

Hand-Object Contact Prediction via Motion-Based Pseudo-Labeling and Guided Progressive Label Correction

Takuma Yagi
tyagi@iis.u-tokyo.ac.jp

Md. Tasnimul Hasan
tasnim@iis.u-tokyo.ac.jp

Yoichi Sato
ysato@iis.u-tokyo.ac.jp

The University of Tokyo
Tokyo, Japan

Abstract

Every hand-object interaction begins with contact. Despite predicting the contact state between hands and objects is useful in understanding hand-object interactions, prior methods on hand-object analysis have assumed that the interacting hands and objects are known, and were not studied in detail. In this study, we introduce a video-based method for predicting contact between a hand and an object. Specifically, given a video and a pair of hand and object tracks, we predict a binary contact state (contact or no-contact) for each frame. However, annotating a large number of hand-object tracks and contact labels is costly. To overcome the difficulty, we propose a semi-supervised framework consisting of (i) automatic collection of training data with motion-based pseudo-labels and (ii) guided progressive label correction (gPLC), which corrects noisy pseudo-labels with a small amount of trusted data. We validated our framework’s effectiveness on a newly built benchmark dataset for hand-object contact prediction and showed superior performance against existing baseline methods. Code and data are available at https://github.com/takumayagi/hand_object_contact_prediction.

1 Introduction

Recognizing how hands interact with objects is crucial to understand how we interact with the world. Hand-object interaction analysis contributes to several fields such as action prediction [10], rehabilitation [28], robotics [68], and virtual reality [17].

Every hand-object interaction begins with contact. In determining which hand-object pairs are interacting, it is important to infer when hands and objects are in contact. However, despite its importance, finding the beginning and the end of hand-object interaction has not received much attention. For instance, prior works on action recognition (*e.g.*, [12]) attempt to recognize different types of hand object interactions at the video clip level, *i.e.*, recognizing one action for each video clip given as input. Some other works on action localization (*e.g.*, [24]) can be used for detecting hand object interactions but localized action segments




Figure 1: **Overlap on Image = Contact?!**: Right hand (masked red region) grabs the onion (middle). Surrounding objects (cutting board, knife) overlaps with hand but not in contact. While it is difficult to determine contact state from single image, we can ease the problem by looking at temporal context (left and right).

are not necessarily related to the beginning and the end of contact between hands and objects. Contact between a hand and an object has been studied in the context of 3D reconstruction of hand object interaction [8, 18]. However, they assumed that hands and objects are already interacting with each other. Only the moment when hands and objects are interacting in a stable grasp was targeted for analysis.

In this work, we tackle the task of predicting contact between a hand and an object from visual input. Predicting contact between a hand and an object from visual input is not trivial. For example, even if the hand area and the bounding box of an object overlap, it does not necessarily mean that the hand and the object are in contact (see Figure 1). In determining whether a hand and an object are in contact, it is essential to consider the spatiotemporal relationship between them. While some methods claim that the hand contact state can be classified by looking at hand shape [62, 40], they did not explicitly predict the contact state between a specific pair of a hand and an object, limiting their utility.

We propose a video-based method for predicting binary contact states (contact or no-contact) between a hand and an object in every frame. We assume tracks of hands and objects specified by bounding boxes (*hand-object tracks*) as input. However, annotating a large number of hand-object tracks and their contact states can become too costly. To overcome this difficulty, we propose a semi-supervised framework consisting of (i) automatic training data collection with motion-based pseudo-labels and (ii) guided progressive label correction (gPLC) which corrects noisy pseudo-labels with a small amount of trusted data.

Given unlabeled videos, we apply off-the-shelf detection and tracking models to form a set of hand-object tracks. Then we assign pseudo-contact state labels to each track by looking at its motion pattern. Specifically, we assign a contact label when a hand and an object are moving in the same direction and a no-contact label when a hand is moving alone.

While generated pseudo-labels can provide valuable information on determining the state of contact states with various types of objects when training a prediction model, the pseudo-labels also contain errors that hurt the model’s performance. To alleviate this problem, we correct those errors by the guidance of an additional model trained on a small amount of trusted data. In gPLC, we train two networks each trained with noisy labels and trusted labels. During the training, we iteratively correct noisy pseudo-labels based on both network’s confidence scores. We use the small-scale trusted data to guide which label to be corrected and yield reliable training labels for automatically extracted hand-object tracks.

