On Representation Redundancy and Disentanglement in Deep Reinforcement Learning for Robot Navigation

Zhiyuan Nan*, Yifang Shi†, Haewoon Nam*†

*Hanyang University, Department of Electrical and Electronic Engineering, Ansan, Korea

†Hangzhou Dianzi University, School of Automation, Hangzhou, China

Emails: namjw@hanyang.ac.kr, syf2008@hdu.edu.cn, hnam@hanyang.ac.kr

†Corresponding author: hnam@hanyang.ac.kr

***Abstract*—Deep reinforcement learning (DRL) for robot navigation relies heavily on the quality of representations extracted from high-dimensional visual and geometric observations. However, RGB, depth, and LiDAR inputs often contain substantial appearance-related redundancy, while navigation behavior primarily depends on geometric semantics. When a shared encoder is jointly optimized by critic objectives and auxiliary representation learning losses, these heterogeneous learning signals tend to mix task-relevant and task-irrelevant factors within the same feature space, leading to representation entanglement. Such entanglement increases the variance of value estimation, reduces policy robustness to appearance changes, and limits generalization across environments.**

This paper provides a conceptual and structural analysis of how redundancy and entanglement emerge in current DRL frameworks for navigation, and why disentangling navigation semantics from appearance factors is fundamentally important. The discussion highlights key challenges arising from shared-encoder training and outlines potential benefits of adopting disentangled representations for improving stability, robustness, and transferability in DRL-based navigation systems.

***Index Terms*—Deep Reinforcement learning, robot navigation, representation learning, Disentanglement learning.**

I. Introduction

With the rapid progress of deep reinforcement learning (DRL) in robotics, vision-based and multimodal autonomous navigation has emerged as a central research topic. Mobile robots operating in real-world environments must process high-dimensional observations collected from RGB cameras, depth sensors, and LiDAR units. Although these sensing modalities provide rich environmental information, they also contain large amounts of task-irrelevant variations, such as changes in illumination, texture, material appearance, and sensor noise. These appearance-related factors significantly increase observational redundancy, thereby making effective policy learning more challenging in navigation tasks.

Recent studies have combined DRL with auxiliary representation-learning techniques—including reconstruction [1], [2], self-supervised contrastive learning [2]–[6], and predictive representation learning [2], [6]–[8]—to improve training stability and sample efficiency. In typical architectures, the encoder is jointly updated by critic losses and auxiliary losses, while the actor consumes a stop-gradient version of the encoder features to avoid destabilizing gradients. This paradigm has shown strong performance on standard visuomotor benchmarks. However, because multiple heterogeneous learning objectives simultaneously update the same encoder, the resulting representations often entangle navigation-relevant geometric semantics with appearance-induced disturbances, a phenomenon known as representation entanglement. Such entanglement can destabilize value estimation and make the learned policy sensitive to appearance changes, ultimately compromising reliability and generalization across environments and platforms.

These representation issues are especially pronounced in mobile robot navigation. Real-world environments exhibit substantial variability in building layouts, lighting conditions, floor materials, dynamic occlusions, and even domain shifts between simulation and reality. If the learned representation fails to separate task-relevant from task-irrelevant factors, the policy may perform well only in the training environment but degrade significantly—or behave unpredictably—under new visual conditions or in unseen scenes.

Motivated by these challenges, this paper provides a conceptual and structural analysis of two fundamental representation issues in DRL-based robot navigation: observational redundancy and representation entanglement. Specifically, we examine (i) why navigation observations inherently possess high redundancy, (ii) how entanglement emerges when the encoder is jointly optimized by critic and auxiliary objectives, and (iii) why incorporating some form of representation disentanglement is structurally beneficial for enhancing robustness and cross-scene transferability. Rather than proposing a new algorithm, this work aims to offer motivation and analytical grounding for future research on task-aware representation design and disentanglement methods in DRL-based navigation.

II. Redundancy in Navigation Observations

Mobile robots operating in real-world environments must process large volumes of data from multiple sensing modalities, including RGB images, depth maps, and LiDAR scans. Although these modalities provide rich environmental cues, a substantial portion of their observed dimensions is not directly relevant to navigation behavior. For example, RGB inputs typically contain appearance-related factors—such as texture, material properties, illumination, and color variations—that do not correspond to traversability or goal-related geometry. Depth and LiDAR observations likewise include numerous distant points, noise artifacts, and fine-grained geometric details that contribute little to local decision-making, thereby introducing structural redundancy into the high-dimensional observation space.

