5 and Figure 6 respectively. As the figure shows, our forces are quite meaningful. For example, in case of plane at $t = 0$ (See Figure 5(a)), initially the yellow and blue arrow is pushing to left and purple arrow is on the other side is pushing to right. This creates a twist motion and as seen at $t = 4$ rotates the plane. Also, in case of skillet (Figure 5(b)), there is a big change in the orientation of the object from $t = 0$ to $t = 4$, therefore a bigger magnitude of force is required. However, after getting the initial

momentum the forces decrease to the minimum needed for maintaining the current state. Refer to project page* for more visualizations.

Next we show few examples of contact point prediction. If contact points are predicted in isolation, small differences can lead to fundamentally different grasps. By enforcing physical meaning to these grasps (via forces), our approach ensures more meaningful predictions. An example is the top of Figure 5, where isolated prediction leads to grasp on the rim of the pitcher; but joint prediction leads to prediction of grasp on the handle.

4.2. Quantitative Evaluation

To quantitatively evaluate our approach, we measure performance by evaluating if application of forces as predicted lead to effects as depicted in the videos. Table 1 shows the performance in terms of CP metric and KP error. We report both the joint optimization (depicted as $\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$) and the case where contact point prediction and force prediction modules are learned in an isolated manner ($\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$). We observe that the joint model is significantly better (low error), both in contact point prediction and keypoint prediction.

Next, we measure what happens if the contact point prediction was perfectly matched with the annotated ground truth. So, instead of predicting them, we use ground truth contact points. Given a video, initial pose, and human contact points, we predict the forces applied to each point of contact to replicate the motion. Table 2 shows the results for unseen videos. Note that the training and test sets of contact points are disjoint, so model needs to be able to generalize to applying forces to the object under novel contact point configurations.

Comparing results in Table 1 and Table 2 shows that the jointly optimized model ($\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$), (even though it is using the predicted contact points which may have er-

