

Object-Centric Action-Enhanced Representations for Robot Visuo-Motor Policy Learning

Nikos Giannakakis^{1†}, Argyris Manetas^{2†}, Panagiotis P. Filntisis², Petros Maragos^{1,2} and George Retsinas²

Abstract—Learning visual representations from observing actions to benefit robot visuo-motor policy generation is a promising direction that closely resembles human cognitive function and perception. Motivated by this, and further inspired by psychological theories suggesting that humans process scenes in an object-based fashion, we propose an object-centric encoder that performs semantic segmentation and visual representation generation in a coupled manner, unlike other works, which treat these as separate processes. To achieve this, we leverage the Slot Attention mechanism and use the SOLV [1] model, pretrained on large out-of-domain datasets, to bootstrap fine-tuning on human action video data. Through simulated robotic tasks, we demonstrate that visual representations can enhance reinforcement and imitation learning training, highlighting the effectiveness of our integrated approach for semantic segmentation and encoding. Furthermore, we show that exploiting models pretrained on out-of-domain datasets can benefit this process, and that fine-tuning on datasets depicting human actions—although still out-of-domain—can significantly improve performance due to close alignment with robotic tasks. These findings show the capability to reduce reliance on annotated or robot-specific action datasets and the potential to build on existing visual encoders to accelerate training and improve generalizability.

I. INTRODUCTION

In recent years, the intersection of machine learning and robotics have transformed our understanding of how machines perceive and interact with their environments. This has encouraged researchers to explore more deeply the underlying principles of human learning, motivation and action, delving into the field of psychology, to often guide their algorithm or system design. One such area investigates how different individuals perceive objects and other parts of the scene to construct internal representations that capture feasible actions on them—these action possibilities constitute what is referred to in literature as affordances. Specifically, according to J.J. Gibson [2], who termed this concept, objects have inherent “values” or “meanings” [3], that observers uncover through interaction and exploration, linking them to potential actions. Children, for example, perceive the shapes, sizes, colors, motions, and other characteristics of objects and

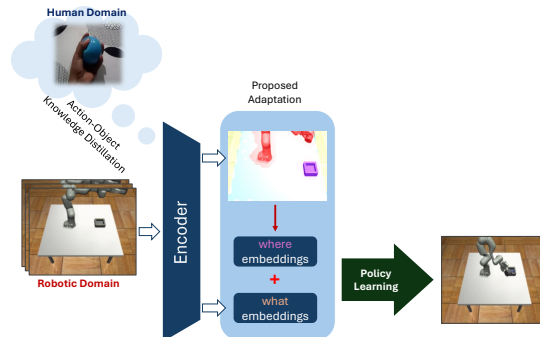


Fig. 1. **Approach Overview:** Our proposed encoder extracts “what” and “where” visual object-centric embeddings, which are combined to form scene representations. These embeddings distill action-based knowledge acquired from pretraining on human action videos to effectively guide robot policy learning.

are able to learn how things can be manipulated, e.g. moved, squeezed, stacked, and others. Although all senses participate in this process, the question of how artificial intelligence can construct *visual* representations to mimic the way humans perceive such cues, has attracted considerable attention, due to the widespread availability of cameras.

These visual features can complement explicit state representations, enhancing visuo-motor policy training in reinforcement and imitation learning scenarios, as demonstrated in previous works [4], [5]. Approaches following this paradigm typically pre-train neural networks in an unsupervised manner—using either in-domain or out-of-domain datasets, such as [6], [7], [8], [9]—before applying them to robotics tasks. This strategy has been shown to reduce training time and improve performance and generalization compared to end-to-end learning methods [10], [11], [12]. The resulting learned representations can subsequently be fine-tuned or directly employed for downstream robot manipulation tasks. Particularly when leveraging video datasets, the sequential structure naturally distills embeddings with information about *potential* dynamic behavior of parts of the scene—objects and others—, closely resembling human cognitive development through everyday interaction with the world. For example, infants intuitively lift, push and drop objects, observing the outcomes and incrementally constructing their world knowledge.

