

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031

HOW NEURAL IS A NEURAL FOUNDATION MODEL?

Anonymous authors

Paper under double-blind review

ABSTRACT

Foundation models have shown remarkable success in fitting biological visual systems; however, their black-box nature inherently limits their utility for understanding brain function. Here, we peek inside a SOTA foundation model of neural activity (Wang et al., 2025) as a physiologist might, characterizing each ‘neuron’ based on its temporal response properties to parametric stimuli. We analyze how different stimuli are represented in neural activity space by building *decoding manifolds*, and we analyze how different neurons are represented in stimulus-response space by building *neural encoding manifolds*. We find that the different processing stages of the model (i.e., the feedforward *encoder*, *recurrent*, and *readout* modules) each exhibit qualitatively different representational structures in these manifolds. The *recurrent* module shows a jump in capabilities over the *encoder* module by “pushing apart” the representations of different temporal stimulus patterns. Our “tubularity” metric quantifies this stimulus-dependent development of neural activity as biologically plausible. The *readout* module achieves high fidelity by using numerous specialized feature maps rather than biologically plausible mechanisms. Overall, this study provides a window into the inner workings of a prominent neural foundation model, gaining insights into the biological relevance of its internals through the novel analysis of its neurons’ joint temporal response patterns. Our findings suggest design changes that could bring neural foundation models into closer alignment with biological systems: introducing recurrence in early encoder stages, and constraining features in the readout module.

1 INTRODUCTION

Viewed in the large, deep neural networks are intriguing models of the mouse visual system, since they learn to predict neural responses directly from visual input (Cowley et al., 2023; Ustyuzhaninov et al., 2022; Huang et al., 2023; Averbeck et al., 2006; Qazi et al., 2025; Li et al., 2023), and recent foundation models can generalize, to some extent, beyond training data (Li et al., 2023). Viewed in the small, Representational Similarity Analysis (RSA) Kriegeskorte et al. (2008) shows that, on average, many units in these networks reflect properties (e.g. orientation selectivity) resembling those found in biology (Conwell et al., 2021; Qazi et al., 2025). However, while this progress has been impressive, questions are arising about whether the pairwise activity of units in artificial networks agrees with biological data (Liscai et al., 2025). Moreover, in the large view the input/output maps are far from complete (normalized response correlation ceilings around 70% (Wang et al., 2025)), raising questions about their robustness. In effect, response correlation measures how well the input drives the system to the correct output; it does not address the inverse question of how ambiguity in the output obscures the input. That is, one must consider both the “forward” and the “inverse” mappings. Such issues are classical in modeling: control theory teaches us that, without a perfect model, one must “look inside the box” to achieve identifiability (cf. (Åström, 2012)). We seek to do just this on the Foundation Neural Network (FNN) (Wang et al., 2025). Without this, we cannot guarantee correct, robust, and generalizable behavior, especially on out-of-distribution data, to confidently build hypotheses about the brain using the FNN. The FNN was selected because it was trained on MICrONS, the largest available functional connectomics dataset of the mouse visual system (Bae et al., 2025) and is based on artificial and naturalistic input videos across multiple animals. The FNN thus provides the SOTA in modeling.

FNN consists of multiple stages (Figs. 1B and 6) and millions of units, so analyses beyond pairwise interactions—such as third- or fourth-order statistics—are computationally prohibitive. To “look inside,” we use three techniques popular in neuroscience. These allow us to: (1) evaluate how the state of the network represents the different stimuli; i.e. how stimuli are related to one another in global

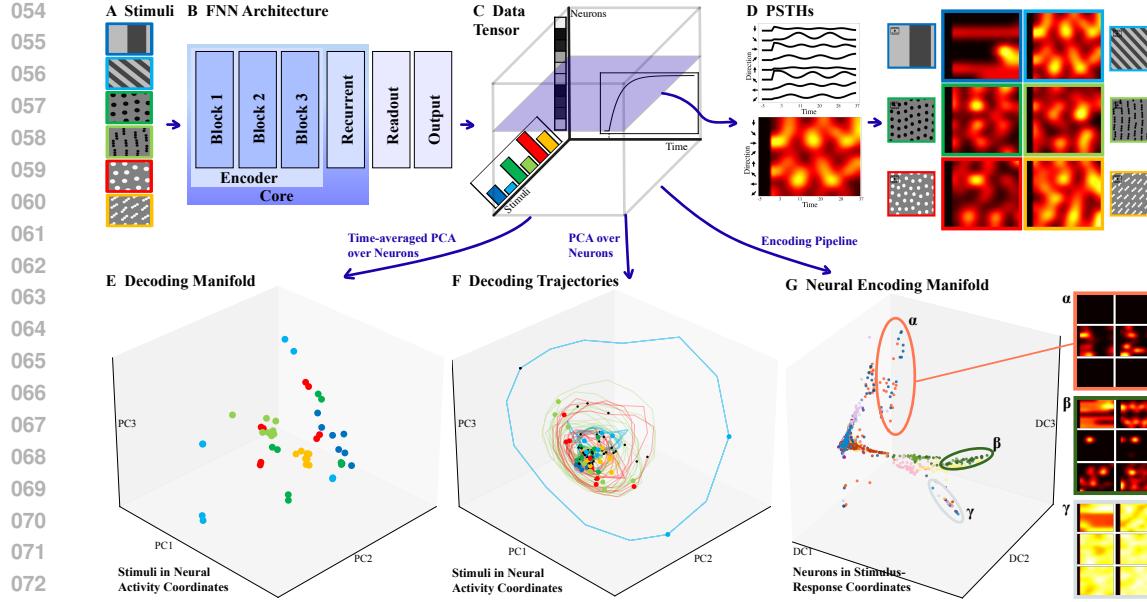


Figure 1: Approach and manifolds analysis **A** Stimulus ensemble provides input. **B** FNN consists of multiple encoding blocks, modeled as convolutional layers, followed by recurrent and readout/interpolation layers. **C** The tensor of data, containing the response (in time) of each sampled unit to the stimulus ensemble. **D** PeriStimulus Time Histogram: The response (instantaneous “firing rate”) of a single unit/neuron to a stimulus pattern drifting in each of 8 different directions. The curves are redrawn as an image, with brightness corresponding to activity. A plane through the data tensor shows the PSTHs for each of the 6 stimulus classes, drifting in all directions. **E** Decoding manifold, plots the total activity for each stimulus in PCA-reduced neural coordinates. Colors correspond to stimulus classes in **A**. **F** The time evolution of each stimulus presentation, plotted in PCA-reduced neural coordinates for the early encoder layer. Note the nested, periodic trajectories indicating a stimulus drifting over a receptive field filter. **G** Encoding manifold plots individual units/neurons in stimulus/response coordinates. Note the clustering of units with similar responses across the ensemble.

neural coordinates (Fig. 1E); (2) to show how all the units in the network are related to each other functionally when driven by the stimuli (Fig. 1G); i.e. how they encode information; and (3) how the dynamics evolve as the network processes the different stimuli; i.e., how the global neural state changes in time during a computation (Fig. 1F). The first two techniques result in manifolds characterizing *forward* and *backward* mappings, respectively, and the third in trajectories over these manifolds; all can then be compared against biology. The result, in brief, is that while the FNN learned a forward map reasonably well, it processes stimuli quite differently from the mouse, and hence is only a partial “digital twin” in the dynamical sense. Importantly, our manifolds identify where the disparities lie.

In more detail, (1) we build *neural decoding manifolds* (Chung and Abbott, 2021), in which trials are embedded in the space of neural activity coordinates (Fig. 1E), then dimensionality-reduced using Principal Component Analysis (PCA) (Cunningham and Yu, 2014). Typically, trials involving the same stimulus cluster together, facilitating a read-out of the brain’s state. (2) To switch from trials to neurons, we build *neural encoding manifolds* (Fig. 1G) (Dyballa et al., 2024a) in which each point is a neuron in the space of stimulus-response coordinates, dimensionality-reduced using tensor factorization (Williams et al., 2018). Proximity between neurons in an encoding manifold denotes similar responses to similar stimuli; i.e., groupings of neurons that are likely to share circuit properties. For a review of classic encoding/decoding in neuroscience, see (Mathis et al., 2024). Finally, (3) the relationship between these two manifolds is captured by the temporal evolution of each neuron’s activity for each stimulus trial. Recalling that a ‘neural computation’ can be viewed as the result of a dynamical system in neural state space (Hopfield, 1984), we plot these both as PeriStimulus Time Histograms (PSTHs, Fig. 1D) and as streamline traces (decoding trajectories, Fig. 1F). While streamline representations have been used previously for decision tasks (Duncker and Sahani, 2021) and the motor system (Churchland et al., 2012; Safaie et al., 2023), we note: (i) the activity integral along such *decoding tra-*

108 *jectories* (Fig. 1F) defines the decoding manifold, while (ii) shared tubular neighborhoods (developed
 109 below) specify position in the encoding manifold. These three perspectives enable us to investigate dif-
 110 ferent aspects of alignment: (1) Decoding manifolds reveal whether the model maintains stimulus sep-
 111 arability like biology; (2) Encoding manifolds reveal whether functional topology of neurons is brain-
 112 like; (3) Trajectories reveal whether the model performs computations through brain-like dynamics.
 113 Critically, a model could succeed at one level while failing at others. We use modeling tools available
 114 online (references in Methods), stimuli similar to those used in FNN’s original training (Wang et al.,
 115 2025), and add naturalistic flow stimuli used in mouse physiology (Dyballa et al., 2018) (Fig. 1A).

116 **Prior Work.** There is an extensive literature on modeling biological neural responses (Averbeck
 117 et al., 2006; Ustyuzhaninov et al., 2022; Qazi et al., 2025), including other foundation models
 118 (Zhang et al., 2025; Azabou et al., 2023; Ryoo et al., 2025; Ye et al., 2023, 2025). We highlight that
 119 compared to these other approaches, the FNN is concerned with predicting neural activity from input
 120 videos. The FNN is an example of a data-driven predictive model (Klindt et al., 2018; Turishcheva
 121 et al., 2024; Nellen et al., 2025) with Gaussian readout (Lurz et al., 2021) that interprets the readout
 122 as per-neuron basis functions with individual readout weights. The readout thus provides an encoding
 123 embedding of biological neurons. For comparability, we use our encoding method to compare
 124 the embeddings of biological neurons and individual readout neurons, investigating not only the
 125 final embedding but also the readout embedding. Different loss functions have been used (Nayebi
 126 et al., 2023; Bakhtiari et al., 2021; Shi et al., 2022), and others have studied decoding manifolds
 127 (Froudarakis et al., 2020; Beshkov and Tiesinga, 2022; Beshkov et al., 2024), focusing on topological
 128 properties. For a recent general review, see Doerig et al. (2023). Some studies are supportive of
 129 modeling brains with deep networks (Kriegeskorte, 2015; Yamins et al., 2014; Margalit et al., 2024),
 130 while others raise questions (Serre, 2019a). For the reasons stated above we focus on the FNN.

131 To our knowledge, this is the first time all three of the encoding and decoding manifold techniques
 132 have been utilized together for analysis of a perceptual system; i.e., toward *interpretability* for a
 133 foundation model. Interpretability is a rapidly evolving field for analyzing large language models
 134 (Elhage et al., 2021; Bricken et al., 2023; Skean et al., 2025), vision models (Simonyan et al., 2014;
 135 Olah et al., 2017), and recurrent models (Krakovna and Doshi-Velez, 2016). This field has been
 136 connected to neuroscience, arguing that both aim to understand complex intelligent black boxes (Kar
 137 et al., 2022; Tolooshams et al., 2025; He et al., 2024; Mineault et al., 2025). It aims to investigate
 138 the function of individual neurons, circuits, and modules in artificial networks, while in neuroscience
 139 it additionally focuses on the alignment between artificial models and biological systems (Kar et al.,
 140 2022). We tackle both challenges by trying to understand what functions the FNN modules fulfill
 141 and by testing alignment with biological representations.

142 Within this framework, we ask: *Do neural decoding and encoding manifolds reveal new insights*
 143 *into how foundation models represent temporal response patterns? Are their representations brain-*
 144 *like?* We hypothesize that each processing stage contributes distinct representational capabilities, all
 145 essential for fitting neural data. In particular, one might expect the *recurrent* module to enrich the
 146 temporal structure of representations, analogously to the cortex, and the encoder layers to resemble
 147 the retina with its limited recurrence. Following a brief description of our methods, we proceed to
 148 develop each of the manifolds in turn.

149 2 METHODS

150 Our work makes novel use of publicly available open-source resources. Specifically, we employed
 151 the pretrained foundation model of neural activity (denoted FNN) provided by Wang et al. (2025),
 152 available here; and the stimulus generation tools and neural encoding manifold construction pipeline
 153 introduced by Dyballa et al. (2024a), accessible at here. Below we briefly outline our methods, and
 154 refer readers to Appendix A for the full details.

155 **Model:** The FNN consists of five modules: perspective, modulation, *encoder*, *recurrent*, and *readout*
 156 (see Fig. 6). The perspective and modulation modules model the mouse’s state and transform the
 157 inputs to approximate the actual visual information received. Thus, only the *encoder*, *recurrent*, and
 158 *readout* modules perform the core computation, and are the focus of this work. The *encoder* module
 159 is a 10-layer DenseNet-style convolutional encoder (Huang et al., 2017). For analysis, we use a
 160 subset of *encoder* layers; we report results from the very first layer and the last block as representative
 161 examples (remaining layers in the appendix). Notably, the encoder includes 3D convolutions, which

in principle enable it to capture temporal patterns for up to 12 timesteps in the last encoder layers. The *recurrent* module is preceded by an attention layer and consists of a convolutional LSTM, followed by a single convolutional layer that produces its output. This feedforward–recurrent combination constitutes the core of the FNN, which is trained on data from all mice combined. Finally, a separate *readout* module is trained on each mouse individually: it performs an interpolation on the recurrent output followed by a linear transformation to produce the FNN output. We included one scan (session 8, scan 5) for readout and output analysis, and validated the findings on other sessions and scans. We claim that comparison across mice on the *population* level, rather than on individual neurons, is valid.

Stimuli: Our stimulus set is composed of drifting square-wave gratings and optical flows with varying spatial frequencies moving in eight directions. This yields 88 unique input sequences with stochastic initial positions and velocities (Fig. 1A). To ensure that these stimuli would drive the network in a representative manner, we compared the output of the network for these stimuli with the output for the original natural movie stimuli used to train the network (Appendix Figs. 9 and 10); we found the results to be quantitatively similar in all measured respects.

PSTH visualization: To visualize the network responses to stimuli concisely, we group together the model’s PeriStimulus Time Histogram responses (PSTH) corresponding to all flow directions of a given stimulus pattern with time on the *x*-axis and flow direction on the *y*-axis (Fig. 1D).

Decoding manifolds & trajectories: Following traditional analysis techniques, we first constructed decoding manifolds by performing PCA on the stimulus-time-averaged activity data. Therefore, the decoding manifold contains 48 points, one for each unique sequence, colored by the corresponding base-stimulus (as shown in Fig. 1A); different spatial frequencies of the same stimulus are summarized with the same color. To construct *decoding trajectories*, we treated each time step as a separate data point rather than averaging across time before applying PCA. We compared with biological decoding trajectories using the experimental data from Dyballa et al. (2024a).

Tubularity: To investigate neural dynamics, we modeled trajectories by bundling them into tubular neighborhoods around a central skeleton (Budanur, 2023). We operationalized this idea for discrete data using the tubular neighborhood theorem (Da Silva, 2008), which guarantees that smooth submanifolds admit non-intersecting neighborhoods diffeomorphic to their normal bundles. Let $\{\gamma_i\}_{i=1}^m \subset \mathbb{R}^D$ denote a set of m trajectories (curves). We define this set as *tubular* if it remains close to a common centerline c and exhibits minimal transverse intersections. Formally, the tube is obtained by expanding c with a radius profile $R(\cdot)$ such that all points at parameter u within distance $R(u)$ of $c(u)$ are included. In practice, curves are first clustered (e.g., via HDBSCAN (Campello et al., 2013) using the Sobolev H^1 metric, or with ground truth) to separate distinct tubes before computing tubularity scores. We introduce *tightness*, which measures how tight a group of curves is around the centerline, and *crossings*, measuring how many transverse crossings occur in each trajectory bundle. Therefore, tubularity is not a trajectory-matching metric but a population-geometry metric: it assesses the structure of collections of trajectories rather than the similarity of individual pairs.

