REPRESENTING SHORT-TERM MOVEMENT TO REVEAL ANIMAL IDENTITY AND LEARNING PROGRESS

Anonymous authors
Paper under double-blind review

ABSTRACT

How learning progresses over time and varies across individuals is a longstanding question in cognitive and behavioral science. Behavioral changes are observable outcomes of learning and have been a long-established bedrock for understanding learning progress. However, a major challenge is identifying observable behavioral measures that faithfully capture learning progression. Existing approaches commonly rely on predefined behavioral features or task-level performance, while fine-scale posture and movement in freely behaving animals remain less explored. In this paper, we ask: *can short-term movement trajectories reveal learning progress and individual identity?* We introduce a self-supervised representation learning framework built on a simple hypothesis: behavioral characteristics associated with identity and learning remain relatively stable over nearby time intervals. Using temporal proximity as supervision, our framework learns compact representations directly from movement trajectories without manually defining behavioral states or kinematic features. Using behavior data from freely behaving rats, we show that the learned representations not only improve prediction of identity and learning-related readouts over existing methods, but also provide a tool for investigating broader scientific questions about when and how these signals appear in behavior. Specifically, trajectories as short as one second contain informative learning and identity signals. Egocentric pose dynamics are more predictive of identity, whereas allocentric motion is more predictive of learning-related variation, and both signals are strongest at decision junctions. We further use a representation-conditioned diffusion model to inspect movement patterns associated with identity and learning progress. Our findings establish short-term movement as an informative behavioral "fingerprint" for studying learning and individual differences.

1 INTRODUCTION

A longstanding question in cognitive science is how learning progresses over time. We can study this progression by characterizing changes in brain networks and behavior. While characterizing neural changes requires recording brain activity, behavioral changes provide a directly observable window into the learning process. A major challenge, however, is identifying which aspects of behavior show learning-related changes and how they evolve. Standard approaches rely on predefined behavioral features, such as speed, latency, or the frequency of a target behavior, as well as task-level performance measures, such as success rate. These measures are interpretable and readily quantifiable, but often summarize behavior as coarse scalar changes over extended periods (Smith et al., 2004). In contrast, fine-scale changes in limb posture and movement can provide a wealth of information about learning progression, as demonstrated in studies of motor learning (Kawai et al., 2015; Bova et al., 2021; Dhawale et al., 2019; Nguyen et al., 2021). Yet, in unconstrained settings, where animals learn to navigate freely in space (Rosenberg et al., 2021), the complexity of movement patterns makes it difficult to determine which behavioral features to measure and how they can be used to understand changes in the underlying cognitive processes. Addressing this challenge could reveal the hidden structures and dynamics of learning.

In addition to behavioral changes over time, inter-individual differences in movement have been reported across species (Godesberg et al., 2024; Chao et al., 2019). Such movement "fingerprints" are increasingly being linked to biologically important variation in genetic and cognitive traits (Levy et al., 2023; Forkosh et al., 2019; Freund et al., 2013; von Ziegler et al., 2024; Sheppard et al., 2022;

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

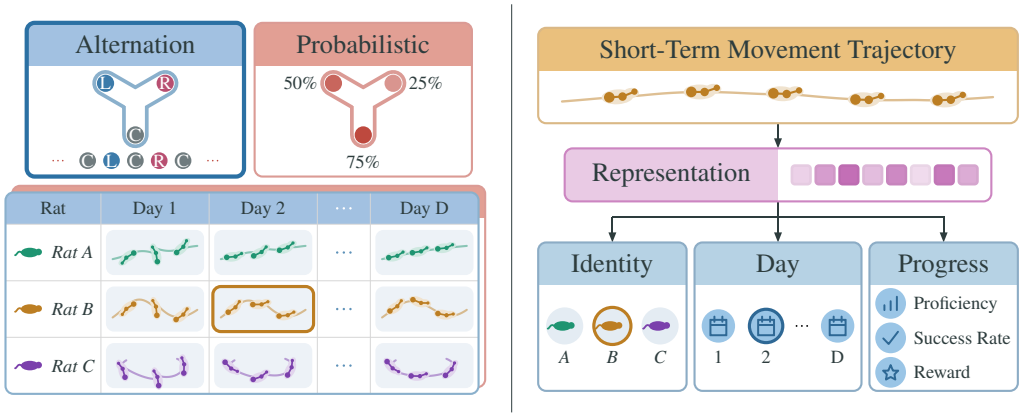

Figure 1: **Dataset and problem setting**. *Left:* Repeated pose recordings span animals and days under two reward rules over the Y-shape spatial learning dataset: alternation and probabilistic rewards. *Right:* A compact representation of short-term movement trajectories is learned and evaluated for inferring identity, recording day, and performance (e.g., proficiency, success rate, and reward).

Ding et al., 2026). If successfully applied to learning-related movement analysis, these fingerprints could add a new dimension to understanding individual learning dynamics. However, current methods vary widely in their performance and the observation periods required to identify individuals, ranging from tens of seconds to minutes (Levy et al., 2023; Forkosh et al., 2019; von Ziegler et al., 2024).

These limitations raise a fundamental question: *can short-term movement trajectories reveal learning progress and individual identity?* We focus on body posture and movement patterns unfolding over only a few seconds—a timescale that is often insufficient to assess task-level performance. Nevertheless, how animals move, coordinate their body parts, or respond to their surroundings within these short intervals may already provide informative signals (Weinreb et al., 2026). We investigate whether a representation learned directly from movement trajectories can capture these signals without manually defining behavioral states or kinematic features.

To this end, we study movement data from rat spatial learning experiments described in (Shridhar et al., 2026). We first evaluate existing approaches for representing movement trajectories, spanning traditional state-based methods (Durbin & Koopman, 2012; Wiltschko et al., 2015; Paraschos et al., 2013; Weinreb et al., 2024), recent machine learning methods (Hinton & Salakhutdinov, 2006; Luxem et al., 2022; Grill et al., 2020; Bai et al., 2018; Azabou et al., 2023; Stoffl et al., 2025), and time-series foundation models (Google Research, 2026; Ansari et al., 2025). Across identity- and learning-related evaluations such as task performance, these approaches have limited predictive performance.

We therefore develop a representation learning framework, TRAIL (TRAjectory-Identity Latent), tailored to capturing information about individual differences and learning progress from short segments of movement data, as illustrated in Figure 1. Our framework combines a trajectory encoder with a self-supervised objective. Inspired by Wayformer (Nayakanti et al., 2022), an architecture developed for short-term motion forecasting in autonomous driving, the encoder maps short movement trajectories into compact latent representations. Our training objective builds on a simple hypothesis: although an animal's movements may vary rapidly, the underlying behavioral characteristics associated with identity and learning are more stable over nearby time intervals. We therefore use temporal proximity as a self-supervisory signal, encouraging representations of nearby trajectory segments to be similar while contrasting them with segments from other recordings or sufficiently distant times within the same recording. This objective enables representation learning without identity or learning-progress labels during training.

We evaluate our framework on the spatial learning data (Shridhar et al., 2026) and find clear improvements over existing methods. Linear probes trained on the learned representations predict learning-related readouts more accurately, demonstrating that such information is accessible from short movement trajectories. Interestingly, the representations also support accurate identity prediction, even though segments from different recordings of the same individual can be contrasted during training. Beyond these learning-focused datasets, our method remains competitive on the established behavioral benchmark MABe22 (Sun et al., 2023), suggesting broader applicability to representation learning from short movement sequences.

The superior performance of our framework not only provides an affirmative answer to our fundamental question, but also offers a tool for investigating broader scientific questions: *At what timescale does behavior vary in relation to learning? How long must a trajectory be to capture identity- and learning-related signals, and under what conditions are these signals strongest? Which movement patterns reveal individual identity and learning progress?*

First, our temporal-window analysis suggests that behaviorally relevant variation occurs on a timescale of approximately 60 seconds: treating segments beyond this window as behaviorally uniform substantially reduces performance. Second, trajectories as short as one second already contain information predictive of both identity and learning-related readouts. Both signals are stronger at locations where animals must choose a direction than along straight paths. Third, local pose dynamics are more predictive of identity, whereas global motion is more predictive of learning-related readouts. Together, these findings identify temporal, spatial, and postural factors that can guide more targeted hypotheses about behavior and learning.

Finally, we use a representation-conditioned diffusion model to explore the movement patterns associated with individual identity and learning progress. By varying the conditioning representation, we generate trajectories that preserve broad motion structure while exhibiting subtle changes in stopping, accelerating, and stretching. These qualitative results make behavioral differences associated with the learned representations directly inspectable, providing a tool for hypothesis generation and subsequent experimental validation.

Taken together, our study shows that short movement trajectories contain observable signatures of both individual identity and learning progress. By combining self-supervised representation learning with controlled analyses and conditional generation, we provide a framework for predicting these properties and examining the movement patterns associated with them.

2 Related Work and Background

Behavioral tasks & datasets. CalMS21 (Sun et al., 2021) and MABe22 (Sun et al., 2023) support behavioral recognition and representation learning across diverse settings, while longitudinal resources such as CRCNS hc-28 (Jadhav & Frank, 2020) track rats learning rewarded alternation across days. In this paper, we use data from two spatial navigation tasks for rats (Shridhar et al., 2026). The alternation and probability reward rules define learning objectives with quantifiable performance metrics over days. Unlike spontaneous behavior, this design allows us to relate locomotion and posture features to learning progress and experience in the environment. Longitudinal recordings of the same animals over consecutive days enable analysis of individual learning trajectories, with body keypoint coordinates available for both conditions. The two reward rules provide complementary evaluation settings, including behavior under reward uncertainty.

Algorithms. Reconstruction, prediction, and representation alignment provide approaches to learning behavioral structure without labels. VAME (Luxem et al., 2022) combines reconstruction with future prediction, while BAMS (Azabou et al., 2023) and hBehaveMAE (Stoffl et al., 2025) capture behavioral structure at multiple temporal scales. Our approach combines contrastive learning with reconstruction to learn a compact representation from short trajectory windows. This representation supports analysis of movement changes during learning and differences across individuals.

3 Proposed Representation Learning Framework: TRAIL

We learn compact representations of short movements without predefined behavioral categories or hand-crafted features, and evaluate their ability to capture identity and learning-related variation.

3.1 Overview

Data format. We use data recorded from each animal over multiple days, with multiple sessions per day (Shridhar et al., 2026). Each trajectory segment $x_i \in \mathbb{R}^{T\times K\times 2}$ contains T frames of K keypoints and is associated with animal r_i, session s_i, start time t_i, and within-animal recording day d_i. We consider three views of movement: local pose, global trajectory, and full movement (Appendix C.2). We denote the selected view by $v_i = V(x_i)$ and optional environmental geometry by c_i.

Design principles. To capture identity-specific patterns and learning-related changes from short-term movements, we consider three design choices: (1) multi-frequency encoding to preserve fine-grained movement details (Mildenhall et al., 2022), (2) a compact latent space learned through contrastive learning, based on the hypothesis that nearby trajectory segments share similar underlying behavioral states, and (3) an auxiliary reconstruction objective. Together, these choices encourage TRAIL to retain informative movement details and learn behavioral relationships without predefined categories or explicit labels.

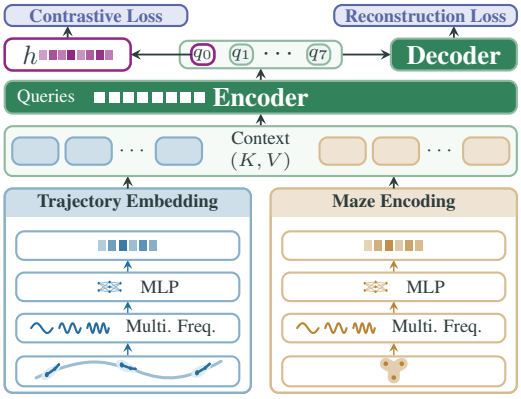

Figure 2: **TRAIL high-level architecture.**

High-level architecture. Figure 2 shows the encoder-decoder architecture of TRAIL. Inspired by Wayformer (Nayakanti et al., 2022), the encoder uses learnable queries to summarize multi-frequency embeddings of movement trajectories and environmental context (e.g., maze layout). The first output token is compressed into a behavioral representation trained through contrastive learning, and then projected back to the original feature dimension and combined with the remaining tokens for trajectory reconstruction. Reconstruction encourages the model to retain fine-grained movement details. We detail each component below.

3.2 Model Architecture and Learning Objectives

Multi-frequency trajectory embedding. We take keypoint trajectories $\mathbf{X} \in \mathbb{R}^{T\times K\times 2}$ and maze coordinates $\mathbf{M} \in \mathbb{R}^{N\times 2}$ as input; the maze provides global context for the trajectory. Inspired by Wayformer (Nayakanti et al., 2022), we anchor each trajectory to its final frame by centering the body point and vertically aligning the body-neck vector in that frame. We then apply the same transformation to all trajectory frames and maze coordinates. Each two-dimensional coordinate receives NeRF-style multi-frequency encoding with $L = 10$ frequencies (Mildenhall et al., 2022) and a linear projection to 64 dimensions. To identify each token's corresponding keypoint and time frame, we introduce two learnable positional embeddings: a keypoint embedding shared across frames and a temporal embedding shared across keypoints within the same frame. Maze coordinates undergo the same encoding procedure, with a learnable additive embedding distinguishing them from trajectory tokens. Together, the $K \times T$ trajectory tokens and N maze tokens form $\mathbf{Z} \in \mathbb{R}^{(KT+N)\times 64}$.