In line with these observations, an emerging direction in machine learning and computer vision is object-centric representation learning, which aims to represent complex environments in terms of discrete objects, rather than treating an entire scene as a single entity. This approach is backed by and aligns closely with the proposed “principles of grouping”

[†]equal contribution

¹School of Electrical & Computer Engineering, National Technical University of Athens, Greece ngiannakakis@mail.ntua.gr

²Robotics Institute, Athena Research Center, 15125 Marousi, Greece {a.manetas, pfilntisis, petros.maragos, george.retsinas}@athenarc.gr

This research work was supported by the project “Applied Research for Autonomous Robotic Systems” (MIS5200632) which is implemented within the framework of the National Recovery and Resilience Plan (NNRP) “Greece 2.0” (Measure: 16618- Basic and Applied Research) and is funded by the European Union-NextGenerationEU.

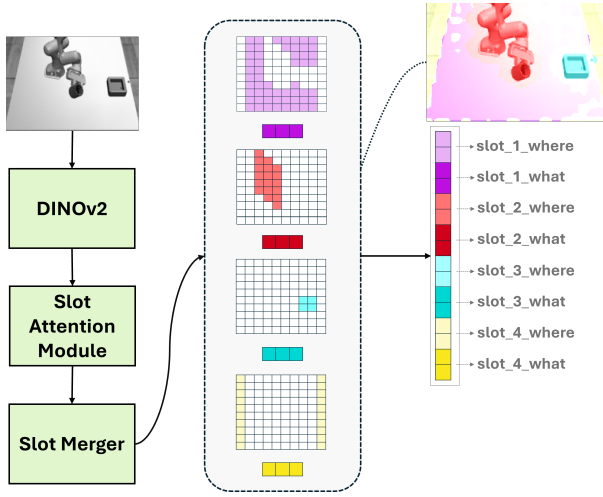


Fig. 2. **Our proposed architecture:** Image embeddings produced by DINOv2 are processed by the Slot Attention Module to generate slots. The Slot Merger then outputs four final object-specific features, referred to as “what” vectors, each corresponding to a distinct semantic region. These are combined with the associated “where” vectors, which encode spatial information derived from the attention masks, to form a unified object-centric scene representation.

from psychology and cognitive science, which describe how humans process visual signals by organizing them into distinct objects [13]. Consistent with this, experiments in developmental psychology show that infants learn very early on action-object associations; even before word associations, which become more important later in development [14]. Systems following this object-centric paradigm have shown promising results [15], [16], while also offering improved interpretability. This underscores the significance of incorporating these human development insights into the design of representation encoding methods.

Nevertheless, many object-centric methods decouple the problems of object segmentation and visual representation generation, potentially introducing propagating errors—in the case of inaccurate semantic segmentation—and requiring that all visible object classes in the frames be known in advance. To overcome these limitations, we propose an encoder capable of producing object-centric representations that capture information about both the characteristics of the objects and their position. Our system leverages the SOLV [1] framework, which builds on the slot attention mechanism of previous works [17], [18] to enable concurrent object scene segmentation and object-specific vector representation generation. Capitalizing on the pretrained nature of SOLV, we fine-tune it on a subset of the Something-Something action video dataset [7] to develop an encoder that outputs object-centric visual features, which are aggregated into a common scene representation encompassing potential actions and spatiality information of the visible objects. A high-level overview of our approach is illustrated in Figure 1. We evaluate the effectiveness of this object-centric scene representation to guide robot visuo-motor policy generation in reinforcement and imitation learning scenarios, comparing them against state-of-the-art encoders, demonstrating that our

encoder, benefits from the object-centric approach, and that knowledge acquired from human action videos transfers well to robotic tasks, significantly enhancing the robot’s efficacy.

Our key contributions can therefore be summarized as:

- We develop an object-centric visual encoder by adapting the SOLV framework for visuo-motor policy learning.
- We fine-tune the encoder on human action videos to distill action-relevant knowledge into the representations.
- We evaluate the effectiveness of these representations on imitation and reinforcement learning scenarios, outperforming state-of-the-art image encoders on the TOTO [12] benchmark.