Alignment metrics: To validate our results against the literature, we calculated scores for Representational Similarity Analysis (RSA) (Kriegeskorte et al., 2008), Canonical Correlation Analysis (CCA) (Raghu et al., 2017), Linear Predictivity (LP) (Yamins et al., 2014), and Dynamic Similarity Analysis (DSA) (Ostrow et al., 2023) (details in Appendix A.12)

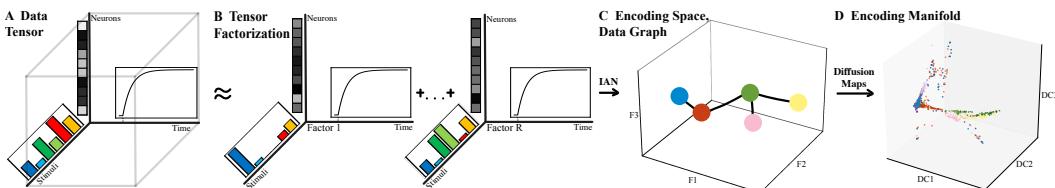


Figure 2: Encoding Manifold Pipeline (A, B) A non-negative tensor factorization of the original data tensor identifies those neural factors that account for part of the stimulus ensemble over comparable time epochs. **(C)** Collecting the neural factors into a linear vector space, an adaptive-neighborhood kernel builds a data graph. **(D)** Diffusion maps yield the encoding manifold.

Encoding manifolds: To understand the response properties of *neurons* with respect to all stimuli (rather than the representation of *stimuli* in the space of all neurons), we finally constructed *encoding manifolds*. At a high level (Fig. 2), these manifolds allow one to examine the global topology of neuronal populations based on their stimulus selectivities and temporal response patterns (Dyballa

et al., 2024a). The neural encoding manifold was constructed in three steps. First, a 3-tensor was built with the temporal responses from each neuron for each stimulus, and decomposed using Nonnegative Tensor Factorization (details in Appendix), with each component comprised of neural, stimulus, and temporal response factors. The neural factors then serve as position coordinates, embedding the neurons into a stimulus-response framework called the neural encoding space. Second, we constructed a data graph in this neural encoding space using the IAN algorithm (Dyballa and Zucker, 2023). Third, applying diffusion maps (Coifman et al., 2005; Coifman and Lafon, 2006) to the data graph yielded the manifold. We followed the methodological choices of Dyballa et al. (2024a), where extensive parameter analysis for biological neural data was conducted.

3 RESULTS

We built encoding and decoding manifolds, as well as decoding trajectories, for all layers of the modules considered in the FNN. Here, we focus on the results that were most informative for interpreting the computational role of each stage of the network and for comparing the FNN representations to biological results (see Appendix for extended results). The **decoding manifolds** assess *stimulus separability*, the **encoding manifolds** capture *global neuronal response similarity and topology*, and the **trajectories** characterize *response dynamics*. Together, these analyses provide complementary perspectives for evaluating brain alignment at the population level.

3.1 DECODING MANIFOLDS

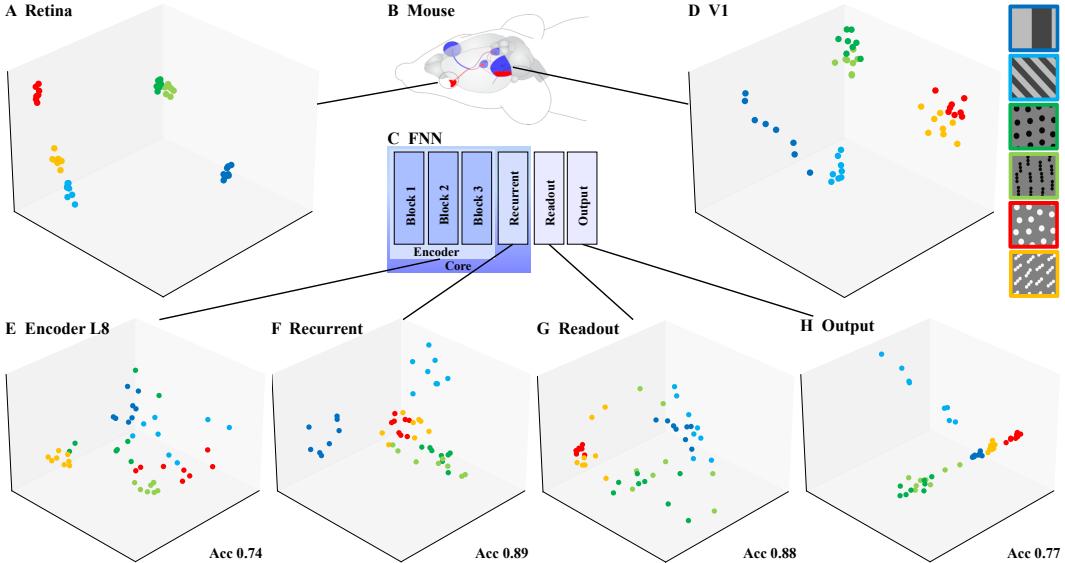


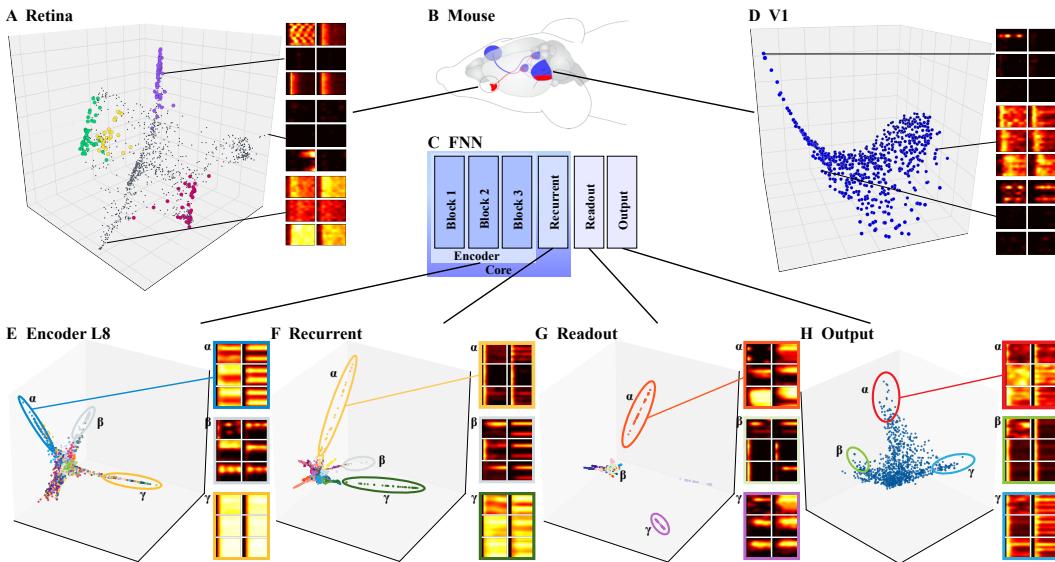
Figure 3: **Decoding Manifolds** for the mouse (**A**) retina and (**D**) visual cortex are highly clustered by stimulus (color labels shown in top-right bar) supporting decoding (i.e., reading out the stimulus from neural responses) in both cases. By contrast, the FNN is most clustered at the recurrent and readout stages (**E–H**). Acc: classification accuracy for that layer (see Table 1). Notice how the encoder (first stage in the FNN) differs significantly from the retina (first stage in the visual system); on the other hand, the recurrent layer is most analogous to V1.

Table 1: Stimulus classification accuracy for Leave-One-Out 3-Nearest Neighbor (3-NN) and Logistic Regression (LR) classifiers trained on each layer’s activations. Methods in Appendix A.

Accuracy	L1	L2	L4	L5	L7	L8	Rec	RecOut	Readout	Out
LR	0.59	0.62	0.66	0.65	0.71	0.74	0.89	0.90	0.88	0.77
3-NN	0.41	0.66	0.58	0.52	0.53	0.61	0.73	0.64	0.63	0.67

270 The biological decoding manifolds (Fig. 3A, D) showed clear clustering by stimulus with some overlap
 271 between the related 1-dot and 3-dot stimuli. It follows that neural responses at both the retina and
 272 cortical levels can be used to “read out” the stimulus. By contrast, the first encoder layer (L1) yielded a
 273 poorly clustered decoding manifold (Fig. 1E) in which stimulus classes were mixed. This implies that
 274 the latent feature representation at this point within the FNN is not sufficient to distinguish between
 275 the different stimuli (indeed, its classification accuracy is lowest; see Table 1). The decoding manifold
 276 for layer 8 (L8) was similar to that for L1, but with greater stimulus-specific clustering. **The recurrent**
 277 **decoding manifold was closest to the biological data, showing more distinct clusters and greater**
 278 **overlap between 1-dot and 3-dot stimuli.** Following this, the readout and output decoding manifolds
 279 showed weaker clustering, suggesting these stages are responsible for fitting neural data rather than
 280 enriching the model’s representations. This aligns with the **classification accuracy being highest**
 281 **for the recurrent stage and dropping again afterwards, rather differently from biology.**

282 3.2 ENCODING MANIFOLDS



302 **Figure 4: Encoding Manifolds** for the mouse **(A)** retina and **(D)** visual cortex differ significantly:
 303 retina is clustered and cortex is continuous. Example PSTHs show how functionality varies smoothly
 304 in cortex but not in the retina. **(E)** The encoder stage showed a distinct arm of orientation-selective
 305 units (α), which are compatible with biological results, and another of intensity-based units (γ),
 306 which are not. **(F)** The recurrent stage showed many direction-selective units, but the following **(G)**
 307 readout stage was the most clustered among all stages. This “bottleneck” layer is then interpolated to
 308 a continuous **(H)** output layer. While the topology of this final layer is similar to that of biological
 309 visual cortex, the responses of individual units (PSTHs) are not.

310 The encoding manifolds were even more revealing about differences between the mouse and FNN.
 311 Replotting data from Dyballa et al. (2024a), we start with the retinal manifold (Fig. 4A). The neurons
 312 form clear clusters, each one with distinct response patterns (PSTHs) that corresponded to known
 313 retinal ganglion cell types. By contrast, the V1 encoding manifold is continuous, with smooth
 314 transitions in response patterns as it is traversed. See Dyballa et al. (2024a) for further discussion.

316 The encoding manifold for L1 (Fig. 1G) revealed that most neurons belonging to the same feature
 317 map (points with the same color label) formed contiguous clusters, or regions, over the manifold;
 318 this was not entirely surprising given the weight-sharing property of these convolutional layers.
 319 Nevertheless, several feature maps were found mixed into the same “arm” (labeled β). Examining
 320 the response patterns (PSTHs) of these neurons in detail, we observed strong, continuous activity
 321 across the entire trial duration with no selectivity for directions or stimulus classes. There was no
 322 biological counterpart to this type of neurons.

323 We now move on to the late-stage encoder layer, L8 (Fig. 4 E). Its encoding manifold again showed
 324 grouping by FNN feature maps, but with more mixing than in L1. This was especially true in the

poorly selective “intensity arm” of neurons, (β) which exhibited strong response (PSTHs) for all stimuli across multiple feature maps. Further investigation revealed that the intensity arm resulted from an FNN technical requirement: padding artifacts at the edges of feature maps. Such artifacts are a well-known issue in convolutional models (Alsallakh et al., 2020), and we also observed them in Du et al. (2025)’s model (Fig. 17). Sampling only from the central regions of feature maps eliminated both the intensity arm and the shared activity pattern seen in the decoding trajectories (see Supplemental Fig. 16). Although these artifacts distort the representation—indeed, the smoothness of the intensity arm reflects how padding-related information propagates across feature maps—they are part of the network’s normal operation. Excluding them would therefore misrepresent the model’s true internal dynamics, so we retained them in our manifold analysis.

We emphasize that the non-selective groups of neurons with high activity (labeled as β in Figs. 1E and 4E) were a significant departure from what is found in biological networks: in the retina, there are no such non-selective neurons. Although low selectivity has been observed in cortex, it is restricted to inhibitory (inter)neurons and continuously mixes with other, more selective responses; they do not segregate into an arm or cluster (Dyballa et al., 2024a).

The *recurrent* module was qualitatively different. Its encoding manifold showed that different regions exhibited distinct selectivity and temporal response patterns, as evidenced by their PSTHs (Fig. 4F). Furthermore, although segregation by feature map was still present, there was no longer a cluster of neurons with no selectivity; instead, the highlighted β group showed selectivity for particular directions or orientations, as is typical in biological visual neurons (e.g., PSTHs in Fig. 4D).

The final stages of the network—the *readout* and *output* layers—were again different. The encoding manifold for the readout layer analyzes the intermediate readout neurons in stimulus-response space, not the final biological output neurons. **It was highly disconnected** (Fig. 4G), with each cluster corresponding almost exclusively to neurons sampled from a single feature map. Each feature map exhibited a distinct response pattern that was invariant across its neurons. Compared to this, the biological results (e.g., Baden et al. (2016); Dyballa et al. (2024a)) showed more variability within functional cell “types”, even in the retina. Curiously, and despite this intra-map uniformity, the large number of feature maps (see PSTHs) and the rich dynamics within each one, somehow enable the *output* to represent the complex behavior of neurons (Fig. 4H). These behaviors are captured in the FNN output via a linear combination of *readout* features. Since classification accuracy has declined slightly at this stage (Supplemental Fig. 8), but orientation and direction selectivity agree (Supplemental Fig. 10), we conjecture that these dynamics interpolate the spiking activity individually for each mouse data used as input. **The smooth manifold aligned most closely with the biological V1 manifold** (Fig. 4D), although **the large number of transient responses in the FNN did not match what was found in V1** (across different animals, scans, and sampling procedures).

3.3 DECODING TRAJECTORIES

The encoding manifolds revealed functional differences between FNN and biology: both in the topology of the neuronal organization, and in the PSTHs i.e. temporal responses for multiple stimulus classes. This motivated a direct analysis of the population response dynamics. **The biological decoding trajectories showed stimulus-dependent development of activity** (Fig. 5A,D). They formed segregated, stimulus-dependent bundles whose temporal dynamics allowed linear separability during much of the trial’s time course. Here, V1 activity showed more bundles and less collinear development of trajectories. This indicates a higher complexity of response patterns in V1 compared to the retina.

Turning to FNN, the decoding trajectories for L1 revealed that periodic stimuli were represented as loops (Fig. 1F). This was likely due to the translation equivariance of the convolutional layers used in the encoder stage, which preserved the circular geometric structure of these stimulus sequences (Cohen and Welling, 2016). However, we saw that these loops could take on many different forms (such as that for the high spatial frequency gratings, shown in light blue), influenced by the responses of particular groups of neurons to each stimulus. Layer 8, by contrast, showed stimulus-independent temporal decoding trajectories (Fig. 5E). Our analysis of removing the intensity arm from the encoding manifold showed that this temporal development of activity could be attributed to an non-selective increase in intensity during the first timesteps (Supplemental Fig. 16). Without the intensity arm, L8 has highly stationary neural activity. Thus, despite temporal convolutions, the **FNN feedforward encoder appears to lack biologically plausible stimulus-dependent temporal patterns** and primarily reports features present in the input, with varying intensities.

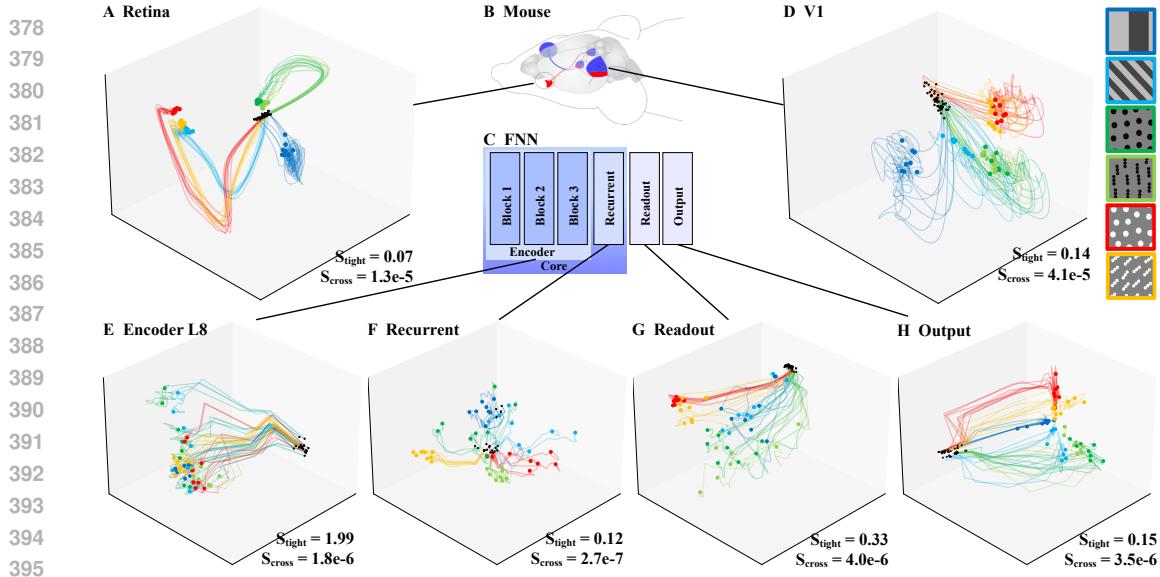


Figure 5: **Decoding Trajectories** in the retina (A) and V1 (D) show the development of neural activity dynamics into stimulus tubes. The encoder (E) shows only a non-selective increase in activity (see also Figure 16) rather than stimulus-dependent tubes. From the recurrent stage onward (F–H), tubular trajectories similar to those seen in biological data are present. The tubularity metrics quantify this phenomenon (S_{tight}), and also highlight a lack of complexity in FNN activity compared to the biological data, reflected in their lower crossings values (S_{cross}).