For reinforcement learning–based navigation systems, such redundancy implies that the encoder must extract behavior-critical semantics from an extremely high-dimensional and largely irrelevant observation space, making representation learning considerably more difficult [9]. Moreover, modern visual DRL frameworks often incorporate reconstruction, contrastive learning, or predictive auxiliary tasks to improve training stability. However, these auxiliary objectives tend to encourage the encoder to preserve appearance details or visual consistency, causing it to retain even more task-irrelevant information and thereby amplifying redundancy in the learned representation. This is particularly evident in RGB-based inputs, where color and material variations occupy significant encoder capacity despite being peripheral to navigation decisions.

In multimodal settings, redundancy may further accumulate across modalities. RGB inputs emphasize appearance patterns, while depth maps provide dense geometric structure—yet many dimensions in both modalities are unrelated to obstacle avoidance or short-horizon path planning. When these modalities are jointly processed by a shared encoder, the model must cope not only with large-scale high-dimensional input, but also with modality-dependent statistical discrepancies that complicate feature fusion. As a result, the learned representation often becomes mixed and highly redundant, increasing susceptibility to noise, slowing convergence, and reducing generalization to new environments.

In summary, the high redundancy present in navigation observations arises from large input dimensionality, substantial appearance variability, repeated geometric information, and pervasive sensor noise. Redundancy is not merely a property of raw perception; it also consumes encoder capacity and fundamentally increases the difficulty of learning task-relevant representations. These factors lay the structural groundwork for the representation entanglement discussed in the following section.

III. Representation Entanglement in DRL for Navigation

As shown in Fig. 1, in visual reinforcement learning systems for navigation, the policy and value networks typically share a common encoder to improve training efficiency. However, this shared structure must simultaneously serve multiple learning objectives: the critic objective drives the encoder to extract reward-relevant geometric semantics, while auxiliary objectives—such as contrastive learning, prediction, or reconstruction—encourage the encoder to preserve appearance-related information that supports visual consistency or reconstructability. When these heterogeneous objectives act on the same encoder parameters, task-relevant and task-irrelevant factors become difficult to separate, giving rise to representation entanglement in the learned feature space.

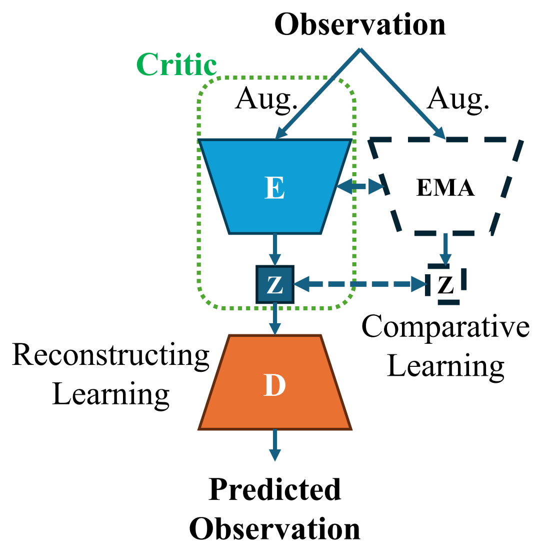

Fig. 1. A common encoder-sharing paradigm in visual DRL navigation, where critic and auxiliary objectives jointly update the encoder. Such multi-objective optimization often leads to representation entanglement.

More concretely, as shown in Eqs. (1) and (2), the feature vector $\mathbf{h}$ in the critic branch is passed through $\mathbf{f_c}$ to produce the value-related latent representation $\mathbf{Z}_c$. Critic supervision relies on geometric cues—such as obstacle boundaries, goal direction, and local traversability—and thus requires the encoder to emphasize these semantics. In contrast, auxiliary tasks promote the retention of textures, colors, illumination patterns, and other appearance attributes to satisfy reconstruction or contrastive matching. When both objectives jointly update the CNN encoder, geometric and appearance information are inevitably mixed within the same representation, preventing the encoder from forming a clean, task-aligned semantic structure.

$$\mathbf{obs} \xrightarrow{f_{\text{cnn}}} \mathbf{h} \xrightarrow{f_c} \mathbf{Z}_c \xrightarrow{f_Q} Q(s,a), \tag{1}$$

$$\mathbf{obs} \xrightarrow{f_{\text{cnn}}} \mathbf{h} \xrightarrow{f_c} \mathbf{Z}_c \rightarrow \mathcal{L}_{\text{aux}}, \tag{2}$$