Compact behavioral representation. We adopt a four-layer Q-Former (Li et al., 2023) as the encoder, which takes eight learned queries to attend to $\mathbf{Z}$, yielding $\mathbf{Q} = [\mathbf{q}_0, \ldots, \mathbf{q}_7] \in \mathbb{R}^{8\times 64}$. We then linearly project $\mathbf{q}_0$ to produce the 16-dimensional latent behavior representation $\mathbf{h}$. To learn the representations, we hypothesize that nearby trajectory segments share similar underlying behavioral states. We therefore use a balanced InfoNCE loss (Oord et al., 2018). Positive pairs consist of segments from the same session within a certain time interval; negatives are drawn from temporally distant segments or different sessions:

$$\mathcal{L}_{\text{contrast}} = -\frac{1}{|\mathcal{P}|}\sum_{\mathbf{h}^+\in\mathcal{P}} \log \frac{\exp\big(s(\mathbf{h},\mathbf{h}^+)\big)}{\frac{1}{|\mathcal{P}|}\sum_{\mathbf{h}'\in\mathcal{P}}\exp\big(s(\mathbf{h},\mathbf{h}')\big) + \frac{1}{|\mathcal{N}|}\sum_{\mathbf{h}^-\in\mathcal{N}}\exp\big(s(\mathbf{h},\mathbf{h}^-)\big)}, \tag{1}$$

where $\mathcal{P}$ and $\mathcal{N}$ denote the positive and negative sets for an anchor representation $\mathbf{h}$, and s is the cosine similarity. The compressed representation $\mathbf{h}$ is then projected back to 64 dimensions and concatenated with the remaining tokens $q_{1:7}$ to form the decoder input for trajectory reconstruction.

Decoder for reconstruction. The decoder follows the same Q-Former architecture as the encoder, with separate parameters and learnable queries that attend to the encoder output. An MLP head maps the first query representation to the parameters of a Gaussian distribution over the trajectory, which we train using a negative log-likelihood reconstruction loss $\mathcal{L}_{\text{recon}}$.

Overall objective. The final objective weights contrastive learning and reconstruction equally:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{contrast}} + \mathcal{L}_{\text{recon}}. \tag{2}$$

We evaluate each component through ablation studies in Appendix C.1.

Table 1: **Performance on the spatial-learning dataset (alternation rewards)**. **Bold**: best; underlined: second best by unrounded values within each context group and column.

	Method	Identity (Acc ↑)	Day (MSE ↓)	Proficiency (MSE ↓)	Success Rate (MSE ↓)	Reward (MSE ↓)
Classic	Direct Trajectory	27.16	1.96	.021	0.14	0.76
	LG-SSM (Durbin & Koopman, 2012)	26.05	2.02	.021	0.14	0.77
	AR-HMM (Wiltschko et al., 2015)	44.09	3.69	.023	0.21	0.69
	ProMPs (Paraschos et al., 2013)	39.80	1.94	.021	0.14	0.76
	Keypoint-MoSeq (Weinreb et al., 2024)	27.43	1.98	.021	0.14	0.77
Foundation Model	TimesFM-3 (Zero-shot) (Google Research, 2026)	37.23	1.68	.019	0.14	0.71
	TimesFM-3 (LoRA)	44.41	1.64	.019	0.14	0.68
	Chronos-2 (Zero-shot) (Ansari et al., 2025)	37.63	1.72	.019	0.14	0.72
	Chronos-2 (LoRA)	46.83	1.61	.018	0.14	0.68
Deep Learning	AutoEncoder (Hinton & Salakhutdinov, 2006)	59.45	1.61	.018	0.13	0.68
	VAME (Luxem et al., 2022)	62.51	1.61	<u>.017</u>	<u>0.13</u>	0.68
	T-BYOL (Grill et al., 2020)	59.63	1.68	.018	0.13	0.70
	TCN (Bai et al., 2018)	65.43	<u>1.51</u>	.018	0.14	0.70
	BAMS (Azabou et al., 2023)	63.74	1.59	.018	0.13	0.67
	hBehaveMAE (Stoffl et al., 2025)	47.81	1.64	.018	**0.13**	<u>0.64</u>
	Multi-task Supervised	<u>78.41</u>	1.69	.018	0.16	0.80
	TRAIL	**84.00**	**1.26**	**.016**	0.14	**0.61**

4 Experiments on Representation Learning

Dataset description. We use data from two spatial-learning cohorts recorded in a Y-shaped maze (Shridhar et al., 2026): six rats learning an *alternation* rule that rewards goal visits in a prescribed sequence, and eight different rats learning a *probabilistic* rule with goal-specific reward probabilities. The maze consists of three goal platforms connected to a central junction by arms. We represent its geometry using 25 waypoints. Recordings span five consecutive days with 3–4 approximately ten-minute sessions per day at 11 Hz, totaling 275 sessions and 44.92 hours. We consider the nose, neck, and body keypoints, which are the most reliably tracked using DeepLabCut (Mathis et al., 2018). We split each session chronologically into 80% training, 5% validation, and 15% testing, and then aggregate them across sessions to form the training, validation, and test sets. (see Appendix A.1).

Evaluation metrics. We consider five evaluation targets: *Identity*, *Day*, trial-level *Proficiency*, and *Success Rate* and *Reward* over one-minute intervals. Binary reward outcomes are recorded for each trial that spans departure from one goal platform to arrival at the next; we then obtain the ground-truth proficiency following (Smith et al., 2004).

Unless stated otherwise, we freeze the encoder and train a separate linear probe for each target. We use logistic regression for identity classification and ridge regression for the remaining tasks. Probes are trained on the training trajectories and evaluated on the test trajectories. We report classification accuracy and mean squared error (MSE) for the regression tasks.

Baselines. We compare three baseline groups. *Classic* methods include linear probing on raw trajectories, Linear-Gaussian State-space Model (Durbin & Koopman, 2012), Autoregressive Hidden Markov Model (Wiltschko et al., 2015), Probabilistic Movement Primitives (Paraschos et al., 2013), and Keypoint-MoSeq (Weinreb et al., 2024). *Deep Learning* methods include AutoEncoder (Hinton & Salakhutdinov, 2006), VAME (Luxem et al., 2022), Temporal ConvNet (Bai et al., 2018), temporal BYOL (Grill et al., 2020), BAMS (Azabou et al., 2023), and hBehaveMAE (Stoffl et al., 2025). We also evaluate time-series *Foundation Models*, Chronos-2 (Ansari et al., 2025) and TimesFM-3 (Das et al., 2024; Google Research, 2026), frozen or with LoRA fine-tuning (Hu et al., 2022).

For our framework TRAIL, we also consider a *Multi-task Supervised* variant that uses the same trajectory embedding and encoder architecture but is trained from scratch jointly on all five target tasks, using cross-entropy for classification and MSE for regression.

Implementation details. For the short-context Y-maze comparison, all methods use 3-second trajectories (i.e., $T = 3 \times 11$), with BAMS restricted to its short-term branch. For each method, including TRAIL, we evaluate tens of hyperparameter configurations and select the best based on validation performance on Identity and Day in the alternation reward setting (Appendix A.2).

4.1 RESULTS

Spatial learning dataset. Table 1 shows that TRAIL outperforms all baselines on Identity and most learning-related tasks, except for Success Rate under the alternation reward rule. Its advantage over the multi-task supervised variant suggests that temporal contrastive learning and reconstruction encourage more effective representations than direct supervision alone. TRAIL also outperforms fine-tuned Chronos-2 and TimesFM-3, highlighting the value of representations tailored to movement trajectories. Similar trends hold under the probabilistic reward rule (Appendix B.1).

Mouse Triplet dataset. We also evaluate TRAIL on MABe22 Mouse Triplets (Sun et al., 2023), a behavioral dataset that does not focus on learning, comprising 5,336 one-minute clips and 13 tasks (4 sequence-level and 9 frame-level). Following BAMS and hBehaveMAE, we use inductive evaluation with short (2s) and long (45s) contexts. We additionally evaluate within-clip *Identity* by training a linear probe on the first 42 seconds and testing on the final 18 seconds. We omit this task in the long-context setting because the temporal splits are shorter than the context window. We use the reported configurations for BAMS and hBehaveMAE and those selected in our main experiments for the other baselines.

Table 2: **Performance on the MABe22 Mouse Triplet dataset with short- and long-term contexts.** MSE values scaled by 100.

	Method	ID↑ Acc	Tasks MSE↓	Tasks F1%↑	Frame F1%↑	Seq. F1%↑
	Short-term Context (2s)					
Classic	LG-SSM	65.32	9.42	21.01	13.17	56.25
	AR-HMM	45.62	9.36	22.50	14.70	57.61
	ProMPs	67.65	9.43	22.59	14.73	57.96
	Keypoint-MoSeq	63.50	9.39	22.42	13.67	61.79
Deep Learning	AutoEncoder	67.53	9.39	25.16	17.06	61.62
	VAME	68.24	9.36	25.93	18.06	61.34
	T-BYOL	69.30	9.43	25.72	17.13	64.36
	TCN	67.26	9.45	24.72	16.00	63.97
	BAMS	63.08	9.34	26.02	16.97	66.77
	hBehaveMAE	73.07	9.31	**29.88**	**20.78**	70.85
	TRAIL	**73.65**	**9.26**	28.64	17.81	**77.39**
	Long-term Context (45s)					
Deep Learning	BAMS	–	9.24	27.74	17.36	74.47
	hBehaveMAE	–	9.32	**29.13**	**20.05**	70.02
	TRAIL	–	**9.19**	29.02	17.90	**79.08**

Table 2 shows that TRAIL remains competitive at both context lengths and achieves the highest on Identity. Its weaker frame-level performance may reflect an emphasis on temporal consistency over rapidly changing movements. These results demonstrate the applicability of TRAIL across experimental settings and timescales.

5 SCIENTIFIC ANALYSIS AND INSIGHT

Having established strong identity and learning-related readouts across two datasets, we use TRAIL to investigate broader scientific questions: *At what timescale does behavior vary in relation to learning? How long must a trajectory be to capture identity- and learning-related signals, and under what conditions are these signals strongest? Which movement patterns reveal individual identity and learning progress?* We conduct systematic, controlled analyses using the six rats performing a Y-maze alternation task. These analyses move beyond predictive performance to explore how brief movements reveal individuality and learning-related changes.

At what timescale does behavior vary in relation to learning? Our contrastive learning objective assumes that learning unfolds gradually over hours to days (Shridhar et al., 2026), while behavioral states remain relatively stable within a local positive-pair window (Markowitz et al., 2023; Weinreb et al., 2026). Figure 3 (*left*) quantifies the effect of window size. Treating segments more than 60

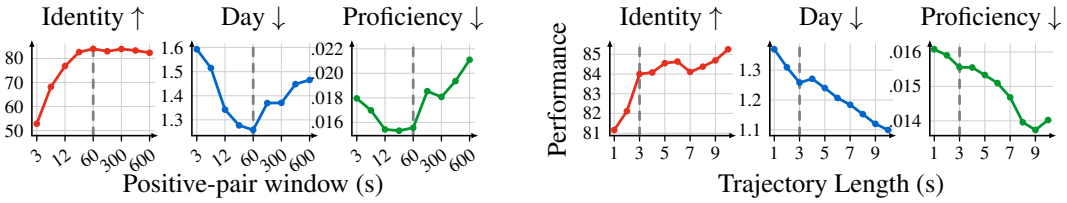

Figure 3: **Positive-pair window (*left*) and trajectory length (*right*).** Dashed line marks the default setting.

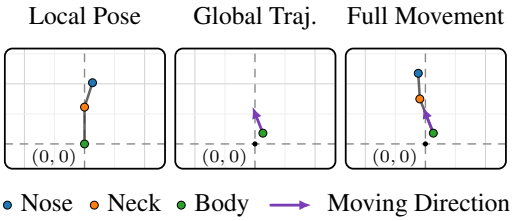

Figure 4: **Decomposition of movement dynamics.** Local Pose captures body posture changes, while Global Traj. captures spatial translocation.

Table 3: **Different movement dynamics carry complementary behavioral information**. Local body dynamics better reveal individuality, while global trajectories better capture learning-related changes.

	Identity↑	Day↓	Proficiency↓
Local Pose	<u>77.1</u>	1.61	0.019
Global Traj.	72.1	<u>1.47</u>	<u>0.017</u>
Full Movement	**84.0**	**1.26**	**0.016**

seconds apart as behaviorally uniform substantially reduces downstream performance, suggesting that behaviorally relevant differences may emerge on a timescale of approximately one minute.

How long must a trajectory be to capture identity- and learning-related signals? We examine the effect of trajectory length. As shown in Figure 3 (*right*), longer segments generally improve identity and proficiency prediction by providing more movement context. Interestingly, TRAIL remains competitive with the default three-second setting even with just one second of input, suggesting that brief movements already contain informative signals, particularly about individual identity.

Which movement dynamic tells more about identity and learning progress? Movement can be examined in egocentric or allocentric reference frames. Here, we decompose *Full Movement* into *Local Pose* and *Global Trajectory*: the former expresses other keypoints relative to the body keypoint at each frame, while the latter tracks the body keypoint over time (Figure 4; Appendix C.2). This allows us to examine how individual identity and learning-related changes are reflected in body posture and movement through space. As shown in Table 3, the two components exhibit different strengths: *Local Pose* is more predictive of identity, whereas *Global Trajectory* is more predictive of learning-related readouts. Combining both yields the strongest overall performance, suggesting that posture dynamics and global movement patterns carry complementary behavioral information.