The recurrent module showed a qualitative change in decoding trajectories compared to the encoder (Fig. 5F). **Similarly to the biological results, tubular temporal patterns were present at the recurrent stage.** Still, the organization of decoding trajectories was noticeably more entangled than both retina and V1 (compare with Fig. 5A,D). This phenomenon was quantified using *tubularity* metrics based on the geometry of the observed decoding bundles (see Methods). The tightness scores were comparable between biological and FNN data from the recurrent stage onward (Fig. 5, Table 4). The retinal trajectories were the tightest, while V1 and FNN trajectories showed more expanded cones of trajectories. In particular, the FNN readout trajectories were less tight because they linearly spread out from the origin. The recurrent trajectories were also spread out, but retained a tight stimulus-dependent organization towards the end of the time frame. The tightness score for trajectories from the output stage was difficult to interpret: the predominance of transient responses caused a convergence towards a common point of low activity, which might bias the tightness score to be lower.

A more pronounced difference was observed in the **crossings scores**. Biological trajectories exhibited more crossings than those of the FNN, despite their tight tubular development ($p < 0.005$, Bonferroni-corrected, for all layers). These crossings occurred predominantly toward the end of the time frame, when the activity seemed to settle into a steady state. Several factors could explain this pattern. One possibility is that biological recordings contain inherently more noise, which could artificially lead to more crossings. However, the noise observed toward the end of the biological trajectories is of similar magnitude as the overall tube diameter. If measurement noise were the only cause, we would expect less coherent (tubular), more erratic (noisier) trajectory development already at earlier time steps, which is not observed. A second possibility is that the crossings reflect genuine neural dynamics captured in the data, suggesting that biological systems exhibit more complex temporal processing than the FNN. Modulatory phenomena such as clique-like interactions (Miller and Zucker, 1999) or traveling wave activity (Pitts and McCulloch, 1947; Milner, 1974; Keller et al., 2024) could generate these apparent fluctuations late in the trial. These results indicate that while parts of the FNN reproduce certain aspects of biological temporal structure (such as tubular structure), it is not yet capable of fitting the full intricacies of neural dynamics observed in real neural populations.

The readout and output stages exhibited tubular trajectories that were less well separated than those observed in retina and V1 (Fig. 5G,H), consistent with the less clustered organization seen in the decoding manifolds. In the output trajectories, the bias towards transient responses was clearly visible

432 as all trajectories originated from a common point (black, high activity) and converged toward a
 433 shared low-activity point via different paths.
 434

435 **3.4 REPRESENTATIONAL ALIGNMENT METRICS**
 436

437 To validate the results of our manifold analysis, we quantified the representational alignment of the
 438 FNN with both V1 and retina using standard alignment metrics from the literature (Kriegeskorte
 439 et al., 2008; Raghu et al., 2017; Yamins et al., 2014; Ostrow et al., 2023). We found that our
 440 result of the recurrent module being most aligned with biology in terms of decoding analysis was
 441 supported by these metrics (see Tables 2, 6). The DSA metric (Ostrow et al., 2023), while correctly
 442 showing higher values for tubular dynamics in the recurrent stage and after, wrongly predicted high
 443 alignment between the FNN’s L1 and the biological data. This is likely due to tubular trajectories
 444 arising for entirely different reasons (i.e., local stimulus periodicity). Moreover, **smoothness and**
 445 **neuronal responses (PSTHs) in the encoding manifold showed a clear misalignment between**
 446 **the FNN’s recurrent stage and V1. This relationship was not captured by the standard metrics**,
 447 underscoring the need for our analysis at the population level.
 448

449 **4 DISCUSSION**
 450

451 **Decoding manifolds and trajectories** allow us to assess whether networks achieve comparable
 452 degrees of stimulus representation and separability. **Encoding manifolds**, on the other hand, evaluate
 453 at a global level how the responses and global organization of individual neurons compare to those in
 454 biological systems; in other words, whether the FNN and biological networks employ similar encod-
 455 ing mechanisms to produce similar outputs. Finally, **decoding trajectories** serve as a surrogate for
 456 *computation*, reflecting the dynamics of activity over the neural state space (cf. (Hopfield, 1984)). Our
 457 analysis of the FNN revealed an increasing richness of representation up to the *recurrent* module (cf.
 458 Hoeller et al. (2024); see also contrasts with Xu et al. (2023); Nayebi et al. (2023); Froudarakis et al.
 459 (2020)). However, most PSTHs lacked the characteristic temporal response profiles observed in bio-
 460 logical recordings Ringach et al. (2016); Ko et al. (2011). Since the FNN was trained to predict neural
 461 spike trains, **classification behavior evolved implicitly** (cf. Table 1)). Thus, it is plausible that the re-
 462 current features are sufficiently complex for robust feature representation and that the subsequent mod-
 463 ules serve to fit the neural data rather than to provide additional biologically meaningful computations.
 464

465 However, the highly clustered topology of the latent representation observed in the *readout* module
 466 was not consistent with that of the retina or cortex (cf. Baden et al. (2016); Dyballa et al. (2024a), nor
 467 with those of higher visual areas (cf. Glickfeld and Olsen (2017); Dyballa et al. (2024b); Yu et al.
 468 (2022)). Nevertheless, the rich dynamics within each feature map (as evident in the PSTHs), together
 469 with their large number, seem to enable the output layer to capture the complex response patterns
 470 of neurons, resulting in the network’s strong performance in predicting neural activity. Still, it is
 471 somewhat surprising that such biologically realistic outputs are produced at the FNN’s output through
 472 a simple linear combination of readout features—one would expect the fitting of neural activity to
 473 occur throughout the entire network, rather than as a separate appendage module.
 474

475 Our analysis pipeline was validated by its overall agreement with commonly used alignment metrics
 476 (Kriegeskorte et al., 2008; Yamins et al., 2014; Raghu et al., 2017; Ostrow et al., 2023) in predicting the
 477 closest alignment at the recurrent stage. However, the reliability of such metrics has been questioned
 478 in the recent literature (Schaeffer et al., 2025; Anonymous, 2025; Bowers et al., 2023; Lampinen
 479 et al., 2025; Dujmovic et al., 2024; Serre, 2019b). Beyond this high-level alignment, our analysis
 480 also exposed some limitations of these alignment approaches, such as with the DSA metric (Ostrow
 481 et al., 2023). This highlights the advantage of our manifold-based framework over simple metric in-
 482

483 **Table 2: Mean representational alignment metrics.** Mean taken over Representational Similarity
 484 Analysis (RSA), Canonical Correlation Analysis (CCA), Linear Predictivity (LP) and Dynamic
 485 Similarity Analysis (DSA) scores. Details in Appendix A.12); individual metric values in Table 6.
 486

Region	Enc L1	Enc L2	Enc L4	Enc L5	Enc L7	Enc L8	Rec	Readout	Output
Retina	0.26	0.26	0.30	0.33	0.28	0.28	0.40	0.34	0.34
V1	0.29	0.21	0.32	0.30	0.30	0.32	0.53	0.38	0.48

486 spection: it provides a deeper understanding of the model’s internal computations and representations.
 487 Tubularity was developed as a descriptive, data-driven characterization of population-level temporal
 488 organization. Rather than constituting an optimality principle for model design, it highlighted a salient
 489 structural property empirically present in biological recordings that was absent in early FNN layers.
 490

491 **Future architecture improvements:** Our findings suggest several actionable insights for bringing
 492 foundation models, such as the FNN, into closer alignment with biological systems. (1) Coupling
 493 feature extraction with temporal dynamics: In biological systems, feature extraction and the develop-
 494 ment of temporal response dynamics occur simultaneously. Enforcing temporal dynamics in the early
 495 layers could enable more adequate modeling of the rich retinal dynamics. The FNN uses two tempo-
 496 rally aware mechanisms in the recurrent module: attention and recurrence. We argue that recurrence,
 497 rather than attention, is the critical mechanism, as the FNN without attention yielded equal or better
 498 performance (Wang et al., 2025). Although our analysis was limited to the published attention-based
 499 version, we propose introducing early-stage recurrence that mimics amacrine cell connectivity in the
 500 retina (Marc et al., 2014). (2) Addressing padding-related artifacts: While padding is not an issue in
 501 biological systems, the resulting intensity artifacts can distort model representations. These artifacts
 502 are well known in convolutional architectures (Alsallakh et al., 2020), and could be addressed through
 503 alternative padding strategies, or tailored regularization, thereby freeing model capacity rather than
 504 requiring the readout to “unlearn” non-biological features. (3) Revising the readout stage: The current
 505 Gaussian readout layer (Lurz et al., 2021) combines a large number of feature maps through a single
 506 linear combination step, producing unrealistically distinct feature representations. Enforcing mixed
 507 features while reducing their number to better reflect biological cell type diversity (Bae et al., 2025)
 508 could push the representation towards smoother and more biologically realistic manifolds.
 509

510 **Limitations:** Our analysis used a single foundation model, due to the limited availability of other
 511 video-based foundation models of neural activity over time. Moreover, we worked with a restricted
 512 set of stimuli (see Methods) to ensure comparability with biological data. However, there is evidence
 513 that these stimuli exercise much of the mouse visual cortex Dyballa et al. (2018), so they provide at
 514 least a necessary component for out-of-sample examination. Moreover, we show that these stimuli
 515 elicit activity patterns in the FNN similar to those evoked by the natural movies on which they
 516 were trained (Appendix Fig. 9), supporting their empirical validity. Finally, the tubularity metrics
 517 introduced here represent a novel approach for quantifying the geometry of neural trajectories. As no
 518 established methodological standards currently exist, further investigation of these metrics would
 519 be valuable. Specifically, systematic investigations of on both biological data and synthetic datasets
 520 would help assess robustness and for obtaining clear baselines. Additionally, incorporating curvature
 521 information could extend the metrics to capture additional characteristics of neural trajectories.
 522

5 CONCLUSION

523 We found a rich diversity of encoding and decoding topologies in the FNN, highlighting its capability
 524 to fit complex neural data. Distinct representation patterns emerged across modules, reflecting its archi-
 525 tecture. First, the *recurrent* module appears to learn generalizable representations of temporal stimuli,
 526 promoting uniformity and alignment, as in general self-supervised foundation models (Wang and
 527 Isola, 2022). Second, the *readout* module accounts for rich biological variability, but does so through
 528 a large number of self-similar feature maps, differing from the heterogeneous organization known in
 529 V1. Finally, the output layer achieves a continuous representation by linearly combining the readout
 530 features, ultimately enabling the network to associate spike trains with input movies *a posteriori*.
 531

532 Using our novel tubularity metrics, we found that biological data exhibited strong stimulus-dependent
 533 structure in both retina and V1, whereas the FNN encoder trajectories lacked such tubularity. Only
 534 from the recurrent module onward did the FNN begin to form bundles of activity, reaching higher-
 535 though still sub-biological-levels of representational cohesion. This emphasizes the role of recurrence
 536 in generating biologically plausible temporal representations, suggesting that models may benefit from
 537 placing recurrence after a more light-weight, local encoder (e.g., emulating the amacrine connectivity
 538 in the retina (Marc et al., 2014)) and that constrain feature dimensionality to reflect biological cell-type
 539 diversity (Bae et al., 2025). While biological fidelity is not a prerequisite for achieving high predictive
 540 accuracy, digital-twin use cases require enough internal alignment to support mechanistic and
 541 interventional inference. Such designs could help bridge the gap between computational performance
 542 and biological plausibility, moving toward truly brain-aligned foundation models.
 543

540 6 ETHICS STATEMENT
541542 There are no ethical concerns for this paper.
543544 7 REPRODUCIBILITY STATEMENT
545546 We provide an overview of our methods in the main text (Section 2) and include further details
547 for reproducing our results in the Appendix A. Upon acceptance, we will make the code for all
548 experiments and figures available on GitHub.
549550 551 REFERENCES
552553 Bilal Alsallakh, Narine Kokhlikyan, Vivek Miglani, Jun Yuan, and Orion Reblitz-Richardson. Mind
554 the Pad – CNNs can Develop Blind Spots, October 2020. URL <http://arxiv.org/abs/2010.02178>. arXiv:2010.02178 [cs].
555556 Anonymous. Only brains align with brains: Cross-region patterns expose limits of normative models.
557 In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025. URL
558 <https://openreview.net/forum?id=cMGJcHHI7d>. under review.
559560 Karl J Åström. *Introduction to stochastic control theory*. Courier Corporation, 2012.
561562 Bruno B. Averbeck, Peter E. Latham, and Alexandre Pouget. Neural correlations, population coding
563 and computation. *Nature Reviews Neuroscience*, 7(5):358–366, May 2006. ISSN 1471-003X, 1471-
564 0048. doi: 10.1038/nrn1888. URL <https://www.nature.com/articles/nrn1888>.
565566 Mehdi Azabou, Vinam Arora, Venkataramana Ganesh, Ximeng Mao, et al. A unified, scalable
567 framework for neural population decoding. In *Thirty-seventh Conference on Neural Information
568 Processing Systems*, 2023. URL <https://openreview.net/forum?id=sw2Y0sirtM>.
569570 Tom Baden, Philipp Berens, Katrin Franke, Miroslav Román Rosón, Matthias Bethge, and Thomas
571 Euler. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586):345,
572 2016.
573574 Brett W. Bader, Tamara G. Kolda, et al. Tensor toolbox for matlab, version 3.6. <https://www.tensortoolbox.org>, 2023.
575576 J. Alexander Bae, Mahaly Baptiste, Maya R. Baptiste, Caitlyn A. Bishop, Agnes L. Bodor, et al.
577 Functional connectomics spanning multiple areas of mouse visual cortex. *Nature*, 640(8058):
578 435–447, April 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-08790-w. URL <https://doi.org/10.1038/s41586-025-08790-w>.
579580 Shahab Bakhtiari, Patrick Mineault, Timothy Lillicrap, Christopher Pack, and Blake Richards.
581 The functional specialization of visual cortex emerges from training parallel pathways with
582 self-supervised predictive learning. *Advances in Neural Information Processing Systems*, 34:
583 25164–25178, 2021.
584585 Kosio Beshkov and Paul Tiesinga. Geodesic-based distance reveals nonlinear topological features in
586 neural activity from mouse visual cortex. *Biological Cybernetics*, 116(1):53–68, 2022.
587588 Kosio Beshkov, Marianne Fyhn, Torkel Hafting, and Gaute T Einevoll. Topological structure of
589 population activity in mouse visual cortex encodes densely sampled stimulus rotations. *Iscience*,
27(4), 2024.
590591 Jeffrey S. Bowers, Gaurav Malhotra, Marin Dujmović, Milton Llera Montero, et al. Deep
592 problems with neural network models of human vision. *Behavioral and Brain Sciences*, 46:
593 e385, 2023. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X22002813. URL https://www.cambridge.org/core/product/identifier/S0140525X22002813/type/journal_article.
594

594 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, et al. Towards monosemanticity:
 595 Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023.
 596 URL <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
 597

598 Nazmi Burak Budanur. Prediction and control of spatiotemporal chaos by learning conjugate tubular
 599 neighborhoods, 2023. URL <https://arxiv.org/abs/2310.12178>.
 600

601 Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. Density-based clustering based on
 602 hierarchical density estimates. In Jian Pei, Vincent S. Tseng, Longbing Cao, Hiroshi Motoda,
 603 and Guandong Xu, editors, *Advances in Knowledge Discovery and Data Mining*, pages 160–172,
 604 Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-37456-2.
 605

606 SueYeon Chung and L. F. Abbott. Neural population geometry: An approach for understanding
 607 biological and artificial neural networks. *Current Opinion in Neurobiology*, 70:137–144, October
 608 2021. ISSN 09594388. doi: 10.1016/j.conb.2021.10.010. URL <http://arxiv.org/abs/2104.07059> [q-bio].
 609