* <https://ehsanik.github.io/forcecvpr2020>

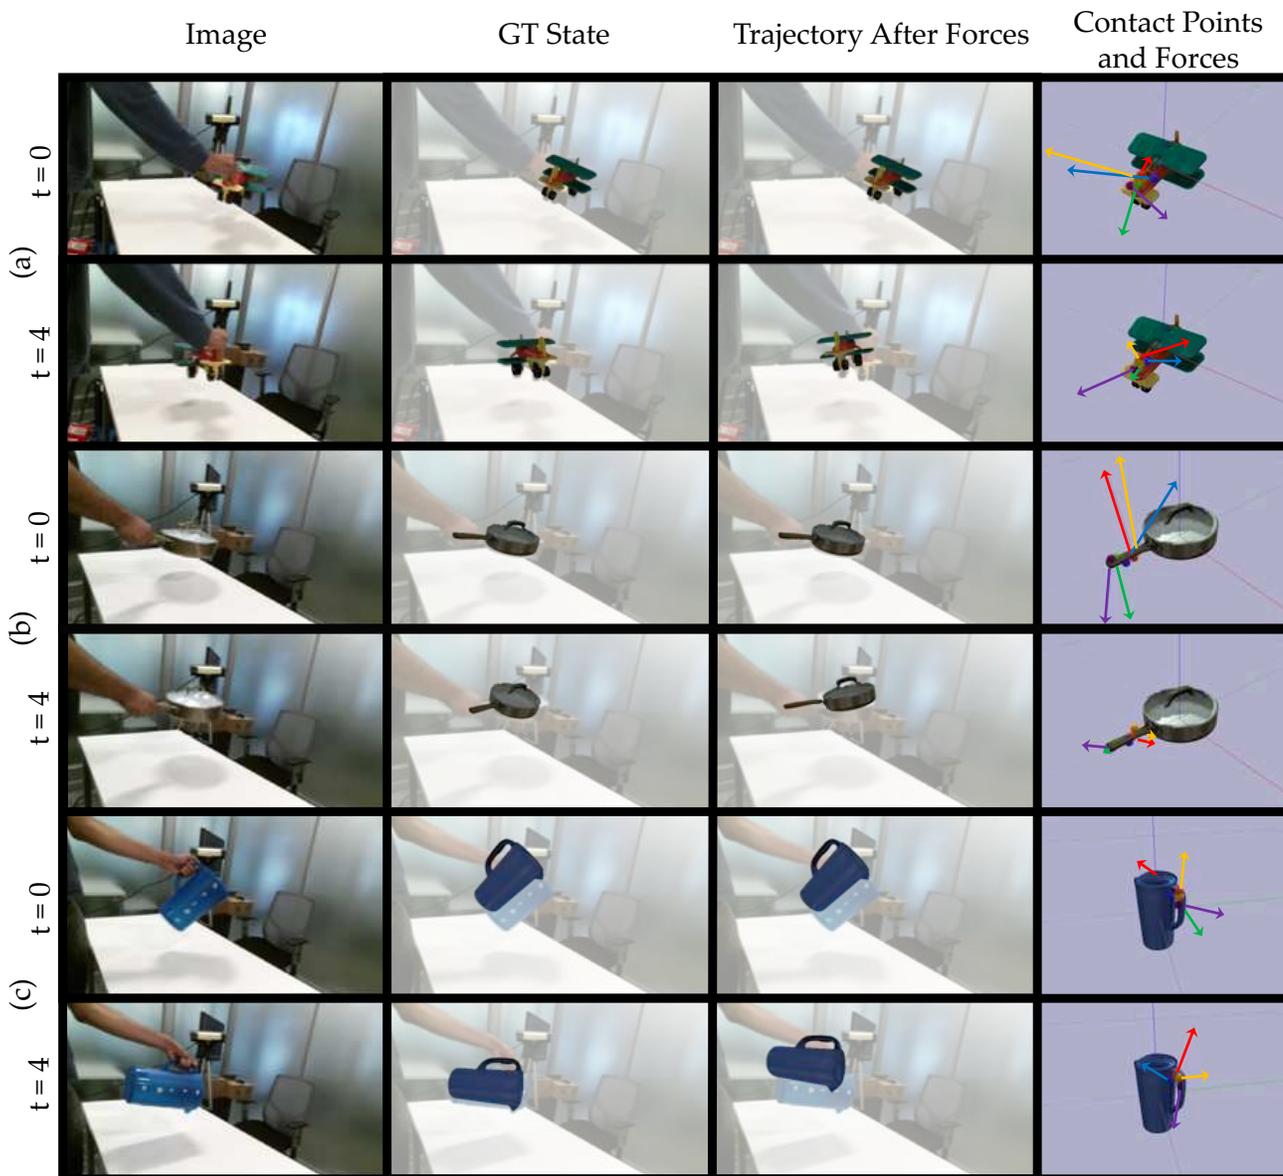


Figure 5: Qualitative Results. We show the results for the model which optimizes for both $\mathcal{L}_{keypoint}$ and \mathcal{L}_{cp} . Due to space limitations only Frames for $t = 0, 4$ are shown. For more videos and contact point visualizations refer to supplementary material.

rors), predicts forces that result in better KP metric in compare to the model that uses ground-truth contact points (and is only trained on $\mathcal{L}_{keypoint}$). This shows that jointly reasoning about where to apply the force and what force to apply can help with a better physical understanding.

Training one model for different objects. We try two training setup for each experiment: (1) training a separate model per object, and (2) sharing weights for one model across all objects (referred to as "All objects"). The common trend of performance improvement for both metrics after joint optimization is observed when training one model for all objects as well.

4.3. Few-shot Generalization to Novel Objects

In order to evaluate the representation we learn from this training schema, we train a model for manipulating plane and use that to predict forces on unseen objects using few shot examples. Figure 7 shows that the representation we learned from predicting forces on one object can generalize to estimating forces for held-out objects using only few examples. Increasing the number of training samples yields better results on the test set. The 10 few shot experiment without pre-training (red) has a significantly lower accuracy than the one with our pre-training (light green) and is comparable with our 1-shot experiment (dark green).

Object	Input	Objective	CP Error	Keypoint Error (in px)	Rotation Error	Translation Error
Plane	Image	$\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$	8.79e-2	109.73	0.220	0.152
	Image	$\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$	7.97e-2	99.89	0.152	0.182
Skillet	Image	$\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$	2.37e-2	70.02	0.085	0.094
	Image	$\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$	2.18e-2	65.25	0.076	0.063
Pitcher	Image	$\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$	7.3e-2	131.62	0.126	0.212
	Image	$\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$	5.61e-2	113.40	0.131	0.129
Drill	Image	$\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$	7.08e-2	85.86	0.192	0.312
	Image	$\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$	6.47e-2	71.00	0.163	0.250
All objects	Image	$\mathcal{L}_{keypoint}/\mathcal{L}_{cp}$	6.83e-2	104.57	0.218	0.250
	Image	$\mathcal{L}_{keypoint} + \mathcal{L}_{cp}$	6.71e-2	99.38	0.183	0.212

Table 1: Contact point prediction and key point projection error on test set, independent vs. joint optimization In each set, the first row shows the results of the model optimizing for contact point and keypoint projection error separately, and the second row represents the joint optimization results. End to end training improves the results for both contact point prediction and keypoint projection.