II. PROPOSED METHOD

An outline of our proposed method can be seen in Figure 2. Our proposed methodology is designed to be domain- and task-agnostic, allowing the encoding of object- and scene-specific information into representations that can be utilized in reinforcement and imitation learning scenarios for relevant robotic task completion. Given our object-centric focus, our work builds on recent advances in slot attention mechanisms and SOLV (Self-Supervised Object-Centric Learning for Videos), a self-supervised framework capable of segmenting and tracking multiple objects in real-world video sequences, without the need for cumbersome annotations or additional modalities such as depth and optical flow. In the following paragraphs, we present the core ideas that shaped our approach and the main components of our architecture.

A key component of our architecture is the slot attention module, introduced in [17] and later extended in [18] to enforce transformation-invariant visual representations. This mechanism assigns regions of an image or video frame to a set of “slots”, which are learnable vectors that represent distinct objects or components in the scene. Each slot is updated iteratively by attending to the input features, allowing it to bind to a specific object or semantic region, thus producing object-centric representations without supervision.

SOLV [1] leverages the slot attention mechanism introduced in [18] to produce spatial slots that are associated across video frames, resulting in temporally-aware object representations. This makes SOLV a natural choice for our setting, where understanding potential object behavior is essential. Specifically, by fine-tuning on videos depicting human actions, we aim to encourage the model to capture action-relevant features, such as possible object dynamics and affordances, thereby making the resulting representations more effective in robot policy learning.

Accordingly, we slightly augment the definition of slots and modify the SOLV architecture to create an encoder that generates suitable object-centric representations for visuo-motor policy learning following the same principles as in [15], [16], for every video frame. Our encoder implementation consists of the following SOLV components¹:

- **Visual Backbone:** The frozen DINOv2 [19] encoder extracts visual embeddings from the input video frames.

¹for more details on the SOLV components and their implementation see the original paper [1]



Fig. 3. **Experimental Environment:** We simulate the Franka Emika Panda robot arm pouring task within the TOTO framework. The goal is to transfer as many spheres as possible from the cup to the container.

- **Slot Attention:** The slot attention mechanism binds slots into semantically different regions for each frame, following the paradigm of [18], resulting in object segmentation. Shared initialization is used across time steps to ensure temporal consistency.
- **Slot Merging:** This module merges slots based on their similarity, through a configured Agglomerative Clustering (AC) algorithm, which outputs 4 final slots — a quantity which through our experimentation ensured the best mask granularity avoiding over- and under-segmentation (see Section III-E).

Unlike SOLV, we do not include a temporary binding step in the final encoder (referred to as the *Temporal Binding Component* in the original SOLV paper), as we argue that the slot vectors produced by the slot attention module already capture non-trivial temporal information, akin to a video-to-image distillation process, and thus some action-object affordance associations, which we find sufficient for the short-horizon target robotic tasks. Since the slot attention module is pretrained in a system that processes temporal information, the resulting slots are optimized for inter-frame attention, and hence inherently capture temporal signals.

We augment the final slots by introducing an explicit spatial component, following [16]. The initial four slots, produced by iterative slot binding with dimension $D_{\text{what}} = 128$, encode objects’ intrinsic feature cues, forming what we term the “what” component. Since these representations will be used to generate the visuo-motor policy of robots, adding explicit information about object position is intuitive. We achieve this by processing the attention masks of each slot, which are resized to dimensions $H'_{\text{att}} \times W'_{\text{att}} = 10 \times 10$ from $H_{\text{att}} \times W_{\text{att}} = 24 \times 36$ via bilinear interpolation. The scaled softmax activation function is applied to the resized attention masks, amplifying regions with initially high values — an adjustment that improved performance in our experiments. The resulting masks are flattened to produce localization vectors, hereinafter referred to as the “where” vectors, which combined with the “what” vectors, form the final object-centric representation of the image, with a shape of $(D_{\text{what}} + D_{\text{where}}) \times \text{“number of slots”} = (128 + 100) \times 4 = 912$.

III. EXPERIMENTAL EVALUATION

To study the ability of our encoder’s object-centric representations in guiding a robotic learning task, we incorporate them into reinforcement and imitation learning scenarios

as part of the state space, and compare their effectiveness against representations from other state-of-the-art image encoders that lack action-object association knowledge.