How is individuality expressed in movement? To investigate whether TRAIL extracts latent representations that reveal individual identity, we project the behavioral latents of all animals into a common t-SNE space (Figure 5); the projection shows that each animal occupies a distinct region. We also fit a separate t-SNE projection for each animal using its latents from all five days. These per-animal views suggest day-related variation, although the extent of separation differs across animals and days can overlap. Together, the visualizations provide qualitative evidence of identity structure and across-day changes, complementing the quantitative decoding results.

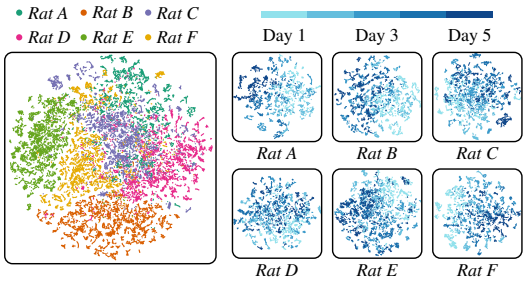

Figure 5: **t-SNE embeddings over 3s trajectory.** *Left*: across animals. *Right*: across days per rat.

Under what conditions are identity and learning-related signals most apparent? Having shown that movements contain signatures of identity and learning-related changes, we next ask

Table 4: **Behavior analysis across maze regions, recording days, and movement motifs.** Day prediction is omitted from the per-day analysis. **Bold** and <u>underlined</u> values indicate the best and second-best results within each subtable, based on unrounded values.

(a) Maze Regions.

	Identity↑	Day↓	Proficiency↓
Arm segment	81.7	1.59	0.021
Arm endpoint	<u>83.0</u>	<u>1.20</u>	<u>0.015</u>
Maze center	**88.7**	**1.20**	**0.014**

(b) Recording Days.

	Identity↑	Proficiency↓
Day 1	79.3	0.019
Day 2	81.9	0.018
Day 3	78.8	0.016
Day 4	**89.5**	<u>0.013</u>
Day 5	<u>87.6</u>	**0.012**

(c) Movement Motifs.

	Identity↑	Day↓	Proficiency↓
Remaining still	<u>87.1</u>	<u>1.11</u>	<u>0.012</u>
Scanning	76.6	1.34	0.020
Going straight	84.0	1.38	0.018
Turning	**98.0**	**1.07**	**0.008**
Turning back	84.2	1.56	0.020

whether these signals are expressed continuously or intermittently in the animal's behavior. We systematically quantify when, where, and during which movements these signals are most apparent (Appendix C.5). Tables 4a to 4c summarize the results. These signals are strongest at the *maze center* and clearer on Days 4–5 than on Days 1–3, suggesting that they become more apparent with experience. *Turning* is the most informative motif, consistent with its occurrence primarily at the maze center, while *remaining still* retains informative signals despite minimal movement.

6 INTERPRETATION

6.1 GOAL

Learning is accompanied by behavior changes that vary across individuals. Our preceding results show that the learned representations carry information about learning progress and individual identity. Here, we explore whether translating these representations back into movement can make these associations visible, offering a qualitative view of the behavioral differences they capture.

6.2 METHOD

To connect variation in representation space to visible changes in behavior, we first make the information encoded in these representations observable through generated movement, then vary the representations along directions associated with learning progression and individual differences.

Making representation observable. We train a conditional diffusion model $G_\psi(u, h)$ to generate movement trajectories from a diffusion latent noise u and a conditioning representation h produced by the frozen movement encoder (Peebles & Xie, 2023). Given an observed trajectory x_o with representation h_o, we obtain its latent noise through approximate DDIM inversion. We then hold u_o and the model parameters ψ fixed while varying the conditioning representation. The unmodified condition yields the reconstruction $\widehat{x}_o = G_\psi(u_o, h_o)$, providing a baseline for inspecting movement changes induced by representation transformation.

Tracing directions of progression. To trace how movement evolves as learning unfolds, we seek directions that organize progression across recording days, alongside variation across individuals. We estimate these directions using the leading principal components of daily representation centers within an animal or animal-specific centers on the same relative recording day. This provides a shared axis along which to follow gradual changes in movement. Given a selected unit direction v, we use an observed representation h_o and a reference representation h_r to define the transformation:

$$\begin{aligned} d &= \big[(h_r - h_o)^\top v\big] v, \\ h(\tau) &= h_o + \tau d, \qquad \tau \in [0, 1], \\ \widehat{x}(\tau) &= G_\psi\big(u_o, h(\tau)\big). \end{aligned} \tag{3}$$

Increasing τ traces a continuous path along this axis toward the reference. Decoding this path with the fixed latent u_o helps preserve the broad motion structure of x_0, such as the overall direction of travel, while allowing finer movement patterns to change. This lets us inspect how movement progressively shifts toward characteristics associated with a later recording day or another individual. Further implementation and training details are described in Appendices E.1 and E.2.

6.3 RESULTS

We examine movement changes from two complementary perspectives: progression across recording days within an animal, and variation between animals on the same relative recording day. Here, we focus on turning, which carries the strongest day and identity signals, as shown in Table 4c, making it an informative setting for both comparisons. Both comparisons generate Full Movement sequences. Our earlier findings in Section 5 (Table 3) show that Global Trajectory predicts day and proficiency more accurately than Local Pose, whereas Local Pose better predicts identity. These findings guide how we inspect the generated movements: through Body Trajectory to follow travel across days, and through Local Pose to follow relative body extension across individuals.

Across days, we examine how movement execution varies toward a later-day reference. The Body Trajectory sequence on Figure 6 (*left*) traces the body keypoint's path through space, allowing us to

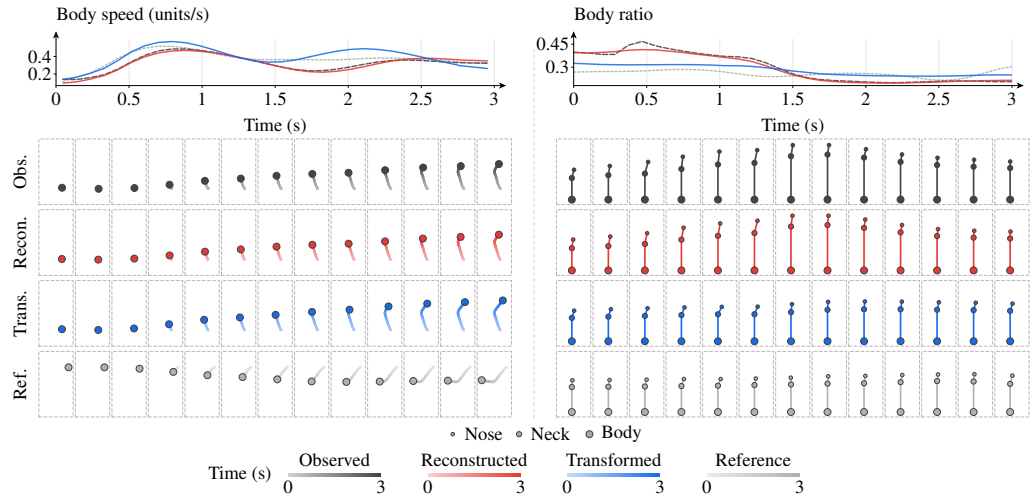

Figure 6: **Turning movement transformations across days and individuals.** *Left:* cross-day Body Trajectory with body speed, the magnitude of body-keypoint velocity. *Right:* cross-individual Local Pose with body ratio, the nose–neck distance divided by the neck–body distance. Obs., Recon., Trans., and Ref. denote observed, reconstructed, transformed, and reference movements, respectively.

follow both where the animal moves and how the turn unfolds. The transformed movement continues in the observed movement's overall direction while taking on aspects of the later-day reference's pace and turning rhythm. The body-speed curve above the sequence shows a corresponding shift toward the reference's speed over time. Together, these observations suggest that the transformation expresses day-related differences in movement execution while preserving the observed movement's broad spatial structure.

Across individuals, we ask how a movement takes on another animal's postural characteristics. The Local Pose sequence on Figure 6 (*right*) shows how the head, neck, and body extend relative to one another during the turn. As the representation shifts toward the reference animal, the transformed sequence adopts aspects of that animal's pattern of extension, accompanied by a body-ratio curve that moves toward the reference. Together, these observations make individual differences visible in how the head, neck, and body extend relative to one another during the same behavior, connecting the identity prediction results in Table 3 to visible local postural dynamics. The complementary views—cross-day Local Pose and cross-individual Body Trajectory—are provided in Appendix E.3, while Appendix E.4 extends both comparisons to the going-straight motif.

7 Discussion and Conclusion

In this paper, we present TRAIL, a framework for learning representations from short-term animal movements. We show that movement "fingerprints" can reliably predict individual identity and learning-related readouts. These results suggest that fine-scale movement patterns in freely behaving animals differ across individuals and change systematically over days of learning.

Beyond predictive performance, TRAIL provides a tool for generating scientific hypotheses. Through systematic, controlled analyses, we examine when, where, and in which movement patterns identity- and learning-related signals are most apparent. We find that their strength varies across recording days, spatial locations, and movement motifs, and that local pose dynamics and global motion carry complementary information. Conditional generation further allows us to inspect movement differences associated with identity and learning progress, providing hypotheses for subsequent experimental validation.

Moving forward, we aim to extend TRAIL to behavioral time-series data from freely moving animals and humans, particularly in experiments where movement patterns evolve over time. By helping researchers identify which aspects of behavior change and how they change, this framework could support more detailed studies of learning. Applied to larger populations, it could also help reveal how learning dynamics vary across individuals.

AI USE STATEMENT

The research and manuscript were primarily developed and written by the authors. Generative AI tools provided assistance with identifying potentially relevant literature, brainstorming aspects of the research, implementing code, editing LaTeX, preparing figures, and polishing the manuscript for clarity. All final scientific and methodological decisions were made by the authors, who take full responsibility for the content of this work.

REFERENCES

Abdul Fatir Ansari, Oleksandr Shchur, Jaris Küken, Andreas Auer, Boran Han, Pedro Mercado, Syama Sundar Rangapuram, Huibin Shen, Lorenzo Stella, Xiyuan Zhang, Mononito Goswami, Shubham Kapoor, Danielle C. Maddix, Pablo Guerron, Tony Hu, Junming Yin, Nick Erickson, Prateek Mutalik Desai, Hao Wang, Huzefa Rangwala, George Karypis, Yuyang Wang, and Michael Bohlke-Schneider. Chronos-2: From univariate to universal forecasting. *arXiv preprint arXiv:2510.15821*, 2025. URL https://arxiv.org/abs/2510.15821.

Mehdi Azabou, Michael Mendelson, Nauman Ahad, Maks Sorokin, Shantanu Thakoor, Carolina Urzay, and Eva Dyer. Relax, it doesn't matter how you get there: A new self-supervised approach for multi-timescale behavior analysis. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 28491–28509. Curran Associates, Inc., 2023. doi: 10.52202/075280-1238. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/5aad86aa2a3c00b70c71e19bc4780319-Paper-Conference.pdf.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv:1803.01271*, 2018.

Alexandra Bova, Kenneth Ferris, and Daniel K. Leventhal. Evolution of gross forelimb and fine digit kinematics during skilled reaching acquisition in rats. *eNeuro*, 8(5):ENEURO.0153–21.2021, 2021. doi: 10.1523/ENEURO.0153-21.2021.

Hanqing Chao, Yiwei He, Junping Zhang, and Jianfeng Feng. GaitSet: Regarding gait as a set for cross-view gait recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 8126–8133, 2019. doi: 10.1609/aaai.v33i01.33018126.

Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10148–10167. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/das24c.html.

Ashesh K. Dhawale, Yohsuke R. Miyamoto, Maurice A. Smith, and Bence P. Ölveczky. Adaptive regulation of motor variability. *Current Biology*, 29(21):3551–3562, 2019. doi: 10.1016/j.cub.2019.08.052.

Yiran Ding, Yuen Gao, Chunqi Qian, and Zijun Cui. GEESE: Genotype-aware end-to-end spatio-temporal embedding for behavioral phenotyping, 2026.

James Durbin and Siem Jan Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, second edition, May 2012. ISBN 9780199641178. doi: 10.1093/acprof:oso/9780199641178.001.0001. URL http://dx.doi.org/10.1093/acprof:oso/9780199641178.001.0001.

Oren Forkosh, Stoyo Karamihalev, Simone Roeh, Uri Alon, Sergey Anpilov, Chadi Touma, Markus Nussbaumer, Cornelia Flachskamm, Paul M. Kaplick, Yair Shemesh, and Alon Chen. Identity domains capture individual differences from across the behavioral repertoire. *Nature Neuroscience*, 22(12):2023–2028, 2019. doi: 10.1038/s41593-019-0516-y.

Julia Freund, Andreas M. Brandmaier, Lars Lewejohann, Imke Kirste, Mareike Kritzler, Antonio Krüger, Norbert Sachser, Ulman Lindenberger, and Gerd Kempermann. Emergence of individuality in genetically identical mice. *Science*, 340(6133):756–759, 2013. doi: 10.1126/science.1235294.

Vincent Godesberg, Till Bockemühl, and Ansgar Büschges. Natural variability and individuality of walking behavior in Drosophila. *Journal of Experimental Biology*, 227(22):jeb247878, 2024. doi: 10.1242/jeb.247878.