610 Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian,
 611 Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487
 612 (7405):51–56, 2012.
 613

614 Taco Cohen and Max Welling. Group equivariant convolutional networks. In Maria Florina Balcan
 615 and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine
 616 Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999, New York,
 617 New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/cohenc16.html>.
 618

619 R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Ge-
 620ometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion
 621 maps. *Proceedings of the National Academy of Sciences*, 102(21):7426–7431, May 2005. ISSN
 622 0027-8424, 1091-6490. doi: 10.1073/pnas.0500334102. URL <https://pnas.org/doi/full/10.1073/pnas.0500334102>. Publisher: Proceedings of the National Academy of
 623 Sciences.
 624

625 Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic
 626 Analysis*, 21(1):5–30, July 2006. ISSN 1063-5203. doi: 10.1016/j.acha.2006.04.006. URL <https://linkinghub.elsevier.com/retrieve/pii/S1063520306000546>. Publisher:
 627 Elsevier BV.
 628

630 Colin Conwell, David Mayo, Michael A. Buice, Boris Katz, George A. Alvarez, and Andrei Barbu.
 631 Neural Regression, Representational Similarity, Model Zoology & Neural Taskonomy at Scale in
 632 Rodent Visual Cortex, June 2021. URL <http://biorxiv.org/lookup/doi/10.1101/2021.06.18.448431>.
 633

635 Benjamin R. Cowley, Patricia L. Stan, Jonathan W. Pillow, and Matthew A. Smith. Compact
 636 deep neural network models of visual cortex. *bioRxiv: The Preprint Server for Biology*, page
 637 2023.11.22.568315, November 2023. ISSN 2692-8205. doi: 10.1101/2023.11.22.568315.
 638

639 John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings.
 640 *Nature neuroscience*, 17(11):1500–1509, 2014.

641 Ana Cannas Da Silva. *Lectures on Symplectic Geometry*, volume 1764 of *Lecture Notes in Mathe-
 642 matics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISBN 978-3-540-42195-5 978-3-
 643 540-45330-7. doi: 10.1007/978-3-540-45330-7. URL <http://link.springer.com/10.1007/978-3-540-45330-7>.
 644

646 Adrien Doerig, Rowan P Sommers, Katja Seeliger, Blake Richards, Jenann Ismael, Grace W Lindsay,
 647 Konrad P Kording, Talia Konkle, Marcel AJ Van Gerven, Nikolaus Kriegeskorte, et al. The
 648 neuroconnectionist research programme. *Nature Reviews Neuroscience*, 24(7):431–450, 2023.

648 Fengtong Du, Miguel Angel Núñez-Ochoa, Marius Pachitariu, and Carsen Stringer. A simplified
 649 minimodel of visual cortical neurons. *Nature Communications*, 16(1):5724, July 2025.
 650 ISSN 2041-1723. doi: 10.1038/s41467-025-61171-9. URL <https://doi.org/10.1038/s41467-025-61171-9>.

652 Marin Dujmovic, Jeffrey S Bowers, Federico Adolfi, and Gaurav Malhotra. INFERRING DNN-
 653 BRAIN ALIGNMENT USING REPRESENTATIONAL SIMILARITY ANALYSES. 2024.

655 Lea Duncker and Maneesh Sahani. Dynamics on the manifold: Identifying computational dynamical
 656 activity from neural population recordings. *Current Opinion in Neurobiology*, 70:163–170, October
 657 2021. ISSN 0959-4388. doi: 10.1016/j.conb.2021.10.014. URL <https://linkinghub.elsevier.com/retrieve/pii/S0959438821001264>. Publisher: Elsevier BV.

659 Luciano Dyballa and Steven W. Zucker. IAN: Iterated Adaptive Neighborhoods for manifold learning
 660 and dimensionality estimation. *Neural Computation*, 35(3):453–524, February 2023. ISSN
 661 0899-7667, 1530-888X. doi: 10.1162/neco_a_01566. URL <http://arxiv.org/abs/2208.09123> [cs].

664 Luciano Dyballa, Mahmood S Hoseini, Maria C Dadarlat, Steven W Zucker, and Michael P Stryker.
 665 Flow stimuli reveal ecologically appropriate responses in mouse visual cortex. *Proc Natl Acad Sci
 666 USA*, 115(44):11304–11309, 2018.

667 Luciano Dyballa, Andra M. Rudzite, Mahmood S. Hoseini, Mishek Thapa, Michael P. Stryker,
 668 Greg D. Field, and Steven W. Zucker. Population encoding of stimulus features along the visual
 669 hierarchy. *Proceedings of the National Academy of Sciences*, 121(4), January 2024a. ISSN
 670 0027-8424, 1091-6490. doi: 10.1073/pnas.2317773121. URL <https://pnas.org/doi/10.1073/pnas.2317773121>. Publisher: Proceedings of the National Academy of Sciences.

673 Luciano Dyballa, Greg D Field, Michael P Stryker, and Steven W Zucker. Functional organization
 674 and natural scene responses across mouse visual cortical areas revealed with encoding manifolds.
 675 *bioRxiv*, 2024b.

677 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, et al. A mathematical frame-
 678 work for transformer circuits. *Transformer Circuits Thread*, 2021. URL <https://transformer-circuits.pub/2021/framework/index.html>.

680 Emmanouil Froudarakis, Uri Cohen, Maria Diamantaki, Edgar Y Walker, Jacob Reimer, Philipp
 681 Berens, Haim Sompolinsky, and Andreas S Tolias. Object manifold geometry across the mouse
 682 cortical visual hierarchy. *BioRxiv*, pages 2020–08, 2020.

684 Lindsey L Glickfeld and Shawn R Olsen. Higher-order areas of the mouse visual cortex. *Annual
 685 review of vision science*, 3:251–273, 2017.

687 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, et al.
 688 Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.

691 Zhonghao He, Jascha Achterberg, Katie Collins, Kevin Nejad, et al. Multilevel Interpretability Of
 692 Artificial Neural Networks: Leveraging Framework And Methods From Neuroscience, August
 693 2024. URL <http://arxiv.org/abs/2408.12664>. arXiv:2408.12664 [cs].

694 Judith Hoeller, Lin Zhong, Marius Pachitariu, and Sandro Romani. Bridging tuning and invariance
 695 with equivariant neuronal representations. *bioRxiv*, pages 2024–08, 2024.

697 John J Hopfield. Neurons with graded response have collective computational properties like those of
 698 two-state neurons. *Proceedings of the national academy of sciences*, 81(10):3088–3092, 1984.

699 Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected
 700 convolutional networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern
 701 Recognition*, pages 4700–4708, 2017.

702 Liwei Huang, Zhengyu Ma, Liutao Yu, Huihui Zhou, and Yonghong Tian. Deep Spiking Neural
 703 Networks with High Representation Similarity Model Visual Pathways of Macaque and Mouse.
 704 *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1):31–39, June 2023. ISSN
 705 2374-3468, 2159-5399. doi: 10.1609/aaai.v37i1.25073. URL <https://ojs.aaai.org/index.php/AAAI/article/view/25073>.

707 J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):
 708 90–95, 2007. doi: 10.1109/MCSE.2007.55.

710 Plotly Technologies Inc. Collaborative data science. <https://plot.ly>, 2015.

712 Kohitij Kar, Simon Kornblith, and Evelina Fedorenko. Interpretability of artificial neural network
 713 models in artificial Intelligence vs. neuroscience. *Nature Machine Intelligence*, 4(12):1065–
 714 1067, December 2022. ISSN 2522-5839. doi: 10.1038/s42256-022-00592-3. URL [http://arxiv.org/abs/2206.03951](https://arxiv.org/abs/2206.03951) [q-bio].

716 T. Anderson Keller, Lyle Muller, Terrence Sejnowski, and Max Welling. Traveling Waves Encode the
 717 Recent Past and Enhance Sequence Learning, March 2024. URL [http://arxiv.org/abs/2309.08045](https://arxiv.org/abs/2309.08045). arXiv:2309.08045 [cs].

720 David A. Klindt, Alexander S. Ecker, Thomas Euler, and Matthias Bethge. Neural system identifica-
 721 tion for large populations separating "what" and "where", 2018. URL <https://arxiv.org/abs/1711.02653>.

723 Ho Ko, Sonja B Hofer, Bruno Pichler, Katherine A Buchanan, P Jesper Sjöström, and Thomas D
 724 Mrsic-Flogel. Functional specificity of local synaptic connections in neocortical networks. *Nature*,
 725 473(7345):87–91, 2011.

727 Viktoriya Krakovna and Finale Doshi-Velez. Increasing the Interpretability of Recurrent Neural
 728 Networks Using Hidden Markov Models, September 2016. URL [http://arxiv.org/abs/1606.05320](https://arxiv.org/abs/1606.05320). arXiv:1606.05320 [stat].

730 Nikolaus Kriegeskorte. Deep neural networks: a new framework for modeling biological vision and
 731 brain information processing. *Annu Rev Vis Sci*, 1:417–446, 2015.

733 Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis-
 734 connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:249,
 735 2008.

736 Nicholas Krämer et al. Tueplots, 2024. URL <https://tueplots.readthedocs.io/en/latest/index.html>.

739 Andrew Kyle Lampinen, Stephanie C. Y. Chan, Yuxuan Li, and Katherine Hermann. Representation
 740 biases: will we achieve complete understanding by analyzing representations?, August 2025. URL
 741 [http://arxiv.org/abs/2507.22216](https://arxiv.org/abs/2507.22216). arXiv:2507.22216 [q-bio].

742 Bryan M. Li, Isabel M. Cornacchia, Nathalie L. Rochefort, and Arno Onken. V1T: large-scale mouse
 743 V1 response prediction using a Vision Transformer, September 2023. URL [http://arxiv.org/abs/2302.03023](https://arxiv.org/abs/2302.03023). arXiv:2302.03023 [cs].

746 Dario Liscai, Emanuele Luconi, Alessandro Marin Vargas, and Alessandro Sanzeni. Beyond single
 747 neurons: population response geometry in digital twins of mouse visual cortex. *Journal of
 748 Statistical Mechanics: Theory and Experiment*, 2025(9):094003, 2025.

749 Konstantin-Klemens Lurz, Mohammad Bashiri, Konstantin Willeke, Akshay Jagadish, Eric Wang,
 750 Edgar Y. Walker, Santiago A Cadena, Taliah Muhammad, Erick Cobos, Andreas S. Tolias,
 751 Alexander S Ecker, and Fabian H. Sinz. Generalization in data-driven models of primary vi-
 752 sual cortex. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Tp7kI90Htd>.

755 Robert E Marc, James R Anderson, Bryan W Jones, Crystal L Sigulinsky, and James S Lauritzen.
 The aii amacrine cell connectome: a dense network hub. *Frontiers in neural circuits*, 8:104, 2014.

756 Eshed Margalit, Hyodong Lee, Dawn Finzi, James J. DiCarlo, Kalanit Grill-Spector, and Daniel L.K.
 757 Yamins. A unifying framework for functional organization in early and higher ventral vi-
 758 sual cortex. *Neuron*, 112(14):2435–2451.e7, July 2024. ISSN 0896-6273. doi: 10.1016/
 759 j.neuron.2024.04.018. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627324002794>. Publisher: Elsevier BV.
 760

761 Mackenzie Weygandt Mathis, Adriana Perez Rotondo, Edward F Chang, Andreas S Tolias, and
 762 Alexander Mathis. Decoding the brain: From neural representations to mechanistic models. *Cell*,
 763 187(21):5814–5832, 2024.
 764

765 Douglas A. Miller and Steven W. Zucker. Computing with Self-Excitatory Cliques: A Model and
 766 an Application to Hyperacuity-Scale Computation in Visual Cortex. *Neural Computation*, 11(1):
 767 21–66, January 1999. ISSN 0899-7667, 1530-888X. doi: 10.1162/089976699300016782. URL
 768 <https://direct.mit.edu/neco/article/11/1/21-66/6227>.

769 Peter M. Milner. A model for visual shape recognition. *Psychological Review*, 81(6):521–535, 1974.
 770 ISSN 1939-1471, 0033-295X. doi: 10.1037/h0037149. URL <https://doi.apa.org/doi/10.1037/h0037149>.

772 Patrick Mineault, Niccolò Zanichelli, Joanne Zichen Peng, Anton Arkhipov, et al. NeuroAI for AI
 773 Safety, April 2025. URL <http://arxiv.org/abs/2411.18526>. arXiv:2411.18526 [cs].
 774

775 Aran Nayebi, Nathan C. L. Kong, Chengxu Zhuang, Justin L. Gardner, Anthony M. Norcia, and
 776 Daniel L. K. Yamins. Mouse visual cortex as a limited resource system that self-learns an
 777 ecologically-general representation. *PLOS Computational Biology*, 19(10):e1011506, October
 778 2023. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1011506. URL <https://dx.plos.org/10.1371/journal.pcbi.1011506>.

779

780 Nina S. Nellen, Polina Turishcheva, Michaela Vystrčilová, Shashwat Sridhar, Tim Gollisch, Andreas S.
 781 Tolias, and Alexander S. Ecker. Learning to cluster neuronal function, 2025. URL <https://arxiv.org/abs/2506.03293>.
 782

783

784 Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*,
 785 2017. doi: 10.23915/distill.00007. URL <https://distill.pub/2017/feature-visualization/>.

786

787 Mitchell Ostrow, Adam Eisen, Leo Kozachkov, and Ila Fiete. Beyond geometry: Comparing the
 788 temporal structure of computation in neural circuits with dynamical similarity analysis, 2023. URL
 789 <https://arxiv.org/abs/2306.10168>.

790

791 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, et al. Pytorch: An
 792 imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703, 2019. URL
 793 <http://arxiv.org/abs/1912.01703>.

794

795 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, et al. Scikit-learn: Machine learning
 796 in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

797

798 Walter Pitts and Warren S. McCulloch. How we know universals the perception of auditory and
 799 visual forms. *The Bulletin of Mathematical Biophysics*, 9(3):127–147, September 1947. ISSN
 800 0007-4985, 1522-9602. doi: 10.1007/BF02478291. URL <http://link.springer.com/10.1007/BF02478291>.

801

802 Ahmed Qazi, Hamd Jalil, and Asim Iqbal. Mice to Machines: Neural Representations from Visual Cor-
 803 tex for Domain Generalization, May 2025. URL <http://arxiv.org/abs/2505.06886>.
 arXiv:2505.06886 [cs].

804

805 Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector
 806 canonical correlation analysis for deep learning dynamics and interpretability, 2017. URL <https://arxiv.org/abs/1706.05806>.

807

808 Dario L Ringach, Patrick J Mineault, Elaine Tring, Nicholas D Olivas, Pablo Garcia-Junco-Clemente,
 809 and Joshua T Trachtenberg. Spatial clustering of tuning in mouse primary visual cortex. *Nat Commun*, 7, 2016.

810 Avery Hee-Woon Ryoo, Nanda H Krishna, Ximeng Mao, Mehdi Azabou, Eva L Dyer, Matthew G
 811 Perich, and Guillaume Lajoie. Generalizable, real-time neural decoding with hybrid state-space
 812 models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
 813 URL <https://openreview.net/forum?id=1i4wNFgHDd>.

814 Mostafa Safaie, Joanna C. Chang, Junchol Park, Lee E. Miller, Joshua T. Dudman, Matthew G
 815 Perich, and Juan A. Gallego. Preserved neural dynamics across animals performing similar
 816 behaviour. *Nature*, 623(7988):765–771, 2023. doi: 10.1038/s41586-023-06714-0. URL <https://doi.org/10.1038/s41586-023-06714-0>.

817 Rylan Schaeffer, Mikail Khona, Sarthak Chandra, Mitchell Ostrow, Brando Miranda, and Sanmi
 818 Koyejo. Position: Maximizing Neural Regression Scores May Not Identify Good Models of the
 819 Brain. 2025.

820 Thomas Serre. Deep learning: the good, the bad, and the ugly. *Annual review of vision science*, 5(1):
 821 399–426, 2019a.

822 Thomas Serre. Deep learning: the good, the bad, and the ugly. *Annual review of vision science*, 5(1):
 823 399–426, 2019b. Publisher: Annual Reviews.

824 Jianghong Shi, Bryan Tripp, Eric Shea-Brown, Stefan Mihalas, and Michael A. Buice. MouseNet: A
 825 biologically constrained convolutional neural network model for the mouse visual cortex. *PLOS
 826 Computational Biology*, 18(9):e1010427, September 2022. ISSN 1553-7358. doi: 10.1371/journal.
 827 pcbi.1010427. URL <https://dx.plos.org/10.1371/journal.pcbi.1010427>.