A. Pretraining and Fine-Tuning

In our approach, we leverage SOLV’s ability to generate object-centric representations, acquired through training on large-scale datasets for semantic segmentation, and adapt it to capture action-object associations by fine-tuning its pretrained weights on video datasets containing such interactions. Specifically, to improve its capacity for constructing object-centric representations relevant to robot-specific tasks — such as rolling, squeezing, and other interactions — we fine-tune it for 100 epochs on a subset of the Something-Something [7] dataset, which includes such actions. The resulting network layers are then integrated into the final encoder described in Section II. Although human videos from the Something-Something dataset may introduce limitations due to the human-robot embodiment gap, we use them for their accessibility and to demonstrate that even out-of-domain data can meaningfully improve robot policy learning.

B. Experimental Environment

We evaluate our approach on the Train Offline, Test Online (TOTO) [12] benchmark, which provides a protocol to assess both visual representations and policy learning. TOTO is a robotics benchmark that addresses the lack of standardization across research centers by providing remote research teams with access to shared robotic hardware and an open-source dataset of tasks for offline training. The dataset includes noise-augmented trajectories collected through robot teleoperation, and trajectories generated by Behavior Cloning (BC) agents. The benchmark focuses on two manipulation tasks: (i) pouring and (ii) scooping, which, despite being commonplace for humans, remain challenging to robots, due to varying initial conditions and object types, thus offering valuable insights into the methods under evaluation.

The TOTO benchmark includes a simulation package, which uses the MuJoCo physics simulation engine [20], and simulates the 7 DoF Franka Emika Panda robot arm [21] (see Figure 3). Joints are restricted within a specified range of positions, offering the two-fold advantage of a smaller control space and improved safety in real-world deployments.

In our test setup, the simulated robot arm is initialized in random positions, holding a cup filled with twelve small

spheres, with the goal of pouring as many of them into another container. A successful manipulation corresponds to experiments where at least one sphere is deposited in the container, while the reward is the percentage of the total spheres being successfully deposited. The training set used consists of 82 successful and 21 failed trajectories. In the following subsections, we present the results obtained from our imitation and reinforcement learning experiments.

C. Imitation Learning

In our first approach, we employ Imitation Learning (IL) to train the robot’s policy, using expert demonstrations. Unlike reinforcement learning, where an agent learns by exploring and interacting with the environment, IL attempts to replicate expert behavior, avoiding the high costs, risks, and time demands of real-world exploration. Among IL methods, we adopt Behavioral Cloning (BC) [22], one of the earliest and simplest approaches. BC formulates policy learning as a supervised learning problem, where the agent directly maps states to expert actions without requiring exploration.

We train separate BC agents for each encoder under evaluation—our own, and those we compare against—using the 82 successful out of the 103 total trajectories (i.e., those with a reward above zero), with inputs produced by the corresponding encoder. During training the BC algorithm trains a neural network to serve as the agent’s policy, receiving as input the encoders’ result concatenated with the robot’s current joint angles. The network predicts joint angle targets, which are fed to the MuJoCo controller to move the simulated robot arm. Therefore, the input dimension is $R_{\text{dim}} + 7$, where R_{dim} is the size of the representation vector, while the output dimension is $\text{out}_{\text{dim}} = 7 \times h$, where h (in our experiments $h = 10$) represents the prediction horizon.

Our encoder, A-SOLV (AE), is compared against state-of-the-art image encoders: BYOL [23], CLIP [24], DINOv2 [19], MoCo [25], and ResNet50 [26]. We also evaluate our adapted encoder architecture using the pretrained SOLV weights (referred to as A-SOLV), trained on the Youtube-VIS 2019 [27] real-world video dataset, to assess the impact of action-object associations introduced by fine-tuning on action videos from a subset of the Something-Something dataset. Each BC agent is trained for 80 epochs with a 0.001 learning rate. For each encoder, we train 5 BC agents and evaluate them on 100 randomly initialized trajectories to reduce stochasticity and obtain more reliable performance estimates. The results are summarized in Table I, showing that our encoder consistently achieves the highest reward and success rate.

