seq-JEPA: Autoregressive Predictive Learning of Invariant-Equivariant World Models

Hafez Ghaemi 1,2,3 * Eilif B. Muller 1,2,3 † Shahab Bakhtiari 1,2 †

¹Université de Montréal, ²Mila - Quebec AI Institute,

Abstract

Joint-embedding self-supervised learning (SSL) commonly relies on transformations such as data augmentation and masking to learn visual representations, a task achieved by enforcing invariance or equivariance with respect to these transformations applied to two views of an image. This dominant two-view paradigm in SSL often limits the flexibility of learned representations for downstream adaptation by creating performance trade-offs between high-level invariance-demanding tasks such as image classification and more fine-grained equivariance-related tasks. In this work, we propose seq-JEPA, a world modeling framework that introduces architectural inductive biases into joint-embedding predictive architectures to resolve this trade-off. Without relying on dual equivariance predictors or loss terms, seq-JEPA simultaneously learns two architecturally segregated representations: one equivariant to specified transformations and another invariant to them. To do so, our model processes short sequences of different views (observations) of inputs. Each encoded view is concatenated with an embedding of the relative transformation (action) that produces the next observation in the sequence. These view-action pairs are passed through a transformer encoder that outputs an aggregate representation. A predictor head then conditions this aggregate representation on the upcoming action to predict the representation of the next observation. Empirically, seq-JEPA demonstrates strong performance on both equivariant and invariant benchmarks without sacrificing one for the other. Furthermore, it excels at tasks that inherently require aggregating a sequence of observations, such as path integration across actions and predictive learning across eye movements.

1 Introduction

Self-supervised learning (SSL) in latent space has made significant progress in visual representation learning, closing the gap with supervised methods across many tasks. Most SSL methods rely on comparing two transformed views of an image and enforcing invariance to the transformations [Misra and van der Maaten, 2020, Chen et al., 2020, He et al., 2020, Dwibedi et al., 2021, HaoChen et al., 2021, Yeh et al., 2022, Caron et al., 2020, 2021, Ermolov et al., 2021, Assran et al., 2022, Zbontar et al., 2021, Bardes et al., 2022]. Another group of methods employ techniques to preserve transformation-specific information, thereby learning equivariant representations [Lee et al., 2021, Xiao et al., 2021, Park et al., 2022, Dangovski et al., 2022, Gupta et al., 2023, Garrido et al., 2023, 2024, Gupta et al., 2023, 2024, Yerxa et al., 2024].

³Centre de Recherche Azrieli du CHU Sainte-Justine

^{*}Correspondence to hafez.ghaemi@umontreal.ca

[†]Equal Contribution

 $^{^{1}}$ Code available at https://github.com/hafezgh/seq-jepa

Equivariance is a crucial representational property for downstream tasks that require fine-grained distinctions. For example, given representations that are invariant to color, it is not possible to distinguish between certain species of flowers or birds [Lee et al., 2021, Xiao et al., 2021]. Moreover, recent work has shown that equivariant representations are better aligned with neural responses in primate visual cortex and could be important for building more accurate models thereof [Yerxa et al., 2024]. While some equivariant SSL approaches have reported minor gains on tasks typically associated with invariance (e.g., classification) [Devillers and Lefort, 2022, Park et al., 2022, Gupta et al., 2023], a growing body of work highlights a fundamental trade-off between learning invariance and equivariance, i.e., models that capture equivariance-related style latents do not fare well in classification and vice versa [Garrido et al., 2023, 2024, Gupta et al., 2024, Yerxa et al., 2024, Rusak et al., 2025]. This trade-off has recently received theoretical support [Wang et al., 2024b], underscoring the need for new architectural and objective designs that can reconcile these competing goals.

In contrast to the two-view paradigm in SSL, humans and other animals rely on a *sequence* of actions and consequent observations (views) for developing appropriate visual representation during novel object learning [Harman et al., 1999, Vuilleumier et al., 2002]. For example, they recognize a 3-D object by changing their viewpoint and examining different sides of the object [Tarr et al., 1998]. Inspired by this, we introduce **seq-JEPA**, a self-supervised world modeling framework that combines joint-embedding predictive architectures [LeCun, 2022] with inductive biases for sequential processing. seq-JEPA simultaneously learns two architecturally distinct representations: one that is equivariant to a specified set of transformations, and another that is suited for invariance-demanding tasks, such as image classification.

Specifically, our framework (Figure 1) processes a short sequence of transformed views (observations) of an image. Each view is encoded and concatenated with an embedding corresponding to the relative transformation (action) that produces the next observation in the sequence. These view-action pairs are passed through a transformer encoder, *a form of learned working memory*, that outputs an aggregate representation of them. A predictor head then conditions this aggregate representation on the upcoming action to predict the representation of the next observation.

Our results demonstrate that individual encoded views in seq-JEPA become transformation/action-equivariant. Through ablations, we show that action conditioning plays a key role in promoting equivariant representation learning in the encoder network. In contrast, the aggregate representation of views, produced at the output of the transformer, becomes largely transformation/action-invariant. This emergent architectural disentanglement of invariant and equivariant representations is central to seq-JEPA's competitive performance compared to invariant and equivariant SSL methods on both categories of tasks (Figure 3). Unlike most prior equivariant SSL methods [Lee et al., 2021, Park et al., 2022, Dangovski et al., 2022, Gupta et al., 2023, Garrido et al., 2023, Gupta et al., 2024, Yerxa et al., 2024], our model does not rely on explicitly crafted loss terms or objectives to achieve equivariance, nor is it instructed to learn the decomposition of two representation types. Instead, the dual representation structure arises naturally from the model architecture and action-conditioned predictive learning.

Beyond resolving the invariance-equivariance trade-off, seq-JEPA further benefits from processing a sequence of observation views; we show that it performs well on tasks requiring integration over sequences of observations. In one scenario, inspired by embodied vision in primates, our model learns image representations without augmentations or masking, solely by predicting across simulated eye movements (saccades). In another setting, it performs path integration over sequences of actions—such as eye movements or 3D object rotations in 3DIEBench [Garrido et al., 2023]. **Our key contributions are as follows:**

- We introduce seq-JEPA, a self-supervised world model that learns architecturally distinct equivariant and invariant representations through sequential prediction over action-observation pairs, without requiring explicit equivariance losses or dual predictors.
- We empirically validate that seq-JEPA matches or outperforms existing invariant and equivariant SSL methods across tasks requiring either representational property.
- We demonstrate that seq-JEPA naturally supports tasks that involve sequential integration of observations, such as predictive learning across saccades and path integration over action sequences.

2 Method

2.1 Invariant and equivariant representations

Before presenting our architecture and training procedure, we briefly define invariance and equivariance in the context of SSL [Dangovski et al., 2022, Devillers and Lefort, 2022]. Let \mathcal{T} denote a distribution over transformations, parameterized by a vector t. These transformations—such as augmentations or masking—can be used to generate multiple views from a single image x. Let x_1 and x_2 be two such views, produced by applying transformations t_1 and t_2 sampled from \mathcal{T} . Additionally, let a denote the *relative* transformation that maps x_1 to x_2 . Additionally, we denote a as a transformation that transforms x_1 to x_2 . We distinguish between t (an individual transformation) and a (an action), where the latter reflects the change from one view to another. Let f be an encoder that maps inputs to a latent space. We say that f is equivariant to t if:

$$\forall t \in \mathcal{T}, \ \exists, u_t; \text{s.t.} \qquad f(t(x)) = u_t(f(x)),$$
 (1)

where u_t is a transformation in latent space corresponding to t. Equivariance can similarly be defined in terms of relative transformations (actions):

$$\forall a \in \mathcal{T}, \exists, u_a; \text{s.t.} \qquad f(x_2) = u_a(f(x_1)). \tag{2}$$

As a special case, if u_t and u_a are identity functions, then f is invariant to the transformation: f(t(x)) = f(x) or $f(x_2) = f(x_1)$.

2.2 Architecture

Figure 1 presents the overall architecture of seq-JEPA. Let $\{x_i\}_{i=1}^{M+1}$ be a sequence of views generated from a sample x via transformations $\{t_i\}_{i=1}^{M+1}$. The relative transformations (actions) $\{a_i\}_{i=1}^{M}$ are defined as $a_i \triangleq \Delta t_{i,i+1}$, i.e., the transformation mapping x_i to x_{i+1} . In our default setting, we use a learnable linear projector to encode these actions.

A backbone encoder, f encodes the first M views, producing representations $\{z_i\}_{i=1}^M$. Except for z_M , each z_i is concatenated with its corresponding action embedding and passed to a transformer encoder g (no MLP projector is used after the encoder), which aggregates the sequence of action-observation pairs. The transformer uses a learnable [AGG] token (analogous to the [CLS] token in ViT [Dosovitskiy et al., 2020]) to generate the aggregate representation:

$$z_{AGG} = g((z_1, a_1), (z_2, a_2)..., (z_{M-1}, a_{M-1}), z_M)$$
(3)

This aggregate representation z_{AGG} is then concatenated with the final action embedding a_M (corresponding to the transformation from x_M to x_{M+1}), and passed to an MLP predictor h to predict the representation of x_{M+1} :

$$\hat{z}_{M+1} = h(z_{AGG}, a_M). \tag{4}$$

The ground truth z_{M+1} is computed using a target encoder—an exponential moving average (EMA) of f. This target representation is passed through a stop-gradient operator (sg) to avoid representational collapse [Grill et al., 2020, Chen and He, 2021, Assran et al., 2023]. The training objective is to maximize the cosine similarity between \hat{z}_{M+1} and z_{M+1} with the loss function:

$$\mathcal{L}_{\text{seq-JEPA}} = 1 - \frac{\hat{z}_{M+1}}{\|\hat{z}_{M+1}\|_2} \cdot \frac{\text{sg}(z_{M+1})}{\|\text{sg}(z_{M+1})\|_2}.$$
 (5)

No additional loss terms or equivariance-specific predictors are used during training.

