

Coupled Particle Filters for Robust Affordance Estimation

Patrick Lowin^{1,2}

Vito Mengers^{1,2,3}

Oliver Brock^{1,2,3}

Abstract—Robotic affordance estimation is challenging due to visual, geometric, and semantic ambiguities in sensory input. We propose a method that disambiguates these signals using two coupled recursive estimators for sub-aspects of affordances: graspable and movable regions. Each estimator encodes property-specific regularities to reduce uncertainty, while their coupling enables bidirectional information exchange that focuses attention on regions where both agree, i.e., affordances. Evaluated on a real-world dataset, our method outperforms three recent affordance estimators (Where2Act, Hands-as-Probes, and HRP) by 308%, 245%, and 257% in precision, and remains robust under challenging conditions such as low light or cluttered environments. Furthermore, our method achieves a 70% success rate in our real-world evaluation. These results demonstrate that coupling complementary estimators yields precise, robust, and embodiment-appropriate affordance predictions.

I. INTRODUCTION

Affordances are the possible actions an agent can perform in its environment [1]. For robots that want to interact with the world, detecting affordances is essential yet challenging, as affordance cues are often ambiguous in the sensory input. Since robots can interact by grasping and moving objects, affordances arise in regions where these properties co-occur. In this paper, we present a robust method for detecting affordances by coupling estimators of graspable and movable regions and using their relationship to resolve ambiguities in the sensor data.

Coupling the estimation processes for graspable and movable regions improves robustness in two ways. First, while we lack explicit affordance priors, we exploit available priors for graspability and movability to reduce uncertainty. For instance, predicted grasps in physically infeasible locations, such as edges caused by depth discontinuities, are rejected. Second, the relationship between the two properties provides an additional prior, allowing the estimators to focus on regions where both agree. This resolves ambiguities that a single estimator cannot.

Initially, each estimator focuses on the regions most likely for its property. Through information sharing, they gradually align attention on regions important to both, i.e., the true affordances (Fig. 1). The strength of this coupling can be adjusted to control the interaction between the estimators. By coordinating in this way, they share complementary




Fig. 1. Robots can interact by grasping an object and then moving it. Consequently, affordances emerge in regions where these properties co-occur. By estimating these properties separately and leveraging their relationship, we can resolve ambiguities in both modalities and estimate robust affordances.

information, reduce uncertainty, and produce more robust affordance estimates.

We evaluate our method on the RBO dataset [2] of articulated objects and interactions, which we extend with manipulable-region annotations. We compare against recent learning-based approaches, outperforming a geometry-based estimator trained in simulation by 308% and two appearance-based methods trained on human interactions by 245% and 257%. At the same time, our approach demonstrates improved robustness under challenging conditions such as low light or clutter. We further evaluate this robustness in real-world manipulation trials, achieving a 70% success rate. These results highlight that leveraging shared structure, complementary information, and coordinated attention through coupled estimation processes enables robust affordance estimation.

Dataset annotations and code are available at <https://github.com/tu-rbo/CoupledParticleFilters>.

II. RELATED WORK

Affordance estimation is typically approached by learning affordances in simulation or from human interaction data within a single, monolithic model. We review these approaches in Sec. II-A, highlighting their struggles with out-of-distribution data and perceptual ambiguities. To address these challenges, we instead estimate affordances with

¹ Robotics and Biology Laboratory, Technische Universität Berlin

² Robotics Institute Germany

³ Science of Intelligence (SCIOI), Cluster of Excellence, Berlin, Germany

This work has been partially funded by the German Federal Ministry of Research, Technology and Space (BMFTR) under the Robotics Institute Germany (RIG) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2002 – 390523135.

two coupled recursive estimators for related properties, an approach that has already proven effective for other complex perceptual problems, as we review in Sec. II-B.

A. Affordance Estimation

Affordances are characterized as an agent’s possible actions in an environment [1]. Their perception thus depends on the agent’s capabilities and tasks. Early works used hand-crafted, often geometry-based rules to predict object affordances [3,4], but defining such rules for complex scenarios is infeasible.

Recent learning-based methods capture general affordance features to enable interaction in simulation [5]–[8], yet they often struggle to generalize to real-world environments due to the sim-to-real gap [9]. Other approaches learn from real-world human interaction videos [10]–[15], but face visual ambiguities or embodiment mismatches. Moreover, they often assume different properties to be known in advance, such as actions [5,6,8], objects [5]–[8,16], movable parts [5,6,8], or segmentation [5]–[8,16].

Our approach avoids these assumptions by factorizing affordances into two simpler properties: graspability and movability. While we also use a model trained on human data [10] to determine movability, we avoid embodiment mismatch by coupling with a robot-specific graspability model [17]. This coupling also aids in disambiguating affordance cues through complementarity of the two estimation processes.