D. Implicit Q-Learning

The Imitation Learning experiments in subsection III-C demonstrated the effectiveness of object-centric representations in improving robot policies learned from expert demonstrations, but to further evaluate their impact in standard reinforcement learning algorithms, we conduct experiments using Implicit Q-Learning (IQL) [28], an offline reinforcement learning algorithm that is designed to improve

TABLE I
Comparison of Pre-trained Visual Representation Models in Training Behavior Cloning Agents for the TOTO Benchmark Simulated Pouring Task. *Adapted SOLV, **Adapted SOLV (Action Enhanced).

Encoder	Dataset	Repr. Size	Success Rate	Mean Reward
BYOL	ImageNet [8]	512	0.46 ± 0.05	17.48 ± 2.40
CLIP	WIT (WebImageText) [24]	512	0.49 ± 0.06	18.61 ± 4.09
DINOv2	LVD-142M [19]	768	0.55 ± 0.04	18.72 ± 1.78
MoCo	ImageNet [8]	2048	0.31 ± 0.04	9.40 ± 2.16
ResNet50	ImageNet [8]	2048	0.56 ± 0.11	21.15 ± 5.60
<hr/>				
A-SOLV*	Youtube-VIS 2019 [27]	912	0.46 ± 0.06	15.07 ± 2.24
A-SOLV (AE)**	Something-Something [7]	912	0.61 ± 0.04	25.25 ± 2.95

TABLE II
Encoder Comparison in the Implicit Q-Learning Scenario. *Adapted SOLV, **Adapted SOLV (Action Enhanced).

Encoder	Representation Size	Success Rate	Mean Reward
BYOL	512	0.34	10.83
ResNet50	2048	0.46	14.5
<hr/>			
A-SOLV*	912	0.43	18.0
A-SOLV (AE)**	912	0.57	24.41

stability and performance by avoiding out-of-distribution Q-value queries. IQL achieves this by employing expectile regression to approximate optimal Q-values, mitigating extrapolation errors from unseen actions present in previous methods, delivering state-of-the-art results, thus rendering this approach particularly well-suited for our offline scenario.

This time, our training dataset consists of all 103 TOTO trajectories, which we use to train an IQL agent. We compare the effectiveness of our encoder against BYOL, ResNet50, and A-SOLV. The results are presented in Table II.

E. Number of Slots

To determine the optimal number of slots output by the slot merger module for achieving the best segmentation resolution—with the least over- or under-segmentation of objects—we conducted experiments with the AC algorithm configured to produce 4, 6, and 8 slots. Based on qualitative analysis of the resulting segmentations (see Figure 4) and a re-evaluation of the imitation learning experiments (see Figure 5), we found that using 4 slots yielded the best results.

F. Contribution of the “where” Components

Although the initial slots encode intrinsic object properties (such as possible dynamics or affordances) distilled from temporal cues observed in the videos during pretraining, they are still produced by the invariant slot attention mechanism. As a result, they may lack crucial spatial information needed for effective reinforcement learning during action execution. We addressed this limitation in Section II by incorporating “where” vectors through concatenation. To evaluate their contribution, we compare imitation learning performance with and without their inclusion. The results, shown in Figure 6, highlight the significance of incorporating spatial information in the slot vectors.



Fig. 4. Simulation frames from the TOTO pouring task, segmented by the Slot Attention masks of the SOLV model, with the Slot Merger module outputting 4, 6, and 8 slots respectively (from left to right).




Fig. 5. Success rates and mean rewards of different number of slots configurations in the TOTO pouring simulation task. Each variant was trained five times and evaluated across 100 trajectories.

G. Results Discussion

In our experimental evaluation, the representations produced by our encoder consistently outperformed those of other state-of-the-art encoders, both in terms of success rate and mean reward. In the imitation learning setting (Table I), it achieved a mean success rate of 0.61 and a mean reward of 25.25, outperforming the closest competitor (ResNet50) by 8.93% and 19.39% in the respective metrics. Additionally, our method exhibited the lowest variance in success rate (alongside DINOv2) and comparable variance in mean reward—much smaller than that of the next best-performing encoder, ResNet50—suggesting consistency across runs. Moreover, our method uses a more compact representation than the second-best model (ResNet50), and a comparable size to the next best-performing models (CLIP and DINOv2). Similar observations hold for the Implicit Q-Learning scenario, where our method achieved a success rate of 0.57 (11.76% improvement over the next best), and a mean reward of 24.41 (19.24% improvement compared to the next best).