Google Research. TimesFM 3.0 (PyTorch). Official model card, 2026. URL https://huggingface.co/google/timesfm-3.0-pytorch.

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 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 21271–21284. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf.

G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, July 2006. ISSN 1095-9203. doi: 10.1126/science.1127647. URL http://dx.doi.org/10.1126/science.1127647.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Shantanu P. Jadhav and Loren M. Frank. Simultaneous extracellular recordings from hippocampal area ca1 and medial prefrontal cortex from rats performing a w-track alternation task., 2020. URL http://crcns.org/data-sets/hc/hc-28.

Risa Kawai, Timothy Markman, Rajesh Poddar, Raymond Ko, Antoniu L. Fantana, Ashesh K. Dhawale, Adam R. Kampff, and Bence P. Ölveczky. Motor cortex is required for learning but not for executing a motor skill. *Neuron*, 86(3):800–812, 2015. doi: 10.1016/j.neuron.2015.03.024.

Dana Rubi Levy, Nigel Hunter, Sherry Lin, Emma Marie Robinson, Winthrop Gillis, Eli Benjamin Conlin, Rockwell Anyoha, Rebecca M. Shansky, and Sandeep Robert Datta. Mouse spontaneous behavior reflects individual variation rather than estrous state. *Current Biology*, 33(7):1358–1364, 2023. doi: 10.1016/j.cub.2023.02.035.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023. URL https://arxiv.org/abs/2301.12597.

Kevin Luxem, Petra Mocellin, Falko Fuhrmann, Johannes Kürsch, Stephanie R. Miller, Jorge J. Palop, Stefan Remy, and Pavol Bauer. Identifying behavioral structure from deep variational embeddings of animal motion. *Communications Biology*, 5(1):1267, Nov 2022. ISSN 2399-3642. doi: 10.1038/s42003-022-04080-7. URL https://doi.org/10.1038/s42003-022-04080-7.

Jeffrey E Markowitz, Winthrop F Gillis, Maya Jay, Jeffrey Wood, Ryley W Harris, Robert Cieszkowski, Rebecca Scott, David Brann, Dorothy Koveal, Tomasz Kula, et al. Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature*, 614(7946):108–117, 2023.

Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21:1281–1289, 2018. doi: 10.1038/s41593-018-0209-y.

Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2022. doi: 10.1145/3503250.

Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S Refaat, and Benjamin Sapp. Wayformer: Motion forecasting via simple & efficient attention networks. *arXiv preprint arXiv:2207.05844*, 2022.

Katrina P. Nguyen, Abhinav Sharma, Mauricio Gil-Silva, Aryn H. Gittis, and Steven M. Chase. Distinct kinematic adjustments over multiple timescales accompany locomotor skill development in mice. *Neuroscience*, 466:260–272, 2021. doi: 10.1016/j.neuroscience.2021.05.002.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

Alexandros Paraschos, Christian Daniel, Jan R Peters, and Gerhard Neumann. Probabilistic movement primitives. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/e53a0a2978c28872a4505bdb51db06dc-Paper.pdf.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4172–4182. IEEE, 2023.

Matthew Rosenberg, Tony Zhang, Pietro Perona, and Markus Meister. Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration. *eLife*, 10:e66175, 2021. doi: 10.7554/eLife.66175.

Keith Sheppard, Justin Gardin, Gautam S. Sabnis, Asaf Peer, Megan Darrell, Sean Deats, Brian Geuther, Cathleen M. Lutz, and Vivek Kumar. Stride-level analysis of mouse open field behavior using deep-learning-based pose estimation. *Cell Reports*, 38(2):110231, 2022. doi: 10.1016/j.celrep.2021.110231.

Sameera Shridhar, Zachary M. Leveroni, Nan Zhou, Ellie Quattrocchi, Audrey Kaye, Jaquelin Gutierrez, Amanda Rodríguez León, Leslie M. Kay, and Jai Y. Yu. Goal location beta oscillations and sharp wave ripples are complementary markers of learning-related changes in hippocampal-neocortical networks. *eNeuro*, 2026. doi: 10.1523/ENEURO.0048-26.2026. URL https://doi.org/10.1523/ENEURO.0048-26.2026.

Anne C. Smith, Loren M. Frank, Sylvia Wirth, Marianna Yanike, Dan Hu, Yasuo Kubota, Ann M. Graybiel, Wendy A. Suzuki, and Emery N. Brown. Dynamic analysis of learning in behavioral experiments. *Journal of Neuroscience*, 24(2):447–461, 2004. doi: 10.1523/JNEUROSCI.2908-03.2004.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

Lucas Stoffl, Andy Bonnetto, Stéphane d'Ascoli, and Alexander Mathis. Elucidating the hierarchical nature of behavior with masked autoencoders. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision – ECCV 2024*, pp. 106–125, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-73039-9. doi: 10.1007/978-3-031-73039-9_7.

Jennifer J Sun, Tomomi Karigo, Dipam Chakraborty, Sharada Mohanty, Benjamin Wild, Quan Sun, Chen Chen, David Anderson, Pietro Perona, Yisong Yue, and Ann Kennedy. The multi-agent behavior dataset: Mouse dyadic social interactions. In J. Vanschoren and S. Yeung (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/7f1de29e6da19d22b51c68001e7e0e54-Paper-round1.pdf.

Jennifer J. Sun, Markus Marks, Andrew Wesley Ulmer, Dipam Chakraborty, Brian Geuther, Edward Hayes, Heng Jia, Vivek Kumar, Sebastian Oleszko, Zachary Partridge, Milan Peelman, Alice Robie, Catherine E Schretter, Keith Sheppard, Chao Sun, Param Uttarwar, Julian Morgan Wagner, Erik Werner, Joseph Parker, Pietro Perona, Yisong Yue, Kristin Branson, and Ann Kennedy. MABe22: A multi-species multi-task benchmark for learned representations of behavior. In Andreas Krause,

Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 32936–32990. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/sun23g.html.

Lukas M. von Ziegler, Fabienne K. Roessler, Oliver Sturman, Rebecca Waag, Mattia Privitera, Sian N. Duss, Eoin C. O'Connor, and Johannes Bohacek. Analysis of behavioral flow resolves latent phenotypes. *Nature Methods*, 2024. doi: 10.1038/s41592-024-02500-6.

Caleb Weinreb, Jonah E. Pearl, Sherry Lin, Mohammed Abdal Monium Osman, Libby Zhang, Sidharth Annapragada, Eli Conlin, Red Hoffmann, Sofia Makowska, Winthrop F. Gillis, Maya Jay, Shaokai Ye, Alexander Mathis, Mackenzie W. Mathis, Talmo Pereira, Scott W. Linderman, and Sandeep Robert Datta. Keypoint-moseq: parsing behavior by linking point tracking to pose dynamics. *Nature Methods*, 21(7):1329–1339, Jul 2024. ISSN 1548-7105. doi: 10.1038/s41592-024-02318-2. URL https://doi.org/10.1038/s41592-024-02318-2.

Caleb Weinreb, Lakshanyaa Thamarai Kannan, Alia Newman-Boulle, Tim Sainburg, Winthrop F. Gillis, Alex Plotnikoff, Sofia Makowska, Jonah E. Pearl, Mohammed Abdal Monium Osman, Scott W. Linderman, and Sandeep Robert Datta. Spontaneous behavior is a succession of self-directed tasks. *Neuron*, 114(5):922–937, 2026. doi: 10.1016/j.neuron.2025.11.021.

Alexander B. Wiltschko, Matthew J. Johnson, Giuliano Iurilli, Ralph E. Peterson, Jesse M. Katon, Stan L. Pashkovski, Victoria E. Abraira, Ryan P. Adams, and Sandeep Robert Datta. Mapping sub-second structure in mouse behavior. *Neuron*, 88(6):1121–1135, Dec 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.11.031. URL http://dx.doi.org/10.1016/j.neuron.2015.11.031.

A Data, Models, and Evaluation

A.1 Spatial Learning Dataset Statistics

Table 5 gives per-animal counts for the cohorts in Section 4. Recording days are indexed from the start of each animal's schedule; later recording phases are excluded.

Frame counts are taken from the processed trajectory files. Each session is counted once across its training, validation, and test partitions. Recording duration is computed from the first-to-last timestamp span of each session and summed across sessions, excluding intervals between recordings.

Table 5: **Statistics for the two reward datasets**. Frame counts include all data partitions; durations are computed from session timestamps.

Identity	Animal-days	Sessions	Frames	Duration (min)
		Alternation Reward		
Rat A	5	15	98,708	148.9
Rat B	5	20	135,646	205.2
Rat C	5	20	131,844	198.9
Rat D	5	20	127,143	191.8
Rat E	5	20	133,732	202.2
Rat F	5	20	128,259	193.5
Subtotal (6 rats)	**30**	**115**	**755,332**	**1,140.5**
		Probabilistic Reward		
Rat G	5	20	130,300	196.6
Rat H	5	20	132,163	199.4
Rat I	5	20	130,566	197.3
Rat J	5	20	92,583	139.7
Rat K	5	20	135,531	204.9
Rat L	5	20	135,655	205.1
Rat M	5	20	136,430	206.3
Rat N	5	20	135,964	205.6
Subtotal (8 rats)	**40**	**160**	**1,029,192**	**1,554.8**
Total (14 rats)	**70**	**275**	**1,784,524**	**2,695.3**

Table 6 gives the processed-frame counts for the chronological partitions described in Section 4.

Table 6: **Frame counts by partition across two datasets**. Rounding at session boundaries produces small deviations from the nominal split proportions.

Reward rule	Train	Validation	Test
Alternation Reward	604,223	37,752	113,357
Probabilistic Reward	823,294	51,436	154,462
Total	**1,427,517**	**89,188**	**267,819**

A.2 Implementation Details

For each method, we selected the configuration with the highest mean of two standardized validation scores on the alternation cohort (6 rats): z-scored Identity accuracy and sign-reversed z-scored Day MSE.

Classical baselines were fitted using their respective procedures: expectation-maximization for LG-SSM and AR-HMM, Gibbs sampling for Keypoint-MoSeq, and closed-form least squares for ProMPs.

The deep baselines (Autoencoder, VAME, T-BYOL, TCN, BAMS, and hBehaveMAE) were trained for 200 epochs, with learning rates selected through validation grid search.

For LoRA fine-tuning of Chronos-2 and TimesFM-3, we train low-rank adapters for 10 epochs using the self-supervised objective of TRAIL. The validation criterion above selects the learning rate and LoRA rank.

A.2.1 PRIMARY MODEL AND LINEAR PROBES

Preprocessing and windows. We use the registered maze coordinates and downsampled keypoint trajectories from the dataset preprocessing pipeline. Frames with nonfinite tracked coordinates are removed before downsampling. Each recording is split chronologically before constructing stride-one windows, so a trajectory never crosses a partition boundary. The primary feature cache requires 33 history frames and 33 following frames for window eligibility; the encoder receives only the history. After the final-frame alignment described in Section 3.2, we apply no random rotation, reflection, scaling, or coordinate noise.

Encoder and optimization. The encoder and decoder are trained from scratch. The encoder uses eight attention heads per layer; the reconstruction decoder has four layers with four attention heads each. Both stacks use pre-layer normalization, feed-forward width 256, and dropout 0.1. We train for 200 epochs with AdamW, batch size 256, learning rate 10^{-3}, $(\beta_1, \beta_2) = (0.9, 0.999)$, weight decay 0.01, and gradient-norm clipping at 1. Cosine decay reduces the learning rate to 3×10^{-5}. The reported primary run uses seed 42, and features are extracted from the final epoch.

Positive and negative sampling. The sampler uses the history endpoint relative to the beginning of its recording partition. If D_s is that partition's duration and $u_i = t_i^{\mathrm{end}}/D_s$, a positive partner is drawn uniformly from the other eligible windows in the same partition satisfying $|u_i - u_j| \le 0.05$. At an endpoint, the interval is truncated to the available windows; if it contains no other window, the nearest one is used. Thus the nominal neighborhood spans 10% of the partition. For the 80% training portion of a ten-minute recording, this is approximately 48 seconds; its duration scales with the partition length. Each epoch shuffles the eligible training windows; incomplete final batches are dropped. Each batch contains 256 anchors and one separately sampled positive for each anchor. For an anchor, the contrastive candidates are its sampled positive and the other batch anchors. Other anchors satisfying the same locality rule are also positives; the remaining candidates are negatives. The anchor itself is excluded. Positive and negative candidate sets receive equal total weight in the balanced loss. Pair construction uses recording membership and time only.

Contrastive objective. For an anchor representation $\mathbf{h}$, let $\mathcal{P}$ and $\mathcal{N}$ be its positive and negative candidate sets. With s denoting cosine similarity, the training logit is $\ell(\mathbf{h}, \mathbf{h}') = s(\mathbf{h}, \mathbf{h}')/\tau$, with $\tau = 0.1$. The implementation assigns total weight $1/2$ to each candidate set. After cancelling this common factor, the per-anchor loss is

$$\mathcal{L}_{\mathrm{contrast}} = -\frac{1}{|\mathcal{P}|} \sum_{\mathbf{h}^+ \in \mathcal{P}} \log \frac{|\mathcal{P}|^{-1} \exp\big(\ell(\mathbf{h}, \mathbf{h}^+)\big)}{|\mathcal{P}|^{-1} \sum_{\mathbf{h}' \in \mathcal{P}} \exp\big(\ell(\mathbf{h}, \mathbf{h}')\big) + |\mathcal{N}|^{-1} \sum_{\mathbf{h}^- \in \mathcal{N}} \exp\big(\ell(\mathbf{h}, \mathbf{h}^-)\big)}. \quad (4)$$

The numerator normalization contributes an additive $\log |\mathcal{P}|$ term. Candidate membership is fixed by the sampled windows, so this term has zero gradient with respect to model parameters. We average the loss over anchors.