B. Coupled Recursive Estimation

Coupled recursive estimation encodes information about perceptual sub-aspects and their relationships, enabling robust interpretation of sensory signals through Bayesian inference. This approach has been applied to the estimation of kinematic structure and other dynamic properties of objects [18] as well as the overall object segmentation [19]. Its semi-decomposed structure allows integration of additional sensor modalities through simple composition, such as proprioception [20], audio signals [21], or hand poses of surrounding humans [22].

Unlike traditional distributed state estimation methods [23, 24], coupled recursive estimation enforces structured interactions between estimators, akin to factor graphs [25] but more specific. This structure constrains system dynamics, enabling complex action generation with simple gradient descent [26], and shares similarities with biological perceptual processes, motivating its use as a model of human vision [27,28].

We adopt this framework for affordance estimation by coupling particle filters that operate on different signals and representations, graspability and movability. While prior work has explored coupling filters over the same state variables [29], our approach coordinates multiple, distinct estimators via cross-modal fusion, as we explain in detail in the following.

III. COUPLED PARTICLE FILTERS FOR AFFORDANCES

We factorize affordances into graspability and movability, estimate these separately, and couple them to obtain

Algorithm 1 Coupled Particle Filters

```

1: Initialize: Two sets of particles  $\{x_1^i\}_{i=1}^N$  and  $\{x_2^i\}_{i=1}^N$ .
2: for each time step  $t$  do
3:   for each Filter  $f \in \{1, 2\}$  do
4:     Prediction Step: Propagate particles according to
       optical flow.
5:     Injection Step: Inject new particles in low-particle-
       density areas.
6:     Update Step: Weight particles using neural net-
       work measurements.
7:   end for
8:   Cross-Modal Fusion: Weigh particles based on the
       other belief’s density.
9:   for each Filter  $f \in \{1, 2\}$  do
10:    Resample: Draw particles from belief according to
       normalized weights.
11:  end for
12: end for

```

an embodiment-appropriate and disambiguated belief over affordances. Both estimation problems are multi-modal, motivating the use of particle filters, where we incorporate additional priors to address property-specific ambiguities. To couple them, we implement a cross-modal information exchange that directs each estimator toward mutually consistent affordances. First, we describe the template for both particle filters in Sec. III-A, then their instantiation for movability in Sec. III-B, for graspability in Sec. III-C, and finally their fusion in Sec. III-D.

A. Shared Particle Filter Procedure

Estimating either graspability or movability requires reasoning over noisy, ambiguous observations across many candidate regions. We measure each property with a learning-based model from the literature, then refine the estimate with a particle filter that enforces additional constraints. Here, we describe the shared procedure before detailing the property-specific instantiations in the next two sections.

Each particle represents a 3D position $x \in \mathbb{R}^3$ corresponding to a graspable or movable region. Rather than initializing particles uniformly, we draw them from the first measurement (Alg. 1, L. 1) to concentrate particles near likely states and improve initial convergence. We then recursively refine beliefs through prediction and update steps (Alg. 1, L. 2–7).

During prediction, particles are propagated according to optical flow with added noise for diversity to account for motion in the scene. Optical flow measurements are often ambiguous in cluttered scenes, so we track particle depth and reject physically inconsistent motion. To recover from local minima and detect new objects, we include an injection step (Alg. 1, L. 4), which samples new particles in low-density regions, where new objects are most likely.

After prediction and injection, each filter updates particles using its own measurement source and then incorporates