Since there was no benchmark suitable for this task, we newly annotated contact states to various types of interactions appearing in the EPIC-KITCHENS dataset [8, 9] which includes in-the-wild cooking activities. We show that our prediction model achieves superior performance against frame-based models [62, 40], and the performance further boosted by using motion-based pseudo-labels along with the proposed gPLC scheme.

Our contributions include: (1) A video-based method of predicting contact between a

hand and an object leveraging temporal context; (2) A semi-supervised framework of automatic pseudo-contact state label collection and guided label correction to complement lack of annotations; (3) Evaluation on newly collected annotation over a real-world dataset.

2 Related Works

Reconstructing hand-object interaction Reconstruction of the spatial configuration of hands and their interacting objects plays a crucial role to understand hand-object interaction. 2D segmentation [30] and 3D pose/mesh estimation [15, 18, 30, 39, 44] of hand-object interaction were studied actively in recent years. However, they assume (1) 3D CAD models exist for initialization (except [18]) (2) the hand is interacting with objects, making the methods inapplicable when hand and object are not interacting with each other. While multiple datasets appear for hand-object interaction analysis [9, 14, 31, 40], no dataset focused on the entire process of interaction including beginning and termination of contact. It is worth mentioning DexYCB [6], which captured sequences of picking up an object. However, the performed action was very simple and their analysis focused on 3D pose estimation rather than contact modeling between hands and objects. We study the front stage of the hand-object reconstruction problem—whether the hand interacts with the object or not.

Hand-object contact prediction Contact prediction is found to be a difficult problem because contact cannot be directly observed due to occlusions. To avoid using intrusive hand-mounted sensors, contact and force prediction from visual input was studied [10, 11, 57, 62]. For example, Pham *et al.* [32] present an RNN-based force estimation method trained on force and kinematics measurements from force transducers. These methods require a careful setup of sensors, making it hard to apply them in an unconstrained environment. Instead of precise force measurement, a few methods study contact state classification (*e.g.*, no contact, self contact, other people contact, object contact) from an image [52, 40]. Shan *et al.* [40] collected a large-scale dataset of hand-object interaction along with annotated bounding boxes of hands and objects in contact. They train a network that detects hands and their contact state from their appearance. Narasimhaswamy *et al.* [32] extends the task into multi-class prediction. While their formulation is simple, they did not take the relationship between hands and objects explicitly and were prone to false-positive prediction. To balance utility and convenience, we take the middle way between the two approaches—binary contact state prediction between a hand and an object specified by bounding boxes.

Learning from noisy labels Since dense labels are often costly to collect, methods to learn from large unlabeled data are studied. While learning features from weak cues are studied in object recognition [53] and instance segmentation [54], it was not well studied in a sequence prediction task. Generated pseudo-labels typically include noise that hurts the model’s performance. Various approaches such as loss correction [19, 35], label correction [43, 45], sample selection [21], and co-teaching [16, 26] are proposed to deal with noisy labels. However, most methods assume feature-independent feature noise which is over-simplified, and only a few works study realistic feature-dependent label noise [1, 45]. Zhang *et al.* [45] propose progressive label correction (PLC) which iteratively corrects labels based on the network’s confidence score with theoretical guarantees against feature-dependent noise patterns. Inspired by PLC [45], we propose gPLC which iteratively corrects noisy labels by not only the prediction model but also with the clean model trained on small-scale trusted labels.

3 Proposed Method

In contrast to prior works [32, 40] we formalize the hand-object contact prediction problem as predicting the contact states between a hand and a specific object appearing in a image sequence. We assume video frames $\mathcal{X} = \{X^1, \dots, X^T\}$, hand instance masks $\mathcal{H} = \{H^1, \dots, H^T\}$, and target object bounding boxes $\mathcal{O} = \{O^1, \dots, O^T\}$ as inputs, forming a hand-object track $\mathcal{T} = (\mathcal{X}, \mathcal{H}, \mathcal{O})$.