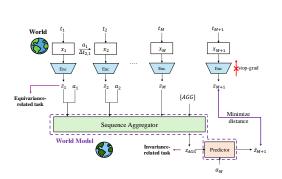


Figure 1: seq-JEPA is a self-supervised world model that leverages a sequence of action-observation pairs to learn architecturally distinct action-invariant and -equivariant representations.

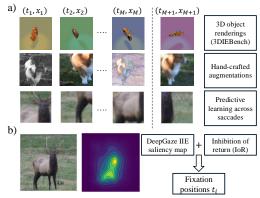


Figure 2: **a.** Transformations and observations used for training **b.** In predictive learning across saccades, image saliencies and inhibition of return help create more informative and less redundant patch sequences.

2.3 Action and observation sets

To evaluate generalization across transformation types, we consider three sets of action-observation pairs (Figure 2). See Appendix A.1 for details of each setup.

3D Invariant Equivariant Benchmark (3DIEBench). The 3DIEBench dataset [Garrido et al., 2023] is designed to evaluate representational invariance and equivariance. It includes 3D object renderings with variations in rotation, floor hue, and lighting. In this benchmark, the action between two views is the relative difference of these three factors of variation. We primarily study equivariance to SO(3) rotations and secondarily to appearance factors of floor and light hue.

Hand-Crafted Augmentations. In this setting, we use transformed views generated via common SSL augmentations (e.g., crop, color jitter, blur), and actions correspond to relative augmentation parameters. We use CIFAR100 and Tiny ImageNet for experiments in this setup, and follow EquiMod's augmentation protocol [Devillers and Lefort, 2022].

Predictive Learning Across Saccades (PLS). Going beyond conventional transformations such as augmentations or 3D rotations, we show that seq-JEPA can learn visual representations from a sequence of partial observations thanks to architectural inductive biases—without relying on any hand-crafted augmentations. Our PLS has a similar flavor to I-JEPA [Assran et al., 2023] but does not require engineered masking strategies. In PLS, we train seq-JEPA on sequences of patches extracted from full-resolution images. For instance, with STL-10 dataset, we use small 32×32 patches to form the observation sequence. In this setting, actions correspond to the relative positions between patch centers, simulating saccadic eye movements and inducing 2-D positional equivariance to representations. To select fixation points, we adopt two biologically inspired techniques that increase informativeness, reduce redundancy, and improve the downstream utility of the aggregate representation (Figure 2):

- Saliency-Based Fixation Sampling. Using DeepGaze IIE [Linardos et al., 2021], we extract saliency maps for each image and use them to probabilistically sample fixation points [Itti et al., 1998, Li, 2002, Zhaoping, 2014]. The maps are pre-computed and introduce no training overhead.
- Inhibition of Return (IoR). To reduce spatial overlap between patches and emulate natural exploration [Posner et al., 1985], we implement IoR by zeroing out the sampling probability of areas surrounding previously sampled fixations.

3 Experimental Setup

3.1 Compared methods and baselines

We compare seq-JEPA against both invariant and equivariant SSL baselines. Invariant methods include SimCLR [Chen et al., 2020], BYOL [Grill et al., 2020], and VICReg [Bardes et al., 2022] and VICReg with trajectory regularization [Wang et al., 2024a]. Equivariant methods include SEN [Park et al., 2022], EquiMod [Devillers and Lefort, 2022], SIE [Garrido et al., 2023], and ContextSSL [Gupta et al., 2024]. For all baselines, architectural details are given in Appendix A.3. We also evaluate two hybrid baselines based on our architecture:

- Conditional BYOL. A two-view version of seq-JEPA with no sequence aggregator, where BYOL's predictor is conditioned on the relative transformation between target and online views. This encourages representations to encode transformation information.
- **Conv-JEPA.** A baseline for the saccades setting. It uses the same sequence of saliency-sampled patches as seq-JEPA and predicts the final patch's representation from each earlier patch individually. These losses are summed across the pairs before backpropagation.

3.2 Training protocol

All models use ResNet-18 [He et al., 2016] as the backbone encoder. For action conditioning, we use a learnable linear projection to learn action embeddings (default action embedding is 128-d). The sequence aggregator in seq-JEPA is a lightweight transformer encoder [Vaswani et al., 2017] with three layers and four attention heads. The predictor is a 2-layer MLP with 1024 hidden units and ReLU activation. In order to control for and eliminate any performance gain resulting from using a transformer encoder in seq-JEPA instead of an MLP projection head, we trained baselines that typically use MLP projectors in two variants: (1) with original MLP projector; and (2) with the MLP replaced by a transformer encoder and a sequence length of one. We did not see any benefit from switching to transformer projectors in any of the baselines, and include the transformer-projector results in Appendix B.3. All models are trained from scratch with a batch size of 512. We use 1000 epochs for 3DIEBench and 2000 epochs for other datasets to obtain asymptotic performance. We use AdamW for models with transformer projectors (including seq-JEPA) due to its stability and improved regularization in transformer training. For ConvNet-only models with MLP heads, we use the Adam optimizer. Full hyperparameters are detailed in Appendix A.

3.3 Evaluation metrics and protocol

To assess equivariance, we follow the protocol of Garrido et al. [2023] and train a regressor on frozen encoder representations to predict the relative transformation (action) between two views. We report the R^2 score on the test set. In addition to action decoding R^2 , we also report retrieval-based metrics including Mean Reciprocal Rank (MRR), Hit@1, and Hit@5 [Kipf et al., 2020, Park et al., 2022, Garrido et al., 2023] to evaluate the quality of the predictor. As a proxy measure of invariance, we use top-1 classification accuracy of a linear probe on top of frozen representations. For all baselines, probes are trained on encoder outputs. For seq-JEPA, we measure accuracies on top of the aggregate representation (z_{AGG} in Figure 1) and report the number of observation views used during training and inference. For completeness, we also report classification performance on top of encoder representation (z_i in Figure 1) for seq-JEPA models in Appendix B.8. Training details of evaluation heads are given in Appendix A.2.

4 Results

4.1 Quantitative evaluation on 3DIEBench

We use the 3DIEBench benchmark to quantitatively compare invariant and equivariant representations in seq-JEPA with baseline methods. This benchmark allows us to measure equivariance through decoding 3D object rotations while enabling invariance measurement through object classification. Table 1 provides a summary of our evaluation on the 3DIEBench where equivariant methods have been conditioned on rotation (a 4-D quaternion representing the relative rotation between two views). In addition to the relative rotation between two views, in the last column we provide the R^2 score for predicting individual transformation parameters from representations of a single view. For seq-JEPA,

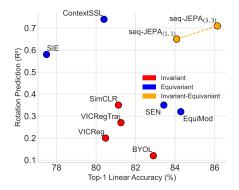


Figure 3: Top-1 linear classification (invariance) vs. rotation prediction (equivariance) performance on 3DIEBench. seq-JEPA learns good representations for both tasks. Subscripts indicate training and inference sequence lengths.

Table 1: Evaluation on 3DIEBench for linear probe classification (invariance) and rotation prediction (equivariance). Equivariant models and seq-JEPA are conditioned on rotation. For seq-JEPA, training and inference sequence lengths are denoted by M_{tr} and M_{val} .

Method	Top-1 Acc. (%)	Rel. Rot. (\mathbb{R}^2)	Indiv. Rot. (R^2)
Invariant			
BYOL	82.90	0.12	0.25
SimCLR	81.13	0.35	0.54
VICReg	80.48	0.20	0.36
VICRegTraj	81.26	0.27	0.43
Equivariant			
SÉN	83.43	0.35	0.57
EquiMod	84.29	0.32	0.55
SIE	77.49	0.58	0.62
Conditional BYOL	82.61	0.31	0.47
ContextSSL ($c = 126$)	80.40	0.74	0.78
Invariant-Equivariant			
seq-JEPA (1, 1)	84.08	0.65	0.69
seq-JEPA (1, 3)	<u>85.31</u>	0.65	0.69
seq-JEPA (3, 3)	86.14	0.71	0.74
seq-JEPA (3, 5)	87.41	0.71	0.74
seq-JEPA (no act cond)	86.05	0.29	0.37

we trained models with varying training sequence lengths (denoted by M_{tr} in the table) and measure the linear classification performance on top of aggregate representations with different inference lengths (denoted by M_{val}).

Among invariant methods, BYOL achieves the highest classification accuracy, yet does not offer a high level of equivariance. Adding a linear trajectory regularization loss to VICReg Wang et al. [2024a] without using ground-truth transformations improves over VICReg, which shows that imposing geometric priors can improve both invariant and equivariant performance—even when these priors do not fully materialize in the environment, e.g., when we have non-smooth angle changes as in 3DIEBench. Among equivariant baselines, SIE and ContextSSL yield strong rotation prediction performance due to their specialized equivariance predictors and loss functions, but underperform in classification. EquiMod and SEN offer better classification performance, yet compromise equivariance. In contrast, seq-JEPA achieves strong performance in both, matching the best rotation R^2 scores while exceeding baselines in classification, with gains increasing with inference sequence length M_{val} . Ablating action conditioning leads to a sharp drop in equivariance but retains classification accuracy, confirming our hypothesis that action-conditioned sequential aggregation enables a segregated invariant-equivariant representation structure. For additional results on 3DIEBench including MRR, Hit@1, and Hit@5 metrics, performance of models with varying training and inference sequence lengths and models conditioned on both rotation and color, and evaluation on an out-of-distribution (OOD) set of 3DIEBench see Appendices B.1 to B.7.

4.2 Qualitative evaluation on 3DIEBench

To visualize equivariance in representational space, we retrieve the three nearest representations of a query image from the validation set of 3DIEBench (Figure 4). While all models retrieve the correct object category, only seq-JEPA and SIE consistently preserve rotation across all retrieved views, consistent with their high R^2 scores. Next, we projected encoder and aggregate representations using 2D UMAP (Figure 5). The left panel shows encoder representations colored by class label, while the middle panel displays the same encoder representations colored by rotation angle. The smooth color gradation across the map within each class cluster in the middle panel suggests that the encoder captures rotation angle as a continuous factor, implying equivariance to rotation (e.g., the red class in the bottom-right corner of the right panel and the corresponding part in the middle panel). The right panel shows aggregate representations colored by class label. Comparing the class-colored plots (left and right panels), we observe that both encoder and aggregate representations contain class information. However, when we aggregate multiple views of a sample, some of the intra-class variability (resulting from transformations such as rotation) is eliminated, causing each class' representational cluster to become more homogeneous. This aggregation procedure likely reduces variation due to rotation and makes the representations more invariant, resulting in decreased intra-class spread and increased inter-class distance. We create a similar UMAP visualization for seq-

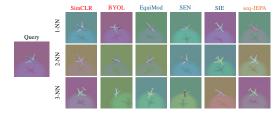


Figure 4: Retrieval of nearest representations; given a query image, we extract the three nearest encoder representations in the validation set of 3DIEBench. The retrieved views of models with the highest quantitative rotation equivariance performance maintain the rotation of the query image across all retrieved views.