828 Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep Inside Convolutional Networks:
 829 Visualising Image Classification Models and Saliency Maps, April 2014. URL <http://arxiv.org/abs/1312.6034> [cs].

830 Oscar Skean, Md Rifat Arefin, Dan Zhao, Niket Patel, Jalal Naghiyev, Yann LeCun, and Ravid
 831 Shwartz-Ziv. Layer by Layer: Uncovering Hidden Representations in Language Models, June
 832 2025. URL <http://arxiv.org/abs/2502.02013>. arXiv:2502.02013 [cs].

833 Bahareh Tolooshams, Sara Matias, Hao Wu, Simona Temereanca, Naoshige Uchida, Venkatesh N.
 834 Murthy, Paul Masset, and Demba Ba. Interpretable deep learning for deconvolutional analysis
 835 of neural signals. *Neuron*, 113(8):1151–1168.e13, April 2025. ISSN 08966273. doi: 10.1016/j.neuron.2025.02.006. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627325001199>.

836 Polina Turishcheva, Max Burg, Fabian H. Sinz, and Alexander Ecker. Reproducibility of predictive
 837 networks for mouse visual cortex, 2024. URL <https://arxiv.org/abs/2406.12625>.

838 Ivan Ustyuzhaninov, Max F. Burg, Santiago A. Cadena, Jiakun Fu, et al. Digital twin reveals
 839 combinatorial code of non-linear computations in the mouse primary visual cortex, February 2022.
 840 URL <https://biorxiv.org/lookup/doi/10.1101/2022.02.10.479884>.

841 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, et al. SciPy 1.0: Fundamental
 842 Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

843 Eric Y. Wang, Paul G. Fahey, Zhuokun Ding, Stelios Papadopoulos, Kayla Ponder, et al. Foundation
 844 model of neural activity predicts response to new stimulus types. *Nature*, 640(8058):470–477,
 845 April 2025. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-025-08829-y. URL <https://www.nature.com/articles/s41586-025-08829-y>. Publisher: Springer Science
 846 and Business Media LLC.

847 Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through align-
 848 ment and uniformity on the hypersphere, 2022. URL <https://arxiv.org/abs/2005.10242>.

849 Alex H. Williams, Tony Hyun Kim, Forea Wang, Saurabh Vyas, Stephen I. Ryu, et al. Unsupervised
 850 Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through
 851 Tensor Component Analysis. *Neuron*, 98(6):1099–1115.e8, June 2018. ISSN 0896-6273. doi: 10.
 852 1016/j.neuron.2018.05.015. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627318303878>. Publisher: Elsevier BV.

864 Aiwen Xu, Yuchen Hou, Christopher Niell, and Michael Beyeler. Multimodal deep learning model
865 unveils behavioral dynamics of v1 activity in freely moving mice. *Advances in neural information*
866 *processing systems*, 36:15341–15357, 2023.

867 Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J
868 DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual
869 cortex. *Proc Natl Acad Sci USA*, 111(23):8619–8624, 2014.

870 Joel Ye, Jennifer L Collinger, Leila Wehbe, and Robert Gaunt. Neural data transformer 2: Multi-
871 context pretraining for neural spiking activity. In *Thirty-seventh Conference on Neural Information*
872 *Processing Systems*, 2023. URL <https://openreview.net/forum?id=CBBtMn1TGq>.

873 Joel Ye, Fabio Rizzoglio, Xuan Ma, Adam Smoulder, et al. A generalist intracortical motor decoder.
874 In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL
875 <https://openreview.net/forum?id=utXSSdD9mt>.

876 Yiyi Yu, Jeffrey N Stirman, Christopher R Dorsett, and Spencer L Smith. Selective representations of
877 texture and motion in mouse higher visual areas. *Current Biology*, 32(13):2810–2820, 2022.

878 Yizi Zhang, Yanchen Wang, Mehdi Azabou, Alexandre Andre, Zixuan Wang, Hanrui Lyu, Interna-
879 tional Brain Laboratory, Eva L Dyer, Liam Paninski, and Cole Lincoln Hurwitz. Neural encoding
880 and decoding at scale. In *Forty-second International Conference on Machine Learning*, 2025.
881 URL <https://openreview.net/forum?id=v0dz3zhSCj>.

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

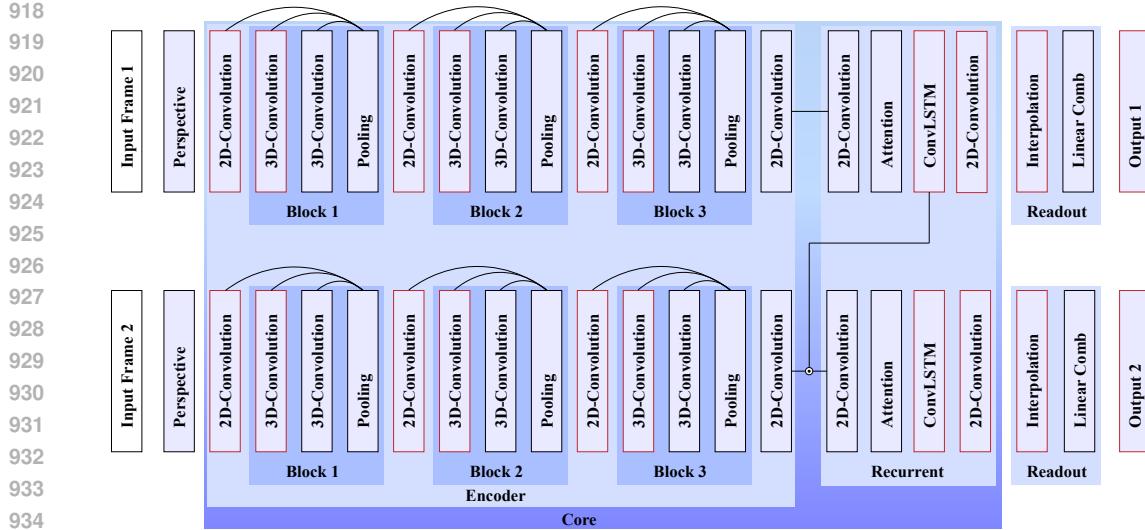


Figure 6: **FNN architecture.** Layers used for sampling are highlighted. Modulation module omitted as it has no effect for our analysis. The FNN used GeLU activations in the convolutional layers, and Tanh activations in the Recurrent module.

Appendix

A METHODS

A.1 ONLINE MATERIAL

Our work made use of publicly available open-source resources. Specifically, we employed the pretrained FNN model provided by Wang et al. (2025), available at <https://github.com/cajal/fnn/tree/main>. For the analysis of this model, we used the stimulus generation tools and neural encoding manifold construction pipeline introduced by Dyballa et al. (2024a), accessible at <https://github.com/dyballa/NeuralEncodingManifolds>.

A.2 FNN

The FNN consists of five modules: perspective, modulation, *encoder*, *recurrent*, and *readout* (Fig. 6). The perspective and modulation modules model the mouse’s state and transform the inputs to approximate the actual visual information received. Thus, only the *encoder*, *recurrent*, and *readout* modules perform the core computation and are the focus of this work.

The *encoder* module is a 10-layer DenseNet-style convolutional encoder. Notably, it includes 3D convolutions, which in principle enable the encoder to capture temporal patterns, for up to 12 time steps into the past for later encoder layers. The *recurrent* module is optionally preceded by an attention layer and consists of a convolutional LSTM, followed by a single convolutional layer that produces its output. This feedforward–recurrent combination constitutes the core of the FNN, which is trained on all data. Finally, the *readout* module is mouse-specific: it performs an interpolation on the recurrent output followed by a linear transformation to produce the FNN output. We used the FNN readout from session 8, scan 5 as this was the exemplary scan used in the authors’ tutorial. We validated findings on several other sessions and scans.

A.3 INPUT VIDEOS

We used the visual stimuli from Dyballa et al. (2024a), consisting of drifting square-wave gratings and optical flows moving in eight directions. The flow stimuli include oriented (lines) and non-oriented (dots) stimuli with spatial frequencies between 0.04 and 0.5 $\frac{\text{cycles}}{\text{deg}}$. This yields 88 unique input

972 sequences with stochastic initial positions and velocities. The stimuli were scaled and cropped to fit
 973 the required FNN input shape of 144×256 pixels. This resulted in an image sequence: $\{\mathbf{x}_0, \dots, \mathbf{x}_T\}$,
 974 where each $\mathbf{x}_i \in \mathbb{R}^{H \times W}$. Stimuli were generated using the tools available at <https://github.com/dyballa/NeuralEncodingManifolds>.
 975

976 The FNN (Wang et al., 2025) processes 2.33-second sequences of 70 frames each, corresponding
 977 to 30 frames per second. Since in Dyballa et al. (2024a) the trials were 1.25 s long, we adapted
 978 the stimuli to contain 37 frames to maintain consistency with the FNN framework. This adaptation
 979 was performed using the hyperparameters of the stimulus generation pipeline, allowing comparable
 980 stimuli dynamically created for different lengths and number of frames.
 981

982 We acknowledge a difference in the experimental setups regarding the visual field: Wang et al.
 983 (2025) used a screen distance of 15 cm, whereas the stimuli from Dyballa et al. (2024a) were
 984 originally designed for a 25 cm viewing distance. This discrepancy potentially affects the visual field
 985 transformations performed by the model’s perspective module, as the visual angle subtended by the
 986 stimuli differs between the two configurations. We applied a global scaling factor of 0.7 to all stimuli
 987 to address this. This adjustment was empirically found to optimize stimulus discriminability across
 988 network layers, effectively bridging the geometric gap between the training and analysis domains.
 989

990 A.4 DATA SAMPLING

991 Neural responses were computed using PyTorch and extracted by sampling activations from 2000
 992 units across selected FNN layers. Within each layer, 40 feature maps were sampled. Then, 50 neurons
 993 were sampled from each feature map. Feature map sampling probabilities were calculated from the
 994 mean maximum response across all neurons within each map, while neuron sampling probabilities
 995 within each selected feature map were based on individual neuron maximum responses, biasing the
 996 sampling to include active neurons. This sampling procedure was chosen to ensure comparability to
 997 the biological results from Dyballa et al. (2024a). This sampling procedure was tested and validated in
 998 Dyballa et al. (2024a); we performed further tests with random sampling to validate this bias does not
 999 filter out relevant structures. One exemplary sampling result, showing qualitative stability of results
 1000 across sampling strategies and sizes can be found in Fig 18. Increasing the sampling rate beyond
 1001 2000 units did not significantly alter manifold topology but hindered cluster separation in diffusion
 1002 map analysis. The resulting tensor data had dimensions $(N \times S \times O \times T)$ with $N = 2000$ neurons,
 1003 $S = 11$ stimulus types, $O = 8$ orientations and $T = 37$ time steps. For manifold construction,
 1004 the optimal spatial frequency was selected (resulting in $S = 6$ stimuli) whereas for classification
 1005 performance all spatial frequencies were kept. We report results from a single random seed per
 1006 layer, as preliminary analysis showed consistent manifold structure across different random activity
 1007 samples. These neural activation tensors served as input for subsequent classification and manifold
 1008 analysis. This sampling procedure was developed by Dyballa et al. (2024a) and tested against other
 1009 sampling methods there. We also experimented with the sampling procedure, finding that random
 1010 sampling and increased sampling rate did not introduce qualitative changes to the manifolds.
 1011

1012 A.5 STIMULUS ADEQUACY

1013 For every FNN layer investigated in this paper, we extracted the activation to the stimulus
 1014 ensemble consisting of gratings and flows (see Section A.3) as well as to a 100-second-long natural input video from the MICrONS functional dataset (Bae et al., 2025),
 1015 downloaded from [s3://bossdb-open-data/iarpa_microns/minnie/functional_data/stimulus_movies/](https://bossdb-open-data.iarpa_microns/minnie/functional_data/stimulus_movies/). Both stimulus sets produced similar activation magnitudes across
 1016 the entire network (see Fig. 9), which shows the adequacy of the stimulus ensemble used for testing
 1017 the FNN.
 1018

1019 For orientation and direction selectivity, we followed Wang et al. (2025)’s procedure: We input
 1020 directional pink noise (16 directions, 37 frames) to the model and record the output activations. Addi-
 1021 tionally, we recorded the outputs for our flow stimuli. For both datasets, we computed the Orientation
 1022 Selectivity Index (OSI) and Direction Selectivity Index (DSI) and compared their distributions.
 1023

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036 A.6 CLASSIFICATION ACCURACY

1038 Classification performance of the stimulus set, measured using activations, serves as a proxy for
1039 representational richness. Logistic regression is employed to assess linear separability, while k-
1040 Nearest Neighbor (k-NN) classification is used to evaluate local geometric structure for comparison
1041 with logistic regression.1042 Stimulus classification accuracy based on individual-layer activities was determined by training
1043 multinomial logistic regression classifiers (solver: L-BFGS) with 5-fold cross-validation (CV). Only
1044 sampled neurons were used to classify the 11 stimuli. For each layer and each time point t, two
1045 feature sets were constructed: (i) the mean activity over frames 0 to t (increasing window) and (ii)
1046 the mean activity over frames t to end (decreasing window). For comparison, K-nearest neighbor
1047 classifiers (K=3) were also evaluated using leave-one-out CV. The value K=3 was selected as the
1048 optimal neighborhood size. Leave-one-out CV was used for k-NN due to its suitability for small
1049 datasets, while 5-fold CV was chosen for logistic regression to reduce computational requirements.
1050 Results are summarized in Table 1 and Fig. ??.

1051

1052 A.7 CONSTRUCTION OF DECODING MANIFOLDS

1053

1054 For building the *decoding manifolds*, we applied PCA (scikit-learn) to the averaged activity data. In to-
1055 tal, the *decoding manifolds* contain 48 points, consisting of 6 stimuli and 8 movement directions each.
1056 The 6 stimuli were obtained from a majority vote of all neurons on the optimal spatial frequency elic-
1057 iting higher responses. The decoding manifolds use different colors for each stimulus, as introduced in
1058 Fig. 1. Different spatial frequencies of the same stimulus are summarized with the same color. To con-
1059 struct *decoding trajectories*, we treated each time step as a separate data point rather than averaging
1060 across time before applying PCA. In both cases, we reduced the dimensionality to three components
1061 for visualization after verifying that further dimensions did not encode qualitatively new information.
1062 We constructed biological *decoding trajectories* using experimental data from Dyballa et al. (2024a),
1063 available at <https://github.com/dyballa/NeuralEncodingManifolds>. For the bio-
1064 logical decoding trajectories, we did not use the additional zero-activity time step since a baseline
1065 activity level was already provided by the inter-stimulus intervals in the experiments.

1066

1067 A.8 TUBULARITY

1068 Before calculating tubularity metrics, we standard-scale the data and apply PCA to obtain a 10-
1069 dimensional embedding, thereby speeding up the computation. While visualizations use only the
1070 first 2–3 dimensions, all metrics are calculated in the 10-dimensional space. To ensure comparability,
1071 we resampled all trajectories to length 100. For statistical analysis, we generated 100 bootstrapped
1072 samples, and using ground-truth clusters, performed Bonferroni-corrected Mann-Whitney U tests on
1073 our hypotheses.1074 We formalize how “tight” a group of curves is around the centerline: We reparameterize each curve
1075 by normalized arc length $u \in [0, 1]$ and resample to $\{u_k\}_{k=1}^M$. Let $x_i(u_k) \in \mathbb{R}^D$ denote the samples
1076 and $\tau_i(u_k)$ their unit tangents. We define the *mean curve* as the pointwise average:

1077

1078

1079

$$OSI = \frac{|\sum_{\theta} \bar{r}_{\theta} e^{i2\theta}|}{\sum_{\theta} \bar{r}_{\theta}} \quad (1)$$

$$DSI = \frac{|\sum_{\theta} \bar{r}_{\theta} e^{i\theta}|}{\sum_{\theta} \bar{r}_{\theta}} \quad (2)$$

Here, \bar{r}_{θ} is the mean response for angle θ . We found comparable OSI and DSI distributions (see Fig. 10).

A.6 CLASSIFICATION ACCURACY

Classification performance of the stimulus set, measured using activations, serves as a proxy for representational richness. Logistic regression is employed to assess linear separability, while k-Nearest Neighbor (k-NN) classification is used to evaluate local geometric structure for comparison with logistic regression.

