Accurate Identification of Communication Between Multiple Interacting Neural Populations

Belle Liu¹ Jacob Sacks² Matthew D. Golub²

Abstract

Neural recording technologies now enable simultaneous recording of population activity across many brain regions, motivating the development of data-driven models of communication between brain regions. However, existing models can struggle to disentangle the sources that influence recorded neural populations, leading to inaccurate portraits of inter-regional communication. Here, we introduce Multi-Region Latent Factor Analysis via Dynamical Systems (MR-LFADS), a sequential variational autoencoder designed to disentangle inter-regional communication, inputs from unobserved regions, and local neural population dynamics. We show that MR-LFADS outperforms existing approaches at identifying communication across dozens of simulations of task-trained multi-region networks. When applied to large-scale electrophysiology, MR-LFADS predicts brain-wide effects of circuit perturbations that were held out during model fitting. These validations on synthetic and real neural data position MR-LFADS as a promising tool for discovering principles of brain-wide information processing.

1. Introduction

Large-scale neural recording technologies, such as highdensity electrophysiology (Steinmetz et al., 2019; Siegle et al., 2021; IBL et al., 2023; Chen et al., 2024; Bennett et al., 2024) and calcium imaging (Sofroniew et al., 2016; Song et al., 2017; Allen et al., 2017), now enable simultaneous recording of neural population activity across many brain regions. These advances have revealed that many sensory, cognitive, and motor processes engage spatially distributed networks in the brain (Makino et al., 2017; Gilad et al., 2018; Stringer et al., 2019; Musall et al., 2019; Allen et al., 2019; Jia et al., 2022). Consequently, there has been growing interest in the design of data-driven *communica-tion models* that seek to infer the pathways and content of communication between the recorded regions.

Accurately identifying inter-regional communication is challenging for at least *four* reasons (Biswas et al., 2020; Kang & Druckmann, 2020; Keeley et al., 2020; Perich & Rajan, 2020; Semedo et al., 2020; Kass et al., 2023). First, communication signals are not directly observed in multi-region recordings. Although some recorded neurons might project to other recorded regions, the identity and targets of these projection neurons are typically unknown. Second, models may need to account for inputs to the recorded brain regions from other regions that were not recorded during the experiment. Third, models should faithfully reconstruct activity within each recorded region by accounting for communication between recorded regions, inputs from unrecorded regions, and local neural population dynamics-capturing complex features such as structured trial-to-trial variability (Goris et al., 2014) and nonlinear, nonstationary, and state-dependent population dynamics (Shenoy et al., 2013; Vyas et al., 2020; Duncker & Sahani, 2021; Durstewitz et al., 2023). Fourth, accurate reconstruction of the recorded data does not guarantee accurate inference of the underlying communication. Many different models may sufficiently explain the recorded data, leading to ambiguities about which, if any, should be trusted for scientific interpretation.

In this work, we introduce Multi-Region Latent Factor Analysis via Dynamical Systems (**MR-LFADS**), a multi-region communication model that directly addresses all of the challenges outlined above. MR-LFADS is a probabilistic model that represents each recorded region with a distinct set of stacked recurrent neural networks (**RNNs**) that capture the region's potentially nonlinear and nonstationary population dynamics. MR-LFADS represents communication between observed regions and inputs from unobserved regions as disentangled sets of latent variables. Structured information bottlenecks encourage the model to infer inputs from unobserved regions only when their effects cannot be explained by communication among the recorded regions. MR-LFADS infers single-trial initial conditions and time-

¹Graduate Program in Neuroscience, University of Washington; ²Paul G. Allen School of Computer Science & Engineering, University of Washington. Correspondence to: Matthew Golub <mgolub@cs.washington.edu>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

varying inputs that together account for trial-to-trial variability in the recorded activity. This automatic inference of inputs eliminates the need to manually specify input signals, thereby avoiding strong, difficult-to-validate assumptions about how external signals influence each region. Finally, MR-LFADS constrains communication to originate from model-reconstructed neural activity, rather than from more flexible latent representations—a design choice that, as we will show, enables more accurate inference of communication without sacrificing the quality of data reconstruction.

To evaluate MR-LFADS, we developed 37 synthetic multiregion datasets that capture real-world challenges in communication modeling across a range of neuroscience-relevant scenarios. On these datasets, MR-LFADS consistently outperforms existing models in recovering the pathways and content of communication. Through targeted ablations of key design features, we demonstrate that these features indeed improve the identification of communication. We then applied MR-LFADS to multi-region electrophysiological recordings in mice performing a decision-making task (Chen et al., 2024). In a subset of trials that were held out during model fitting, photoinhibition was applied to the anterior lateral motor cortex. MR-LFADS predicted the brainwide effects of these circuit perturbations, suggesting that MR-LFADS inferred an accurate account of inter-regional communication. Moreover, MR-LFADS infers consistent communication across multiple training runs from different random initializations, demonstrating its robustness and reliability in real-data settings.