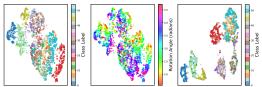


Figure 5: 2-D UMAP projections of seq-JEPA's encoder and aggregate representations trained on 3DIEBench and conditioned on rotation with $M_{tr}=3$ and $M_{val}=5$. Encoder representations for each view observation, color-coded by class (**left**) and rotation angle (**middle**). Aggregate representation for $M_{val}=5$, color-coded by class (**right**).

Table 2: Evaluation with hand-crafted augmentations on CIFAR100 and Tiny ImageNet; equivariance is measured by predicting relative transformation parameters associated with crop, color jitter, or blur augmentations. For all seq-JEPA models, $M_{val}=5$.

			CIFAR100)			Tiny ImageN		
Conditioning	Method	Classification (top-1)	Crop (R^2)	Jitter (R^2)	Blur (R^2)	Classification (top-1)	Crop (R^2)	Jitter (R^2)	Blur (R^2)
	Invariant								
	SimCLR	61.72	0.56	0.18	0.04	34.29	0.30	0.08	0.13
-	BYOL	62.17	0.39	0.05	-0.01	35.70	0.17	0.01	0.16
	VICReg	61.35	0.49	0.11	-0.06	35.29	0.31	0.05	0.19
	VICRegTraj	61.07	0.46	0.14	0.02	34.95	0.28	0.09	0.16
	Equivariant								
	SEN	61.94	0.65	0.51	0.85	36.01	0.24	0.48	0.87
Crop+Jitter+Blur	EquiMod	61.80	0.59	0.49	0.74	36.75	0.38	0.46	0.86
Стортящегтыш	SIE	58.81	0.34	0.26	0.53	31.37	0.26	0.56	0.88
	Conditional BYOL	60.63	0.56	0.46	0.73	35.49	0.34	0.54	0.87
	SEN	61.56	0.66	0.15	0.10	35.95	0.24	0.11	0.49
Crop	EquiMod	61.77	0.62	0.15	0.01	36.83	0.26	0.12	0.30
Стор	SIE	57.55	0.69	0.11	0.25	32.38	0.34	0.05	0.14
	Conditional BYOL	60.17	0.55	0.10	-0.01	37.07	0.22	0.01	0.24
	SEN	61.78	0.50	0.52	0.02	36.59	0.26	0.50	0.21
Color jitter	EquiMod	61.53	0.44	0.50	-0.02	37.18	0.20	0.52	0.29
Color Jitter	SIE	59.29	0.48	0.59	0.06	34.37	0.39	0.62	0.35
	Conditional BYOL	61.30	0.36	0.52	0.02	37.93	0.24	0.62	0.30
	SEN	61.47	0.43	0.15	0.84	34.62	0.16	0.08	0.79
Blur	EquiMod	62.72	0.41	0.15	0.74	36.12	0.24	0.10	0.91
Diui	SIE	57.66	0.40	0.07	0.71	31.00	0.26	0.05	0.85
	Conditional BYOL	60.74	0.36	0.11	0.69	36.17	0.19	0.02	0.85
	Invariant-Equivariant								
Crop+Jitter+Blur	seq-JEPA ($M_{tr} = 1$)	52.90	0.77	0.52	0.23	37.10	0.64	0.49	0.89
Crop+Jitter+Blur	seq-JEPA $(M_{tr} = 2)$	60.17	0.78	0.64	0.88	35.56	0.69	0.42	0.93
Crop+Jitter+Blur	seq-JEPA ($M_{tr} = 3$)	58.33	0.79	0.63	0.92	34.85	0.67	0.64	0.96
Crop	seq-JEPA $(M_{tr} = 2)$	59.32	0.78	0.01	0.10	35.74	0.70	0.12	0.46
Color Jitter	seq-JEPA ($M_{tr} = 3$)	58.62	0.68	0.68	0.29	35.21	0.60	0.66	0.62
Blur	seq-JEPA ($M_{tr} = 3$)	56.82	0.71	0.15	0.74	35.79	0.58	0.22	0.97
-	seq-JEPA ($M_{tr} = 2$)	58.37	0.64	0.14	0.16	35.97	0.52	0.18	0.47

JEPA with ablated rotation conditioning in Appendix B.9 to highlight the role of action conditioning in achieving equivariance to rotation.

4.3 Evaluation with Hand-Crafted Augmentations

We assess invariance and equivariance under hand-crafted augmentations by training on CIFAR100 and Tiny ImageNet (Table 2). Models with action conditioning are trained conditioned on crop, color jitter, blur, or all three (indicated in the first column of the table). seq-JEPA consistently achieves higher equivariance than both invariant and equivariant baselines across all transformations. Notably, except for the model trained on CIFAR-100 and conditioned on blur, the best equivariance performance for a given augmentation is achieved when the model is specialized and conditioned only on that augmentation. Furthermore, ablating actions (last row in the table) causes seq-JEPA to lose its equivariance across transformations compared to action-conditioned models. Overall, our model outperforms both invariant and equivariant families in terms of equivariance, while being competitive in terms of classification performance. For additional results with varying training and inference sequence lengths, see Appendix B.7.

4.4 Predictive Learning across Saccades and Path Integration

In our third action-observation setting, we consider predictive learning across simulated eye movements to exhibit seq-JEPA's ability in leveraging a sequence of partial observations to learn visual

Table 3: Evaluation of predictive learning across saccades on STL-10. Equivariance is measured with respect to fixation coordinates. Unless stated otherwise, $M_{tr} = M_{val} = 4$.

Conditioning		Top-1 Acc. (%)	Position (\mathbb{R}^2)
-	Invariant SimCLR (augmentations)	85.23	-0.06
position	Equivariant Conv-JEPA ($M_{tr,val} = 4$)	80.04	0.80
-	Invariant-Equivariant seq-JEPA	70.45	0.38
position position	seq-JEPA seq-JEPA $(M_{val} = 6)$	83.44 84.12	0.80 0.80
position position	seq-JEPA (w/o saliency & IoR) seq-JEPA (w/o IoR)	79.85 77.97	0.88 0.85

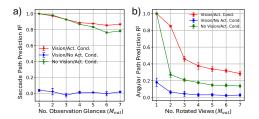


Figure 6: a) Visual path integration across eye movements, b) Angular path integration across object rotations (results over three random seeds)

representations. In Table 3, seq-JEPA reaches 83.44% top-1 accuracy on STL-10, comparable to SimCLR (85.23%) trained with full-resolution images and strong augmentations. This gap narrows further when increasing the inference sequence length from $M_{val}=4$ to $M_{val}=6$. Ablating action conditioning causes the accuracy on top of the aggregate representations to drop sharply, indicating that 2-D positional awareness is essential to forming semantic representations across simulated eye movements. Compared to Conv-JEPA—which accumulates prediction losses pairwise—seq-JEPA performs better in classification, highlighting the importance of sequence aggregation when dealing with partial observations in SSL. Further ablations show that saliency-driven sampling and IoR are critical for forming informative, non-overlapping patch sequences and subsequently a high-quality aggregate representation. Interestingly, while random uniform patch sampling negatively impacts classification accuracy due to lower semantic content, it results in the highest positional equivariance as the model samples patches and corresponding saccade actions from a more diverse set of positions across the entire image, not just the salient regions. The UMAP projections for PLS with and without action conditioning in Appendix B.9 further underscore the role of action conditioning in enabling positional equivariance.

Path Integration. In the context of eye movement-driven or any sequential observations, an ability that naturally arises from predictive learning is *path integration* [McNaughton et al., 2006], i.e., predicting the cumulative transformation/action from a sequence of actions. We evaluate this task in both eye movements in PLS (visual path integration) and object rotations in 3DIEBench (angular path integration). As shown in Figure 6, seq-JEPA demonstrates strong performance in both settings, with performance degrading gracefully as sequence length increases. Ablating action conditioning causes path integration to fail, whereas ablating the visual stream has only a minor impact—highlighting that action information is the dominant signal for this task. Full details of the path integration setup are provided in Appendix B.10.

4.5 Action Conditioning Ablations

To better understand the mechanisms underlying seq-JEPA's invariant-equivariant representation learning and role of action conditioning, we perform a set of ablation experiments on 3DIEBench. Specifically, we study: (i) the role of action conditioning in the transformer and predictor; (ii) the impact of action embedding dimensionality. Table 4 summarizes our ablation results. Removing action conditioning entirely causes a significant drop in equivariance (R^2 from 0.71 to 0.29), although classification accuracy remains high thanks to sequence aggregation and segregated invariance-equivariance in our model. Conditioning only the transformer or only the predictor leads to intermediate results, with predictor conditioning proving more critical for equivariance. We also vary the dimensionality of the learnable action embeddings: performance saturates around the default size of 128, with smaller sizes (e.g., 16 or 64) already sufficient to capture the rotation structure.

4.6 Scaling Properties: Role of Training and Inference Sequence Lengths

We study the effect of both training and inference sequence lengths on performance across tasks, i.e., scalability of seq-JEPA in terms of context length (Figure 7). We draw inspiration from recent findings in foundation models, where increased training and inference context—whether in text [Brown et al., 2020, Touvron et al., 2023], vision [Zellers et al., 2021, Chen et al., 2021], or video [Bain et al., 2021, Arnab et al., 2021]—consistently leads to stronger representations.