As shown in Eq. 3, although the actor branch does not propagate gradients back to the encoder, the actor still consumes $\mathbf{h_{sg}}$, a stop-gradient version of the same shared representation shaped jointly by the critic and auxiliary losses. As a result, policy learning inherently inherits the biases introduced by representation entanglement. For instance, the policy may inadvertently learn to rely on texture or lighting patterns as decision cues—features that can change drastically across environments—leading to unstable behavior during cross-scene evaluation. Meanwhile, the entangled representation also disrupts value estimation, making the critic more sensitive to appearance noise and increasing training instability.

$$\mathbf{obs} \xrightarrow[\text{stopgrad}]{f_{\text{cnn}}} \mathbf{h}_{\text{sg}} \xrightarrow{f_a} \mathbf{Z}_a \xrightarrow{f_\pi} \pi(a|s), \tag{3}$$

Navigation tasks are particularly vulnerable to representation entanglement because their core decisions depend on stable geometric semantics, whereas appearance variations, sensor noise, and real-world physical inconsistencies are widespread and often unavoidable [10]. When such disturbances leak into the encoder's representation, both the policy and the critic may suffer degradation manifested as slower learning, heightened sensitivity to visual changes, poor generalization across environments, and potentially unsafe or unpredictable behavior during real-world deployment.

Thus, representation entanglement is not an incidental byproduct of training details but a structural consequence of jointly optimizing a shared encoder with multiple, and often conflicting, objectives. It fundamentally limits the stability of DRL training in navigation tasks and constrains the reliability of learned policies in real-world applications. These issues provide strong motivation for the disentanglement-based perspective introduced in the next section.

IV. Why Disentanglement is Desirable for Navigation

The previous section examined how representation entanglement arises when a shared encoder is jointly optimized under multiple objectives. In robot navigation, the negative effects of such entanglement are far more pronounced than in general visual tasks, as reliable navigation requires not only stable geometric representations but also robustness to appearance variations, cross-scene differences, and sensor noise. These challenges highlight the need to reconsider representation learning from a structural perspective, ensuring that task-relevant and task-irrelevant factors are appropriately separated within the feature space [2]. This conceptual disentanglement structure is illustrated in Fig. 2.

Navigation inherently depends on geometric semantics—such as the topology of traversable regions, the layout of obstacles, and the relative spatial relationship between the robot and its goal. These factors directly determine action selection. However, visual observations (e.g., RGB and depth) simultaneously contain abundant appearance information, including textures, materials, illumination, colors, and sensor artifacts. When the shared encoder is jointly shaped by critic and auxiliary objectives, these appearance factors are often embedded into the same feature representation as navigation-relevant geometry. As a result, value estimation becomes susceptible to interference, yielding higher variance or even spurious correlations. If the policy branch further relies on such appearance-biased features, its performance may degrade significantly under illumination shifts or scene changes, compromising robustness.

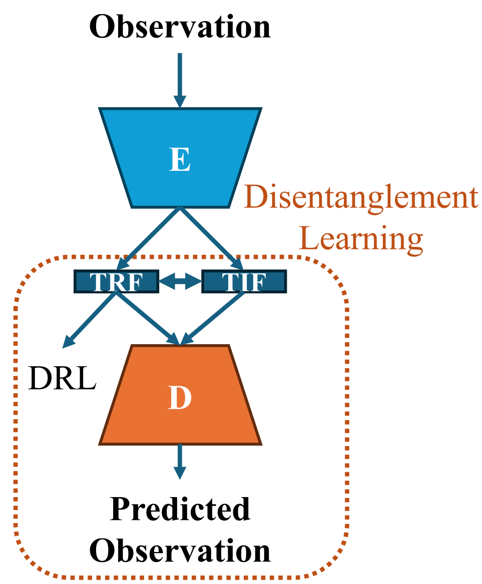

Fig. 2. A conceptual illustration of disentangled representation structure for navigation, where geometry-related semantics and appearance-related factors occupy partially separated feature subspaces. Such separation helps reduce representation entanglement and improves robustness to appearance variations.