Our goal is to predict a sequence of a binary contact state (“no contact” or “contact”) $\mathbf{y} = \{y^1, \dots, y^T\} (y \in \{0, 1\})$ given a hand-object track \mathcal{T} . If any physical contact between the hand and the object exists, the binary contact state y is set to 1, otherwise 0. Although we do not explicitly model two-hands manipulation, we consider the presence of another hand as side information (see Section 3.3 for details).

However, collecting a large number of hand-object tracks and contact states for training can become too costly. We deal with this problem by automatic pseudo-label collection based on motion analysis and a semi-supervised label correction scheme.

3.1 Pseudo-Label Generation from Motion Cues

We automatically detect hand-object tracks and assign pseudo-labels to them based on two critical assumptions. (i) When a hand and an object are in contact, they exhibit similar motion pattern. (ii) When a hand and an object are not in contact, the hand moves while the object remains static (see Figure 2 left for illustration). Because these assumptions are simple yet applicable regardless of object appearance and motion direction, we can use these motion-based pseudo-labels for training to achieve generalization against novel objects.

Hand-object track generation Given a video clip, we first use the hand-object detection model [40] to detect bounding boxes of hands and candidate objects appearing in each frame. Note that the detected object’s contact state is unknown and objects which overlap with hands are detected. For each hand detection, we further apply a hand segmentation model trained on EGTEA dataset [27] to each hand detection to obtain segmentation masks.

Next, we associate adjacent detections using a Kalman Filter-based tracker [3]. However, since [40] does not detect objects away from the hand, we extrapolate object tracks one second before and after using a visual tracker [25], producing \mathcal{H} and \mathcal{O} . Finally, we construct the hand-object track \mathcal{T} by looking for pairs of hand and object tracks which include a spatial overlap between hand mask and object bounding box.

Contact state assignment We find contact (and no-contact) moments by looking at the correlation between hand and object motion. First, we estimate optical flows between adjacent frames. Since we are interested in relative movement of hands and objects against backgrounds, we obtain background motion-compensated optical flow and its magnitude M by homography estimation. Specifically, we sample flow vectors outside detected bounding boxes as matches between frames and estimate the homography using RANSAC [13].

Let $F = (f_{ij}) = \mathbb{I}_{(M > \sigma)}$ be a binary mask of foreground moving region its magnitude larger than a certain threshold σ . For each hand and object binary region mask $H = (h_{ij})$ and $O = (o_{ij})$, we calculate the ratio of moving region within each region: $h_r = \frac{\sum_{ij}(h_{ij} \cdot f_{ij})}{\sum_{ij} h_{ij}}$, $o_r = \frac{\sum_{ij}(o_{ij} \cdot f_{ij})}{\sum_{ij} o_{ij}}$. We assign a label to a frame if $\text{IoU}(H, O) > 0$ and h_r and o_r above certain




Figure 2: (Left) Pseudo-label generation from motion cues. (Right) Pseudo-label extension based on hand-object distance.




Figure 3: **Example of generated pseudo-labels:** (Top) Gray and dark gray bar indicates no-contact/contact labels otherwise no labels assigned. (Bottom) Representative frames. Red, blue, and green regions denote moving hand, object, and background regions, respectively. In rightmost frame, no label is assigned because of abrupt background motion.

thresholds. Similarly, we assign a no-contact label if $\text{IoU}(H, O) = 0$ or h_r above threshold but o_r below threshold. However, the above procedure may wrongly assign contact labels if the motion direction of hand and object are different (e.g., the object handled by the other hand). Thus we calculate the cosine similarity between the average motion vector of hand and object region and assign a contact label if above threshold otherwise a no-contact label. To deal with errors in flow estimation, we cancel the assignment if the background motion ratio $b_r = \frac{\sum_{ij}(b_{ij}f_{ij})}{\sum_{ij}b_{ij}}$ ($B = (b_{ij})$ denotes background mask other than H and O) is above threshold.

Pseudo-label extension The above procedure assigns labels on hand-moving frames, but it does not assign labels when hands are moving slowly or still. To assign labels also on those frames, we extend the assigned contact states if the relationship between hands and objects does not change from the timing when pseudo-labels are assigned (see Figure 2 right).

To track hand-object distance, we find point trajectories from hand and object region which satisfy forward-backward consistency [40]. We then calculate the distance d between each hand-object point pair and compare the average distance of them in each frame. We extend the last contact state if the average distance is within a certain range of that of the starting frame. Figure 3 shows an example of the generated pseudo-labels.