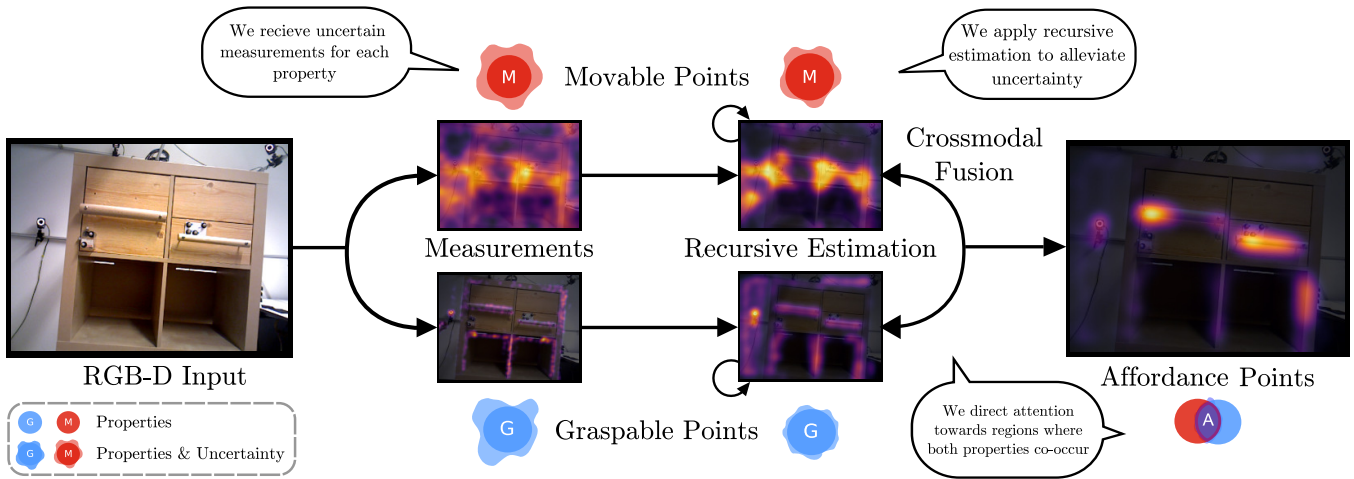


Fig. 2. Our approach estimates affordances as a combination of the graspable and movable regions of our environment. Since our measurement sources provide only uncertain measurements for these properties, we recursively estimate a belief for them and apply additional priors to resolve ambiguities. This aggregates a strong belief in regions with high measurements for their individual property. However, affordances need to be both graspable and movable. Therefore, we exchange information between our estimators, focusing each estimator’s attention on relevant parts of the scene that satisfy both properties, i.e., affordances.

information from the coupled belief, before resampling the particle set.

B. Specific Instantiation for Movability

Movability is measured using Hands-as-Probes [10], which predicts a heatmap of movable regions from RGB images of human interactions. To weight particles according to this measurement, we project their 3D positions into the heatmap and assign the corresponding value as weight $\tilde{w}_{\text{movability}}$.

Because Hands-as-Probes [10] was trained on human interactions, its predictions do not directly transfer to robots. To address this mismatch, we restrict the movable estimator to regions that a robot can interact with by incorporating information about robotic grasps through its coupling with the graspability estimator. This adjustment focuses the movable estimator on affordances for robots, resolving further noise and ambiguities.

C. Specific Instantiation for Graspability

Graspability requires modeling not only location but also grasp pose. Each particle thus represents not just a 3D position $x \in \mathbb{R}^3$ but also an associated gripper pose $\nu \in \mathbb{SE}(3)$.

As a measurement source, we use GraspNet [17], which predicts candidate grasps with contact points and success likelihoods from a scene point cloud. To thus ensure that the tracked particle set contains graspable points with appropriate grasps, we determine particle weights by comparing both grasp likelihood and pose similarity to nearby GraspNet predictions.

Let $\mathbb{Z}_t = (y_j \in \mathbb{R}^3, v_j \in \mathbb{SE}(3), p_j \in [0, 1])$ be the set of candidate grasps predicted by GraspNet at time t , where y_j is a 3D grasp position, v_j is a gripper pose, and p_j is the associated predicted grasp success likelihood. To weight a

particle with position $x \in \mathbb{R}^3$ and gripper pose $\nu \in \mathbb{SE}(3)$, we first ignore gripper orientation and only consider spatial proximity, encoded through a Gaussian kernel for distance $K(x, y)$, and predicted grasp success p_j to determine a grasp success likelihood:

$$v = \max_{(y_j, v_j, p_j) \in \mathbb{Z}_t} p_j K(x, y_j). \quad (1)$$

Then we furthermore consider the similarity of grasps, using cosine similarity $S(\nu, v_j)$ of gripper approach vectors, while considering their overall likelihood p_j and proximity:

$$u = \max_{(y_j, v_j, p_j) \in \mathbb{Z}_t} p_j K(x, y_j) S(\nu, v_j). \quad (2)$$

The final weight is then the product of these two similarity measures, i.e., $\tilde{w}_{\text{graspability}} = u \cdot v$. Multiplying both ensures that particles are weighted highly only if they are both near a high-confidence grasp and have a compatible orientation. To further reduce false positives at geometric ambiguities (e.g., depth discontinuities), we detect such regions in the depth image and down-weight particles accordingly.

These tailored filters allow each estimator to handle its modality-specific ambiguities and adapt to dynamic environments through targeted particle injections. However, each focuses only on its own property, not on the joint affordances of interest. To achieve this coordination, we introduce coupling between the filters, described next.

D. Coupling Estimators for Cross-Modal Fusion

To align both estimators on manipulable regions consistent with the robot’s embodiment, we couple them through cross-modal fusion. This biases resampling toward particles supported by both beliefs, effectively shifting each estimator’s attention to mutually consistent regions.