Stimulus classification accuracy based on individual-layer activities was determined by training multinomial logistic regression classifiers (solver: L-BFGS) with 5-fold cross-validation (CV). Only sampled neurons were used to classify the 11 stimuli. For each layer and each time point t, two feature sets were constructed: (i) the mean activity over frames 0 to t (increasing window) and (ii) the mean activity over frames t to end (decreasing window). For comparison, K-nearest neighbor classifiers (K=3) were also evaluated using leave-one-out CV. The value K=3 was selected as the optimal neighborhood size. Leave-one-out CV was used for k-NN due to its suitability for small datasets, while 5-fold CV was chosen for logistic regression to reduce computational requirements. Results are summarized in Table 1 and Fig. ??.

A.7 CONSTRUCTION OF DECODING MANIFOLDS

For building the *decoding manifolds*, we applied PCA (scikit-learn) to the averaged activity data. In total, the *decoding manifolds* contain 48 points, consisting of 6 stimuli and 8 movement directions each. The 6 stimuli were obtained from a majority vote of all neurons on the optimal spatial frequency eliciting higher responses. The decoding manifolds use different colors for each stimulus, as introduced in Fig. 1. Different spatial frequencies of the same stimulus are summarized with the same color. To construct *decoding trajectories*, we treated each time step as a separate data point rather than averaging across time before applying PCA. In both cases, we reduced the dimensionality to three components for visualization after verifying that further dimensions did not encode qualitatively new information. We constructed biological *decoding trajectories* using experimental data from Dyballa et al. (2024a), available at <https://github.com/dyballa/NeuralEncodingManifolds>. For the biological decoding trajectories, we did not use the additional zero-activity time step since a baseline activity level was already provided by the inter-stimulus intervals in the experiments.

A.8 TUBULARITY

Before calculating tubularity metrics, we standard-scale the data and apply PCA to obtain a 10-dimensional embedding, thereby speeding up the computation. While visualizations use only the first 2–3 dimensions, all metrics are calculated in the 10-dimensional space. To ensure comparability, we resampled all trajectories to length 100. For statistical analysis, we generated 100 bootstrapped samples, and using ground-truth clusters, performed Bonferroni-corrected Mann-Whitney U tests on our hypotheses.

We formalize how “tight” a group of curves is around the centerline: We reparameterize each curve by normalized arc length $u \in [0, 1]$ and resample to $\{u_k\}_{k=1}^M$. Let $x_i(u_k) \in \mathbb{R}^D$ denote the samples and $\tau_i(u_k)$ their unit tangents. We define the *mean curve* as the pointwise average:

$$c(u_k) = \frac{1}{m} \sum_{i=1}^m x_i(u_k), \quad r_i(u_k) = \|x_i(u_k) - c(u_k)\|.$$

1080 The tightness score is calculated by averaging quantile tube radii across bins $\{I_b\}_{b=1}^B$ that partition
 1081 $[0, 1]$, using a high quantile $q \in [0.8, 0.95]$ to ensure robustness to noise. We normalize each tube's
 1082 tightness score by tube length.
 1083

$$1084 \quad S_{\text{tight}} = \frac{1}{B} \sum_{b=1}^B \text{quantile}_q \{ r_i(u) : u \in I_b \text{ over all curves} \}.$$

1088 The second quantity assessed is the uniformity of the tubes relative to one another. That is, the degree
 1089 to which crossings occur in our defined bundle of curves. Tubes are considered disorganized when
 1090 distinct curves pass near each other with *transverse* directions. Let $d_{ij}(u, v) = \|x_i(u) - x_j(v)\|$
 1091 and $\phi_{ij}(u, v) = 1 - \langle \tau_i(u), \tau_j(v) \rangle^2 \in [0, 1]$ (large for near-orthogonal tangents). Using a Gaussian
 1092 kernel $K_\varepsilon(\rho) = \exp(-\rho^2/(2\varepsilon^2))$, we softly count encounters:
 1093

$$1094 \quad \mathcal{X}_\varepsilon = \frac{2}{m(m-1)} \sum_{i < j} \int_0^1 \int_0^1 K_\varepsilon(d_{ij}(u, v)) \phi_{ij}(u, v) du dv.$$

1097 S_{tight} and S_{cross} only depend on distance, unit-tangent inner product, and arc-length. Therefore, they
 1098 are invariant to translations, rotations, and re-timing. We emphasize that, for both scores, smaller
 1099 values indicate more tubular curve bundles, while larger values indicate fewer tubular curve bundles.
 1100

1101 A.9 CONSTRUCTION OF NEURAL ENCODING MANIFOLDS

1103 At a high level, the motivation for constructing *neural encoding manifolds* is to find a space in which
 1104 one can examine the global topology of neuronal populations based on their stimulus selectivities
 1105 and temporal response patterns (Dyballa et al., 2024a). The neural encoding manifold is constructed
 1106 in a three-step procedure. First, a 3-tensor is built with the temporal responses from each neuron
 1107 for each stimulus, and decomposed using Nonnegative Tensor Factorization (details below); each
 1108 component is comprised of neural, stimulus, and temporal response factors. The neural factors then
 1109 serve as position coordinates, embedding the neurons into a stimulus-response framework called the
 1110 neural encoding space. Second, we construct a data graph in this neural encoding space using the
 1111 IAN algorithm (Dyballa and Zucker, 2023). Third, applying diffusion maps (Coifman et al., 2005;
 1112 Coifman and Lafon, 2006) to the data graph yields the manifold.

1113 The methodological choices in our manifold construction procedure are made in accordance with
 1114 Dyballa et al. (2024a), where extensive parameter analysis for biological neural data was conducted.
 1115 Since *neural encoding manifolds* computed with these specific parameters represent the only available
 1116 comparison for biological data from the visual system, we maintained their parameter settings to
 1117 ensure direct comparability between artificial and biological neural representations. We further
 1118 conducted analysis for FNN-specific parameters, such as the sampling procedure, by adapting their
 1119 code to fit the FNN requirements.

1120 A.9.1 PREPROCESSING

1122 The input tensor of neuronal activity (see above) was preprocessed in several steps (using NumPy
 1123 and SciPy). First, the individual responses were smoothed along the time dimension using a one-
 1124 dimensional Gaussian kernel with $\sigma = 3$. Next, we grouped the stimuli into *medium* versus *high*
 1125 spatial frequencies and selected the one exhibiting higher response magnitudes. The temporal
 1126 responses for the 8 directions of motion were then concatenated together into a single vector. Finally,
 1127 we normalized each response and rescaled it by the relative activations of the neuron. The resulting
 1128 tensor \mathbf{T} had shape $((N = 2000) \times (S = 6) \times (O * T = 296))$.
 1129

1130 A.9.2 NONNEGATIVE TENSOR FACTORIZATION

1131 Next, Nonnegative Tensor Factorization (see (Williams et al., 2018) for an overview and applications
 1132 to neuroscience) was applied to our tensor \mathbf{T} . It was decomposed into typically 10–15 rank-1
 1133 tensors which are obtained from the outer product of three vectors each. We selected the number
 of components separately for each data sample based on changes in explained variance and noise,

following the procedure in Dyballa et al. (2024a). The factors in each component are scaled to unit length, and their magnitudes absorbed by a scalar λ_r :

$$\tilde{\mathbf{T}} = \sum_{r=1}^R \lambda_r \mathbf{v}_r^{(1)} \circ \mathbf{v}_r^{(2)} \circ \mathbf{v}_r^{(3)} = [\lambda; \mathbf{X}^{(1)}; \mathbf{X}^{(2)}; \mathbf{X}^{(3)}] \quad (3)$$

For the second equality, the factor matrices $\mathbf{X}^{(k)}$ are constructed using the factor vectors $\mathbf{v}_r^{(k)}$ as columns, and the vector λ contains all individual λ_r s.

Decomposing the tensor \mathbf{T} into these components is an optimization problem with the following objective function and non-negativity constraints:

$$\min_{\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \mathbf{X}^{(3)}} \frac{1}{2} \|\mathbf{T} - \tilde{\mathbf{T}}\|^2 \quad (4)$$

$$\text{such that } \mathbf{X}^{(k)} \geq 0, \forall k \quad (5)$$

The resulting decomposition is interpretable: the third group of vectors, $\mathbf{v}_r^{(3)}$, describes different temporal response patterns; $\mathbf{v}_r^{(2)}$ contain information about which stimuli exhibit these response patterns; and $\mathbf{v}_r^{(1)}$ are the neuronal factors determining which neurons exhibit the response patterns characterized by $\mathbf{v}_r^{(2)}$ and $\mathbf{v}_r^{(3)}$. During decomposition, circular permutations were applied to detect patterns irrespective of the preferred orientations of specific neurons (again, this is necessary to ensure compatibility with the biological results from (Dyballa et al., 2024a)).

Using the OPT method from Tensor Toolbox (Bader et al., 2023)), we ran the decomposition 50 times (different initializations) for each number of components and dataset to ensure robust decomposition results and the choice of the number of factors, R . The manifolds were robust to small changes in R , therefore the heuristic for choosing R based on the explained variance of the decomposition outlined in Dyballa et al. (2024a) proved sufficient. For building the manifolds, we used the result with smallest reconstruction error among the 50 initializations.

A.9.3 NEURAL ENCODING SPACE

Following Dyballa et al. (2024a), we now reformulate the above decomposition to construct the neural encoding space. By defining the diagonal matrix Λ with $\Lambda_{rr} = \lambda_r$, we obtain:

$$\tilde{\mathbf{T}} = \mathbf{X}^{(1)} \Lambda (\mathbf{X}^{(2)} \circ \mathbf{X}^{(3)}) \quad (6)$$

Since the first matrix, $\mathbf{X}^{(1)}$, represents the neuronal factors, we denote it by \mathcal{N} . Now, define a matrix \mathbf{B} with columns $\mathbf{b}_{:,r}$:

$$\mathbf{b}_{:,r} = \text{vec}(\mathbf{v}_r^{(2)} \circ \mathbf{v}_r^{(3)}) \quad (7)$$

Finally, we obtain a matrix representation of \mathbf{T} with respect to neuronal factors as $\mathbf{X}_{\mathcal{N}}$:

$$\mathbf{X}_{\mathcal{N}} = \mathbf{B} \Lambda \mathcal{N}^T \quad (8)$$

This reformulation constructs the neural encoding space. The unit-norm basis vectors of this space are given by the columns of \mathbf{B} . We define the neural matrix containing the positions of all neurons in this space as $\mathcal{N}_{\lambda} = \mathcal{N} \Lambda$. The distances between any two neurons in this space reflect their similarity in stimulus-selective temporal response patterns. Intuitively, neurons with similar selectivity profiles and temporal dynamics should be positioned close together, while neurons with dissimilar response characteristics should be farther apart.

1188 A.9.4 ITERATED ADAPTIVE NEIGHBORHOODS (IAN)
1189

1190 Within this neural encoding space, we construct a weighted graph of the data by inferring a similarity
1191 kernel. This is achieved using the Iterated Adaptive Neighborhoods (IAN) algorithm (Dyballa and
1192 Zucker, 2023), which infers an adaptive local kernel without the need for pre-specifying a fixed
1193 neighborhood size.

1194 IAN first constructs the unweighted Gabriel graph for the data points. In addition, a weighted
1195 graph is constructed using a multiscale Gaussian kernel based on the discrete neighborhood graphs.
1196 Subsequently, the graph is iteratively pruned by ensuring consistency between the discrete and
1197 continuous neighborhoods. The resulting weighted graph is represented by the adjacency (kernel)
1198 matrix \mathbf{K} . This matrix contains similarities computed using locally tuned Gaussian kernels.

1199 A.9.5 DIFFUSION MAPS
1200

1201 Diffusion Maps (Coifman et al., 2005; Coifman and Lafon, 2006) are a dimensionality reduction
1202 technique that retain distances and preserve the intrinsic geometry of the manifold. The diffusion
1203 process is based on graph Laplacian normalization from spectral graph theory.

1204 In detail, we use the weighted graph obtained from IAN as the weighted adjacency matrix \mathbf{K} . The
1205 first step is to normalize and symmetrize it to produce \mathbf{M}_s :

$$1208 \mathbf{d}_i = \sqrt{\sum_j \mathbf{K}_{ij} + \epsilon} \quad (9)$$

$$1211 \mathbf{M}_s = \frac{\mathbf{K}}{\mathbf{d}\mathbf{d}^T} \quad (10)$$

1213 This normalization ensures that nodes of high degree do not dominate the analysis. We then calculate
1214 the spectral decomposition of \mathbf{M}_s with eigenvalues $\lambda_0 = 1 \geq \lambda_1 \geq \lambda_2 \dots$ and eigenvectors ψ_i for
1215 $t = 1$ diffusion steps using $L = 20$ eigenvalues:

$$1218 \mathbf{M}_{s,ij}^t = \sum_{l=0}^L \lambda_l^{2t} \psi_l(i) \psi_l(j) \quad (11)$$

1221 Finally, from the spectral decomposition, we obtain the diffusion map with diffusion coordinates:

$$1224 \Psi_t(i) = \begin{pmatrix} \lambda_0^t \psi_0(i) \\ \lambda_1^t \psi_1(i) \\ \vdots \\ \lambda_{L-1}^t \psi_{L-1}(i) \end{pmatrix} \quad (12)$$

1228 Plotting the data using these diffusion coordinates yields the *neural encoding manifold*.

1230 A.9.6 ENCODING MANIFOLD VISUALIZATION
1231

1232 For visualization purposes, we optionally applied metric multidimensional scaling (MDS) to the
1233 diffusion map coordinates. This was done by computing pairwise squared Euclidean distances
1234 using the first diffusion coordinates, constructing the corresponding Gram matrix $\mathbf{G} = -0.5 * \mathbf{D}^2$,
1235 and applying kernel PCA to obtain a lower-dimensional embedding. This preserves the distance
1236 relationships from the diffusion map while combining multiple diffusion coordinates, enabling a
1237 clearer visualization of the manifold structure.

1238 Based on the manifold topology, we selected groups of neurons to investigate via their Peristimulus
1239 Time Histograms (PSTH). We averaged their activity across trials and constructed the PSTHs as a
1240 2-D heatmap, where each row contains the temporal activity in response to a particular direction of
1241 motion (as displayed in Fig. 1). Additionally, we calculated the average response intensity over time
for these groups and reported the s.e.m. using the shaded regions (see insets in Fig. 2A,D).

1242 A.10 VISUALIZATIONS

1243

1244 Interactive three-dimensional plots of the manifolds were computed using Plotly. Other plots were
1245 created with Matplotlib and TUEplots.

1246

1247 A.11 MINIMODELS

1248

1249 For our additional analysis in Fig. 17, we used the convolutional model introduced in Du et al. (2025).
1250 We downloaded model checkpoints from [https://github.com/MouseLand/minimodel/](https://github.com/MouseLand/minimodel/tree/main)
1251 tree/main. We left the manifold pipeline unchanged for this experiment and sampled activations
1252 from layer 2.

1253

1254 A.12 ALIGNMENT METRICS

1255

1256 For comparability, the biological data was downsampled to 37 time steps for all alignment metric
1257 calculations. Except for DSA, all metrics were calculated on the individual time steps and the
1258 averaged.

1259

A.12.1 REPRESENTATIONAL SIMILARITY ANALYSIS (RSA)

1260

1261 RSA (Kriegeskorte et al., 2008) is computed by obtaining the Representational Dissimilarity Matrices
1262 (RDMs) for every time step individually via $RDM = 1 - PearsonCorrelation$. Then, based on
1263 the upper triangular values (excluding diagonals), the RSA scores are obtained from the Spearman's r
1264 (using `scipy stats`) between biological and artificial data.

1265

1266 A.12.2 CANONICAL CORRELATION ANALYSIS (CCA)

1267

1268 For CCA (Raghu et al., 2017), the data was first dimensionality reduced using PCA (3 components).
1269 Then, using `sklearn`'s CCA function, the first 3 canonical vectors were obtained and their correlations
1270 between brain and model were averaged, yielding CCA.

1271

A.12.3 LINEAR PREDICTIVITY (LP)

1272

1273 Linearly predicting individual biological neurons from FNN data using Ridge Regression did not
1274 yield adequate scores due to the high amount of noise. Therefore, we again used PCA to obtain
1275 3 components for brain data and 20 components for artificial data. We then fit Ridge Regression
1276 ($\alpha = 1$ to predict individual components of biological data using 40 random stimuli, and measured
1277 the average performance on the 8 heldout stimuli via R^2 (Yamins et al., 2014).

1278

A.12.4 DYNAMICAL SIMILARITY ANALYSIS (DSA)

1279

1280 For DSA (Ostrow et al., 2023), we again simplified data using PCA (10 components) and computed
1281 DSA scores using the DSA function from the authors. We reported inverted DSA scores ($1 - DSA$)
1282 to compare with other metrics, and Z-scores compared to a null distribution of 50 samples where the
1283 time steps of FNN data were randomly shuffled before comparing to biology.