3.2 Guided Progressive Label Correction (gPLC)

While generated pseudo-labels include useful information in determining contact states, they also include errors induced by irregular motion patterns. The model may overfit to noise if we simply train it based on these noisy labels. To utilize reliable labels from noisy pseudo-labels, we propose a semi-supervised procedure called guided progressive label correction (gPLC), which works with a small number of trusted labels (see Figure 4 for overview).

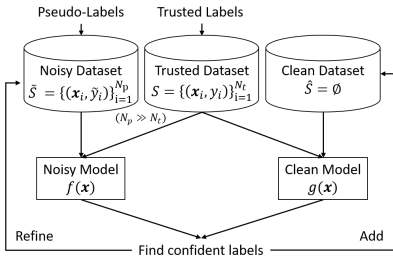


Figure 4: Overview of guided progressive label correction (gPLC).




Figure 5: Architecture of contact prediction model.

Algorithm 1 Guided Progressive Label Correction (gPLC)

Input: Noisy dataset $\tilde{S} = \{(\mathbf{x}_i, \tilde{\mathbf{y}}_i)\}_{i=1}^{N_p}$, trusted dataset $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N_t}$, clean dataset $\hat{S} = \emptyset$, noisy model $f(\mathbf{x})$, clean model $g(\mathbf{x})$, initial and end thresholds (δ_0, δ_{end}) , correction threshold $\delta = \delta_0$, flip ratio α , step size β , supervision interval m , total round N

Output: Trained Model $f(\mathbf{x})$

$\hat{S} = \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^{N_p}$ where $\hat{\mathbf{y}}_i$ is a list with empty elements same size as $\tilde{\mathbf{y}}_i$ // Initialize clean dataset

for $n \leftarrow 1, \dots, N$ **do**

for $i \leftarrow 1, \dots, N_p$ **do**

$\mathbf{z}_i \leftarrow \tilde{\mathbf{y}}_i$ // Keep previous labels

for $t \leftarrow 1, \dots, |\tilde{\mathbf{y}}_i|$ **do**

if $\tilde{y}_i^t \in \{0, 1\}$ **and** $|f(\mathbf{x}_i^t) - \frac{1}{2}| \geq \frac{1}{2} - \delta$ **and** $\mathbb{I}_{\{f(\mathbf{x}_i^t) \geq \frac{1}{2}\}} = \mathbb{I}_{\{g(\mathbf{x}_i^t) \geq \frac{1}{2}\}}$ **then**

$\tilde{y}_i^t, \hat{y}_i^t \leftarrow \mathbb{I}_{\{f(\mathbf{x}_i^t) \geq \frac{1}{2}\}}$ // Refine or add label by confident prediction

end if

end for

 Train $f(\mathbf{x})$ on $(\mathbf{x}_i, \tilde{\mathbf{y}}_i)$ and $g(\mathbf{x})$ on $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$ // Update models

if #iterations % $m = 0$ **then**

 Train $f(\mathbf{x})$ and $g(\mathbf{x})$ on S // Fine-tune on trusted set

end if

end for

if $\sum_{i,t} \mathbb{I}_{\{\tilde{y}_i^t = \hat{y}_i^t\}} < \alpha \cdot \sum_{i,t} \mathbb{I}_{\{\tilde{y}_i^t \in \{0,1\}\}}$ **then**

$\delta \leftarrow \min(\delta + \beta, \delta_{end})$ // Loosen threshold if number of flipped labels are small enough

end if

end for

We assume a noisy dataset \tilde{S} with generated pseudo-labels and a trusted dataset S with manually annotated trusted labels. We train two identical networks, each called noisy model and clean model. The noisy model f is trained on both \tilde{S} and S while the clean model g is trained on S and a clean dataset \hat{S} which is introduced later. We perform label correction against generated pseudo-labels in \tilde{S} using the prediction of both models.