For this coupling, we want to weight each particle depending on its likelihood according to *both* beliefs. For each set,

we first determine its measurement-specific weight \tilde{w} in its respective estimator. Then we determine its likelihood given the particle set of the other estimator. For each estimator with particle set \mathbb{X} , we determine the cross modal weight for each particle by comparing it to the density of the other estimator’s particle set \mathbb{Y} (Alg. 1, L. 8). To do so, we represent both beliefs in 3D and estimate densities using a Gaussian kernel with $\sigma = 0.05$ to the cross-modal weight \hat{w} for particle $x \in \mathbb{X}$:

$$\hat{w} = \sum_{y \in \mathbb{Y}} \exp\left(-\frac{\|x - y\|^2}{2\sigma^2}\right). \quad (3)$$

The final particle weight in each estimator then combines the original measurement-based weight \tilde{w} with the cross-modal weight \hat{w} for each particle:

$$w = f_u(\tilde{w}) f_c(\hat{w}), \quad (4)$$

where f_u and f_c clip values to predefined ranges tuned on a single scene. For graspability, we clip \tilde{w} to $(0.2, 0.25)$ and \hat{w} to $(0.3, 0.6)$. For movability, we clip \tilde{w} to $(0.2, 0.3)$ and \hat{w} to $(0.45, 0.55)$.

We assign broader ranges to the cross-modal weights, giving them higher influence since beliefs are typically more reliable than raw measurements. This ensures particles are more likely resampled if supported by both estimators, enforcing mutual agreement. The result is a robust joint belief over affordances that integrates complementary priors and, in the case of graspability, already encodes executable grasps for direct interaction.

IV. EXPERIMENTS

To evaluate our approach, we compare against three recent learning-based methods on a real-world dataset. First, we describe the experimental setup, including how we extend an existing dataset with affordance annotations, (Sec. IV-A). Next, we demonstrate that fusing graspable and movable information improves performance over recent approaches (Sec. IV-B), while also providing robustness under challenging conditions (Sec. IV-C). Lastly, we show how this effectiveness stems not from our measurement sources but the coupling of estimators (Sec. IV-D), and how cross-modal coupling is beneficial compared to other fusion strategies (Sec. IV-E).

A. Experimental Setup

To evaluate the robustness of affordance estimation in real-world scenarios, we need a dataset for such affordances given real sensor data. To our knowledge, existing datasets lack explicit manipulable region annotations [11,30,31], motivating our extension of the RBO dataset [2] (Sec. IV-A.1). On this dataset, we compare against three recent works on affordance estimation (Sec. IV-A.2), using appropriate metrics (Sec. IV-A.3).

1) *Dataset*: We evaluate our approach using the RBO dataset of articulated objects and interactions [2], which contains RGB-D videos of interactions with various articulated objects in challenging environmental conditions. Additionally, it includes ground-truth models and poses for the articulated objects in the scene, enabling precise evaluation. To assess performance, we annotate the manipulable parts of these object models with interaction volumes (3D regions indicating where a robot can interact) and project them into the sequences.

To evaluate performance and robustness against adverse conditions, we further split the dataset into three parts: well-lit, dark, and cluttered environments. For cluttered environments, we only show qualitative examples, since the dataset does not include object models for the distractor objects.

2) *Learning-Based Approaches*: We compare against three recent learning-based affordance estimation techniques to show both the performance and robustness of our approach. Where2Act [5] predicts grasps and success likelihood from point cloud data and an action primitive, while Hands-as-Probes [10] and a vision transformer trained with HRP [15] predict heatmaps for possible interactions given human interaction data. Since the vision transformer predicts 12 heatmaps, we average their values to a single heatmap.

3) *Metrics*: To quantify the performance of the different approaches, we compute the mean precision across all sequences for each object. For this, we sample 200 affordance points from the prediction of each approach to ensure comparable coverage across methods. We then count the proportion of samples within 2cm of the nearest manipulable region. This tolerance accounts for inaccuracies between shape models and transformations. For objects with multiple manipulable parts, we additionally analyze the distribution across these regions to show if the models detect all relevant affordances.

B. Coupling Reduces Uncertainty and Improves Precision

Our method couples estimators for graspable and movable areas and aligns their attention on regions where these properties co-occur, i.e., affordances. In these regions, the estimators share complementary information, resolving ambiguities that single modalities cannot handle.

We demonstrate the improved performance of our approach by comparing it against three recent affordance estimators, Where2Act [5], Hands-as-Probes [10], and HRP [15], on well-lit and uncluttered scenes from the RBO dataset [2]. As shown in Fig. 4, our approach achieves the highest mean precision across all objects, outperforming Where2Act [5] by 308%, Hands-as-Probes [10] by 245%, and the HRP model [15] by 257%.