Table 4: Ablation results for action conditioning (3DIEBench). All models use $M_{tr}=3$, $M_{val}=5$ (results over three random seeds).

Variant	Top-1 Acc. (%)	Rotation (R^2)
Act. conditioning		
None	87.36 ± 0.7	0.29 ± 0.04
No predictor cond.	87.17 ± 0.3	0.37 ± 0.06
No transformer cond.	86.33 ± 0.1	0.53 ± 0.05
Act. embedding dim.		
$a_{\text{dim}} = 16$	86.29 ± 0.4	0.70 ± 0.01
$a_{\text{dim}} = 64$	87.11 ± 0.2	0.70 ± 0.02
$a_{\text{dim}} = 128 \text{ (default)}$	87.41 ± 0.5	0.71 ± 0.02
$a_{\text{dim}} = 256$	87.26 ± 0.6	0.72 ± 0.00

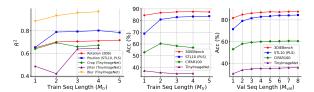


Figure 7: Effect of training and inference sequence length on seq-JEPA's performance; **left:**: Equivariant performance (R^2) versus training sequence length; **middle:** Classification accuracy versus training sequence length; **right:** Classification accuracy versus inference sequence length.

We observe that equivariance generally improves with longer training sequences (left panel). A possible explanation for this observation is the presence of more transitions (z_i, a_i, z_{i+1}) in working memory, which means the predictor has access to a richer context $z_{\rm agg}$. This enables the predictor to more accurately approximate the transition $p(z_{i+1} \mid z_{\rm agg,i}, a_i)$. Because accurate prediction requires the model to preserve and utilize information about a_t , the encoder is implicitly encouraged to learn structured, more equivariant representations.

Classification performance on 3DIEBench and STL-10 (middle panel) benefits from a longer training sequence. In contrast, on CIFAR100 and Tiny ImageNet with synthetic augmentations, longer training sequences slightly decrease classification performance. We hypothesize that leveraging and aggregating a sequence of action-observation pairs, i.e. seq-JEPA's architectural inductive bias, is most effective in settings where the downstream task benefits from sequential observations. In the case of object rotations in 3DIEBench, seeing an object from multiple angles is indeed beneficial in recognizing the object's category, which explains the improved classification accuracy with increased training sequence length. Similarly, in the case of predictive learning across saccades, each eye movement and its subsequent glance provides additional information that can be leveraged for learning a richer aggregate representation.

At inference time, all datasets benefit from longer context lengths (M_{val}) , confirming that richer aggregate representations yield stronger performance (right plot). This scalability via sequence lengths opens avenues for efficient representation learning with small foveated patches in lieu of full-frame inputs, mirroring how foundation models scale with input tokens at test time. Together with our transfer learning results on ImageNet-1k (Appendix B.11), these findings suggest that seq-JEPA's architectural inductive bias enables graceful scaling via longer sequence lengths.

5 Related Work

Non-Generative World Models and Joint-Embedding Predictive Architectures. Non-generative world models predict the effect of transformations or actions directly in latent space, avoiding reconstruction in pixel space. This includes contrastive SSL methods that model transformed views from context representations [van den Oord et al., 2019, Gupta et al., 2024], as well as approaches in model-based reinforcement learning (RL) to improve sample efficiency, generate intrinsic rewards, or capture environment transitions [Schwarzer et al., 2021, Khetarpal et al., 2025, Ni et al., 2024, Tang et al., 2023, Guo et al., 2022]. Joint-embedding predictive architectures [LeCun, 2022] form a subclass of non-generative world models. They introduce an asymmetric predictor conditioned on transformation parameters to infer the outcome of an action applied to a latent view. Examples include I-JEPA[Assran et al., 2023], which predicts masked regions from positional cues, and IWM [Garrido et al., 2024], which conditions on augmentation parameters. JEPAs have been recently extended to physical reasoning [Garrido et al., 2025] and offline planning [Sobal et al., 2025], illustrating the framework's versatility in representation learning and world modeling.

Equivariant SSL. Equivariant SSL methods aim to retain transformation-specific information in the latent space, typically by augmenting invariant objectives with an additional equivariance term. Some approaches directly predict transformation parameters [Lee et al., 2021, Scherr et al., 2022, Gidaris et al., 2018, Gupta et al., 2024, Dangovski et al., 2022]. Methods such as EquiMod [Devillers and Lefort, 2022] and SIE [Garrido et al., 2023] predict the effect of a transformation in latent space via a predictor in addition to their invariant objective. SEN [Park et al., 2022] similarly predicts

transformed representations but omits the invariance term. Xiao et al. [2021] use contrastive learning with separate projection heads for each augmentation, treating same-augmentation pairs as negatives. ContextSSL [Gupta et al., 2024] conditions representations on both current actions and recent context and employs a dual predictor for transformation prediction to avoid collapsing to invariance. Other approaches [Shakerinava et al., 2022, Gupta et al., 2023, Yerxa et al., 2024] do not require explicit transformation parameters but instead enforce equivariance by applying the same transformation to multiple view pairs and minimizing a distance-based loss. Action-conditioned JEPAs incorporate augmentation parameters and mask positions by conditioning the predictor to induce equivariance without additional objectives [Garrido et al., 2024]. Chavhan et al. [2023] use an ensemble of heads trained on top of a pre-trained SSL encoder to span a diverse spectrum of transformation sensitivities across the latent spaces of each head. Downstream probes then learn to linearly combine these feature heads depending on the desired invariance-equivariance trade-off.

Positioning of Our Work. seq-JEPA belongs to the family of joint-embedding predictive architectures and is a non-generative world model. Unlike most equivariant SSL methods, seq-JEPA does not rely on an equivariance loss or transformation prediction objective, nor does it require view pairs with matched transformations. Instead, it leverages action conditioning and architectural inductive biases to learn two separate invariant and equivariant representations. In contrast to ContextSSL, which extends the two-view contrastive setting of contrastive predictive coding van den Oord et al. [2019] using a transformer decoder projector conditioned on previous views via key-value caching, seq-JEPA operates on sequences of action-observation pairs in an online end-to-end manner by incorporating a sequence model (e.g. a transformer encoder) as a *learned working memory* during both training and inference. Moreover, while ContextSSL aims to adapt equivariance to recent transformations by dynamically modifying the training distribution, seq-JEPA is designed to explicitly learn both invariant and equivariant representations with respect to a specified set of transformations, and is also well-suited for downstream tasks that require multi-step observation aggregation.

6 Limitations and Future Perspectives

We have validated viability of seq-JEPA across a range of transformations in the image domain. Here, we discuss limitations and possible future directions. First, seq-JEPA is capable of autoregressive prediction in time, and therefore, can be leveraged for autoregressive latent planning in control tasks. Second, the transformer-based aggregator in seq-JEPA could support multi-modal fusion across language, audio, or proprioceptive inputs—enabling multi-modal world modeling and generalization. Third, while our results show clear benefits from longer sequences, we have experimented with relatively short training and inference sequence lengths as the ResNet backbone is also trained end-to-end. The sequence scaling trends observed in the paper suggest that seq-JEPA can benefit from longer context windows over representations of pre-trained backbones as in [Pang et al., 2023, Lin et al., 2024]. Fourth, our method assumes access to a known transformation group (e.g., SO(3)) for 3D rotations). Designing group-agnostic or learned transformation models [Finzi et al., 2021] without access to transformation parameters or pairs of same transformation is an open challenge in equivariant SSL that future work may tackle. Finally, our preliminary ImageNet-1k transfer results (Appendix B.11) point to potential for broader generalization. Scaling seq-JEPA to larger foveated image settings or video and multi-modal datasets such as Song et al. [2023] could support the development of lightweight, saliency-driven agents capable of learning efficiently from partial observations in embodied settings with a limited field of vision.

Acknowledgments and Disclosure of Funding

This project was supported by funding from NSERC (Discovery Grants RGPIN-2022-05033 to E.B.M., and RGPIN-2023-03875 to S.B.), Canada CIFAR AI Chairs Program and Google to E.B.M., Canada Excellence Research Chairs (CERC) Program, Mila - Quebec AI Institute, Institute for Data Valorization (IVADO), CHU Sainte-Justine Research Centre, Fonds de Recherche du Québec–Santé (FRQS), and a Canada Foundation for Innovation John R. Evans Leaders Fund grant to E.B.M. This research was also supported in part by Digital Research Alliance of Canada (DRAC) and Calcul Québec.

References

- Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. ViViT: A Video Vision Transformer. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 6816–6826, Montreal, QC, Canada, October 2021. IEEE. ISBN 978-1-6654-2812-5. doi: 10.1109/ICCV48922.2021.00676. URL https://ieeexplore.ieee.org/document/9710415/.
- Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked Siamese Networks for Label-Efficient Learning. In *Proceedings of the 17th European Conference on Computer Vision, ECCV 2022*, April 2022. URL http://arxiv.org/abs/2204.07141. arXiv:2204.07141 [cs, eess].
- Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 15619–15629, Vancouver, BC, Canada, June 2023. IEEE. ISBN 979-8-3503-0129-8. doi: 10.1109/CVPR52729.2023.01499. URL https://ieeexplore.ieee.org/document/10205476/.
- Max Bain, Arsha Nagrani, Gul Varol, and Andrew Zisserman. Frozen in Time: A Joint Video and Image Encoder for End-to-End Retrieval. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 1708–1718, Montreal, QC, Canada, October 2021. IEEE. ISBN 978-1-6654-2812-5. doi: 10.1109/ICCV48922.2021.00175. URL https://ieeexplore.ieee.org/document/9711165/.
- Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=xm6YD62D1Ub.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *Advances in Neural Information Processing Systems*, volume 33, pages 9912–9924. Curran Associates, Inc., 2020.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Caron_Emerging_Properties_in_Self-Supervised_Vision_Transformers_ICCV_2021_paper.
- Ruchika Chavhan, Henry Gouk, Da Li, and Timothy Hospedales. Quality diversity for visual pretraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5384–5394, 2023.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1597–1607. PMLR, November 2020. URL https://proceedings.mlr.press/v119/chen20j.html. ISSN: 2640-3498.

- Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. URL https://openaccess.thecvf.com/content/CVPR2021/html/Chen_Exploring_Simple_Siamese_Representation_Learning_CVPR_2021_paper.html.
- Xinlei Chen, Saining Xie, and Kaiming He. An Empirical Study of Training Self-Supervised Vision Transformers, August 2021. URL http://arxiv.org/abs/2104.02057. arXiv:2104.02057 [cs].
- Rumen Dangovski, Li Jing, Charlotte Loh, Seungwook Han, Akash Srivastava, Brian Cheung, Pulkit Agrawal, and Marin Soljacic. Equivariant Self-Supervised Learning: Encouraging Equivariance in Representations. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=gKLAAfiytI.
- Alexandre Devillers and Mathieu Lefort. EquiMod: An Equivariance Module to Improve Visual Instance Discrimination. In *The Eleventh International Conference on Learning Representations*, September 2022. URL https://openreview.net/forum?id=eDLwjKmtYFt.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, October 2020. URL https://openreview.net/forum?id=YicbFdNTTy.
- Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a Little Help From My Friends: Nearest-Neighbor Contrastive Learning of Visual Representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9588–9597, 2021.
- Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for Self-Supervised Representation Learning. In *Proceedings of the 38th International Conference on Machine Learning*, pages 3015–3024. PMLR, July 2021. URL https://proceedings.mlr.press/v139/ermolov21a.html. ISSN: 2640-3498.
- Marc Finzi, Max Welling, and Andrew Gordon Wilson. A Practical Method for Constructing Equivariant Multilayer Perceptrons for Arbitrary Matrix Groups. In *Proceedings of the 38th International Conference on Machine Learning*, pages 3318–3328. PMLR, July 2021. URL https://proceedings.mlr.press/v139/finzi21a.html. ISSN: 2640-3498.
- Quentin Garrido, Laurent Najman, and Yann Lecun. Self-supervised learning of Split Invariant Equivariant representations. In *Proceedings of the 40th International Conference on Machine Learning*, pages 10975–10996. PMLR, July 2023. URL https://proceedings.mlr.press/v202/garrido23b.html. ISSN: 2640-3498.
- Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and Leveraging World Models in Visual Representation Learning, March 2024. URL http://arxiv.org/abs/2403.00504. arXiv:2403.00504 [cs].
- Quentin Garrido, Nicolas Ballas, Mahmoud Assran, Adrien Bardes, Laurent Najman, Michael Rabbat, Emmanuel Dupoux, and Yann LeCun. Intuitive physics understanding emerges from self-supervised pretraining on natural videos, February 2025. URL http://arxiv.org/abs/2502.11831. arXiv:2502.11831 [cs].
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. February 2018. URL https://openreview.net/forum?id=S1v4N210-.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap Your Own Latent A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020.

- Zhaohan Guo, Shantanu Thakoor, Miruna Pislar, Bernardo Avila Pires, Florent Altché, Corentin Tallec, Alaa Saade, Daniele Calandriello, Jean-Bastien Grill, Yunhao Tang, Michal Valko, Remi Munos, Mohammad Gheshlaghi Azar, and Bilal Piot. BYOL-Explore: Exploration by Bootstrapped Prediction. In *Advances in Neural Information Processing Systems*, volume 35, pages 31855–31870, December 2022.
- Zhaohan Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-Bastien Grill, Florent Altché, Remi Munos, and Mohammad Gheshlaghi Azar. Bootstrap Latent-Predictive Representations for Multitask Reinforcement Learning. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3875–3886. PMLR, November 2020. URL https://proceedings.mlr.press/v119/guo20g.html. ISSN: 2640-3498.
- Sharut Gupta, Joshua Robinson, Derek Lim, Soledad Villar, and Stefanie Jegelka. Structuring Representation Geometry with Rotationally Equivariant Contrastive Learning. In *The Twelfth International Conference on Learning Representations*, October 2023. URL https://openreview.net/forum?id=lgaFMvZHSJ.
- Sharut Gupta, Chenyu Wang, Yifei Wang, Tommi Jaakkola, and Stefanie Jegelka. In-Context Symmetries: Self-Supervised Learning through Contextual World Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024. URL https://openreview.net/forum?id=etPAH4xSUn.
- Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable Guarantees for Self-Supervised Deep Learning with Spectral Contrastive Loss. In Advances in Neural Information Processing Systems, volume 34, pages 5000–5011. Curran Associates, Inc., 2021.
- Karin L. Harman, G.Keith Humphrey, and Melvyn A. Goodale. Active manual control of object views facilitates visual recognition. *Current Biology*, 9(22):1315–1318, November 1999. ISSN 09609822. doi: 10.1016/S0960-9822(00)80053-6. URL https://linkinghub.elsevier.com/retrieve/pii/S0960982200800536.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, Las Vegas, NV, USA, June 2016. IEEE. ISBN 978-1-4673-8851-1. doi: 10.1109/CVPR.2016.90. URL http://ieeexplore.ieee.org/document/7780459/.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 9726–9735, Seattle, WA, USA, June 2020. IEEE. ISBN 978-1-72817-168-5. doi: 10.1109/CVPR42600.2020.00975. URL https://ieeexplore.ieee.org/document/9157636/.
- L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, November 1998. ISSN 1939-3539. doi: 10.1109/34.730558. URL https://ieeexplore.ieee.org/document/730558/?arnumber=730558. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Khimya Khetarpal, Zhaohan Daniel Guo, Bernardo Avila Pires, Yunhao Tang, Clare Lyle, Mark Rowland, Nicolas Heess, Diana L. Borsa, Arthur Guez, and Will Dabney. A Unifying Framework for Action-Conditional Self-Predictive Reinforcement Learning. In *The 28th International Conference on Artificial Intelligence and Statistics*, February 2025. URL https://openreview.net/forum?id=5DypCUsMg4.
- Thomas Kipf, Elise van der Pol, and Max Welling. Contrastive Learning of Structured World Models, January 2020. URL http://arxiv.org/abs/1911.12247. arXiv:1911.12247 [cs, stat].
- Yann LeCun. A Path Towards Autonomous Machine Intelligence Version 0.9.2, 2022-06-27, 2022. OpenReview.
- Hankook Lee, Kibok Lee, Kimin Lee, Honglak Lee, and Jinwoo Shin. Improving Transferability of Representations via Augmentation-Aware Self-Supervision. In *Advances in Neural Information Processing Systems*, volume 34, pages 17710–17722. Curran Associates, Inc., 2021.

- Zhaoping Li. A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, 6(1): 9-16, January 2002. ISSN 1364-6613, 1879-307X. doi: 10.1016/S1364-6613(00)01817-9. URL https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(00)01817-9. Publisher: Elsevier.
- Han Lin, Tushar Nagarajan, Nicolas Ballas, Mido Assran, Mojtaba Komeili, Mohit Bansal, and Koustuv Sinha. VEDIT: Latent Prediction Architecture For Procedural Video Representation Learning, October 2024. URL http://arxiv.org/abs/2410.03478. arXiv:2410.03478 [cs].
- Akis Linardos, Matthias Kummerer, Ori Press, and Matthias Bethge. DeepGaze IIE: Calibrated prediction in and out-of-domain for state-of-the-art saliency modeling. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pages 12899–12908, Montreal, QC, Canada, October 2021. IEEE. ISBN 978-1-6654-2812-5. doi: 10.1109/ICCV48922.2021.01268. URL https://ieeexplore.ieee.org/document/9711473/.
- Bruce L. McNaughton, Francesco P. Battaglia, Ole Jensen, Edvard I. Moser, and May-Britt Moser. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, 7 (8):663–678, August 2006. ISSN 1471-0048. doi: 10.1038/nrn1932. URL https://www.nature.com/articles/nrn1932. Publisher: Nature Publishing Group.
- Ishan Misra and Laurens van der Maaten. Self-Supervised Learning of Pretext-Invariant Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Misra_Self-Supervised_Learning_of_Pretext-Invariant_Representations_CVPR_2020_paper.html.
- Tianwei Ni, Benjamin Eysenbach, Erfan SeyedSalehi, Michel Ma, Clement Gehring, Aditya Mahajan, and Pierre-Luc Bacon. Bridging State and History Representations: Understanding Self-Predictive RL. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ms0VgzSGF2.
- Ziqi Pang, Ziyang Xie, Yunze Man, and Yu-Xiong Wang. Frozen Transformers in Language Models Are Effective Visual Encoder Layers. October 2023. URL https://openreview.net/forum?id=t0FI3Q66K5.
- Jung Yeon Park, Ondrej Biza, Linfeng Zhao, Jan-Willem Van De Meent, and Robin Walters. Learning Symmetric Embeddings for Equivariant World Models. In *Proceedings of the 39th International Conference on Machine Learning*, pages 17372–17389. PMLR, June 2022. URL https://proceedings.mlr.press/v162/park22a.html. ISSN: 2640-3498.
- Michael I. Posner, Robert D. Rafal, Lisa S. Choate, and Jonathan Vaughan. Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, 2(3):211–228, August 1985. ISSN 0264-3294, 1464-0627. doi: 10.1080/02643298508252866. URL http://www.tandfonline.com/doi/abs/10.1080/02643298508252866.
- Evgenia Rusak, Patrik Reizinger, Attila Juhos, Oliver Bringmann, Roland S. Zimmermann, and Wieland Brendel. InfoNCE: Identifying the Gap Between Theory and Practice, April 2025. URL http://arxiv.org/abs/2407.00143. arXiv:2407.00143 [cs].
- Franz Scherr, Qinghai Guo, and Timoleon Moraitis. Self-Supervised Learning Through Efference Copies. In *Advances in Neural Information Processing Systems*, October 2022. URL https://openreview.net/forum?id=DotEQCtY67g.
- Max Schwarzer, Ankesh Anand, Rishab Goel, R. Devon Hjelm, Aaron Courville, and Philip Bachman. Data-Efficient Reinforcement Learning with Self-Predictive Representations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=uCQfPZwRaUu&fbclid=IwAR3FMvlynXXYEMJaJzPki1x1wC9jjA3aBDC_moWxrI91hLaDvtk7nnnIXT8.
- Mehran Shakerinava, Arnab Kumar Mondal, and Siamak Ravanbakhsh. Structuring Representations Using Group Invariants. In *Advances in Neural Information Processing Systems*, volume 35, pages 34162–34174, December 2022.