As the training of f proceeds, f will give high confidence against some samples. Similar to PLC [14], we try to correct labels on which f gives high confidence. Note that we correct labels in a frame-wise manner, assuming output contact probability is produced per frame. In gPLC, we correct labels only when f has high confidence and does not contradict the clean network g 's prediction. Because \tilde{S} is generated from motion cues, the decision boundary of f may be different from that of the optimal classifier. Thus the label correction on f alone would not converge to the desired decision boundary. Therefore, we guide the correction process by using g . Starting with a strict threshold on δ , we iteratively correct labels upon training. When the number of corrected labels gets small enough, we increase δ to loosen

the threshold and continue the same procedure. However, since g is trained on a small-scale data, it has the risk of overfitting to S . To prevent overfitting, we iteratively add data that f gives high confidence to another dataset called clean dataset \hat{S} and feed them to g so that g also grows through training. Initially \hat{S} will not contain labels, but high-confident labels will be added over time. See Algorithm 1 for detail. In implementation, $f(\mathbf{x})$ and $g(\mathbf{x})$ are trained beforehand by \hat{S} and S before starting the gPLC iterations.

3.3 Contact Prediction Model

Contact states are predicted by using an RNN-based model that takes RGB images, optical flow, and mask information as input (see Figure 5). For each modality, we crop the input by taking the union of the hand region and the object bounding box. The foreground mask is a four-channel binary mask that tells the presence of a target hand instance mask, a target object bounding box, other detected hand instance masks, and other detected object bounding boxes. The latter two channels prevent confusion when the target hand or object interacts with other entities. RGB and flow images are fed into each encoder branch, concatenated at the middle, and then passed to another encoder. Both encoders consist of several convolutional blocks, each consisted of 3×3 convolution followed by a ReLU and a LayerNorm layer [2], and a 2×2 max-pooling layer. The foreground mask encoder consists of three convolutional layers each followed by a ReLU layer, producing a 1×1 feature map encoding the positional relationship between the target hand, the target object, and the other hands and objects. After concatenating the features extracted from the foreground mask, contact probability is calculated through four bi-directional LSTM layers and three layers of MLP.

We train the network by a standard binary cross-entropy loss weighted by the ratio of the amount of labels in the training data. We did not propagate the error for non-labeled frames.

4 Experiments

Since there was no benchmark suitable for our task, we newly annotated hand-object tracks and contact states between hands and objects against videos in EPIC-KITCHENS dataset [8]. We collected tracks with various objects (*e.g.*, container, pan, knife, sink). The amount of the annotation was 1,200 tracks (67,000 frames) in total. We split the data into a training set (240 tracks), validation set (260 tracks), and test set (700 tracks). For the noisy dataset, we have generated 96,000 tracks with motion-based pseudo-labels.

4.1 Implementation Details

We used FlowNet2 [20] for optical flow estimation. We used Adam [22] for optimization with a learning rate of $3e-4$. We trained the network for 500,000 iterations with a batch size of one and selected the best model by frame accuracy on the validation set. The hyperparameters were set to $\delta_0 = 0.05$, $\delta_{end} = 0.25$, $\alpha = 0.01$, $\beta = 0.025$, $m = 2500$.

4.2 Evaluation Metrics

We prepared several metrics to evaluate the performance. **Frame Accuracy:** Frame-wise accuracy balanced by the ground truth label ratio; **Boundary Score:** F-measure of boundary detection. Performs bipartite graph matching between ground truth and predicted bound-

Method	Frame Acc.	Boundary Score	Peripheral Acc.	Edit Score	Correct Ratio
Fixed	0.500	0.394	0.534	0.429	0.166
IoU	0.642	0.505	0.613	0.678	0.259
ContactHands [33]	0.555	0.440	0.596	0.468	0.136
Shan-Contact [40]	0.608	0.516	0.656	0.507	0.180
Shan-Bbox [40]	0.688	0.435	0.639	0.631	0.189
Shan-Full [40]	0.746	0.477	0.687	0.583	0.193
Supervised (train)	0.770	0.563	0.649	0.718	0.394
Supervised (train+val)	0.816	0.636	0.695	0.793	0.487
Proposed	0.836	0.681	0.730	0.793	0.519

Table 1: Results of hand contact state prediction performance.