Temporal representation learning.

Joint supervised baseline. The five task losses are summed with unit coefficients, and regression targets retain their original units. The trained encoder is evaluated with the probes below.

Probe fitting. Identity is decoded by one pooled logistic-regression probe across animals, with an ℓ_2 penalty, $C = 1$, and at most 300 L-BFGS iterations. Each continuous target has a separate ridge probe for each animal. Feature means and standard deviations are fitted on the corresponding training data. Ridge penalties are selected by leave-one-out cross-validation on the training rows from $\{10^{-3}, 10^{-2.5}, \ldots, 10^{3}\}$. Encoder configurations are selected using held-out validation scores.

Proficiency and Success Rate predictions are clipped to $[0,\ 1]$; Day and Reward predictions use the observed target range of the evaluated dataset, pooled across its partitions. Regression MSE is pooled over all valid test windows, weighting each animal by its window count.

A.2.2 Behavioral target construction

Trial-level proficiency. The supplied state-space labels follow Smith et al. (2004): a Gaussian random walk models the latent learning state, and a logistic observation model links it to the binary trial outcomes. We use the posterior modes of response probability as trial-level labels.

We retain entries for all trials, joining each posterior mode to its trial by animal, recording date, session, maze, and trial identifier. Per-day estimates are used where available; trial keys missing from the probabilistic cohort's per-day table use the supplied estimate fitted across recording days. These are retrospective behavioral labels: state-space smoothing uses outcomes later in the fitted sequence as well as earlier ones.

Let a_j be the start time of trial j and $\hat{p}_j$ its supplied posterior mode. Between consecutive trial starts, every frame inherits the preceding trial's estimate:

$$P(t) = \hat{p}_j \quad \text{for } a_j \le t < a_{j+1}. \tag{5}$$

The first and last trial estimates extend to the recording boundaries. An unmatched trial produces a missing value throughout its interval. For the sampled frame times $\mathcal{T}_i$ covered by history i, the probe target is

$$y_i^{\text{prof}} = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} P(t). \tag{6}$$

A window containing a missing proficiency value is excluded from that target's probe fit and error calculation. Its other targets remain available.

Reward and success rate. Let m_j be the midpoint of trial j and $b_j \in \{0, 1\}$ its reward outcome. At frame time t, we count all trials with $m_j \in [t - 30\,\text{s}, t + 30\,\text{s})$. Reward is the number of rewarded trials in this interval; Success Rate divides that count by the total number of trials and is zero for an empty interval. Each history receives the mean of its frame-level values. Counts use the complete recording's trial table, including trials across a trajectory-split boundary. These centered targets describe the reward context around the movement.

A.2.3 Hyperparameter search

Table 7 lists the hyperparameters considered during model selection. The TRAIL search crosses three learning rates with five reconstruction/contrastive weight pairs, giving 15 configurations; the selected weights are $(1, 1)$. Encoder depth is fixed at four layers in this search. The separate architecture sweeps examine latent dimensions $\{1, 4, 8, 16, 32, 40, 48, 56, 64\}$, history lengths from 1 to 10 seconds, and positive neighborhoods spanning $\{0.5, 1, 2, 5, 10, 25, 50, 75, 100\}\%$ of a recording partition. The architecture sweeps are reported in the ablation experiments.

A.3 Encoder settings for the additional experiments

The generalization, window-selection, repeated-training, movement-descriptor, and maze-context experiments use the supplementary encoder in Table 8. It shares the primary model's trajectory alignment, temporal sampling, contrastive objective, and optimization schedule. Runs use seed 42 unless a table specifies otherwise. The temporal-neighborhood comparison uses the primary model. Both encoders are evaluated below on the same within-session test windows.

Within-session experiments use the primary data split. Session- and day-held-out experiments keep complete recordings together, as specified in Appendix B.3. We freeze the encoder before fitting probes, standardize features using the training partition, and select ridge penalties from $\{10^{-3}, 10^{-2.5}, \ldots, 10^{3}\}$ on the validation partition. Identity uses regularized logistic regression. Proficiency, rewarded count, and success rate follow the target definitions in Appendix A.2.2.

Table 7: **Hyperparameter search.** Candidate values are grouped by model. Loss weights are listed as (reconstruction, contrastive); Keypoint-MoSeq iteration counts are listed as (AR, full model). Hidden width and HoA weight are fixed at the values shown.

Model	Hyperparameter	Candidate values
TRAIL	Learning rate	10^{-4}, 3×10^{-4}, 10^{-3}
	Loss weights	$(1, 1)$, $(0.5, 1)$, $(1, 0.5)$, $(0.25, 0.75)$, $(0.75, 0.25)$
Classic methods		
LG-SSM	Latent dimension	4, 8, 16, 32, 64, 128
	EM iterations	100, 200
AR-HMM	Number of states	32, 64, 128
	Autoregressive lags	1, 3
	EM iterations	100, 200
ProMPs	Basis weights	2, 3, 5, 8, 10, 12, 15, 18, 20, 24, 28, 33, 36, 40, 50
Keypoint-MoSeq	Number of states	20, 30, 50
	Latent dimension	6, 10
	AR / full iterations	$(100, 100)$, $(75, 125)$, $(50, 150)$
Deep learning methods		
Autoencoder	Layers	2, 4, 6, 8
	Learning rate	10^{-4}, 3×10^{-4}, 10^{-3}
TCN / T-BYOL	Channels per layer	16, 32, 64, 128, 256 (three equal-width layers)
	Learning rate	10^{-4}, 3×10^{-4}, 10^{-3}
VAME	Latent dimension	64, 128
	Hidden width	256
	β	1, 2
	Learning rate	10^{-4}, 3×10^{-4}, 5×10^{-4}, 10^{-3}
BAMS	HoA weight	500
	HoA bins	16, 32
	BYOL weight	0.5, 1
	Learning rate	10^{-4}, 3×10^{-4}, 10^{-3}
Foundation-model Fine-tuning		
LoRA	Adapter rank	4, 8, 16
	Learning rate	10^{-5}, 3×10^{-5}, 10^{-4}, 3×10^{-4}, 10^{-3}

Table 8: **Encoder configurations used in the additional experiments.** Both use multi-frequency coordinate features, 64-dimensional tokens, eight queries, four encoder layers, and contrastive temperature 0.1. The performance rows use the same 105,882 test windows and state-space Proficiency labels. Identity is percentage accuracy; the remaining targets are MSE.

Setting	Primary	Supplementary
Coordinate projection	Linear	Two-layer MLP
Maze tokens	Waypoints	Vertices and edge descriptors
First-query key normalization	None	Per-head layer normalization
Reconstruction input	Restored representation and seven queries	Full encoder query bank
Reconstruction coefficient	1	0.5
Identity	84.00	84.07
Day	1.258	1.252
Proficiency	0.01556	0.01552
Reward	0.609	0.593

A.4 UNCERTAINTY ESTIMATES

We quantify uncertainty by resampling animals, keeping all test windows from a selected animal together. For the primary cohort, movement-view, and reconstruction comparisons, we recompute pooled accuracy or MSE from the selected animals' window counts and score sums. The generalization tables use equal-animal averages, as indicated in their captions. The 2.5th and 97.5th percentiles give the reported 95% confidence intervals. For the primary six-animal alternation estimates, we enumerate all 6^6 bootstrap resamples. For the eight-animal probabilistic cohort, we draw 20,000 resamples with random seed 42. The encoder and probes remain fixed, so these intervals describe variation across animals under the fitted models. For paired comparisons, each bootstrap draw uses the same animals for both methods.

B Prediction Across Datasets and Partitions

B.1 Performance across reward conditions

Table 9 gives the full comparison for probabilistic rewards. Table 10 reports animal-bootstrap intervals for TRAIL in both cohorts. The paired comparisons in Table 11 use the primary alternation runs from Table 1.

Table 9: **Performance on the spatial-learning dataset (probabilistic rewards)**.

	Method	Identity (Acc ↑)	Day (MSE ↓)	Proficiency (MSE ↓)	Success Rate (MSE ↓)	Reward (MSE ↓)
Classic	Direct Trajectory	35.18	1.89	.011	0.07	0.91
	LG-SSM	19.40	1.88	.011	0.05	0.85
	AR-HMM	35.40	1.72	.010	0.07	0.83
	ProMPs	29.96	1.84	.010	0.05	0.83
	Keypoint-MoSeq	33.58	1.94	.011	0.07	0.94
Foundation Model	TimesFM-3 (Zero-shot)	31.73	1.87	.011	0.08	0.91
	TimesFM-3 (LoRA)	36.42	1.82	.010	0.08	0.89
	Chronos-2 (Zero-shot)	33.33	1.81	.010	0.08	0.89
	Chronos-2 (LoRA)	41.06	1.73	.010	0.08	0.86
Deep Learning	AutoEncoder	55.63	1.67	<u>.009</u>	<u>0.05</u>	0.73
	VAME	59.49	1.72	.009	0.05	0.75
	T-BYOL	56.72	1.78	.010	0.05	0.79
	TCN	66.10	<u>1.61</u>	.009	0.05	0.81
	BAMS	62.40	1.72	.009	0.05	0.73
	hBehaveMAE	46.58	1.73	.009	**0.05**	<u>0.71</u>
	Multi-task Supervised	<u>82.26</u>	1.75	.009	0.06	0.91
	TRAIL	**85.87**	**1.33**	**.007**	0.05	**0.68**

Table 10: **TRAIL estimates and confidence intervals in both spatial-learning cohorts.** Estimates and animal-bootstrap 95% intervals are computed from the held-out predictions used in the cohort comparisons. Identity is percentage accuracy; the remaining targets are MSE.

Cohort	Target	Estimate	95% CI
Alternation	Identity (%)	84.00	[76.24, 91.10]
Alternation	Day	1.2579	[0.9569, 1.5937]
Alternation	Proficiency	0.0156	[0.0109, 0.0223]
Alternation	Success Rate	0.1345	[0.0888, 0.1798]
Alternation	Reward	0.6090	[0.4488, 0.8493]
Probabilistic	Identity (%)	85.87	[82.59, 88.81]
Probabilistic	Day	1.3329	[1.0275, 1.6831]
Probabilistic	Proficiency	0.0071	[0.0054, 0.0087]
Probabilistic	Success Rate	0.0511	[0.0425, 0.0600]
Probabilistic	Reward	0.6824	[0.6094, 0.7657]

Table 11: **Paired comparisons in the primary alternation cohort.** Values are pooled MSE reductions relative to each baseline, with animal-bootstrap 95% confidence intervals. All methods use the same 105,882 test windows.

Baseline	Day reduction [95% CI]	Proficiency reduction [95% CI]
TCN	0.248 [0.169, 0.330]	0.00264 [0.00202, 0.00325]
VAME	0.354 [0.226, 0.466]	0.00175 [0.00059, 0.00315]
BAMS	0.331 [0.167, 0.478]	0.00223 [0.00062, 0.00422]

B.2 LONG-CONTEXT PREDICTION

Table 12 compares TRAIL with BAMS and hBehaveMAE in the long-context setting. BAMS uses both 3s (33-frame) short-term and 10s (110-frame) long-term windows while TRAIL and hBehaveMAE take 10s trajectories as input.

Table 12: **Long-context prediction in both reward groups**. **Bold**: best; underlined: second best by unrounded values in each column.

Method	Identity (↑)	Day (↓)	Proficiency (↓)	Success Rate (↓)	Reward (↓)
Alternation Reward					
BAMS	73.68	1.32	.017	**0.12**	0.54
hBehaveMAE	50.87	1.55	.017	**0.12**	0.59
TRAIL	**85.26**	**1.10**	**.014**	0.13	**0.51**
Probabilistic Reward					
BAMS	77.64	1.52	.008	**0.04**	0.62
hBehaveMAE	48.65	1.68	.009	**0.04**	0.64
TRAIL	**87.87**	**1.29**	**.006**	0.04	**0.58**

B.3 GENERALIZATION ACROSS SESSIONS, DAYS, AND ANIMALS

B.3.1 HELD-OUT SESSIONS

For each animal and recording day, complete sessions are sorted in acquisition order. Sessions before the final two are used for training, the penultimate session for validation, and the final session for testing. All original temporal shards of a session remain together. This gives 55 training, 30 validation, and 30 test sessions, with all six animals and five recording days represented in the test set.

We compare the supplementary encoder with BAMS using 26-channel local kinematics and maze geometry in the same body-centered frame. BAMS uses 33-frame and 110-frame histories, its HoA and BYOL objectives, seed 42, and 200 training epochs. Both methods are evaluated at the same endpoints eligible for the longer history. Table 13 shows lower error for day, proficiency, and rewarded count.

Table 13: **Prediction on held-out sessions.** MSE is averaged equally across animals. MSE reductions measure the improvement with TRAIL over BAMS, with paired animal-bootstrap 95% CIs. The comparison uses 183,884 windows, except proficiency, which has 181,639 complete cases. BAMS uses local coordinates and maze context.