- Vlad Sobal, Wancong Zhang, Kynghyun Cho, Randall Balestriero, Tim G. J. Rudner, and Yann LeCun. Learning from Reward-Free Offline Data: A Case for Planning with Latent Dynamics Models, February 2025. URL http://arxiv.org/abs/2502.14819. arXiv:2502.14819 [cs].
- Yale Song, Gene Byrne, Tushar Nagarajan, Huiyu Wang, Miguel Martin, and Lorenzo Torresani. Ego4D Goal-Step: Toward Hierarchical Understanding of Procedural Activities. November 2023. URL https://openreview.net/forum?id=3BxYAaovKr¬eId=IoK9NhxnlM.
- Yunhao Tang, Zhaohan Daniel Guo, Pierre Harvey Richemond, Bernardo Avila Pires, Yash Chandak, Remi Munos, Mark Rowland, Mohammad Gheshlaghi Azar, Charline Le Lan, Clare Lyle, András György, Shantanu Thakoor, Will Dabney, Bilal Piot, Daniele Calandriello, and Michal Valko. Understanding Self-Predictive Learning for Reinforcement Learning. In *Proceedings of the 40th International Conference on Machine Learning*, pages 33632–33656. PMLR, July 2023. URL https://proceedings.mlr.press/v202/tang23d.html. ISSN: 2640-3498.
- Michael J. Tarr, Pepper Williams, William G. Hayward, and Isabel Gauthier. Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience*, 1(4):275–277, August 1998. ISSN 1546-1726. doi: 10.1038/1089. URL https://www.nature.com/articles/nn0898_275. Publisher: Nature Publishing Group.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and Efficient Foundation Language Models, February 2023. URL http://arxiv.org/abs/2302.13971. arXiv:2302.13971 [cs].
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding, January 2019. URL http://arxiv.org/abs/1807.03748.arXiv:1807.03748 [cs, stat].
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- P. Vuilleumier, R. N. Henson, J. Driver, and R. J. Dolan. Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, 5(5):491–499, May 2002. ISSN 1546-1726. doi: 10.1038/nn839. URL https://www.nature.com/articles/nn839. Publisher: Nature Publishing Group.
- Jiayun Wang, Yubei Chen, and Stella X Yu. Pose-aware self-supervised learning with viewpoint trajectory regularization. In *European Conference on Computer Vision*, pages 19–37. Springer, 2024a.
- Yifei Wang, Kaiwen Hu, Sharut Gupta, Ziyu Ye, Yisen Wang, and Stefanie Jegelka. Understanding the Role of Equivariance in Self-supervised Learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024b. URL https://openreview.net/forum?id=NLqdudgBfy.
- Tete Xiao, Xiaolong Wang, Alexei A. Efros, and Trevor Darrell. What Should Not Be Contrastive in Contrastive Learning. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=CZ8Y3NzuVzO.
- Chun-Hsiao Yeh, Cheng-Yao Hong, Yen-Chi Hsu, Tyng-Luh Liu, Yubei Chen, and Yann LeCun. Decoupled Contrastive Learning. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision ECCV 2022*, pages 668–684, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-19809-0. doi: 10.1007/978-3-031-19809-0_38.
- Thomas Edward Yerxa, Jenelle Feather, Eero P. Simoncelli, and SueYeon Chung. Contrastive-Equivariant Self-Supervised Learning Improves Alignment with Primate Visual Area IT. November 2024. URL https://openreview.net/forum?id=AiMs8GPP5q.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stephane Deny. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *Proceedings of the 38th International Conference on Machine Learning*, pages 12310–12320. PMLR, July 2021. URL https://proceedings.mlr.press/v139/zbontar21a.html. ISSN: 2640-3498.

Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. MERLOT: Multimodal Neural Script Knowledge Models, October 2021. URL http://arxiv.org/abs/2106.02636.arXiv:2106.02636 [cs].

Li Zhaoping. The V1 hypothesis—creating a bottom-up saliency map for preattentive selection and segmentation. In *Understanding Vision*, pages 189–314. Oxford University PressOxford, 1 edition, May 2014. ISBN 978-0-19-956466-8 978-0-19-177250-4. doi: 10.1093/acprof:oso/9780199564668.003.0005. URL https://academic.oup.com/book/8719/chapter/154784147.

A Implementation Details

A.1 Data preparation

3DIEBench. The original 256×256 images are resized to a 128×128 resolution for all experiments. Normalization is done using the means and standard deviations in Garrido et al. [2023], i.e., $\mu = [0.5016, 0.5037, 0.5060]$ and $\sigma = [0.1030, 0.0999, 0.0969]$ for the three RGB channels, respectively.

CIFAR100. We use 32×32 images with normalization parameters typically used in the literature, i.e., $\mu = [0.4914, 0.4822, 0.4465]$ and $\sigma = [0.247, 0.243, 0.261]$. For data augmentation, we follow EquiMod's augmentation strategy.

Tiny ImageNet. The training set consists of 100000 ImageNet-1k images from 200 classes (500 for each class) downsized to 64×64 . The validation set has 50 images per class. We use normalization parameters $\mu = [0.4914, 0.4822, 0.4465]$ and $\sigma = [0.247, 0.243, 0.261]$, for the three RGB channels, respectively. For data augmentation, we use the same augmentation parameters as CIFAR100 with the kernel size of Gaussian blur adapted to the 64×64 images.

STL10. In order to extract the saliencies, we resize images to 512×512 , feed them to the pretrained DeepGaze IIE [Linardos et al., 2021], resize the output saliencies back to 96×96 , and store them alongside original images. We use normalization parameters $\mu = [0.4467, 0.4398, 0.4066]$, $\sigma = [0.2241, 0.2215, 0.2239]$ for the three RGB channels, respectively. After sampling fixations from saliencies, the patches that are extracted from the image to simulate foveation are 32×32 (compared to the full image size of 96×96). For IoR, we zero-out a circular area with radius of 16 around each previous fixation.

Transformation parameters. For 3DIEBench, we use the rotation and color parameters provided with images for action conditioning as done in Garrido et al. [2023]. For the augmentation setting, we use the parameters corresponding to each of the three augmentations and form the action as the relative augmentation vector between two images. For crop, we use four variables, i.e., vertical and horizontal coordinate, and height and width. For color jitter, we use four variables: brightness, contrast, saturation, and hue. For blur, we use one variable: the standard deviation of the blurring kernel. In predictive learning across saccades, the action is a 2-d vector, i.e., the normalized relative (x,y) coordinate between two patches.

A.2 Training and evaluation details

Additional Training Details. We used the PyTorch framework for training all models. For experiments that use CIFAR100 and low-resolution STL-10 patches in predictive learning across saccades, we use the CIFAR variant of ResNet-18. For models trained with AdamW, we used with default β_1 and β_2 , a weight decay of 0.001, and a learning rate of 4×10^{-4} with a linear warmup for 20 epochs starting from 10^{-5} followed by a cosine decay back to 10^{-5} . For models trained with Adam, we use the Adam optimizer with a learning rate of 10^{-3} , default β_1 and β_2 , and no weight decay.

Protocols for training evaluation heads. For linear probing, we follow a common SSL protocol and train a linear classifier on top of frozen representations with a batch size of 256 for 300 epochs

using the Adam optimizer with default hyperparameters. For action prediction, we follow a similar protocol as SIE [Garrido et al., 2023]. Specifically, for rotation, color jitter, and crop, we train an MLP regressor with a hidden dimension of 1024 and ReLU activation for 300 epochs. For color (in 3DIEBench), blur (in the augmentation setting), and position (in predictive learning across saccades), we use a linear regressor and train it for 50 epochs. For path integration experiments, the same regressor architectures as the equivariance evaluation heads are used, i.e., the MLP for angular path integration and the linear regressor for saccade path integration. All regression heads are trained using the Adam optimizer with default hyperparameters.

Hardware. Each experiment was run on a single NVIDIA A100 GPU with 40GB of accelerator RAM.

Compute cost and FLOP analysis. We report pretraining compute time across methods (see Table 5). All experiments are run under the same configuration: a single A100 GPU, 128×128 resolution, batch size 512, for 1000 epochs on 3DIEBench. seq-JEPA with a sequence length of one has similar runtime to other baselines (e.g., BYOL, SimCLR), while seq-JEPA with sequence length of three incurs a moderate increase in wall-clock time (15.1 GPU-hours) due to encoding additional action-view pairs.

Table 5: GPU-Hour comparison across methods

11 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1					
Method	A100 pretraining GPU hours				
SimCLR	11.4				
BYOL	11.1				
SIE	12.0				
VICReg	10.9				
EquiMod	11.8				
seq-JEPA (train seq $len = 1)$	12.3				
seq-JEPA (train seq $len = 3)$	15.1				

seq-JEPA's inference cost is primarily governed by two factors: sequence length and input resolution. Below, we report detailed FLOP counts (in Gigaflops) per forward pass of a single datapoint across a range of configurations (Note: these reflect inference-time cost).

Table 6: FLOPs for different seq-JEPA configurations (in gigaflops)

Resolution / Config	Encoder only	Post-aggregator				
		seq len = 1	seq len = 2	seq len = 3	seq len = 4	seq len = 8
Each view is 128×128	0.60G	0.63G	1.24G	1.85G	2.46G	4.9G
Views of 64×64 vs. full image (224×224)	1.82G (224) / 0.15G (64)	0.18G	0.34G	0.51G	0.67G	1.32G
Views of 84×84 , vs. full image (224×224)	1.82G (224) / 0.30G (84)	0.33G	0.64G	0.95G	1.26G	2.51G

We observe that for single-view inference at 128×128 resolution, the encoder requires only 0.60G FLOPs, which matches the compute cost of standard SSL baselines such as SimCLR and BYOL. As sequence length increases, the additional compute cost grows sub-linearly as additional views should be encoded and aggregated. For example, lengths 2 and 3—where invariant performance improves substantially—incur \sim 1.2G and \sim 1.85G FLOPs, respectively. The cost can be offset by reducing input resolution. In our saccade-based setup, we can feed low-resolution foveal glimpses (e.g., 64×64 or 84×84 crops sampled via saliency maps), which require only 0.15G–0.30G FLOPs per view. This enables the use of longer sequences at an overall compute cost comparable to full-resolution, single-view pipelines (e.g., compare the FLOPs of aggregating eight 64×64 crops with encoding a single 224×224 image in the second row of Table 6).