Object	Input	Objective	Keypoint Error (in px)	Rotation Error	Translation Error
Plane	Image + CP_{gt}	$\mathcal{L}_{keypoint}$	104.33	0.134	0.199
Skillet	Image + CP_{gt}	$\mathcal{L}_{keypoint}$	67.47	0.075	0.089
Pitcher	Image + CP_{gt}	$\mathcal{L}_{keypoint}$	124.70	0.130	0.176
Drill	Image + CP_{gt}	$\mathcal{L}_{keypoint}$	91.25	0.236	0.352
All objects	Image + CP_{gt}	$\mathcal{L}_{keypoint}$	102.40	0.237	0.273

Table 2: Predicting the forces using ground truth contact points We first train a model to predict the forces on the ground-truth contact points. The first four rows show the quantitative results for training separate model per object using the $\mathcal{L}_{keypoint}$ objective and the last row shows the result for training one shared model for all 8 objects.

4.4. Additional Ablation Analysis

Regressing the force without simulation gradients. One alternative approach to solving the force inference problem is to predict the forces without the gradients from simulation. However, this requires having ground-truth labels for forces, which is almost impossible to obtain. So, we try to optimize for a set of pseudo-ground truth forces.

The goal of this experiment is to investigate how keeping the physics simulation in the training loop can help with understanding physics of the environment and generalizing to unseen trajectories.

To obtain pseudo-labels, we optimize a set of valid forces per training example which minimizes the error in keypoint projection. Then we train a model that given the sequence of images and ground truth contact points regresses these forces. The objective is defined as follows:

$$\mathcal{L}_{force}(f_{gt}, \hat{f}) = \|f_{gt} - \hat{f}\|^2$$

Table 3 shows that even though the error for training is relatively low, it still fails to generalize to unseen trajectories and gets a high test error. The intuition is that the gradients the network gets from interaction helps with learning a

Set	Objective	KP (px)	Rotation	Translation
Train	$\mathcal{L}_{keypoint}$	55.96	0.129	0.063
Test	$\mathcal{L}_{keypoint}$	105.60	0.138	0.198
Train	\mathcal{L}_{force}	59.86	0.120	0.065
Test	\mathcal{L}_{force}	185.33	0.156	0.307

Table 3: Pseudo-ground truth force regression. We trained a model to predict the pseudo-ground truth forces for toy airplane from video, initial pose and contact points as input, by directly optimizing for the force prediction (\mathcal{L}_{force}). Although the training error is similar to the model trained on keypoint projection loss, it fails to generalize to unseen images and trajectories.

more generalizable representation.

Predicting initial state. We want to evaluate the necessity of giving the initial state as the input to the network. Thus, we try to predict the initial pose instead of using the ground truth. We do so by training a model that given a video as input and contact points in object coordinates (independent of object state), predicts the initial state of the object as well as the forces that are applied on each contact point (Table 4).

Adding noise to initial state during inference. We also



Figure 6: Improvements in contact point prediction after joint optimization. We qualitatively show some examples for which the model makes better contact point predictions when it is trained using both $\mathcal{L}_{keypoint}$ and \mathcal{L}_{cp} . Fingers are color-coded (Thumb: orange, Index: Red, Middle: Blue, Ring: Green, Pinky: Purple).

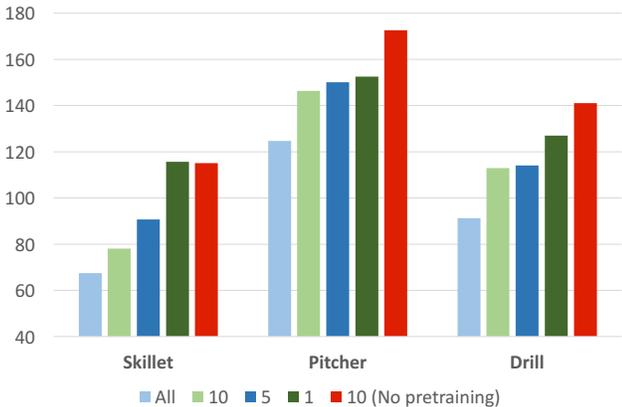


Figure 7: Few shot experiment. Showing the keypoint projection error for the model that is trained on planes and fine-tuned on few examples for held out objects (Lower is better). The results are shown for 1, 5, or 10 examples, as well as the model trained on the whole training set. The red bar shows errors for training on 10 examples without plane pre-training.

want to investigate the effects of adding noise to initial state on performance. This experiment evaluates the robustness of the performance of the model to the initial state estimation. We used the trained model on toy airplane from Table 2, and added noise to the input initial state during inference time. Figure 8 shows the changes in KP metric with respect to tweaks in the rotation and translation of the initial state.