All three learning-based approaches struggle with out-of-distribution data and modality-specific ambiguities in their input representation: Where2Act [5] with geometric and Hands-as-Probes [10] and the HRP model [15] with visual ambiguities (Fig. 3). Our method addresses these challenges by filtering noisy measurements, fusing complementary modalities, and incorporating priors. Grasps encode




Fig. 3. Our method can handle dynamic and complex scenes. While the non-recursive learning-based approaches detect new affordances directly, we rely on our injection step to introduce new affordances, such as the object inside the shelf, into our belief. Our estimates are more precise, focusing on manipulable affordances.

geometric cues that disambiguate flat or textured regions, while movability estimates and priors suppress grasps in non-movable areas.

When the estimators align, they exchange refined information, yielding localized and reliable affordance predictions. Next, we evaluate how this robustness extends to more challenging environments.

C. Coupling Improves Robustness in Adverse Conditions

We further evaluate robustness under adverse conditions, showing that our method remains reliable:

1) *Dark Environments*: Dark environments introduce visual ambiguities that affect the appearance-based Hands-as-Probes [10], which also acts as our movable measurement source. While appearance-based methods struggle under these conditions, our coupled system resolves ambiguities by leveraging grasp information. As a result, we outperform Hands-as-Probes [10] by 235% and the HRP model [15] by 207% in dark environments. Because grasp estimates remain unaffected by lighting, we also outperform Where2Act [5] by 197% (Fig. 4, bottom).

By fusing modalities, our method produces precise and robust affordances despite reduced visual information (Fig. 5). This robustness to visual ambiguity motivates evaluating our method in cluttered environments, which introduce additional geometric and visual challenges.

2) *Cluttered Environments*: Clutter introduces occlusions and texture complexity that hinder both geometry- and

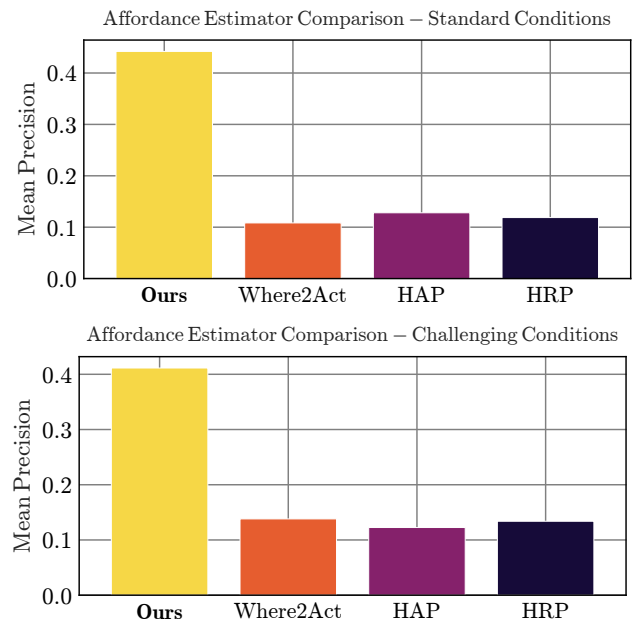


Fig. 4. Our approach outperforms recent learning-based approaches in both standard, i.e., well-lit and clutter-free (top), and more challenging, i.e., dark or cluttered (bottom), settings. The monolithic learned models struggle with real-world ambiguities, while coupling resolves them by aligning estimator attention and filtering noise.

appearance-based estimators. Qualitative results (Figs. 3 and 6) show our estimates remain localized on manipulable parts of the target object, while the predictions of recent learning-based approaches often spill into irrelevant regions.

Our coupled estimators drive attention toward relevant affordances, filtering out spurious grasps such as edges of tables. The cluttered IKEA scene (Fig. 3) highlights robustness in dynamic, complex environments and underscores the method’s potential for real-world robotic interaction. Next, we will show in two ablations that this performance and robustness does not stem from our measurement sources or filtering by themselves, but our coupling approach.

D. Ablation: Our Cross-Modal Coupling Aligns Estimators

Our method aligns the estimators for graspable and movable regions on affordances. We show this by analyzing our individual measurement sources and recursive estimators without coupling. The individual measurement sources struggle to estimate affordances as they only account for one aspect of robotic manipulation (Fig. 7). Consequently, filtering the grasp prediction results in a performance drop because highly graspable areas, such as table edges, are not interactable. Movability, on the other hand, is a stronger indicator for interaction, and filtering improves performance by 13% under good lighting conditions. However, both measurement sources are susceptible to perceptual ambiguities, which we cannot resolve using data from one modality alone.

By enabling the estimators to inform each other, they can resolve these ambiguities and align their attention on manipulable regions. As a result, our coordinated attention estimates precise and robust affordances even under adverse




Fig. 5. Our affordance estimates remain precise in dark environments. Hands-as-Probes [10] and HRP [15] suffer from visual ambiguities, and Where2Act [5] struggles with object complexity, while our coupled system yields reliable affordances.