Target	TRAIL	BAMS	MSE reduction [95% CI]	Animals improved
Day	1.765	1.932	0.167 [0.071, 0.255]	5/6
Proficiency	0.02032	0.02460	0.00428 [0.00124, 0.00897]	5/6
Rewarded count	0.760	1.179	0.419 [0.131, 0.788]	6/6

We also compare the held-out-session encoder with an animal-specific training mean and a ridge probe using 14 kinematic features. The representation gives lower day and proficiency error than either baseline in five of the six animals (Table 14).

B.3.2 HELD-OUT DAY

We train encoders and probes on days 1–3, select probe regularization on day 4, and test on day 5, keeping complete sessions and their temporal shards together. On day 5, TRAIL achieves higher identity accuracy than BAMS and simple kinematics, and lower rewarded-count error than BAMS (Table 15).

Table 14: **Day and proficiency readouts on held-out sessions.** MSE is averaged equally across six animals. The lower block reports the reduction in MSE achieved by TRAIL relative to each baseline, with paired animal-bootstrap 95% confidence intervals. Day uses 190,814 test windows and Proficiency uses 188,261 complete cases.

Readout	**Day**	**Proficiency**
Animal-specific mean	2.012	0.02571
14 kinematic features	1.950	0.02541
TRAIL	1.756	0.02046
MSE reduction with TRAIL [95% CI]		
Training mean	0.256 [0.051, 0.430]	0.00524 [0.00143, 0.00922]
Kinematic features	0.194 [0.054, 0.316]	0.00495 [0.00230, 0.00753]

Table 15: **Identity and rewarded-count prediction on a held-out day.** Days 1–3 train the encoders and probes; day 4 selects regularization. Results use 139,875 day-5 windows. TRAIL improves identity accuracy in all six animals; gains are measured against each listed baseline in percentage points. Rewarded-count MSE is averaged equally across animals; five of six animals improve. Intervals are paired animal-bootstrap 95% CIs.

Representation	**Balanced Accuracy**	**Accuracy gain (pp)**	**95% CI**
Chance	16.67%	58.21 pp	[44.08, 72.33]
Kinematic features	25.90%	48.98 pp	[33.91, 62.35]
BAMS (local + maze)	38.89%	35.99 pp	[20.80, 55.76]
TRAIL	**74.87%**	–	–
Target	**TRAIL MSE**	**BAMS MSE**	**Reduction [95% CI]**
Rewarded count	0.948	1.608	0.660 [0.015, 1.790]

B.3.3 Within-recording temporal alignment

We keep the day-5 encoders and day-4-selected rewarded-count probes fixed, then permute features within each test shard. The analysis uses 4,275 nonoverlapping windows from 69 shards and six animals, with 1,000 permutations. Shuffling increases prediction error for both representations (Table 16), indicating that their features track changes in rewarded count within a recording.

Table 16: **Within-recording time alignment.** Errors are mean per-animal rewarded-count MSE. The reduction compares correctly aligned features with permuted features; intervals are animal-bootstrap 95% CIs. Each method improves in five of six animals (one-sided permutation $p = 0.001$).

Representation	**Aligned**	**Permuted**	**Reduction [95% CI]**
BAMS (local + maze)	1.593	1.661	0.068 [0.0031, 0.1604]
TRAIL	0.948	0.980	0.032 [0.0029, 0.0683]

B.3.4 Transfer to new animals

An encoder trained and selected on ten animals is frozen before evaluation on four different animals. Each new animal supplies a chronological training partition for fitting its own ridge probe; later windows are held out for testing. TRAIL has lower error than the kinematic features on every target in all four animals (Table 17).

B.4 Window selection and repeated training

These experiments use the supplementary encoder. We remove windows within 30 seconds of split boundaries and retain stride-33 windows within each session. The frozen-feature probes use identical training and test windows for TRAIL and BAMS. TRAIL retains lower day, proficiency, and rewarded-count MSE in all six animals (Table 18). The same table reports the mean and standard deviation from three training seeds under the dense-window protocol.

Table 17: **Transfer with calibration on four new animals.** MSE is averaged equally across the four animals using 74,768 test windows. MSE reductions measure the improvement with TRAIL over kinematic features, with paired animal-bootstrap 95% CIs.

Target	TRAIL	Kinematics	MSE reduction [95% CI]
Day	1.621	1.861	0.241 [0.182, 0.284]
Proficiency	0.00944	0.01112	0.00167 [0.00073, 0.00261]
Rewarded count	0.860	1.025	0.165 [0.082, 0.249]
Success rate	0.0503	0.0689	0.0186 [0.0138, 0.0234]

Table 18: **Window selection and training variability.** Identity is accuracy; the remaining readouts are MSE. The nonoverlap comparison uses 17,102 training and 2,004 test windows. MSE reductions are paired across animals, with animal-bootstrap 95% CIs. The final row uses seeds 7, 42, and 2026 with validation-selected checkpoints.

Setting	Identity	Day	Proficiency	Rewarded
BAMS, nonoverlap	64.97%	1.847	0.02117	1.032
TRAIL, nonoverlap	84.98%	1.407	0.01634	0.789
MSE reduction	–	0.441	0.00478	0.242
95% CI	–	[0.289, 0.593]	[0.00337, 0.00659]	[0.141, 0.357]
Dense-window protocol: mean ± SD over three seeds				
TRAIL, three seeds	$83.72\% \pm 0.45\%$	1.273 ± 0.020	0.01606 ± 0.00063	0.607 ± 0.013

B.5 MABe22 Mouse Triplets

The short-context TRAIL model uses 61 frames at the native 30 Hz and produces a 64-dimensional representation for each mouse. For the official tasks, we concatenate the coordinate-wise mean and range of the three mice's representations at each frame, yielding 128 features. The encoder is fitted on `UserTrain` and frozen before feature extraction. Task probes use all of `SubmissionTrain`, ridge penalty 1, and seed 42; classification uses class-balanced ridge regression. The reported benchmark scores use the official private test split and its task-specific scoring rules.

In Table 2, Tasks MSE averages the experiment-day and time-of-day regression errors and multiplies the result by 100. Tasks F1 averages the 11 classification scores. Frame F1 averages the nine interaction tasks, and Seq. F1 averages strain and lighting classification. Each F1 column is expressed as a percentage.

For within-clip Identity, a separate class-balanced ridge classifier predicts the three mouse slots in each test clip. Feature standardization uses that clip's first 42 seconds, and the ridge penalty is 1. We score valid frames in its final 18 seconds and average accuracy equally across clips. Table 19 compares the short-context TRAIL and BAMS runs on the same 1,058 clips, with identical numbers of valid training and test frames per mouse. Confidence intervals use 20,000 paired clip-bootstrap resamples with seed 42, keeping the encoders and fitted probes fixed.

Table 19: **Within-clip identity on MABe22 Mouse Triplets.** Estimates average the accuracies of the 1,058 test clips. Intervals describe variation across clips; mouse slots are defined separately in each clip. The final row reports TRAIL's accuracy improvement over BAMS in percentage points.

Representation	Accuracy (%)	95% CI
BAMS, short context	63.08	[61.72, 64.46]
TRAIL, short context	73.65	[72.28, 75.01]
Accuracy improvement (pp)	10.56	[9.68, 11.45]

C Model Design and Movement Analysis

C.1 Model design

Table 20: **Ablation studies on loss function and multi-frequency modeling. Bold**: best based on unrounded values.

(a) Loss function.

	ID↑	Day↓	Prof.↓	Succ.↓	Rew.↓
w/o Recon.	81.1	1.32	0.017	0.14	0.63
w. Recon.	**84.0**	**1.26**	**0.016**	**0.1345**	**0.61**

(b) Multi-frequency modeling.

	ID↑	Day↓	Prof.↓	Succ.↓	Rew.↓
w/o Multi-freq.	83.2	1.28	0.016	**0.13**	0.61
w. Multi-freq.	**84.0**	**1.26**	**0.016**	0.1345	**0.61**

Reconstruction loss We analyze TRAIL with and without reconstruction loss in Table 20a. The model trained with reconstruction has higher identity accuracy and lower day and proficiency errors. Table 21 reports the reduction in prediction error.

Multi-frequency modeling We evaluate the contribution of multi-frequency modeling in Table 20b. The multi-frequency features give modest improvements in identity, day, and proficiency prediction in this comparison.

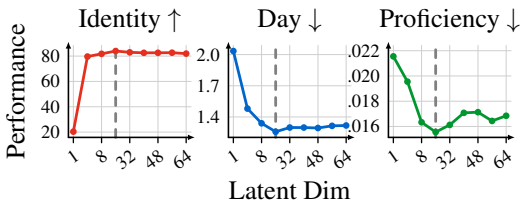

Figure 7: **Latent representation trade-off.** The default representation has 16 dimensions.

Latent representation We present the effect of latent dimensionality on identity and learning-related prediction in Figure 7. The default uses 16 dimensions; the eight-dimensional representation also performs well across the three readouts.

Table 21 gives paired error reductions for the movement-view and reconstruction comparisons. Each row compares the complete primary model with the indicated change.

Table 21: **Movement inputs and reconstruction.** Values are pooled MSE reductions relative to each ablated model, with animal-bootstrap 95% confidence intervals. All methods use the same 105,882 test windows. The comparisons use the runs in Tables 3 and 20a.

Model variant	Day reduction [95% CI]	Proficiency reduction [95% CI]
Local Pose only	0.347 [0.280, 0.408]	0.00330 [0.00131, 0.00556]
Global Trajectory only	0.215 [0.037, 0.382]	0.00169 [0.00032, 0.00297]
Without reconstruction	0.058 [0.022, 0.093]	0.00175 [0.00050, 0.00308]

C.1.1 Temporal neighborhoods

Using the primary encoder and the sampler in Appendix A.2.1, we compare neighborhoods spanning 10% and 100% of a recording partition. Both runs use seed 42 and the same test windows. The narrower neighborhood gives lower day and proficiency error in all six animals (Table 22).

C.1.2 Reconstruction inputs

On a two-animal subset, we train the decoder for 100 epochs with either all history queries or only the restored representation, using seed 42. Table 23 reports validation performance for both inputs.

C.2 Movement views

Full Movement uses the final-frame alignment in Section 3.2.

For *Local Pose*, we apply this standardization independently to each frame, ensuring that the body keypoint remains at the origin and the body–neck segment remains vertically aligned. This removes

Table 22: **Effect of positive-neighborhood width.** Identity is accuracy; Day and Proficiency are MSE. Results use 105,882 windows. The final rows give mean per-animal MSE reductions and paired animal-bootstrap 95% CIs.

Positive sampling	Identity	Day	Proficiency
Broad 100%	82.37%	1.465	0.0211
Local 10%	84.00%	1.258	0.0156
MSE reduction	–	0.211	0.00541
95% CI	–	[0.096, 0.370]	[0.00144, 0.00947]

Table 23: **Reconstruction-context comparison on two animals.** Identity is accuracy; the four regression columns report MSE. The last column is reconstruction negative log-likelihood. Results use 5,250 validation windows.

Context	Identity	Day	Proficiency	Success Rate	Reward	NLL
All queries	87.98%	2.0573	0.0008230	0.10713	0.42644	−6.731
Representation only	87.09%	2.0578	0.0008240	0.10632	0.42615	−6.733

global translation and rotation while retaining changes in relative body configuration. For *Global Trajectory*, we retain only the body coordinates from the Full Movement representation for global position, along with the body–neck unit vector as the orientation over time.

Table 24: **Uncertainty for the movement-view comparison.** Pooled estimates and animal-bootstrap 95% confidence intervals use the same 105,882 test windows as Table 3. Identity is percentage accuracy; Day and Proficiency are MSE.

View	Identity [95% CI]	Day [95% CI]	Proficiency [95% CI]
Local Pose	77.1 [66.2, 87.4]	1.6047 [1.3303, 1.9455]	0.0189 [0.0124, 0.0277]
Global Trajectory	72.1 [61.4, 82.3]	1.4733 [1.3063, 1.6259]	0.0172 [0.0128, 0.0231]
Full Movement	84.0 [76.2, 91.1]	1.2579 [0.9569, 1.5937]	0.0156 [0.0109, 0.0223]

C.3 Movement descriptors and maze geometry

C.3.1 Movement descriptors

We compare the supplementary three-second representation with mean body–neck and neck–nose lengths, and with 14 kinematic features describing displacement, speed, acceleration, scanning, reversals, backtracking, lateral deviation, and head–body decoupling. Linear regularization is selected on validation data. The regression constant is each animal's training-target mean. TRAIL improves identity recognition over both bone lengths and kinematics, and predicts day and reward frequency more accurately than the kinematic features and training-mean predictor (Table 25).

C.3.2 Maze geometry

Adding maze geometry to the body-centered trajectories improves identity recognition and lowers day and reward-frequency error (Table 26).

C.4 Variation across animals

Table 27 reports each animal's test performance using the primary encoder and probes.

C.5 Recording days, maze regions, and movement motifs

Maze regions. We divide the Y-maze into three regions: *center*, where the three arms intersect; *arm segment*, the straight portion of each arm; and *endpoint*, the distal end of each arm.

Table 25: **Movement-descriptor comparison.** Identity is accuracy; the remaining readouts are MSE. Results use 105,882 held-out-time windows. The final rows compare TRAIL with the constant predictor: R^2 uses aggregate MSE, while the mean MSE reductions and bootstrap 95% CIs are paired across animals.