We attribute this overall performance improvement to two main factors. First, our encoder, having been trained on videos, can capture temporally-infused features that result in richer scene representations, enhancing policy learning. Second, fine-tuning on videos depicting human-performed actions—although still out-of-domain—provides a significant advantage, as evidenced by the performance gap with the SOLV-adapted encoder, which was trained on generic out-of-

domain videos. These two conditions enabled the constructed representations to capture action-relevant object features, such as affordances and potential dynamics, thus improving the robot’s visuo-motor policy decisions. It should be noted that the design choices discussed in Sections III-E and III-F had a tangible impact on the results, as the chosen number of slots and the addition of spatial components ensured, through our experimentation, the most optimal segmentation and the best task performance, respectively.

IV. CONCLUSIONS

In this work, we have demonstrated that image representations, derived from aggregating action-object-centric features—produced via a slot attention mechanism, pre-trained in a self-supervised manner on an action database—enhance the effectiveness of reinforcement and imitation learning algorithms. Using our method, robots can construct action-enhanced representations without additional training, enabling better goal-directed behavior, thus significantly improving robots’ open-ended capabilities in a way that aligns with human cognitive development. Furthermore, our results highlight the potential of large-scale, directly-available action videos as valuable resources for visual encoder training.

In future work, we seek to explore segmentation methods that scale well to a larger number of regions, and also to evaluate our approach in real-world robotic scenarios and on more complex, longer-horizon tasks.



Fig. 6. Success rates and mean rewards comparing “what” and unified “what” + “where” representations. Each configuration was trained five times and evaluated across 100 trajectories.