Method	Frame Acc.	Boundary Score	Peripheral Acc.	Edit Score	Correct Ratio
Noisy Label only	0.780	0.569	0.703	0.687	0.344
Noisy + Trusted Label	0.811	0.624	0.708	0.759	0.453
Noisy + Trusted Label w/ PLC [43]	0.821	0.636	0.730	0.768	0.480
Pseudo-Labeling [42]	0.784	0.590	0.703	0.737	0.417
RGB	0.787	0.546	0.681	0.709	0.363
Flow	0.833	0.672	0.725	0.789	0.519
Proposed (RGB+Flow)	0.836	0.681	0.730	0.793	0.519

Table 2: Ablations on input modality and other robust learning methods.

ary [36]. Count as correct if the predicted boundary within six frames from the ground truth boundary; **Peripheral Accuracy**: Frame-wise accuracy within six frames from the ground truth boundary; **Edit Score**: Segmental metric using Levenshtein distance between segments [23]. We assume both contact and no-contact labels are foreground; **Correct Track Ratio**: The ratio of tracks which gives frame accuracy above 0.9 and boundary score of 1.0.

4.3 Baseline Methods

We evaluated several baseline methods: **Fixed**: Predicts always as “contact”; **IoU**: Calculate the mask IoU between the input hand mask and object bounding box. If the score is larger than zero predicts as contact, otherwise no-contact; **ContactHands [33]**: Predicts as a contact if the detected hand’s contact state is “object”; **Shan-Contact [40]**: Predicts as a contact if corresponding hand’s contact state prediction is “portable”; **Shan-Bbox [40]**: Predicts as contact if there is enough overlap between the detected object bounding box and input object bounding box; **Shan-Full [40]**: Combines predictions of Shan-Contact and Shan-Bbox; **Supervised**: Our proposed prediction model, trained by trusted data alone. We note that for the **Shan-*** baselines, the *100k+ego* pre-trained model provided by the authors was used, which is trained on egocentric video datasets including the EPIC-KITCHENS dataset.

4.4 Results

Quantitative results We report the performance in Table 1. Our proposed method consistently outperforms the baseline models on all the metrics, achieving a double correct track ratio compared to **IoU** based on the overlap between hand and object bounding boxes. The frame-based methods (**ContactHands**, **Shan-***) performed equal or worse than **IoU**, producing many false positive predictions. These results suggest that previous methods claiming contact state prediction fails to infer physical contact between hands and objects. While **Supervised** performed well, gPLC further boosted the performance by leveraging diverse motion-based cues with label correction, especially on boundary score.

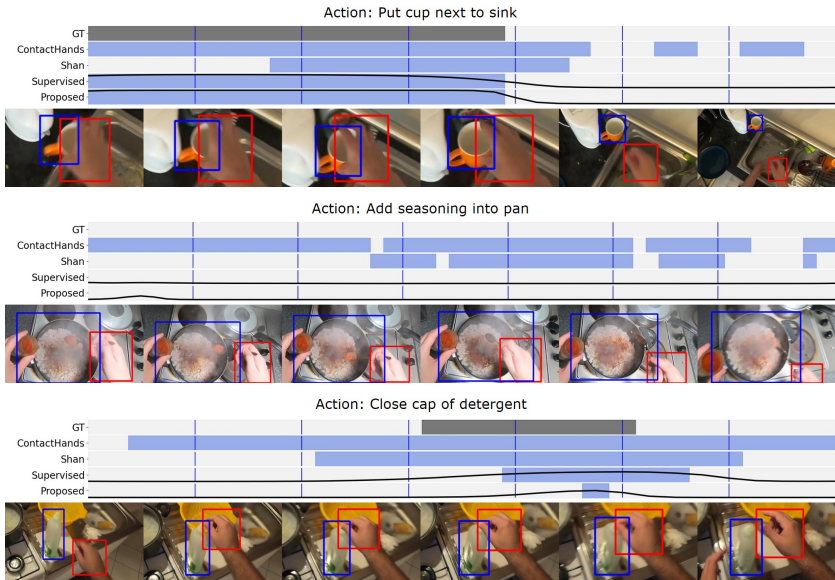


Figure 6: Qualitative examples. Upper chart shows ground truth contact state and prediction of each model (gray and blue region indicates contact, otherwise no-contact) with contact probability in black line. Lower images correspond to blue vertical lines in chart from left to right and red and blue boxes represents input hand and object bounding box.