Fig. 6. Due to the coupling, our affordance estimates remain focused on manipulable object regions even in clutter. Other approaches are less precise, spreading across distractor surfaces such as the table edge.

conditions (Fig. 7). Next, we will show that this cannot be achieved with other fusion methods.

E. Ablation: Coupling of Particle Filters Outperforms Standard Sensor Fusion

In contrast to standard sensor fusion, our approach enables the two estimators to influence and inform each other through cross-modal coupling. To demonstrate the benefits of this method, we compare it to standard sensor fusion approaches (Fig. 8).

Our approach outperforms all other sensor fusing methods, as it accounts for uncertainty and the relationship between the properties. Naive fusion directly combines the estimates for graspability and movability, but this does not account for the uncertainty in both measurement sources. Early and late fusion, on the other hand, alleviate measurement uncertainty through recursive estimation but do not model the relationship between the properties. Late fusion processes the signals independently and thus loses the shared structure. By combining the signals before we apply recursive estimation, we keep the shared structure in the signal but do not explicitly leverage it.

Our method utilizes the relationship between the properties in the interconnection between estimators and addresses uncertainty, and can therefore estimate more robust and precise affordances.

F. Real World Experiments

We conduct real-world experiments to evaluate whether affordances that score highly on benchmark datasets also correspond to actionable grasps in physical scenes. We again compare against Hands-as-Probes [10] and HRP [15], and exclude Where2Act [5] because obtaining the required point-cloud segmentation for real-world scenes is a challenging

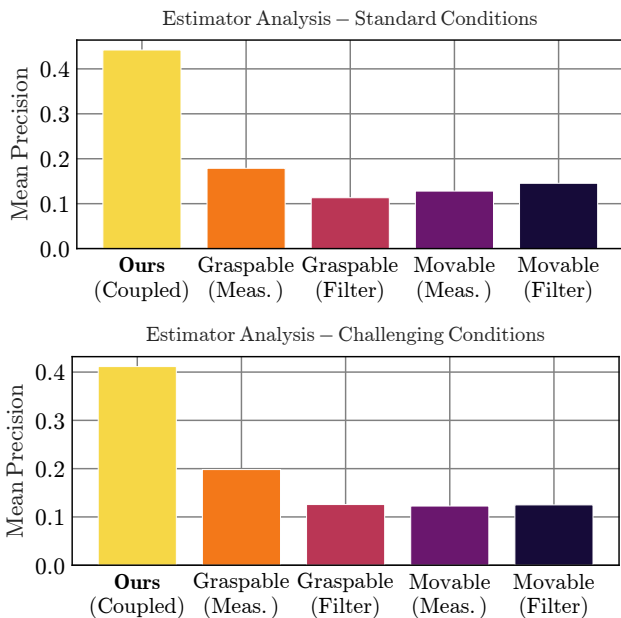


Fig. 7. Our cross-modal fusion aligns the estimators’ attention on regions that satisfy both properties, that is the affordances. The individual estimators suffer from their respective ambiguities, but by sharing information they can resolve these, focus on relevant regions and estimate robust affordances even under adverse conditions.

perception problem by itself. For each trial, we compute the highest-scoring point in the heatmap or the densest region in our belief after 10 iterations and use this as our affordance estimate. If a robot can grasp and move the estimated affordance, we count the trial as a success.

On cluttered tabletops, our method achieves an 80% success rate, outperforming Hands-as-Probes and HRP, which each achieve 40%. In the more challenging cluttered IKEA shelf scenes, our method achieves a 60% success rate, compared to 20% for Hands-as-Probes and 40% for HRP. These results represent up to a threefold improvement over prior methods in difficult scenes.

The appearance-based models often struggle with visual ambiguities and predict affordances at non-manipulable regions. In contrast, our approach resolves many of these ambiguities by incorporating grasp information into the affordance estimation process. Failures in the shelf scenes typically occur due to systemic errors in both measurement sources.

Across all 10 trials, our method achieves a 70% success rate, outperforming the learned baselines and demonstrating the robustness of our coupled estimators in challenging real-world environments.

V. LIMITATIONS & DISCUSSION

Our approach proved remarkably robust throughout our experiments, but is still limited by the accuracy and reliability of its measurement models. In particular, systematic errors that occur in both graspability and movability estimators cannot be resolved through coupling alone. This could be