Representation	**Identity** (↑)	**Day** (↓)	**Proficiency** (↓)	**Reward** (↓)
Constant / chance	16.67%	2.040	0.02163	0.805
Mean bone lengths	29.32%	–	–	–
14 kinematic features	28.81%	1.895	0.02031	0.753
TRAIL	84.07%	1.252	0.01552	0.593
R^2	–	0.386	0.282	0.264
MSE reduction	–	0.800	0.00622	0.200
95% CI	–	[0.429, 1.115]	[0.00268, 0.01014]	[0.0509, 0.371]

Table 26: **Effect of maze context.** Identity is accuracy; the remaining readouts are MSE. Both runs use seed 42 and the same 105,882 held-out-time windows.

Maze context	**Identity** (↑)	**Day** (↓)	**Proficiency** (↓)	**Reward** (↓)
Removed	70.88%	1.599	0.01994	0.648
Included	84.07%	1.252	0.01552	0.593

Table 27: **Per-animal performance in the alternation cohort.** Values use the same held-out predictions as Table 10. Identity is percentage accuracy; the remaining targets are MSE.

Rat ID	**Identity** (↑)	**Day** (↓)	**Proficiency** (↓)	**Success Rate** (↓)	**Reward** (↓)
Rat A	72.70	0.72	0.0119	0.22	0.53
Rat B	95.92	0.91	0.0112	0.12	0.56
Rat C	76.45	1.49	0.0155	0.18	0.38
Rat D	74.40	1.97	0.0313	0.16	0.44
Rat E	95.68	1.33	0.0087	0.04	1.15
Rat F	85.15	1.02	0.0145	0.11	0.55

Movement motifs. We assign five motifs from the Local Pose and Global Trajectory views: *remaining still*, with little translation or postural change; *scanning*, with head or body reorientation while staying in place; *going straight*, with forward travel along an arm; *turning*, with a change of direction between arms through the center; and *turning back*, with a reversal near an arm endpoint.

The subgroup analysis uses nonoverlapping 33-frame histories from the primary test recordings. Of the 3,376 annotated histories, 3,261 also have the 33 subsequent frames required by the feature cache. Table 28 gives their prediction errors and animal-bootstrap intervals. Table 29 reports the sample count for each animal and the spread of the Proficiency target in each group. All groups use the same fitted probes, with a separate regression probe for each animal.

Table 28: **Uncertainty in the primary subgroup analysis.** Values are pooled estimates with animal-bootstrap 95% confidence intervals. Identity is percentage accuracy; Day and Proficiency are MSE. Day prediction is omitted when day defines the group.

Group	n	Identity [95% CI]	Day [95% CI]	Proficiency [95% CI]
Maze region				
Arm segment	558	81.7 [72.4, 90.1]	1.59 [1.04, 2.11]	0.0211 [0.0109, 0.0310]
Arm endpoint	2377	83.0 [74.1, 91.0]	1.20 [0.92, 1.50]	0.0149 [0.0113, 0.0203]
Maze center	326	88.7 [81.1, 95.6]	1.20 [0.89, 1.51]	0.0135 [0.0073, 0.0225]
Recording day				
Day 1	700	79.3 [68.5, 87.6]	–	0.0192 [0.0107, 0.0290]
Day 2	635	81.9 [72.6, 91.2]	–	0.0183 [0.0112, 0.0269]
Day 3	636	78.8 [66.5, 90.8]	–	0.0164 [0.0106, 0.0224]
Day 4	645	89.5 [84.3, 94.4]	–	0.0128 [0.0063, 0.0200]
Day 5	645	87.6 [79.5, 95.4]	–	0.0123 [0.0062, 0.0218]
Movement motif				
Remaining still	1221	87.1 [78.8, 93.6]	1.11 [0.85, 1.36]	0.0115 [0.0083, 0.0164]
Scanning	1080	76.6 [67.1, 86.0]	1.34 [0.99, 1.78]	0.0195 [0.0151, 0.0261]
Going straight	543	84.0 [76.1, 92.8]	1.37 [0.87, 1.82]	0.0184 [0.0102, 0.0271]
Turning	151	98.0 [93.5, 99.5]	1.07 [0.92, 1.44]	0.0083 [0.0055, 0.0145]
Turning back	266	84.2 [72.0, 94.0]	1.56 [1.22, 1.97]	0.0202 [0.0111, 0.0305]

Table 29: **Composition and target spread in the primary subgroup analysis.** Counts include only scored windows. Rat labels follow Table 27. The final column gives the standard deviation of Proficiency across these windows.

Group	Rat A	Rat B	Rat C	Rat D	Rat E	Rat F	Prof. SD
Arm segment	63	97	101	122	96	79	0.174
Arm endpoint	320	449	396	376	404	432	0.155
Maze center	42	40	72	52	77	43	0.157
Day 1	82	158	110	114	120	116	0.105
Day 2	84	107	114	109	113	108	0.165
Day 3	88	97	115	109	116	111	0.111
Day 4	87	112	116	107	114	109	0.080
Day 5	84	112	114	111	114	110	0.093
Remaining still	151	189	227	180	219	255	0.145
Scanning	175	239	171	217	131	147	0.164
Going straight	71	83	128	99	90	72	0.170
Turning	6	31	6	10	70	28	0.100
Turning back	22	44	37	44	67	52	0.149

D REPRESENTATION VISUALIZATIONS

Figure 8 compares three-second representations from the primary evaluation. Each left panel is a joint t-SNE projection colored by rat; the six right panels are independently fitted per-rat projections colored by recording day. Each projection has its own coordinate system.

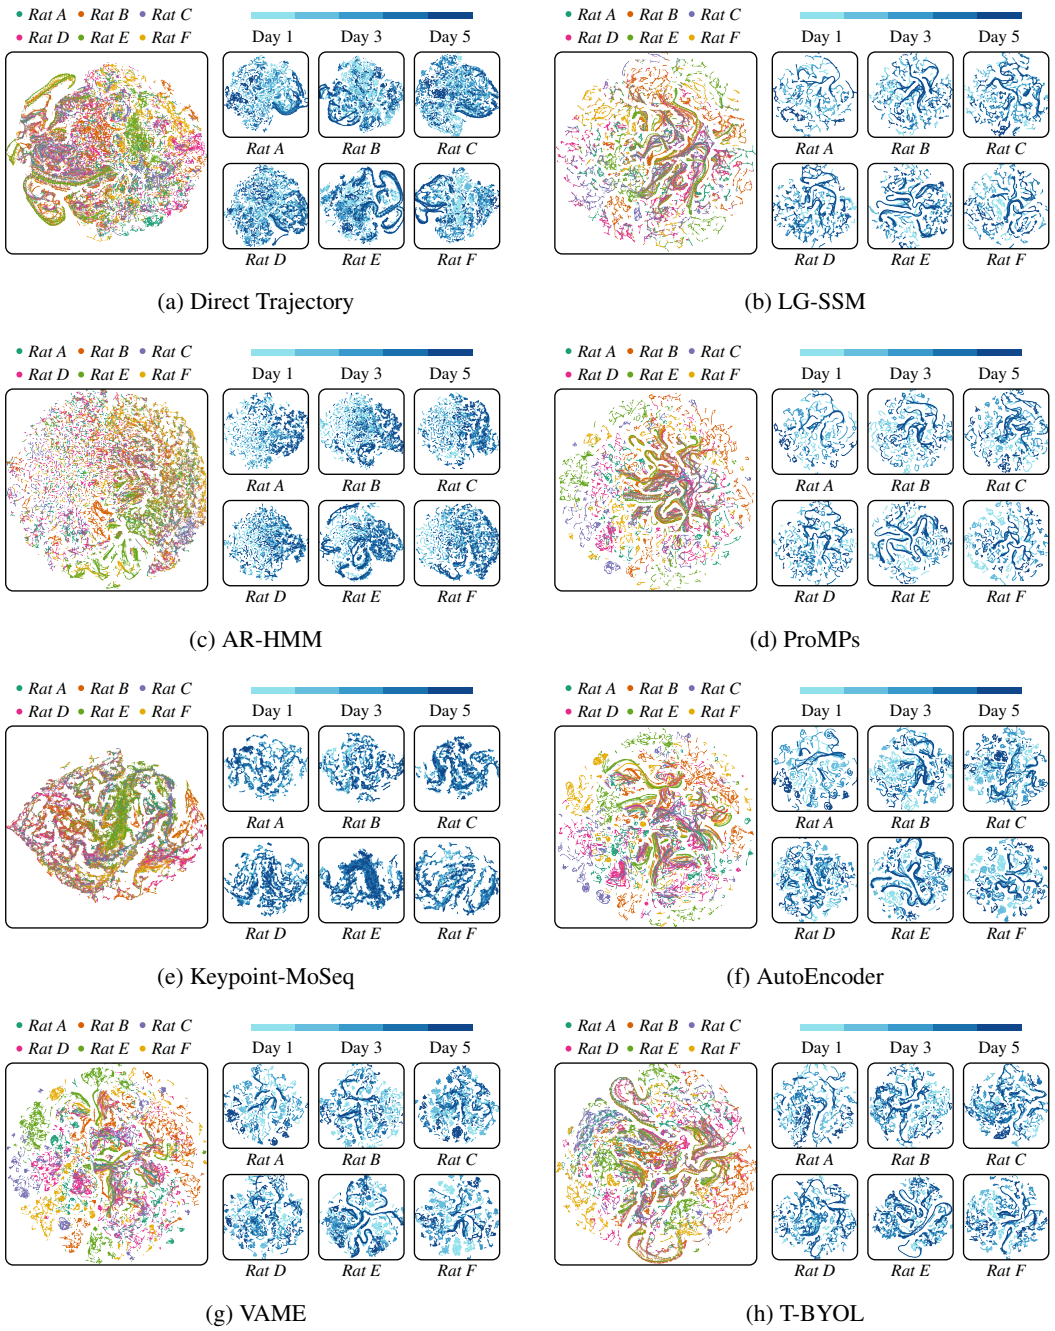

(a) Direct Trajectory

(b) LG-SSM

(c) AR-HMM

(d) ProMPs

(e) Keypoint-MoSeq

(f) AutoEncoder

(g) VAME

(h) T-BYOL

Figure 8: **t-SNE visualizations of 3-second baseline representations.** Each method shows a joint projection colored by rat (left) and six per-rat projections colored by day (right). The figure continues on the next page.

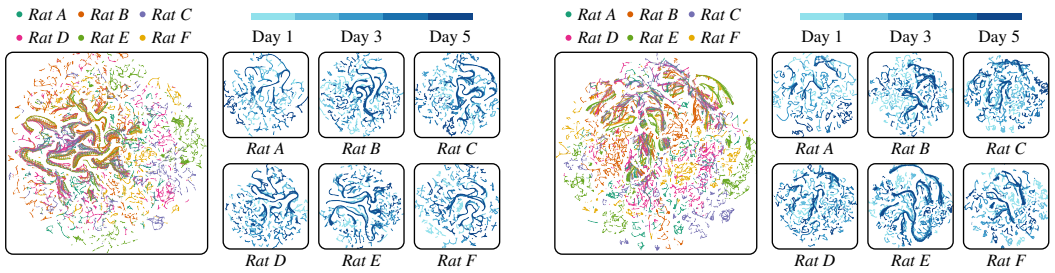

(i) TCN (j) BAMS

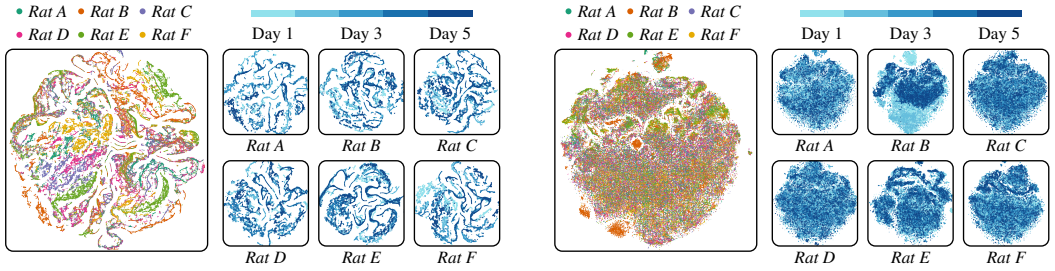

(k) hBehaveMAE (l) TimesFM-3 (zero-shot)

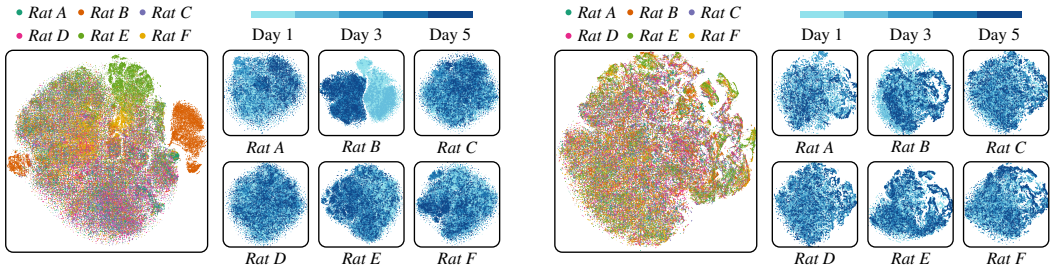

(m) TimesFM-3 (LoRA) (n) Chronos-2 (zero-shot)

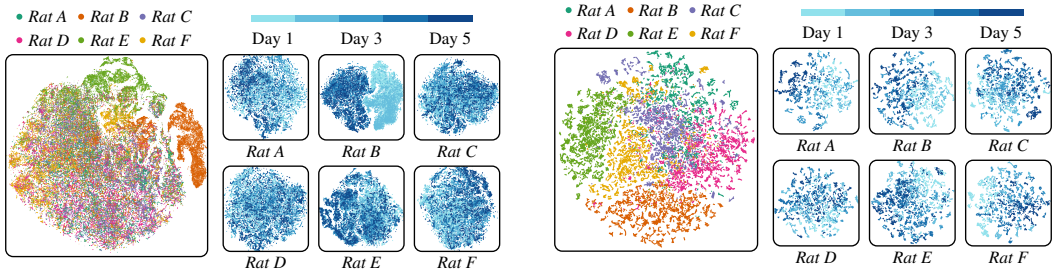

(o) Chronos-2 (LoRA) (p) TRAIL

Figure 8: **t-SNE visualizations of 3-second representations (continued).** The final panel repeats TRAIL from Figure 5 for direct comparison. Colors follow the same rat and day conventions as above.