1284

A.12.5 CRITIQUE

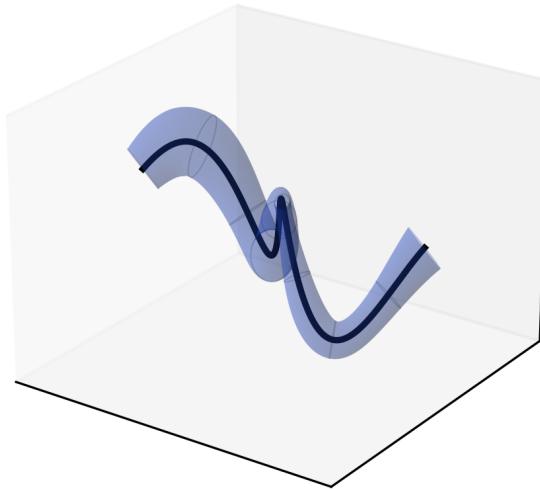
1285

1286 The validity of these methods for comparing brains and machines has been questioned (Serre, 2019b).
1287 Schaeffer et al. (2025) argue that lower LP scores may correspond to less brain-like models, as, instead
1288 of selecting biological models, LP overfits biases in linear regression. They claim that the same
1289 holds for overfitting to other representational similarity metrics. It is unclear what brain alignment
1290 means and what the different alignment metrics truly measure (Anonymous, 2025). Metric variability
1291 can fall within individual subject variability, making clear conclusions difficult (Anonymous, 2025).
1292 Bowers et al. (2023) question the assumption of biological visual systems being optimized to classify
1293 objects. Also, they claim that differences in the features used in DNNs compared to those in the brain
1294 can lead to high similarity. Moreover, simple features can be overrepresented compared to complex
1295 features, biasing the similarity scores (Lampinen et al., 2025). Finally, Dujmovic et al. (2024) find
1296 that metrics like RSA are not robust with respect to input perturbations.

1296 A.13 SOFTWARE
12971298 All software (Table 3) is used in accordance with its respective license.
12991300 Table 3: Software packages used in this work.
1301

1302 Package	1303 Version	1304 License
MATLAB Tensor Toolbox (Bader et al., 2023)	3.6	BSD-2
IAN (Dyballa and Zucker, 2023)	1.1.2	BSD-3
NeuralEncodingManifolds (Dyballa et al., 2024a)	N/A	BSD-2
NumPy (Harris et al., 2020)	1.25.0	BSD-3
SciPy (Virtanen et al., 2020)	1.15.3	BSD-3
scikit-learn (Pedregosa et al., 2011)	1.7.1	BSD-3
PyTorch (Paszke et al., 2019)	2.6.0	MIT
Matplotlib (Hunter, 2007)	3.10.1	PSF-based (BSD-compatible)
Plotly (Inc., 2015)	6.0.0	MIT
TUEplots (Krämer et al., 2024)	0.2.0	MIT

1313
1314 A.14 COMPUTE
13151316 The experiments were conducted on an HPC cluster. FNN sampling uses randomly selected GPUs
1317 (RTX 2080 Ti, or better). All other experiments were performed on CPU. All experiments required
1318 less than 30 GB memory. In total, 10 tensor decomposition experiments were run on CPU, each
1319 taking 2 days on a single CPU. Preliminary results not included in the paper required another 50
1320 tensor decomposition experiments.1321
1322 A.15 LANGUAGE MODEL USAGE1323 At the level of individual words or partial sentences, language models were used to fix language errors.
1324 Minor code sections were produced by language models and used only after careful inspection.
13251326 B DATA AND CODE AVAILABILITY
13271328 Upon acceptance, we will publish a GitHub repository with the full code necessary to reproduce
1329 all experiments and figures in this paper. We will also provide rotating video animations of three-
1330 dimensional visualizations to aid interpretation.1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

1350
1351 C SUPPLEMENTAL FIGURES
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
13721373 Figure 7: A tubular neighborhood around a centerline $c(u)$ with radius profile $R(u)$.
13741375 Table 4: **Tubularity metrics for biological and FNN data.** Low tightness and crossings values
1376 indicate high tubularity. Aligning with the visualization in Fig. ??, the biological trajectories show
1377 highly tubular organizations compared to FNN. Method details in Appendix ??.
1378

1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389	Layer	Ground Truth Labels		HDBSCAN Labels		Clusters
		S _{tight}	S _{cross}	S _{tight}	S _{cross}	
Retina	0.0688	1.29×10^{-5}	0.1017	1.06×10^{-5}		4
V1	0.1357	4.06×10^{-5}	0.1859	3.50×10^{-5}		4
Enc1	0.2018	2.87×10^{-4}	0.7680	1.66×10^{-4}		1
Enc13	1.9885	1.77×10^{-6}	4.3461	1.09×10^{-6}		3
Rec	0.1228	2.65×10^{-7}	0.1697	1.53×10^{-7}		4
RecOut	0.1209	5.72×10^{-7}	0.1650	5.34×10^{-7}		4
Readout	0.3307	3.96×10^{-6}	0.4320	5.40×10^{-6}		4
Output	0.1483	3.53×10^{-6}	0.2784	1.12×10^{-6}		3

1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

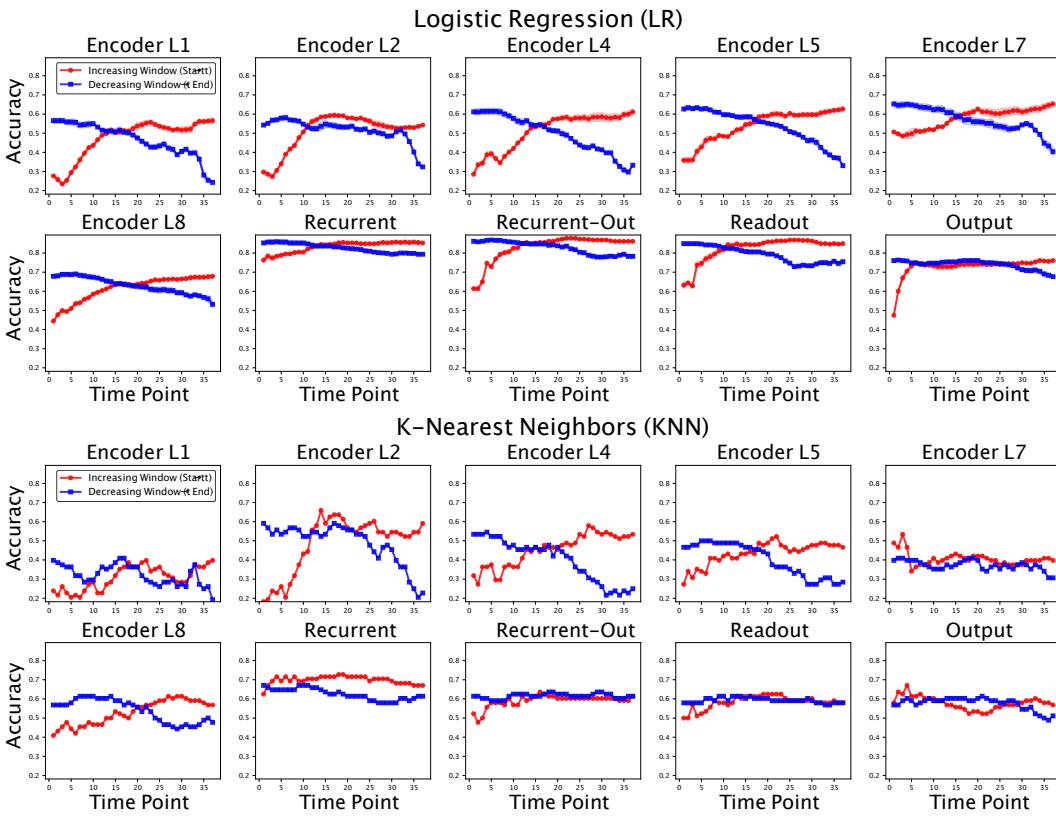


Figure 8: **Logistic regression (LR, top) and K-Nearest Neighbor (KNN, K=3, bottom) classifier accuracy** for each layer. We use increasing time windows (timesteps $0 \rightarrow t$, red) or decreasing time windows ($t \rightarrow 37$, blue) to calculate the accuracies. Shaded regions for LR show the s.e.m. The maxima across panels are summarized in Table 1.

Table 5: Representational Similarity Analysis (RSA), Canonical Correlation Analysis (CCA), Linear Predictivity (LP) and Dynamic Similarity Analysis (DSA) scores and DSA Z-scores to a time-shuffled baseline. High values indicate closer alignment for all metrics. (see Appendix A.12)

Region	Metric	L1	L2	L4	L5	L7	L8	Rec	Readout	Output
Retina	RSA	-.04	-.03	-.03	-.01	-.01	0.03	0.03	0.05	-.01
Retina	CCA	0.19	0.25	0.25	0.30	0.27	0.26	0.34	0.25	0.32
Retina	LP	0.05	0.06	0.17	0.29	0.26	0.28	0.43	0.24	0.26
Retina	DSA	0.84	0.77	0.83	0.73	0.58	0.56	0.80	0.81	0.80
Retina	DSA-Z	5.90	4.77	5.29	4.55	2.35	1.87	2.00	1.79	4.45
V1	RSA	-.11	-.27	-.16	-.22	0.10	0.08	0.46	0.08	0.41
V1	CCA	0.29	0.31	0.33	0.39	0.29	0.35	0.39	0.29	0.38
V1	LP	0.05	0.04	0.18	0.29	0.23	0.27	0.40	0.22	0.24
V1	DSA	0.91	0.76	0.93	0.76	0.59	0.57	0.88	0.92	0.90
V1	DSA-Z	7.03	5.49	6.52	5.91	5.46	4.40	4.72	3.50	5.27

Table 6: Explained Variance (EV, in %) of decoding manifold PCA, and number of tensors (R) and Error Percentage (EP) of tensor factorization (in %)

Metric	L1	L2	L4	L5	L7	L8	Rec	RecOut	Readout	Output
EV	47.91	57.66	59.52	48.20	48.53	42.77	57.73	53.43	53.99	55.23
R	8	12	11	11	13	11	9	12	17	14
EP	37.08	25.50	23.02	23.35	22.34	23.51	15.68	16.45	13.81	9.26

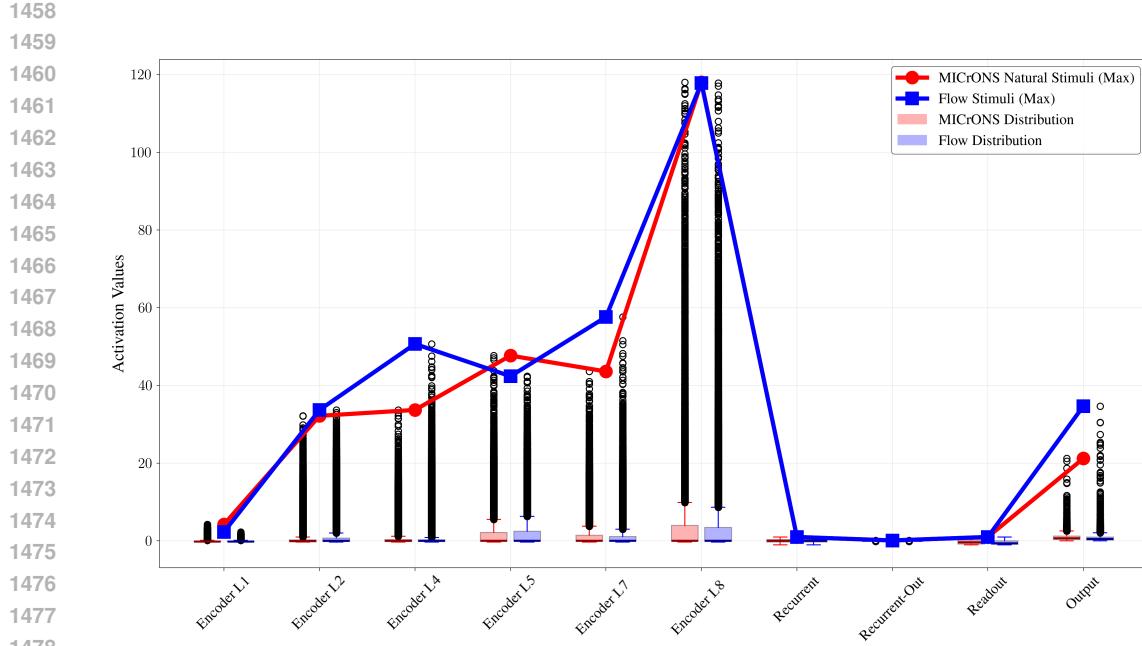


Figure 9: **Activation function output distributions and maxima** for natural MICrONS (Bae et al., 2025) input videos and the flow stimulus ensemble (Dyballa et al., 2024a). The comparable activity across network layers shows the adequacy of investigating the FNN with flow stimuli. The differences in magnitudes across layers are explained by the activations functions (GELU in the *encoder*, Tanh in the *recurrent* and *readout* modules).

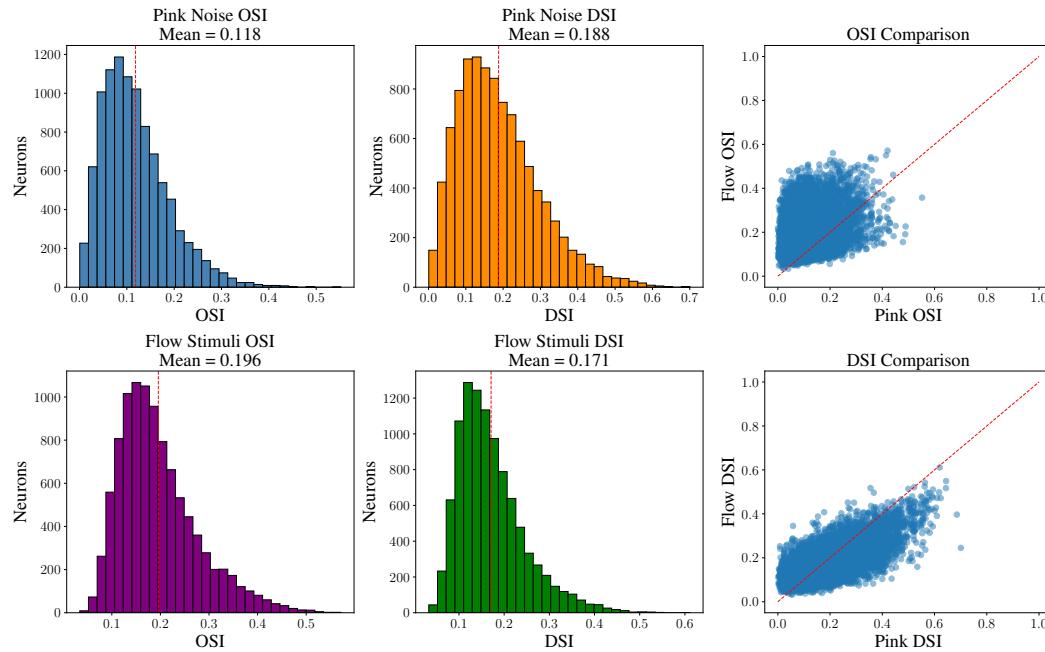


Figure 10: **OSI and DSI of FNN output** for pink noise (as used in Wang et al. (2025)) and for the stimulus ensemble from Dyballa et al. (2024a), meaned over the different stimuli.

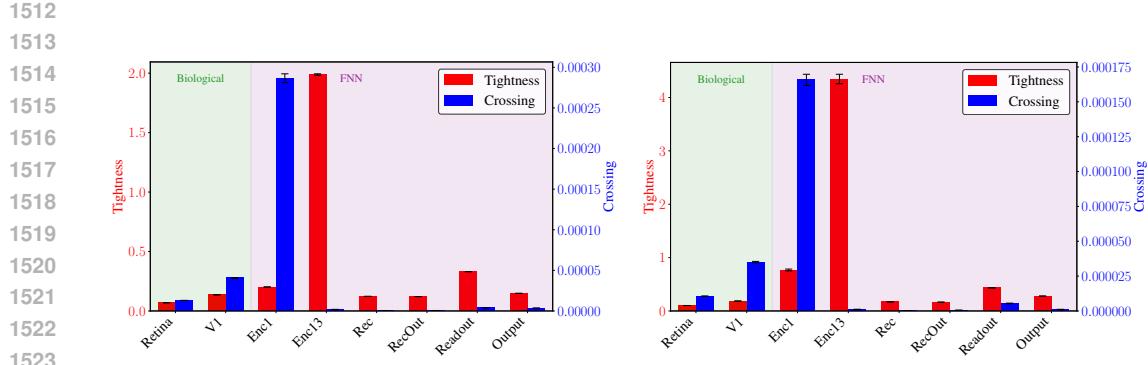


Figure 11: **Tubularity comparison between biological and FNN data.** Tightness, measuring how close trajectories within a bundle are to their centerline, and crossings, measuring the amount of transverse crossings in a bundle, are scores for biological and FNN data. Left: Using ground-truth stimulus class labels. Right: Using HDBSCAN (Campello et al., 2013) labels.