Qualitative results Figure 6 shows the qualitative results. As shown in the top, our method distinguish contact and no-contact states by looking at the interaction between hands and objects while baseline methods yield false positive predictions by looking at box overlaps. The middle shows a typical no-contact case of a hand floating above an object. Our proposed model trained on motion-based pseudo-labels avoid producing false positive prediction.

Comparison against other robust learning methods To show the effectiveness of the proposed gPLC, we report ablations on other robust learning/semi-supervised learning methods (see Table 2 top). As expected, training using motion-based pseudo-labels performed worse due to labeling errors. Joint training with noisy and trusted labels gives marginal gain against the supervised model, but the boundary score remains low since it overfits against pseudo-label noise. We also applied the existing label correction method [45] on a single network with fine-tuning on trusted labels, but its performance was almost equal to joint training, suggesting that label correction on a single network does not yield good correction. We also tried a typical pseudo-labeling [24] without motion-based labels. However, it showed only a marginal improvement over the supervised baseline, suggesting that our motion-based pseudo-labels are necessary for better generalization.

Effect of input modality The bottom of Table 2 reports the ablation results of changing the input modalities. We observed that using RGB images alone impacts the boundary score, suggesting the difficulty of determining the contact state change without motion information. In contrast, the optical flow-based model achieved nearly the same performance as the full model, suggesting that motion information is crucial for accurate prediction.





Figure 7: Accuracy of noisy labels when initialized by corrupted ground-truth labels. Horizontal axis shows elapsed epochs (“0” denotes initial labels). Vertical axis shows frame accuracy (solid) and boundary score (dashed). Note that non-labeled frames are ignored.

Figure 8: Accuracy of noisy labels when initialized by motion-based pseudo-labels. Horizontal axis shows elapsed epochs and vertical axis shows mean frame accuracy per track.

Error analysis While our method can better predict contact states by utilizing the rich supervision from motion-based pseudo-labels, we observed several failure patterns. As shown in Figure 6 bottom, our method often ignored contacts when a person instantly touched objects without yielding apparent object motion. We also observed failures due to unfamiliar grasps, complex in-hand motion, and failure in determining object regions (see supplemental for more results). These errors indicate the limitation of the motion-based pseudo-labels which assigns labels only when clear joint motion is observed. To better deal with subtle/complex hand motions, additional supervision or rules on such patterns may be required.

How does gPLC correct noisy labels? To understand the behavior of gPLC, we measured how gPLC corrects labels during training. We included the validation set into the training data with two patterns of initial labels: (i) randomly corrupted labels from ground truth (with three different corruption ratios $c_r = 0.1/0.2/0.5$) (ii) motion-based pseudo-labels. We trained the full model and measured the accuracy of the labels for every epoch.

First, gPLC succeeded to correct randomly corrupted label even in the case of high corruption ratio of 0.5 (see Figure 7). However, in the case of a small corruption ratio of 0.1, gPLC made wrong corrections which means that both the noisy model and clean model got the prediction wrong. Improved boundary scores showed that gPLC can iteratively suppress inconsistent boundary errors. In the more realistic case of motion-based pseudo-labels, pseudo-labels were assigned to around 44% of the total frames, and achieved initial mean frame accuracy of 91.4% for the labeled frames. While gPLC reduced the error rate by 20% PLC wrongly flipped the contact state, which may have harmed the final performance (see Figure 8). These results indicate that gPLC effectively corrects noisy labels during training.

5 Conclusion

We have presented a simple yet effective method of predicting the contact state between hands and objects. We have introduced a semi-supervised framework of motion-based pseudo-label generation and guided progressive label correction that corrects noisy pseudo-labels guided by a small amount of trusted data. We have newly collected annotation for evaluation and showed the effectiveness of our framework against several baseline methods.

Acknowledgements

This work was supported by JST AIP Acceleration Research Grant Number JPMJCR20U1, JSPS KAKENHI Grant Number JP20H04205 and JP21J11626, Japan. TY was supported by Masason Foundation. We are grateful for the insightful suggestions and feedback from the anonymous reviewers.