A.3 Architectural details

Below, we describe the architectural details and hyperparameters specific to each baseline.

BYOL. We use a a projection head with 2048-2048-2048 intermediate dimensions. The predictor has a hidden dimension of 512-d with ReLU activation. We use the same EMA setup outlined in the original paper [Guo et al., 2020], i.e., the EMA parameter τ starts from $\tau_{\rm base}=0.996$ and is increased following a cosine schedule.

SimCLR. We use a temperature parameter of $\tau = 0.5$ with a projection MLP with 2048-2048 intermediate dimensions.

VICReg. We use $\lambda_{\text{inv}} = \lambda_V = 10$, $\lambda_C = 1$, and a projection head with 2048-2048-2048 intermediate dimensions.

VICRegTraj. We use the same architecture as VICReg and add a trajectory loss (Equation 2 in Wang et al. [2024a]) with a coefficient $\lambda_{traj} = 0.01$.

Conditional BYOL. The architecture is the same as BYOL, except that the predictor also receives the normalized relative transformation parameters. We use a linear action projector of 128-d and the same EMA setup as BYOL.

SIE. For both invariant and equivariant projection heads, we use intermediate dimensions of 1024-1024-1024. For the loss coefficients, we use $\lambda_{\text{inv}} = \lambda_V = 10$, $\lambda_{\text{equi}} = 4.5$, and $\lambda_C = 1$. We use the hypernetwork architecture for all experiments.

SEN. We use a temperature parameter of $\tau=0.5$ with a projection MLP with 2048-2048-2048 intermediate dimensions.

EquiMod. We use the version based on SimCLR (both invariant and equivariant losses are contrastive with $\tau=0.1$ and have equal weights). The projection head has 1024-1024-128 intermediate dimensions. We use a linear action projector of 128-d.

ContextSSL. We use the pre-trained weights provided by the authors (trained for 1000 epochs) and follow their evaluation protocol on 3DIEBench.

seq-JEPA. For the sequence aggregator, we use a transformer encoder with three layers, four attention heads, and post-normalization. For the predictor, we use an MLP with a hidden layer of 1024-d and ReLU activation. The linear action projector in our default setting is 128-d. We use the same EMA setup as BYOL.

B Additional Experimental Results

B.1 Evaluation results on 3DIEBench for models conditioned on rotation and color

Table 7 reports performance of seq-JEPA and several equivariant baselines when conditioned on both rotation and color in 3DIEBench. All methods suffer a drop in classification performance and become highly sensitive to color. Similar performance degradations have been previously observed with 3DIEBench [Garrido et al., 2023, Gupta et al., 2024], though without a clear explanation. One plausible explanation aligned with Principle II in Wang et al. [2024b], is that color in 3DIEBench, i.e., floor and light hue, are weakly correlated with class labels in 3DIEBench (low class relevance). Therefore, forcing the encoder to encode color information would cause class information to be lost, resulting in degradation of classification accuracy.

Table 7: Evaluation on 3DIEBench for rotation and color prediction (equivariance) and linear probe classification (invariance). All models are conditioned on both rotation and color.

Method	Classification (top-1)	Rotation pred. (R^2)	Color pred. (R^2)
SEN	82.17	0.29	0.96
EquiMod	82.58	0.27	0.95
Conditional BYOL	81.95	0.38	0.94
SIE	75.34	0.46	0.97
seq-JEPA ($M_{tr} = 4$, $M_{val} = 5$)	79.31	0.52	0.97

B.2 Measuring invariance of aggregate representations

In order to have a more direct measure of invariance of the aggregate representation, we trained a model on 3DIEBench with $M_{val}=3$ and rotation conditioning, then attempted to predict the individual rotation parameters of each input view from $z_{\rm AGG}$. To prevent leakage of rotation information

during prediction, all action embeddings were zeroed out. The resulting mean R^2 score obtained by predicting individual rotations of the three views using z_{AGG} is 0.141 ± 0.002 ; indicating that the aggregate representation contains little to no rotation-specific information.

B.3 Results with a transformer projector instead of an MLP projector for baselines

By using a transformer projector for baselines that normally use an MLP projector, we tested whether transformer-based projectors alone account for any performance gains. We did not see any benefit from switching to transformer projectors in any of the baselines. Table 8 shows the transformer-projector results for the baselines on 3DIEBench (plus sign indicates a transformer projector).

Table 8: Comparison of baselines trained with transformer projectors on 3DIEBench.

Setting\Metric	Top-1 Acc	R^2 (rel rot)	R^2 (abs rot)
SimCLR+	77.92	0.32	0.48
BYOL+	81.05	0.09	0.16
VICReg+	75.10	0.15	0.28
VICRegTraj+	77.82	0.24	0.36
SEN+	81.74	0.31	0.44
Cond. BYOL+	73.62	0.27	0.36

B.4 Out-of-distribution results

We created an OOD set for 3DIEBench to enable evaluation on unseen transformations. The original dataset samples rotations from $(-\pi/2,\pi/2)$. Our OOD test set instead uses the disjoint range $(-\pi,-\pi/2)\cup(\pi/2,\pi)$, ensuring no angular overlap with the training set. Rendering this dataset took around five hours on a single A6000 GPU. From Table 9, we see that all methods fail at OOD rotation decoding (with R^2 s even reaching near -1.0). The sharp drop in R^2 values may be attributed to a domain shift in transformation geometry: while representations are well-aligned with in-distribution rotation trajectories (inside the range $(-\pi/2,\pi/2)$), they are not geometrically structured to extrapolate beyond this range to $(-\pi,-\pi/2)\cup(\pi/2,\pi)$. Despite the failure in OOD equivariance generalization, seq-JEPA exhibits a graceful degradation in invariant classification accuracy compared to other baselines. We hypothesize that this robustness stems from its sequence-level aggregation mechanism: even when latent representations become less equivariant under OOD transformations, aggregation across multiple views still filters out transformation-specific variability and recovers shared semantic content. Thus, seq-JEPA maintains object identity better under distribution shift, indicating that aggregation over input views promotes invariance, even when equivariance is imperfect.

Table 9: Unseen Transformation Generalization (OOD Rotations)

Setting\Metric	Top-1 Acc (drop)	OOD R^2 (rel rot)	OOD R^2 (abs rot)
SimCLR	63.86 (-17.27)	-0.41	-0.201
BYOL	55.68 (-27.22)	-0.25	-0.198
VICReg	61.28 (-19.20)	-0.31	-0.206
SEN	60.03 (-23.40)	-0.41	-0.201
SIE	60.19 (-17.30)	-0.63	-0.200
EquiMoD	58.54 (-25.75)	-0.45	-0.206
Cond. BYOL	57.91 (-24.70)	-0.42	-0.202
VICRegTraj	62.94 (-18.32)	-0.36	-0.214
seq-JEPA (1,1)	61.53 (-22.55)	-0.69	-0.199
seq-JEPA (3,3)	65.03 (-21.11)	-0.71	-0.211

B.5 Equivariance predictor evaluation metrics

Following the protocol in Garrido et al. [2023], we compute and report MRR, Hit@1, and Hit@5 metrics to evaluate predictor quality on the 3DIEBench validation set for seq-JEPA and two predictor-based baselines (Table 10). These results confirm that seq-JEPA achieves strong equivariance performance not only in terms of R^2 , but also in top-rank retrieval metrics.

Table 10: Predictor evaluation metrics on 3DIEBench validation set.

Setting	MRR	H@1	H@5
SIE	0.319	0.215	0.404
EquiMod	0.136	0.037	0.186
seq-JEPA $(1,1)$	0.340	0.241	0.442
seq-JEPA (3,3)	0.388	0.273	0.468

B.6 Training with more compute

To examine convergence under a high-compute regime, we trained five models using 2000 epochs, 256×256 resolution, and batch size 1024. Each of these experiments was run on 4 A100 GPUs (our main experiments were run on a single A100). The evaluation results (Table 11) confirm that seq-JEPA achieves a strong performance in the high-compute regime without suffering from a trade-off between invariance and equivariance. Importantly, our method achieves near-saturated performance already in the low-compute regime, indicating that it requires fewer steps for convergence and is less sensitive to input resolution and batch size than competing methods.

Table 11: Evaluation results under the high-compute regime.

Setting	Top-1 Acc	\mathbb{R}^2 (rel rot)	\mathbb{R}^2 (abs rot)	MRR	H@1	H@5
SIE	82.652	0.721	0.764	0.411	0.287	0.490
SimCLR	85.961	0.473	0.609	_	_	_
EquiMod	86.833	0.492	0.625	0.154	0.048	0.201
seq-JEPA (1,1)	85.370	0.661	0.713	0.365	0.263	0.447
seq-JEPA (3,3)	87.581	0.736	0.781	0.419	0.282	0.483

B.7 Complete evaluation results for linear probing on top of aggregate representations

We provide our complete evaluation results for linear probing on top of seq-JEPA's aggregate representations for our three transformation settings with different training and inference sequence lengths. Figure 8 shows the top-1 accuracy on 3DIEBench models conditioned on rotation. Figure 9 shows top-1 accuracy on STL-10 for models trained via predictive learning across saccades. Figures 10 and 11 show top-1 accuracy on CIFAR100 and Tiny ImageNet, respectively, with different types of action conditioning (crop, color jitter, blur, or all three). These heatmaps reflect the same trends observed in Figure 7 and discussed in Section 4.6, illustrating the consistent effect of sequence length on representation quality.

B.8 Comparison of evaluation results on encoder representations and aggregate representations

For completeness, we provide linear probe classification on encoder representations for different transformation settings in Table 12 and compare them with accuracy on aggregate representations for different inference evaluation lengths. The aggregate representation generally achieves a much higher classification performance thanks to the architectural inductive bias in seq-JEPA.