2. Related Work

Existing communication models can be broadly categorized as either static or dynamic. Static methods predict each timestep of neural activity in a target region from a corresponding timestep of activity in one or more source regions and then interpret predictive source activity as inter-regional communication (Kaufman et al., 2014; Perich et al., 2018; Ruff & Cohen, 2019; Veuthey et al., 2020). Reduced-rank regression (**RRR**) is a prominent static technique that models target-region activity as a low-rank linear function of the source-region activity (Semedo et al., 2019; MacDowell et al., 2025). While static methods are straightforward to fit and interpret, they are typically limited to capturing instantaneous, linear dependencies. Consequently, they do not readily account for nonlinear or nonstationary relationships, or temporal structure that may arise due to neural population dynamics (Vyas et al., 2020).

Dynamic methods explicitly model temporal dependencies using, for example, switching linear dynamical systems (**SLDS**) (Linderman et al., 2016), RNNs (Perich et al., 2020), or Gaussian processes (Yu et al., 2008; Gokcen et al., 2022; 2024). We will pay particular attention to two such techniques that, like MR-LFADS, are coupled nonlinear dynamical systems, with each dynamical system representing one recorded region: multi-population sticky recurrent SLDS (**mp-srSLDS**) (Glaser et al., 2020), and Multi-Region Switching Dynamical Systems (**MR-SDS**) (Karniol-Tambour et al., 2024). In mp-srSLDS each region is modeled by an SLDS, while in MR-SDS each is modeled as a switching nonlinear dynamical system.

While these existing approaches to communication modeling address some of the challenges outlined in Section 1, none, to our knowledge, address all four challenges. In particular, none of these existing methods support inferring inputs from unobserved brain regions. This functionality is crucial because inputs from unobserved regions might influence the target region's population dynamics and modes of communication. Some approaches attempt to account for unobserved inputs by explicitly providing task-related signals as inputs to each region or by removing condition averages and modeling the residual single-trial neural activity. However, these manual strategies impose strong assumptions about the content and targets of input signals. As we will show, misspecifying such inputs risks confounding inferred population dynamics and communication, leading to models that accurately reconstruct neural activity through incorrect mechanisms.

To address such unobserved inputs, Pandarinath et al. (2018) introduced Latent Factor Analysis via Dynamical Systems (LFADS), a sequential variational autoencoder (sVAE) for modeling single-trial neural population dynamics within a single recorded brain region. LFADS jointly identifies a nonlinear dynamical system, implemented as an RNN, along with the single-trial initial conditions and time-varying unobserved inputs needed to drive the system to reconstruct single-trial neural population recordings. We henceforth use SR-LFADS to refer to a single-region LFADS model. MR-LFADS builds on this foundation to support multi-region modeling with explicit disentangling of communication, unobserved inputs, and local population dynamics.

3. Multi-Region LFADS (MR-LFADS)

MR-LFADS is composed of a set of SR-LFADS modules (Fig. 1a) that interact through constrained communication channels (Fig. 1b). At a high level, MR-LFADS is a coupled set of driven nonlinear dynamical systems that are jointly trained to reconstruct all single trials of a multiregion dataset. Each recorded brain region *i* is modeled as a dynamical system that attempts to reconstruct the region-*i* recorded neural activity x_t^i at each time $t = 1, \ldots, T$. Each region-*i* dynamical system evolves from a single-trial initial state g_0^i and is driven by (1) single-trial time-varying communication messages $m_t^{j \to i}$ from other recorded brain regions *j* and (2) single-trial time-varying inferred inputs



Figure 1. MR-LFADS architecture. (a) Single-region LFADS, as adapted for this work. (b) MR-LFADS with N regions. KL penalties from SR-LFADS in panel (a) are included in MR-LFADS, but are omitted in the diagram for clarity.

 u_t^i , representing input from unobserved brain regions.

Notation. All time-indexed variables and parameters are also indexed by trial, though we omit trial indices for notational simplicity. We use $t_1:t_2$ to denote the inclusive sequence of integers $\{t_1, t_1 + 1, \ldots, t_2\}$. We use $W^i(x):=W^ix+b^i$ to denote an affine transformation with weights W^i and offsets b^i .