Input	KP (px)	Rotation	Translation
Image	135.80	0.160	0.496
Image+ S_0^{gt}	128.55	0.174	0.329

Table 4: Predicting initial state. The result for the model trained on toy plane object set, predicting the initial pose as well as the forces on ground truth contact points. To see how the error in object initial pose estimation is affecting the performance, we input the ground-truth object pose to the model as input during inference to compare. This model is trained with $\mathcal{L}_{keypoint}$ objective.

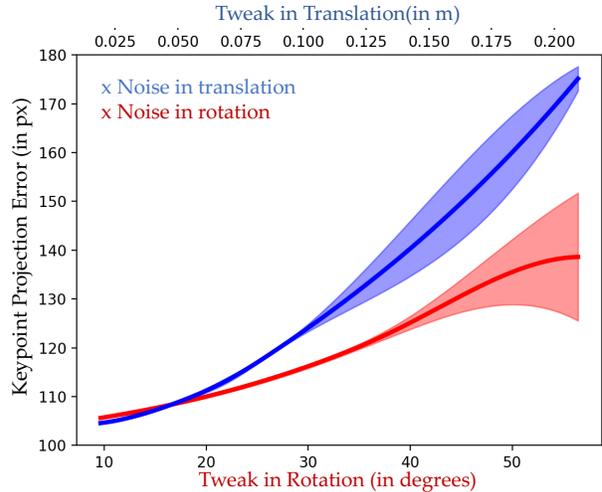


Figure 8: Robustness to initial pose. The changes in the keypoint projection error with respect to the magnitude of the noise added to the initial state. We also calculated the error bars for 5 runs with different random seeds. The labels on the bottom of the chart show the magnitude of the noise in rotation in degrees and the one on the top shows the magnitude of noise in translation in meters.

5. Conclusion

We have proposed a model that given a video depicting human-object interaction, predicts the contact points and the corresponding forces such that it replicates the observed motion. We demonstrate that jointly optimizing for both contact point prediction and keypoint projection error improves the results on both tasks in comparison to training models in isolation. We also show that our model learns a meaningful physical representation that can generalize to novel objects using few examples. Since our approach needs textured non-symmetric objects we were able to show these results on 8 objects from the YCB set [3], but we conjecture that if the keypoint labels in camera frame can be estimated, this method can generalize further than this object set. We believe our work takes step towards integrating action and perception in one common framework – bringing it closer to real-world robotics. A feasible future work would be to investigate how our model’s prediction can speed up the robotic imitation learning procedure.

Acknowledgments. We would like to thank Dhiraj Gandhi for his help with the hardware setup.