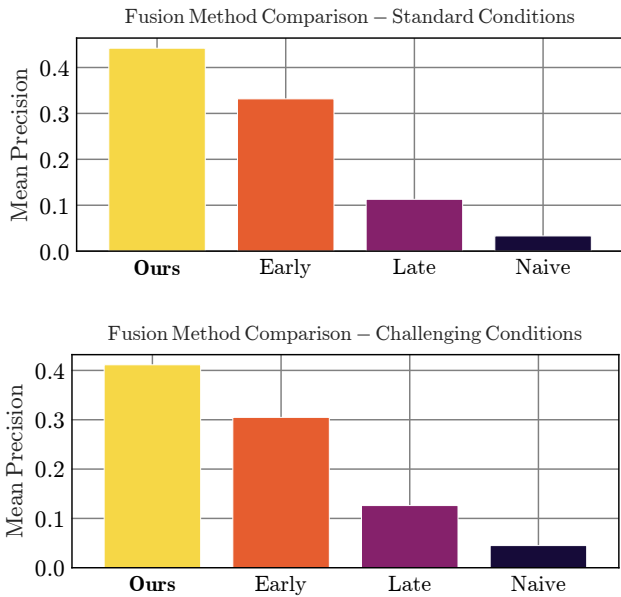


Fig. 8. Coupling particle filters with our cross-modal fusion outperforms other sensor fusion schemes, in particular, naive and late fusion. Early fusion combines signals before recursive estimation, which leads to good estimates, but does not leverage the shared structure of the two properties to further disambiguate the signals.

addressed by introducing additional complementary estimators, either based on additional visual properties or based on other sensors, such as tactile or audio cues. Alternatively, we could make the system interact with its environment to resolve ambiguities by gathering information from different viewing directions or eliminate non-manipulable regions by interacting with the estimated affordances.

A second limitation is computational cost. The coupled particle filters themselves add little overhead, but the measurement models rely on large neural networks, particularly for movability prediction with Hands-as-Probes [10]. Replacing these with faster models would allow real-time affordance estimation without altering the overall pipeline.

Despite these limitations, our approach has several strengths: it substantially improves robustness over raw measurements, is lightweight once measurements are available, and as a modular system it can easily be extended with additional estimators. Parameter tuning proved straightforward, optimized manually on only two sequences, yet generalized across the dataset thanks to the strong inductive bias of the coupled estimation framework and Bayesian inference. This suggests a promising route to enhance the generalizability of modern learning-based perception methods beyond their training assumptions, especially as an increasing number of estimators is incorporated.

VI. CONCLUSION

We propose a method to estimate actionable affordances in real-world environments by identifying regions that are both graspable and movable. To address uncertainty of our learning-based measurement models, we combine recursive

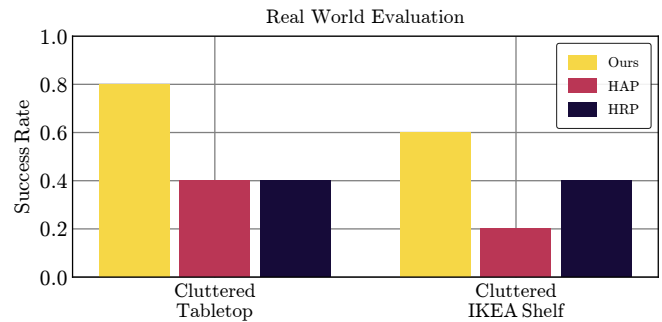


Fig. 9. Our approach estimates robust affordances that enable robots to interact with their environment. We evaluate this in real-world trials on cluttered tabletop and shelf scenes with four objects (cup, banana, toothpaste, and Cheez-it box). Our estimated affordances correspond to precise and actionable grasps, even in cluttered environments.

estimation with coupling between graspability and movability estimators. This alignment allows the estimators to exchange complementary information, resolve ambiguities, and yield robust affordance predictions fitting the robotic embodiment.

In experiments on the RBO dataset [2] of articulated objects, our method outperforms recent learning-based affordance estimators by large margins and demonstrates robustness under adverse conditions such as low light and clutter. The resulting affordances are precise, localized, and probabilistic, supporting reasoning about uncertainty. Together, these properties make our approach well suited for robotic exploration and manipulation, as seen in our real-world experiments, and point toward scalable, modular perception systems built from coupled estimators, capable of reasoning over multiple complementary sensory cues and models.