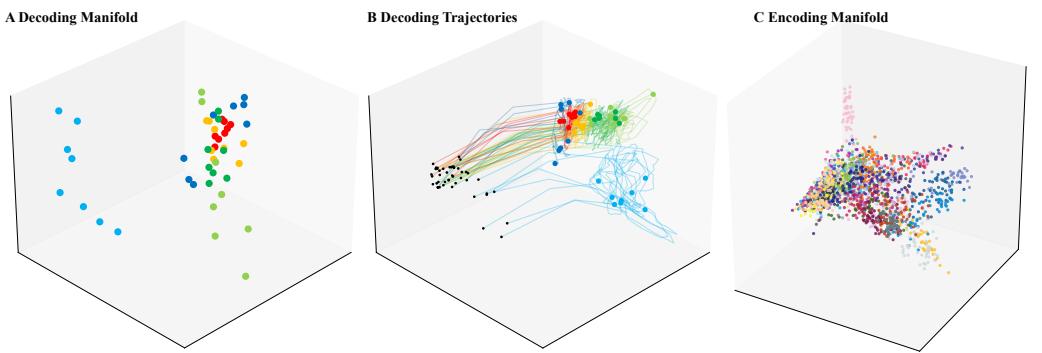


Figure 12: **Encoder L2.** Note the encoding manifold smoothness here results from the early layer only capturing simple features and the dominance of intensity discussed for L8. We therefore do not interpret this as a V1-like smooth encoding manifold.

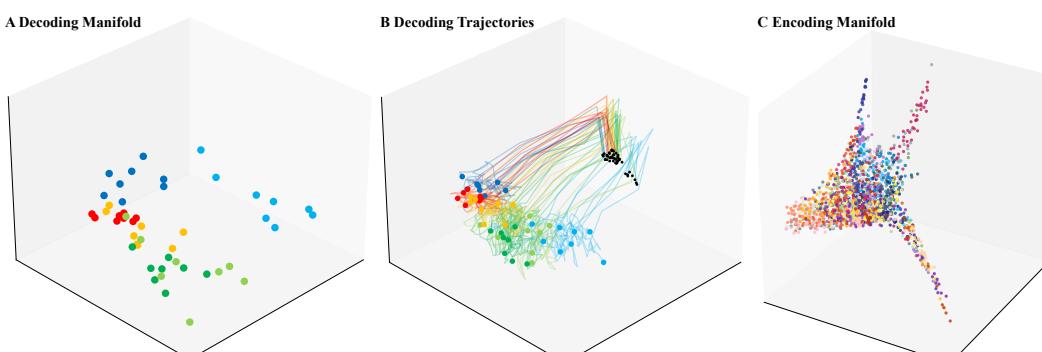
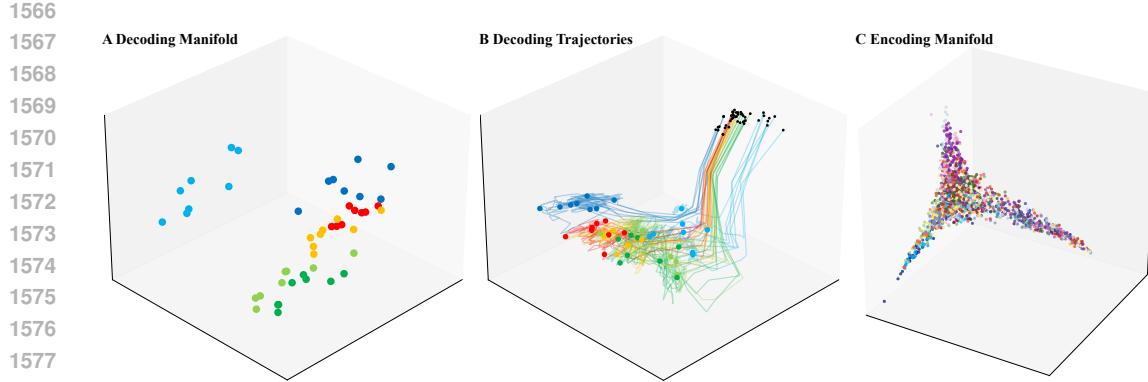


Figure 13: **Encoder L4.** Note the encoding manifold smoothness here results from the early layer only capturing simple features and the dominance of intensity discussed for L8. We therefore do not interpret this as a V1-like smooth encoding manifold.



1579 Figure 14: **Encoder L5.** Note the encoding manifold smoothness here results from the early layer
1580 only capturing simple features and the dominance of intensity discussed for L8. We therefore do not
1581 interpret this as a V1-like smooth encoding manifold.

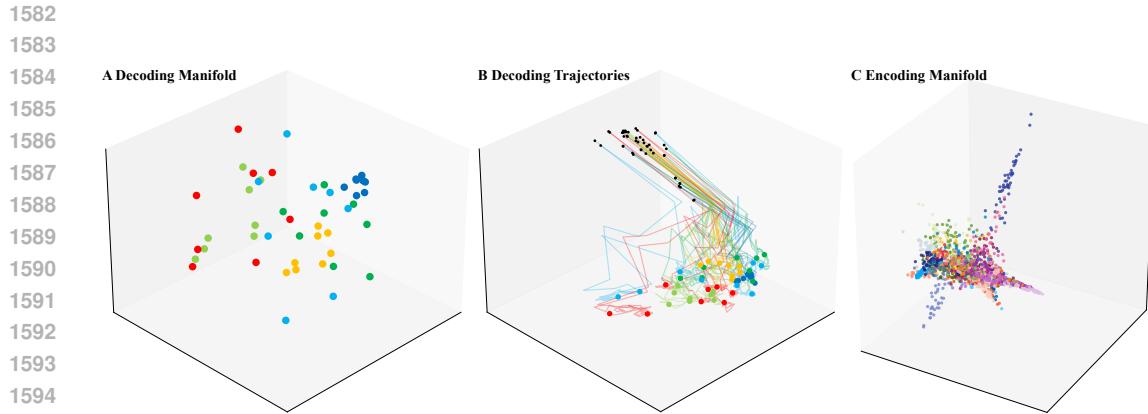
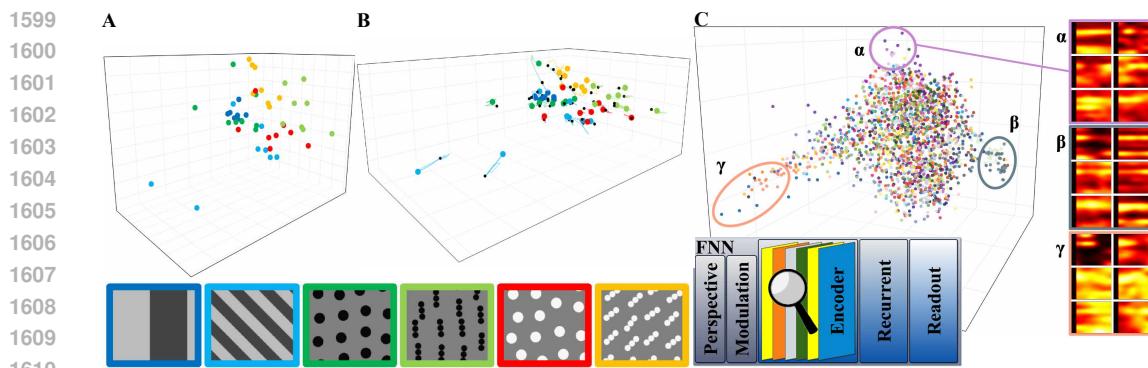


Figure 15: **Encoder L7**



1611 Figure 16: **Encoder L8 decoding manifold, trajectories and encoding manifold without intensity**
1612 **artifacts.** Without the intensity artifacts there is no temporal development at all in the decoding
1613 trajectories (comparable to encoder L1) apart from the jump after the 0-th step. The non-selective
1614 high intensity neurons are padding artifacts at the edges of the image. In the encoder, due to spatial
1615 convolutions, the effect of these artifacts spreads out across the feature maps. This is supported by the
1616 intensity smoothly organizing the manifold with a transition from intensity-only neurons to selective
1617 responses. In the recurrent stage, the function of the attention layer is capable of filtering exactly
1618 those artifacts out. The artifacts are reintroduced by the recurrent-output convolution, but then filtered
1619 out by the readout interpolation from central neurons only.

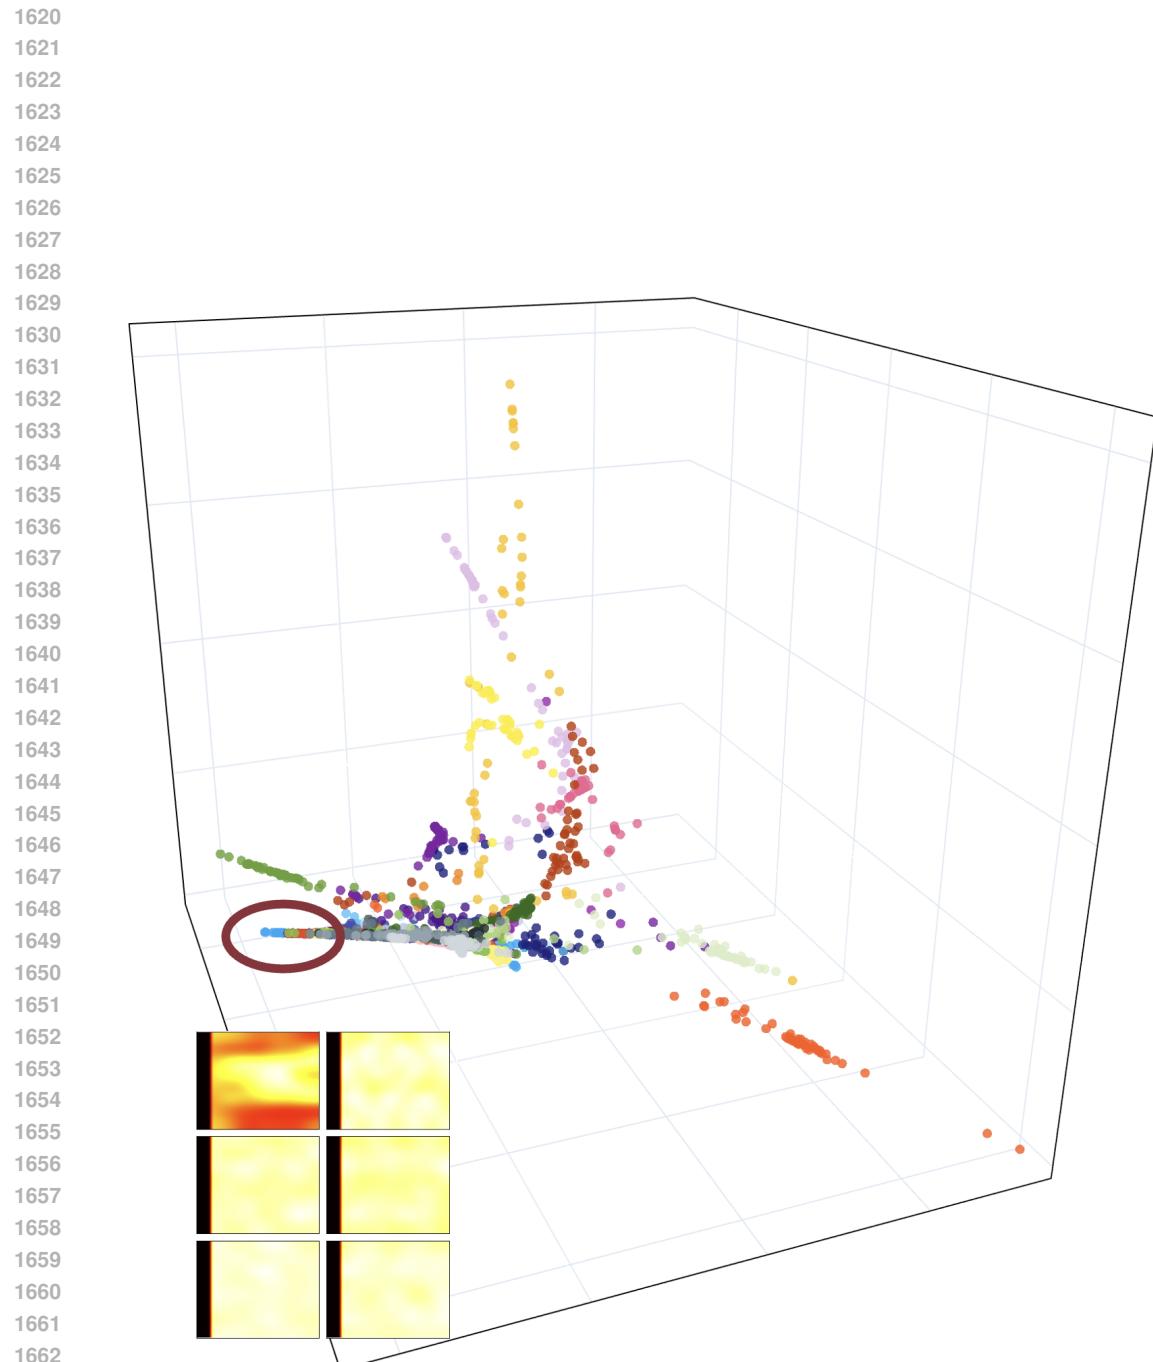
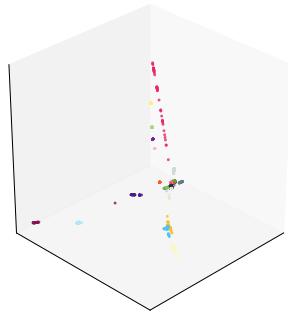


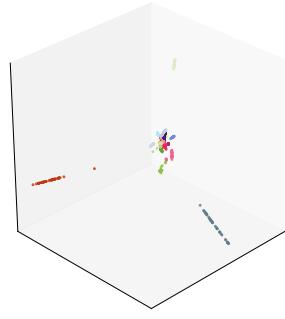
Figure 17: **Minimodel encoding manifold with intensity arm.** The intensity artifacts are also present in the border regions of feature maps in the model from Du et al. (2025).

1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699

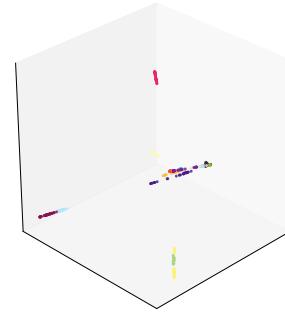
a) 2000 Neurons, 40 feature maps,
 50 neurons each, intensity-based
 feature map and neuron sampling.



b) 2000 Neurons, 40 feature maps,
 50 neurons each, random feature
 map, and intensity-based neuron
 sampling.

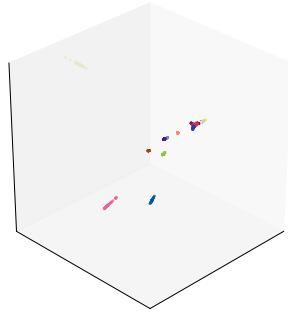


c) 2000 Neurons, 40 feature maps,
 50 neurons each, intensity-based
 feature map, and random neuron
 sampling.

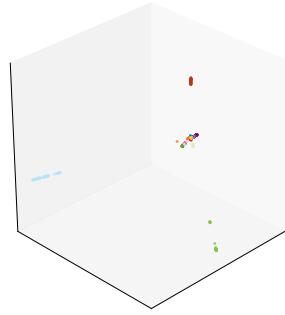


1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714

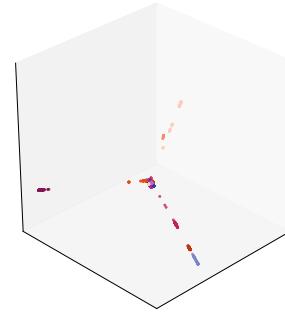
d) 2000 Neurons, 40 feature maps,
 50 neurons each, fully random fea-
 ture map and neuron sampling.



e) 4000 Neurons, 80 feature maps,
 50 neurons each, intensity-based
 feature map and neuron sampling.



f) 5000 Neurons, 50 feature maps,
 100 neurons each, intensity-based
 feature map and neuron sampling.



1715 Figure 18: **Sampling tests for readout encoding manifolds.** The encoding manifold for all sampling
 1716 conditions look similar, having clusters by feature maps.

1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727