B.9 Additional UMAP Visualizations

In Figure 12, we visualize the UMAP projections of seq-JEPA representations trained on 3DIEBench without action conditioning. Similarly, Figure 14 shows projections for models trained on STL-10

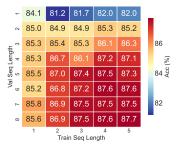


Figure 8: seq-JEPA's performance on 3DIEBench with rotation conditioning; the heatmap shows linear probe accuracy on top of aggregate representations for different training and inference sequence lengths.

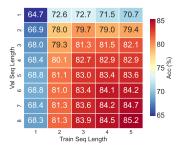


Figure 9: seq-JEPA's performance on STL-10 with predictive learning across saccades; the heatmap shows linear probe accuracy on top of aggregate representations for different training and inference sequence lengths.

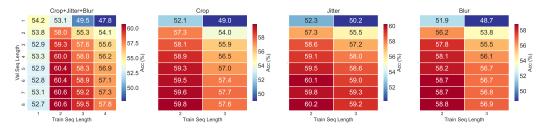


Figure 10: seq-JEPA's performance on CIFAR100 with different types of action conditioning (crop, color jitter, blur, or all three); the heatmap shows linear probe accuracy on top of aggregate representations for different training and inference sequence lengths.

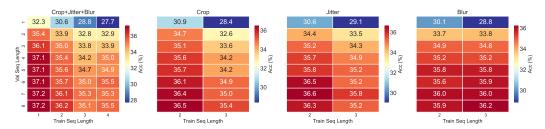


Figure 11: seq-JEPA's performance on Tiny ImageNet with different types of action conditioning (crop, color jitter, blur, or all three); the heatmap shows linear probe accuracy on top of aggregate representations for different training and inference sequence lengths.

via predictive learning across saccades, also without action conditioning. Compared to the action-conditioned counterparts (Figures 5 and 13), these projections exhibit weaker or no smooth color

Table 12: Comparison of seq-JEPA classification performance across datasets and conditioning. Top-1 classification accuracy is reported for z_{res} and z_{agg} , with varying inference lengths M_{eval} .

Dataset	Conditioning	M_{tr}	z_{res}		z_{agg}	
				$\overline{M_{eval}} = 1$	$M_{eval} = 3$	$M_{eval} = 5$
3DIEBench	None	3	80.91	81.61	86.05	87.36
3DIEBench	Rotation	1	84.88	84.08	85.34	85.31
3DIEBench	Rotation	3	82.49	81.72	85.32	87.41
3DIEBench	Rotation + Color	4	74.88	71.14	75.97	79.31
CIFAR100	None	2	53.00	51.60	57.05	58.37
CIFAR100	Crop + Jitter + Blur	1	56.23	54.24	52.90	52.92
CIFAR100	Crop + Jitter + Blur	2	52.07	53.07	59.34	60.35
CIFAR100	Crop + Jitter + Blur	3	46.31	49.48	57.60	58.33
CIFAR100	Crop	2	52.62	52.06	58.07	59.32
CIFAR100	Color Jitter	3	54.92	50.21	57.20	58.62
CIFAR100	Blur	3	51.41	48.69	55.54	56.72
Tiny ImageNet	None	2	32.84	30.48	35.03	35.97
Tiny ImageNet	Crop + Jitter + Blur	1	33.03	32.34	36.14	37.07
Tiny ImageNet	Crop + Jitter + Blur	2	27.20	30.57	34.99	35.56
Tiny ImageNet	Crop + Jitter + Blur	3	24.74	28.84	33.78	34.68
Tiny ImageNet	Crop	2	31.13	30.89	35.07	35.69
Tiny ImageNet	Color Jitter	3	31.85	29.05	34.27	35.21
Tiny ImageNet	Blur	3	27.20	28.83	34.82	35.79
STL-10	None	4	61.38	62.21	69.06	70.45
STL-10	Position	4	81.20	71.45	81.53	83.44
STL-10	Position (no saliency)	4	79.29	63.14	76.93	79.85
STL-10	Position (no IoR)	4	72.49	68.95	76.84	77.97

gradients in their corresponding transformation-colored UMAP—indicating reduced equivariance to transformation parameters.

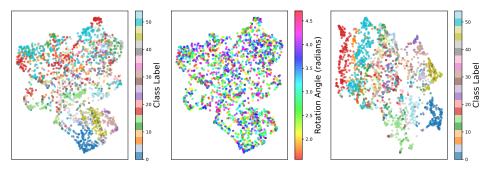


Figure 12: 2-D UMAP projections of seq-JEPA representations on 3DIEBench without action conditioning $(M_{tr} = 3, M_{val} = 5)$. Encoder outputs are color-coded by class (**left**) and rotation angle (**middle**); aggregate token representations are color-coded by class (**right**).

B.10 Details of Path Integration Experiments.

While an agent executes a sequence of actions in an environment, transitioning from an initial state to a final state, it should be capable of tracking its position by integrating its own actions. This is also a crucial cognitive ability that enables animals to estimate their current state in their habitat [McNaughton et al., 2006]. Here, we evaluate whether seq-JEPA is capable of path integration. Given the sequence of observations $\{x_i\}_{i=1}^{M+1}$ generated from transformations $\{t_i\}_{i=1}^{M+1}$ and the corresponding relative action embeddings $\{a_i\}_{i=1}^{M}$, we define the task of path integration over the sequence of actions as predicting the relative action that would directly transform x_1 to x_{M+1} given z_{AGG} and a_M . In other words, given the aggregate representation of a sequence of action-observation pairs and the next action, we would like to predict the overall position change from the starting point (x_1) to the end point (x_{M+1}) . We consider path integration for rotation angles in 3DIEBench and

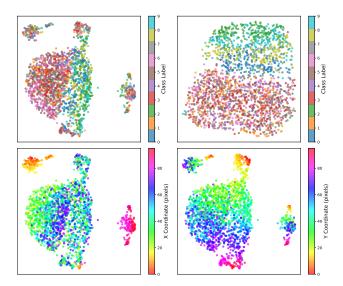


Figure 13: 2-D UMAP projections of seq-JEPA representations on STL-10 with action conditioning ($M_{tr} = M_{val} = 4$). Encoder outputs are color-coded by class (top-left) and by X/Y fixation coordinates (bottom); the aggregate token is color-coded by class (top-right).

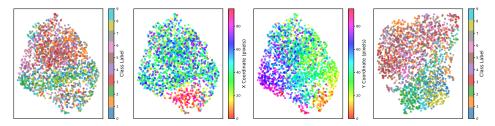


Figure 14: 2-D UMAP projections of seq-JEPA representations on STL-10 without action conditioning $(M_{tr} = M_{val} = 4)$. Encoder outputs are color-coded by class (**left**) and fixation coordinates (**middle**); the aggregate token is color-coded by class (**right**).

across eye movements with STL-10. For rotations, the task is integrating a series of object rotations from the first view to the last, i.e. angular path integration. For eye movements, the task is integrating the eye movements from the first saccade to the last, i.e. visual path integration. To measure path integration performance for inference sequence length M, we train a regression head on top of the concatenation of z_{AGG} and a_M to predict the transformation from x_1 to x_{M+1} . Figure 6 shows that seq-JEPA performs well in both angular and visual path integration. The red curve corresponds to the performance of the original seq-JEPA. The blue curve corresponds to experiments in which the action embeddings are ablated (zeroed-out during inference for all views). The green curve corresponds to experiments in which the encoder (visual) representations are ablated during inference. As expected, path integration becomes more difficult as the number of observations increases (red curves). Ablating action conditioning (blue curves) results in failure of path integration. On the other hand, ablating the visual representations (green curves) results only in a small performance drop compared to the original model, indicating that action conditioning is the key factor that enables path integration.

B.11 Transfer learning results on ImageNet-1k

To evaluate generalization of our model beyond STL-10, we assess transfer performance of the model trained on STL-10 patches via predictive learning across saccades on ImageNet-1k. We follow the same linear probing protocol as in-distribution evaluations: we freeze the ResNet and transformer encoder and train a linear classifier on aggregate representations generated from foveated patches.

We extract saliency maps for ImageNet-1k images using DeepGaze IIE, resize them to 224×224 , and sample sequences of patches sized 32×32 or 84×84 for this out-of-distribution (OOD) evaluation setting. Figure 15 shows top-1 linear probe accuracy on ImageNet-1k validation set across varying inference sequence lengths (M_{val}). For both patch sizes, performance improves with longer inference sequences, validating seq-JEPA's ability to benefit from extended context even in this difficult OOD ImageNet-1k setting. These results echo findings in the main experiments for different training and inference sequence lengths and confirming the possibility of model's scalability in terms of data, parameter count, and compute.

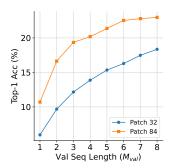


Figure 15: Linear probe transfer learning accuracy on ImageNet-1k for two different patch sizes; the model is pre-trained on STL-10 via predictive learning across saccades.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: All the claims in the abstract and introduction (summarized by the items at the end of introduction) are supported by controlled experimental results on different benchmarks.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The main limitations are discussed in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The work does not contain theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All details required to reproduce our results even without access to the code provided or are included in Sections 3.2 and 3.3 and Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code, date, and checkpoints have been made publicly available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/ public/quides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https: //nips.cc/public/quides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- · At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Ouestion: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The most important experimental details necessary to understand the results are discussed in the main text (Section 3). Other details are provided in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The main experimental results in Tables 1, 2, and 3 are averaged over three independent random seeds. For additional experiments, we pre-trained one model per setting due to compute limitations. Standard deviations are reported in Table 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The hardware setup and approximate wall-clock time are provided in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The paper complies with NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The paper proposes a method and architecture that is not closely tied to any societal impact.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper is not releasing such data or models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All data used in our experiments are open-source. We credit and cite the datasets and methods in the paper when necessary.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We use STL-10 and ImageNet-1K saliency maps and an out-of-distribution set of 3DIEBench. We have made all the three publicly available

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.