References

- [1] Shuichi Akizuki and Yoshimitsu Aoki. Tactile logging for understanding plausible tool use based on human demonstration. In *Proceedings of British Machine Vision Conference*, page 334, 2019.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *Computing Research Repository*, [abs/1607.06450](https://arxiv.org/abs/1607.06450), 2016.
- [3] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3464–3468, 2016.
- [4] Samarth Brahmabhatt, Cusuh Ham, Charles C Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8709–8719, 2019.
- [5] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. *Proceedings of the IEEE international conference on computer vision*, pages 12417–12426, 2021.
- [6] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, et al. Dexycb: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9044–9053, 2021.
- [7] Pengfei Chen, Junjie Ye, Guangyong Chen, Jingwei Zhao, and Pheng-Ann Heng. Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European Conference on Computer Vision*, pages 720–736, 2018.
- [9] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, , Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *International Journal of Computer Vision*, 2021.
- [10] Eadom Dessalene, Chinmaya Devaraj, Michael Maynard, Cornelia Fermuller, and Yiannis Aloimonos. Forecasting action through contact representations from first person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

- [11] Kiana Ehsani, Shubham Tulsiani, Saurabh Gupta, Ali Farhadi, and Abhinav Gupta. Use the force, luke! learning to predict physical forces by simulating effects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 224–233, 2020.
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, October 2019.
- [13] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [14] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 409–419, 2018.
- [15] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3196–3206, 2020.
- [16] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Neural Information Processing Systems*, 2018.
- [17] Shangchen Han, Beibei Liu, Randi Cabezas, Christopher D Twigg, Peizhao Zhang, Jeff Petkau, Tsz-Ho Yu, Chun-Jung Tai, Muzaffer Akbay, Zheng Wang, et al. Megatrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics*, 39(4):87–1, 2020.
- [18] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11807–11816, 2019.
- [19] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [20] Eddy Ilg, Nikolaus Mayer, T Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2462 – 2470, 2017.
- [21] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pages 2304–2313, 2018.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

- [23] Colin Lea, Austin Reiter, René Vidal, and Gregory D Hager. Segmental spatiotemporal cnns for fine-grained action segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 36–52, 2016.
- [24] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013.
- [25] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4282–4291, 2019.
- [26] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2020.
- [27] Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *Proceedings of the European Conference on Computer Vision*, pages 619–635, 2018.
- [28] Jirapat Likitlersuang, Elizabeth R Sumitro, Tianshi Cao, Ryan J Visée, Sukhvinder Kalsi-Ryan, and José Zariffa. Egocentric video: a new tool for capturing hand use of individuals with spinal cord injury at home. *Journal of Neuroengineering and Rehabilitation*, 16(1):83, 2019.
- [29] T. Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3888–3897, 2019.
- [30] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3d hand-object poses estimation with interactions in time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14687–14697, 2021.
- [31] Supreeth Narasimhaswamy and Saif Vazir. Workinghands: A hand-tool assembly dataset for image segmentation and activity mining. In *Proceedings of the British Machine Vision Conference*, page 258, 2019.
- [32] Supreeth Narasimhaswamy, Trung Nguyen, and Minh Hoai. Detecting hands and recognizing physical contact in the wild. In *Advances in Neural Information Processing Systems*, 2020.
- [33] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2701–2710, 2017.
- [34] Deepak Pathak, Yide Shentu, Dian Chen, Pulkit Agrawal, Trevor Darrell, Sergey Levine, and Jitendra Malik. Learning instance segmentation by interaction. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018.

- [35] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1944–1952, 2017.
- [36] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 724–732, 2016.
- [37] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A Argyros, and Abderrahmane Kheddar. Hand-object contact force estimation from markerless visual tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2883–2896, 2017.
- [38] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems*.
- [39] Javier Romero, Hedvig Kjellström, and Danica Kragic. Hands in action: real-time 3d reconstruction of hands in interaction with objects. In *2010 IEEE International Conference on Robotics and Automation*, pages 458–463, 2010.
- [40] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F. Fouhey. Understanding human hands in contact at internet scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9869–9878, 2020.
- [41] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *Proceedings of the European Conference on Computer Vision*, pages 438–451, 2010.
- [42] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: a dataset of whole-body human grasping of objects. In *Proceedings of the European Conference on Computer Vision*, pages 581–600, 2020.
- [43] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5552–5560, 2018.
- [44] Dimitrios Tzionas and Juergen Gall. 3d object reconstruction from hand-object interactions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 729–737, 2015.
- [45] Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature dependent label noise: a progressive approach. In *International Conference on Learning Representations*, 2021.