Introducing a degree of representation disentanglement—such that geometric semantics and appearance factors occupy partially separated regions of the feature space—can substantially mitigate these issues. First, disentanglement allows the critic to perform value estimation on a cleaner and more stable representation, reducing gradient variance and improving convergence. Second, if the representation consumed by the actor emphasizes geometry rather than appearance, the resulting policy becomes less sensitive to texture or lighting variations, thereby improving cross-scene generalization. Moreover, in multimodal navigation systems where RGB, depth, and LiDAR data exhibit high dimensionality and substantial redundancy, disentanglement helps reduce cross-modality interference and enables different sensing modalities to contribute to the representation in a more structured manner.

Overall, representation disentanglement is not intended to increase network capacity or introduce additional supervision, but to address the structural bottlenecks inherent in navigation tasks—namely representation redundancy and entanglement. By constructing a representation space that is more explicitly aligned with navigation semantics, DRL-based navigation systems can achieve greater stability, robustness, and transferability in real-world environments, forming a foundation for tackling more complex navigation challenges.

V. Conclusion

This paper has analyzed key representation challenges faced by deep reinforcement learning in robot navigation, with a particular focus on observational redundancy and representation entanglement arising from shared-encoder architectures.

Visual and geometric observations in navigation often contain substantial task-irrelevant variation, making the encoder susceptible to absorbing appearance noise and environmental perturbations when jointly optimized by multiple objectives. Such entanglement destabilizes critic value estimation and further weakens policy robustness and cross-scene generalization. Through this systematic examination, we highlight the importance and necessity of incorporating representation disentanglement into DRL-based navigation systems.

Representation disentanglement is not intended to increase model complexity, but to provide a cleaner and more stable semantic space at the structural level, allowing reinforcement learning to focus more effectively on geometry that is truly relevant to decision-making. Such disentangled representations reduce critic variance, improve training stability, and enhance policy resilience to appearance shifts and domain changes, thereby strengthening the reliability of navigation systems deployed in real-world environments.

Future research may further explore disentanglement mechanisms tailored to multimodal sensory inputs, generalization strategies for cross-scene navigation, and lightweight representation models that balance efficiency with expressiveness. We believe that addressing representation redundancy and entanglement will remain a key direction for advancing the performance and robustness of visual DRL in real-world robotic navigation.

Acknowledgment

This work was supported in part by the Ministry of Science and Information and Communication Technology (MSIT), South Korea, through the Information Technology Research Center (ITRC) Support Program supervised by the Institute for Information and Communications Technology Planning and Evaluation (IITP) under Grant IITP-2023-RS-2023-00258639.

References

[1] D. Yarats, A. Zhang, I. Kostrikov, B. Amos, J. Pineau, and R. Fergus, "Improving Sample Efficiency in Model-Free Reinforcement Learning from Images," Jul. 2020, arXiv:1910.01741 [cs].

[2] H. Xu, P. Peng, G. Tan, Y. Li, X. Xu, and Y. Tian, "DMR: Decomposed Multi-Modality Representations for Frames and Events Fusion in Visual Reinforcement Learning," 2024, pp. 26 508–26 518.

[3] A. Srinivas, M. Laskin, and P. Abbeel, "CURL: Contrastive Unsupervised Representations for Reinforcement Learning," Sep. 2020, arXiv:2004.04136 [cs].

[4] T. Yu, Z. Zhang, C. Lan, Y. Lu, and Z. Chen, "Mask-based Latent Reconstruction for Reinforcement Learning," Oct. 2022, arXiv:2201.12096 [cs].

[5] I. Kostrikov, D. Yarats, and R. Fergus, "Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels," Mar. 2021, arXiv:2004.13649 [cs].

[6] J. Jiang, Y. Yang, Y. Deng, C. Ma, and J. Zhang, "BEVNav: Robot Autonomous Navigation via Spatial-Temporal Contrastive Learning in Bird's-Eye View," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 10 796–10 802, Feb. 2024.

[7] T. Yu, C. Lan, W. Zeng, M. Feng, Z. Zhang, and Z. Chen, "PlayVirtual: Augmenting Cycle-Consistent Virtual Trajectories for Reinforcement Learning," Oct. 2021, arXiv:2106.04152 [cs].

[8] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. Courville, and P. Bachman, "Data-Efficient Reinforcement Learning with Self-Predictive Representations," Oct. 2020.

[9] W. Huang, Y. Zhou, X. He, and C. Lv, "Goal-Guided Transformer-Enabled Reinforcement Learning for Efficient Autonomous Navigation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 2, pp. 1832–1845, Feb. 2024, arXiv:2301.00362 [cs].

[10] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: a Survey," in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, Dec. 2020, pp. 737–744, arXiv:2009.13303 [cs].