REFERENCES

- [1] J. J. Gibson, *The ecological approach to visual perception: classic edition*. Psychology press, 2014.
- [2] R. Martín-Martín, C. Eppner, and O. Brock, “The RBO dataset of articulated objects and interactions,” *The International Journal of Robotics Research (IJRR)*, vol. 38, no. 9, pp. 1013–1019, 2019.
- [3] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos, “Affordance detection of tool parts from geometric features,” *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1374–1381, 2015.
- [4] M. Hassanin, S. Khan, and M. Tahtali, “Visual affordance and function understanding: A survey,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 3:47, pp. 1–35, 2021.
- [5] K. Mo, L. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani, “Where2Act: From pixels to actions for articulated 3D objects,” *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6793–6803, 2021.
- [6] C. Ning, R. Wu, H. Lu, K. Mo, and H. Dong, “Where2Explore: Few-shot affordance learning for unseen novel categories of articulated objects,” *International Conference on Neural Information Processing Systems (NeurIPS)*, art. no. 203, 2023.
- [7] Z. Xu, H. Zhanpeng, and S. Song, “UMPNet: Universal manipulation policy network for articulated objects,” *IEEE Robotics and Automation Letters (RAL)*, 2022.
- [8] R. Wu, K. Cheng, Y. Shen, C. Ning, G. Zhan, and H. Dong, “Learning environment-aware affordance for 3D articulated object manipulation under occlusions,” *International Conference on Neural Information Processing Systems (NeurIPS)*, art. no. 2664, 2023.
- [9] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 23–30, 2017.
- [10] M. Goyal, S. Modi, R. Goyal, and S. Gupta, “Human hands as probes for interactive object understanding,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3283–3293, 2022.
- [11] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, “Affordances from human videos as a versatile representation for robotics,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13 778–13 790, 2023.
- [12] T. Nagarajan, Y. Li, C. Feichtenhofer, and K. Grauman, “Ego-Topo: Environment affordances from egocentric video,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 163–172, 2020.
- [13] S. Liu, S. Tripathi, S. Majumdar, and X. Wang, “Joint hand motion and interaction hotspots prediction from egocentric videos,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3282–3292, 2022.
- [14] Y. Ye, X. Li, A. Gupta, S. De Mellon, S. Birchfield, J. Song, S. Tulsiani, and S. Liu, “Affordance diffusion: Synthesizing hand-object interactions,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22 479–22 489, 2023.
- [15] M. K. Srirama, S. Dasari, S. Bahl, and A. Gupta, “Hrp: Human affordances for robotic pre-training,” *arXiv*, no. 2407.18911, 2024.
- [16] Y. Ju, K. Hu, G. Zhang, G. Zhang, M. Jiang, and H. Xu, “Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation,” *European Conference on Computer Vision (ECCV)*, pp. 222–239, 2024.
- [17] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-GraspNet: Efficient 6-DoF grasp generation in cluttered scenes,” *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 13 438–13 444, 2021.
- [18] R. Martín-Martín and O. Brock, “Coupled recursive estimation for online interactive perception of articulated objects,” *The International Journal of Robotics Research (IJRR)*, vol. 41, no. 8, pp. 741–777, 2022.
- [19] V. Mengers, A. Battaje, M. Baum, and O. Brock, “Combining motion and appearance for robust probabilistic object segmentation in real time,” *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 683–689, 2023.
- [20] R. Martín-Martín and O. Brock, “Cross-modal interpretation of multi-modal sensor streams in interactive perception based on coupled recursion,” *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3289–3295, 2017.
- [21] M. Baum, A. Froessl, A. Battaje, and O. Brock, “Estimating the motion of drawers from sound,” *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 697–703, 2023.
- [22] A. Pfisterer, X. Li, V. Mengers, and O. Brock, “A helping (human) hand in kinematic structure estimation,” *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [23] S. Das and J. M. F. Moura, “Consensus+Innovations distributed kalman filter with optimized gains,” *IEEE Transactions on Signal Processing (TSP)*, vol. 65, no. 2, pp. 467–481, 2017.
- [24] O. Hlinka, O. Slučiak, F. Hlawatsch, P. M. Djurić, and M. Rupp, “Likelihood consensus and its application to distributed particle filtering,” *IEEE Transactions on Signal Processing (TSP)*, vol. 60, no. 8, pp. 4334–4349, 2012.
- [25] F. Dellaert, M. Kaess *et al.*, “Factor graphs for robot perception,” *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [26] V. Mengers and O. Brock, “No plan but everything under control: Robustly solving sequential tasks with dynamically composed gradient descent,” *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 90–96, 2025.
- [27] V. Mengers, N. Roth, O. Brock, K. Obermayer, and M. Rolfs, “A robotics-inspired scanpath model reveals the importance of uncertainty and semantic object cues for gaze guidance in dynamic scenes,” *Journal of Vision*, vol. 25, no. 2, art. no. 6, 2025.
- [28] A. Battaje, A. Godinez, N. M. Hanning, M. Rolfs, and O. Brock, “An information processing pattern from robotics predicts unknown properties of the human visual system,” *bioRxiv*, no. 2024.06.20.599814, 2024.
- [29] P. E. Jacob, F. Lindsten, and T. B. Schön, “Coupling of particle filters,” *arXiv*, no. 1606.01156, 2016.
- [30] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price *et al.*, “The epic-kitchens dataset: Collection, challenges and baselines,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 11, pp. 4125–4141, 2020.
- [31] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu *et al.*, “Ego4d: Around the world in 3,000 hours of egocentric video,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 18 995–19 012.