REFERENCES

- [1] G. Aydemir, W. Xie, and F. Guneş, “Self-supervised Object-Centric Learning for Videos,” in *Proc. of Advances in Neural Information Processing Systems*, vol. 36, New Orleans (LA), USA, Dec. 2023, pp. 32 879–32 899.
- [2] J. J. Gibson, *The Senses Considered as Perceptual Systems*. Boston (MA), USA: Houghton Mifflin, 1966.
- [3] L. K. Le Goff, O. Yaakoubi, A. Coninx, and S. Doncieux, “Building an Affordances Map with Interactive Perception,” *Frontiers in Neuro-robotics*, vol. 16, p. 504459, May 2022.
- [4] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang, “VIP: Towards Universal Visual Reward and Representation via Value-Implicit Pre-Training,” *arXiv:2210.00030*, Sep. 2022.
- [5] Z. Cui, H. Pan, A. Iyer, S. Haldar, and L. Pinto, “DynaMo: In-Domain Dynamics Pretraining for Visuo-Motor Control,” in *Proc. of Advances in Neural Information Processing Systems*, Vancouver (BC), Canada, Dec. 2024, pp. 33 933–33 961.
- [6] 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 *Proc. of IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, New Orleans (LA), USA, Jun. 2022, pp. 18 995–19 012.
- [7] R. Goyal, S. Ebrahimi Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, F. Hoppe, C. Thureau, I. Bax, and R. Memisevic, “The “something something” video database for learning and evaluating visual common sense,” in *Proc. of IEEE Intl. Conf. on Computer Vision*, Venice, Italy, Oct. 2017, pp. 5843–5851.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, Miami (FL), USA, Jun. 2009, pp. 248–255.
- [9] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn, “RoboNet: Large-Scale Multi-Robot Learning,” *arXiv:1910.11215*, Oct. 2019.
- [10] A. Iosifidis and A. Tefas, *Deep Learning for Robot Perception and Cognition*. Cambridge (MA), USA: Academic Press, 2022.
- [11] A. Sax, J. O. Zhang, B. Emi, A. Zamir, S. Savarese, L. Guibas, and J. Malik, “Learning to Navigate Using Mid-Level Visual Priors,” *arXiv:1912.11121*, Dec. 2019.
- [12] G. Zhou, V. Dean, M. K. Srirama, A. Rajeswaran, J. Pari, K. Hatch, A. Jain, T. Yu, P. Abbeel, L. Pinto, C. Finn, and A. Gupta, “Train Offline, Test Online: A Real Robot Learning Benchmark,” in *Proc. of IEEE Intl. Conf. on Robotics and Automation*, London, UK, May-Jun. 2023, pp. 9197–9203.
- [13] S. E. Palmer, *Vision Science: Photons to Phenomenology*. Cambridge (MA), USA: MIT press, 1999.
- [14] S. F. Eiteljoerge, M. Adam, B. Elsner, and N. Mani, “Word-object and action-object association learning across early development,” *PLOS One*, vol. 14, no. 8, p. e0220317, Aug. 2019.
- [15] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, “VIOLA: Imitation Learning for Vision-Based Manipulation with Object Proposal Priors,” in *Proc. of Conf. on Robot Learning*, Atlanta (GA), USA, Nov. 2023, pp. 1199–1210.
- [16] J. Shi, J. Qian, Y. J. Ma, and D. Jayaraman, “Composing Pre-Trained Object-Centric Representations for Robotics From “What” and “Where” Foundation Models,” in *Proc. of IEEE Intl. Conf. on Robotics and Automation*, Yokohama, Japan, May 2024, pp. 15 424–15 432.
- [17] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf, “Object-Centric Learning with Slot Attention,” in *Proc. of Advances in Neural Information Processing Systems*, vol. 33, Virtual, Dec. 2020, pp. 11 525–11 538.
- [18] O. Biza, S. Van Steenkiste, M. S. Sajjadi, G. F. Elsayed, A. Mahendran, and T. Kipf, “Invariant slot attention: Object discovery with slot-centric reference frames,” in *Proc. of Intl. Conf. on Machine Learning*, Honolulu (HI), USA, Jul. 2023, pp. 2507–2527.
- [19] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, “DINOv2: Learning Robust Visual Features without Supervision,” *arXiv:2304.07193*, Apr. 2023.
- [20] E. Todorov, T. Erez, and Y. Tassa, “MuJoCo: A physics engine for model-based control,” in *Proc. of IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, Vilamoura, Portugal, Oct. 2012.
- [21] S. Haddadin, S. Parusel, L. Johannsmeier, S. Golz, S. Gabl, F. Walch, M. Sabaghian, C. Jähne, L. Hausperger, and S. Haddadin, “The Franka Emika Robot: A Reference Platform for Robotics Research and Education,” *IEEE Robotics & Automation Magazine*, vol. 29, no. 2, pp. 46–64, Jun. 2022.
- [22] F. Torabi, G. Warnell, and P. Stone, “Behavioral Cloning from Observation,” in *Proc. of Intl. Joint Conf. on Artificial Intelligence*, Stockholm, Sweden, Jul. 2018, pp. 4950–4957.
- [23] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. D. Guo, M. Gheshlaghi Azar, B. Piot, M. Valko, and R. Munos, “Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning,” in *Proc. of Advances in Neural Information Processing Systems*, Virtual, Dec. 2020, pp. 21 271–21 284.
- [24] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning Transferable Visual Models from Natural Language Supervision,” in *Proc. of Intl. Conf. on Machine Learning*, Virtual, Jul. 2021, pp. 8748–8763.
- [25] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum Contrast for Unsupervised Visual Representation Learning,” in *Proc. of IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, Seattle (WA), USA, Jun. 2020, pp. 9729–9738.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, Las Vegas (NV), USA, Jun. 2016, pp. 770–778.
- [27] L. Yang, Y. Fan, and N. Xu, “Video Instance Segmentation,” in *Proc. of IEEE/CVF Intl. Conf. on Computer Vision*, Seoul, Korea, October 2019, pp. 5188–5197.
- [28] I. Kostrikov, A. Nair, and S. Levine, “Offline Reinforcement Learning with Implicit Q-Learning,” *arXiv:2110.06169*, Oct. 2021.