References

- [1] Shuichi Akizuki and Yoshimitsu Aoki. Tactile logging for understanding plausible tool use based on human demonstration. In *British Machine Vision Conference (BMVC)*, 2018. 2
- [2] Samarth Brahmabhatt, Cusuh Ham, Charles C Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *Conference on Computer Vision and Pattern Recognition*, pages 8709–8719, 2019. 1, 2
- [3] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *International Conference on Advanced Robotics (ICAR)*, pages 510–517. IEEE, 2015. 2, 3, 8
- [4] Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *International Conference on Robotics and Automation (ICRA)*, pages 8973–8979. IEEE, 2019. 2
- [5] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019. 4
- [6] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end differentiable physics for learning and control. In *Advances in Neural Information Processing Systems*, pages 7178–7189, 2018. 2
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition*, pages 248–255. IEEE, 2009. 4
- [8] Andreas Doumanoglou, Rigas Kouskouridas, Sotiris Malasiotis, and Tae-Kyun Kim. Recovering 6d object pose and predicting next-best-view in the crowd. In *Conference on Computer Vision and Pattern Recognition*, pages 3583–3592, 2016. 2
- [9] Katerina Fragkiadaki, Pulkit Agrawal, Sergey Levine, and Jitendra Malik. Predictive visual models of physics for playing billiards. In *International Conference on Learning Representations*, 2016. 2
- [10] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *Conference on Computer Vision and Pattern Recognition*, pages 409–419, 2018. 2
- [11] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. In *Conference on Computer Vision and Pattern Recognition*, pages 8359–8367, 2018. 2
- [12] Saurabh Gupta, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Aligning 3d models to rgb-d images of cluttered scenes. In *Conference on Computer Vision and Pattern Recognition*, pages 4731–4740, 2015. 2
- [13] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015. 2
- [14] Henning Hamer, Juergen Gall, Thibaut Weise, and Luc Van Gool. An object-dependent hand pose prior from sparse training data. In *Computer Society Conference on Computer Vision and Pattern Recognition*, pages 671–678. IEEE, 2010. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778. IEEE, 2016. 4
- [16] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *International Conference on Robotics and Automation (ICRA)*, pages 6265–6271. IEEE, 2019. 2
- [17] De-An Huang, Minghuang Ma, Wei-Chiu Ma, and Kris M Kitani. How do we use our hands? discovering a diverse set of common grasps. In *Conference on Computer Vision and Pattern Recognition*, pages 666–675, 2015. 2
- [18] Wonjun Hwang and Soo-Chul Lim. Inferring interaction force from visual information without using physical force sensors. *Sensors*, 17(11):2455, 2017. 2
- [19] Wadim Kehl, Fausto Milletari, Federico Tombari, Slobodan Ilic, and Nassir Navab. Deep learning of local rgb-d patches for 3d object detection and 6d pose estimation. In *European Conference on Computer Vision*, pages 205–220. Springer, 2016. 2
- [20] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *European Conference on Computer Vision (ECCV)*, pages 683–698, 2018. 2
- [21] Zongmian Li, Jiri Sedlar, Justin Carpentier, Ivan Laptev, Nicolas Mansard, and Josef Sivic. Estimating 3d motion and forces of person-object interactions from monocular video. In *Conference on Computer Vision and Pattern Recognition*, pages 8640–8649, 2019. 1, 2
- [22] Tao Liu, Yoshio Inoue, and Kyoko Shibata. Design of low-cost tactile force sensor for 3d force scan. In *Sensors*, pages 1513–1516. IEEE, 2008. 2
- [23] Roozbeh Mottaghi, Hessam Bagherinezhad, Mohammad Rastegari, and Ali Farhadi. Newtonian scene understanding: Unfolding the dynamics of objects in static images. In *Conference on Computer Vision and Pattern Recognition*, pages 3521–3529, 2016. 2
- [24] Roozbeh Mottaghi, Mohammad Rastegari, Abhinav Gupta, and Ali Farhadi. “what happens if...” learning to predict the effect of forces in images. In *European Conference on Computer Vision*, pages 269–285. Springer, 2016. 2
- [25] Tu-Hoa Pham, Abderrahmane Kheddar, Ammar Qammaz, and Antonis A Argyros. Towards force sensing from vision: Observing hand-object interactions to infer manipulation forces. In *Conference on Computer Vision and Pattern Recognition*, pages 2810–2819, 2015. 2
- [26] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A Argyros, and Abderrahmane Kheddar. Hand-object contact force estimation from markerless visual tracking. *Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2883–2896, 2017. 2
- [27] Senthil Purushwalkam, Abhinav Gupta, Danny Kaufman, and Bryan Russell. Bounce and learn: Modeling scene dy-

- namics with real-world bounces. In *International Conference on Learning Representations*, 2019. 2
- [28] Caner Sahin and Tae-Kyun Kim. Category-level 6d object pose recovery in depth images. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [29] Shuran Song and Jianxiong Xiao. Sliding shapes for 3d object detection in depth images. In *European conference on computer vision*, pages 634–651. Springer, 2014. 2
- [30] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning*, pages 306–316, 2018. 2
- [31] He Wang, Sören Pirk, Ersin Yumer, Vladimir G Kim, Ozan Sener, Srinath Sridhar, and Leonidas J Guibas. Learning a generative model for multi-step human-object interactions from videos. In *Computer Graphics Forum*, volume 38, pages 367–378. Wiley Online Library, 2019. 2
- [32] Jiajun Wu, Ilker Yildirim, Joseph J Lim, Bill Freeman, and Josh Tenenbaum. Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. In *Advances in Neural Information Processing Systems*, pages 127–135, 2015. 2
- [33] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems (RSS)*, 2018. 2
- [34] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. Situation recognition: Visual semantic role labeling for image understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5534–5542, 2016. 2