E MOVEMENT TRANSFORMATIONS

E.1 DIFFUSION MODEL AND TRAINING DETAILS

Data and backbone. The diffusion model uses the alternation cohort's training and validation partitions, with trajectories in preprocessed absolute coordinates and conditions from the frozen movement encoder. The backbone follows DiT-B (Peebles & Xie, 2023), with 12 transformer blocks, hidden dimension 768, 12 attention heads, and an MLP expansion ratio of 4. Each keypoint at each frame forms a token. The representation h and diffusion timestep are incorporated through adaLN-Zero, and the model predicts the added noise and reverse-process variance.

Denoising and training objective. For each training trajectory x_0, we uniformly sample a diffusion timestep t and construct a noisy trajectory:

$$x_t = \sqrt{\bar{\alpha}_t}\, x_0 + \sqrt{1-\bar{\alpha}_t}\, \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I). \tag{7}$$

We use 1000 diffusion steps with a linear variance schedule from 10^{-4} to 0.02, where $\bar{\alpha}_t$ is the cumulative signal-retention factor. Conditioned on h, the model learns to recover the added noise from x_t while estimating the reverse-process variance:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{x_0,h,t,\epsilon}\left[\text{MSE}\big(\epsilon, \epsilon_\psi(x_t, t, h)\big) + \mathcal{L}_{\text{var}}\right], \tag{8}$$

$$\mathcal{L}_{\text{var}} = \frac{1}{\ln 2}\, \text{mean}_j\; D_{\text{KL}}(q(x_{t-1,j} \mid x_t, x_0) \,\|\, p_\psi(x_{t-1,j} \mid x_t, h)), \qquad t > 0. \tag{9}$$

The MSE term measures noise-prediction error, while the KL term trains the reverse-process variance by comparing the forward-process posterior q with the learned reverse transition p_ψ. The factor $1/\ln 2$ expresses the divergence in bits. Both terms are averaged over coordinates j in frames with no missing keypoints. For variance learning, gradients do not propagate through the predicted noise, and $\mathcal{L}_{\text{var}}$ is set to zero at $t = 0$.

Optimization. With the movement encoder frozen, we train the diffusion model from random initialization for 100 epochs using AdamW, a learning rate of 10^{-4}, batch size 256 per device, and zero weight decay. We maintain an exponential moving average (EMA) of model parameters with decay 0.9999 and select the EMA checkpoint with the lowest validation denoising loss.

DDIM inversion and generation. We use deterministic DDIM (Song et al., 2020) with 100 steps and $\eta = 0$. With the original representation h_o fixed, each step estimates the clean trajectory and advances to a higher noise level:

$$\begin{aligned} \widehat{x}_0 &= \frac{x_t - \sqrt{1-\bar{\alpha}_t}\, \epsilon_\psi(x_t, t, h_o)}{\sqrt{\bar{\alpha}_t}}, \\ x_{t'} &= \sqrt{\bar{\alpha}_{t'}}\, \widehat{x}_0 + \sqrt{1-\bar{\alpha}_{t'}}\, \epsilon_\psi(x_t, t, h_o), \qquad t' > t. \end{aligned} \tag{10}$$

Here, t' is the next timestep on the inversion grid; iterating to the final noise level yields u_o.

E.2 ESTIMATING DIRECTIONS OF PROGRESSION

Centers and axes. Let $\mathcal{S}_{a,d}$ denote the recording sessions for animal a on day d, and let $h_{s,i}$ be the representation of window i in session s. We first average the N_s representations within each session, then give session means equal weight:

$$\mu_s = \frac{1}{N_s} \sum_{i=1}^{N_s} h_{s,i}, \qquad c_{a,d} = \frac{1}{|\mathcal{S}_{a,d}|} \sum_{s \in \mathcal{S}_{a,d}} \mu_s. \tag{11}$$

Cross-day PCA uses the five centers $\mathcal{C}_a^{\text{day}} = \{c_{a,d}\}_{d=1}^{5}$ of one animal; cross-individual PCA uses the six centers $\mathcal{C}_d^{\text{individual}} = \{c_{a,d}\}_{a=1}^{6}$ on the same relative recording day. For either set $\mathcal{C}$, we define the unit axis by

$$v(\mathcal{C}) = \arg\max_{\|w\|_2=1} \sum_{c \in \mathcal{C}} \big[w^\top (c - \bar{c})\big]^2, \qquad \bar{c} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} c. \tag{12}$$

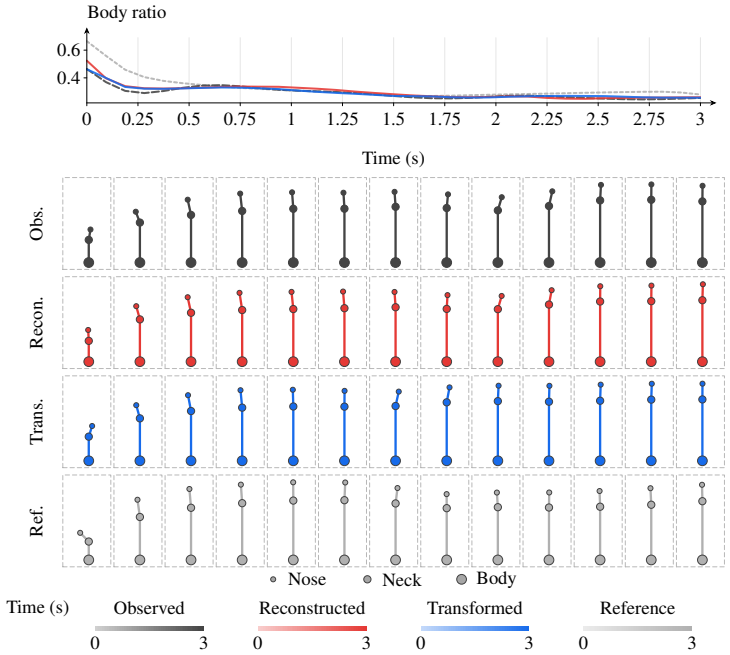

Figure 9: **Cross-day transformation of turning in local pose.** Skeleton sequences and body ratio for the same transformation as the left panel of Figure 6.

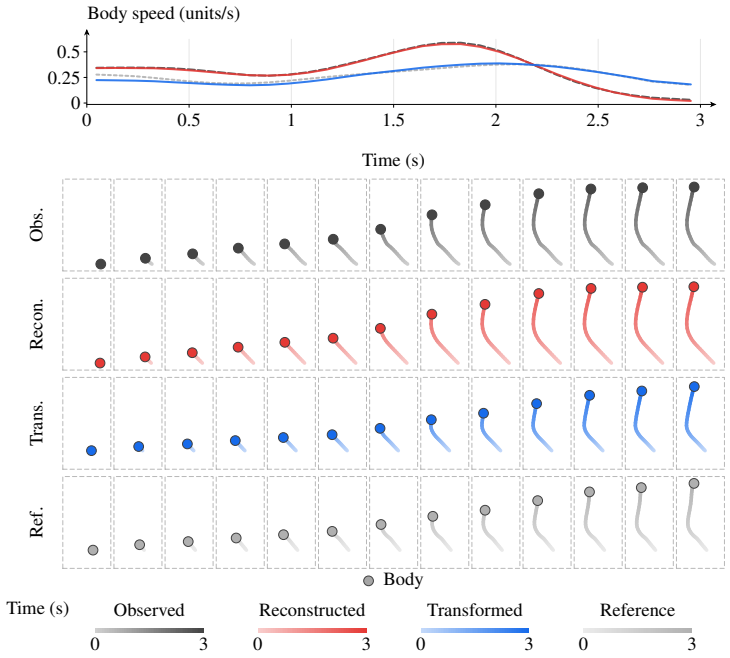

Figure 10: **Cross-individual transformation of turning in body trajectory.** Body-keypoint trajectories and body speed for the same transformation as the right panel of Figure 6.

This is the first principal component of the mean-centered daily centers, without dimension-wise standardization. Each animal has its own cross-day axis and each day its own cross-individual axis; these remain fixed across trajectory-pair comparisons.

Interpolation grid. We evaluate the transformation at $\tau \in \{0, 0.1, \ldots, 1\}$. At $\tau = 1$, the condition matches the reference's projection onto the axis while retaining the observed representation's orthogonal components.

E.3 Complementary Views of Turning Transformations

Figures 9 and 10 show the complementary views of the turning examples in Figure 6: Local Pose for the cross-day transformation and Body Trajectory for the cross-individual transformation.

Cross-day transformation of postural timing. In Figure 9, the transformed sequence reaches similar postures earlier than the reconstruction. The head's rightward inclination and subsequent realignment both occur sooner. The body-ratio curves remain close, so the visible change is chiefly in the timing of these postures.

Cross-individual transformation of movement pace. The cross-individual transformation in Figure 10 produces a lower, later speed peak and sustains forward movement near the end of the window. Both changes follow the reference animal's trajectory.

E.4 Movement Transformations in Going Straight

Figures 11 and 12 show the corresponding transformations for going straight. Each figure pairs Body Trajectory and body speed on the left with Local Pose and body ratio on the right.

In the cross-day example, the observed and reconstructed movements start slowly and accelerate later. The transformed movement starts faster, peaks early, and then slows, following the later-day reference. Its body-ratio curve remains close to the reconstruction, while the speed profile changes across the window.

In the cross-individual example, the transformed speed curve moves toward the reference's earlier peak and subsequent decline. The nose–neck to neck–body ratio also rises through the middle of the sequence, approaching the reference. Here, the transformation changes both movement pace and relative body extension.

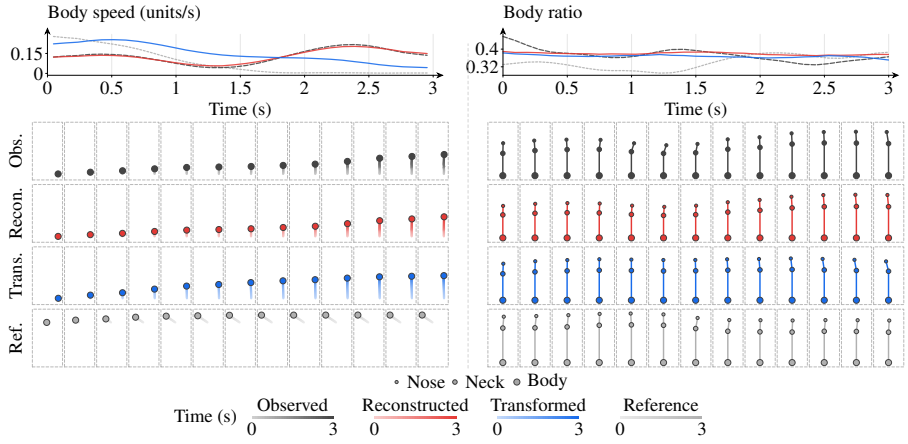

Figure 11: **Cross-day transformation of going straight.** *Left:* Body Trajectory with body speed. *Right:* Local Pose with body ratio.

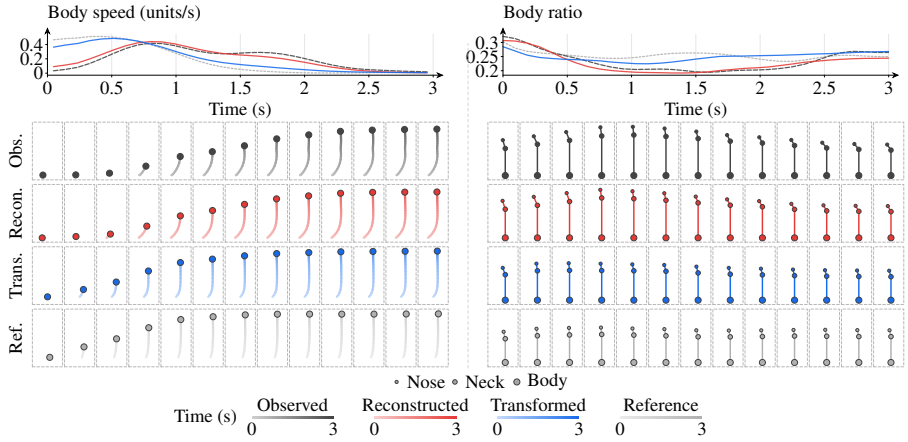

Figure 12: **Cross-individual transformation of going straight.** *Left:* Body Trajectory with body speed. *Right:* Local Pose with body ratio.