Generative Model. MR-LFADS treats all initial states g_0^i , communications $m_t^{j \to i}$, and inferred inputs u_t^i as latent variables. The prior distributions over these latent variables are modeled as:

$$g_0^i, u_t^i, m_t^{j \to i} \sim \mathcal{N}(0, \sigma^2 I) \tag{1}$$

Given these quantities, the neural population dynamics in region *i* are modeled by a "generator" gated recurrent unit (**GRU**) network, GRU_{gen}^i , with internal states g_t^i that evolve according to:

$$g_t^i = \operatorname{GRU}_{\operatorname{gen}}^i \left(g_{t-1}^i, \left[\{ m_t^{j \to i} \}_{j \neq i}; u_t^i \right] \right)$$
(2)

A set of region-*i* factors f_t^i are defined as an affine readout from the corresponding generator states:

$$f_t^i = W_f^i(g_t^i) \tag{3}$$

These factors are then transformed into parameters of timevarying output distributions in a manner dependent on the nature of the neural recordings. For continuous-valued observations, as in calcium imaging, a Gaussian or zeroinflated Gamma distribution may be appropriate (Zhu et al., 2022). In the synthetic-data experiments of Section 4, we apply a Gaussian output distribution:

$$P(x_t^i \mid g_0^i, u_{1:t}^i, \{m_{1:t}^{j \to i}\}_{j \neq i}) = \mathcal{N}(r_t^i, \Sigma_{r,t}^i)$$
(4)

$$r_t^i = W_r^i(f_t^i) \qquad \Sigma_{r,t}^i = \operatorname{diag}\left(\exp(W_{\sigma_r}^i(f_t^i))\right) \tag{5}$$

where r_t^i and $\Sigma_{r,t}^i$ are the region-*i* predicted mean and covariance of the time-*t* recorded neural activity, respectively, and are each computed via separate affine transformations, W_r^i and $W_{\sigma_r}^i$. For spike count observations, as modeled in the electrophysiology experiments of Section 5, we apply a Poisson output distribution:

$$P(x_t^i \mid \cdot) = \text{Poisson}(r_t^i) \qquad r_t^i = \exp\left(W_r^i(f_t^i)\right) \quad (6)$$

where the exponential nonlinearity ensures non-negative predicted firing rates r_t^i .

Inference Model. Following VAE conventions (Kingma & Welling, 2013), MR-LFADS approximates the intractable true posterior distributions over the latent variables using variational posteriors, denoted $q(\cdot|\cdot)$.

MR-LFADS defines the approximate posteriors over communication messages from observed regions j to i as Gaussian distributions:

$$q(m_t^{j \to i} \mid x_{1:t}^j) = q(m_t^{j \to i} \mid r_t^j) = \mathcal{N}(\mu_{m,t}^{j \to i}, \Sigma_{m,t}^{j \to i})$$
(7)

with parameters derived from the region-*j* predicted firing rates r_t^j :

$$\mu_{m,t}^{j \to i} = W_{\mu_m}^{j \to i}(r_t^j)$$

$$\Sigma_{m,t}^{j \to i} = \operatorname{diag}\left(\exp\left(W_{\sigma_m}^{j \to i}(r_t^j)\right)\right)$$
(8)

Constraining communication to be derived from r_t^j anchors it to the neural recordings and in doing so reduces ambiguity in system identification. We refer to this rate-based communication model as **MR-LFADS(R)**. In Section 4, we also explore generator-based **MR-LFADS(G)** and factor-based **MR-LFADS(F)** communication models, which replace all instances of r_t^j in Eq. 8 with g_t^j and f_t^j , respectively. Approximate posteriors over the region-*i* initial generator states g_0^i and inferred inputs (from unobserved brain regions) u_t^i are defined as the following Gaussian distributions:

$$q(g_0^i \mid x_{-\tau:0}^i) = \mathcal{N}(\mu_{g_0}^i, \Sigma_{g_0}^i)$$
(9)

$$q(u_t^i \mid x_{1:t}^i) = \mathcal{N}(\mu_{u,t}^i, \Sigma_{u,t}^i)$$
(10)

where the mean and covariance parameters are computed from the corresponding conditioning recorded neural activity x^i via a set of region-*i*-specific "encoder" and "controller" GRU networks (see Appendix A.1). Our approach here slightly modifies the original SR-LFADS specification, which allows acausal inference via a bidirectional encoder network that processes each entire *T*-timestep neural recording $x_{1:T}$ to infer g_0 and each element of $u_{1:T}$. In contrast, we infer g_0^i (Eq. 9) using a bidirectional encoder applied only to past neural activity $x_{-\tau:0}^i$, preserving causality, and we infer u_t^i (Eq. 10) using a unidirectional encoder RNN that processes $x_{1:t}^i$ in a strictly forward, causal manner. This formulation ensures that all predicted firing rates r_t^i , and thus all derived communication signals $m_t^{j \to i}$, are inferred causally from neural activity recorded up to time t.

Model Fitting. Following VAE conventions, MR-LFADS is trained by maximizing the evidence lower bound (**ELBO**), a variational lower bound on the data log-likelihood. The ELBO is a sum of two terms: (1) the expected log-likelihood

$$\sum_{t=1}^{T} \mathbb{E}_q \left[\log P(\{x_t^i\} \mid \{g_0^i\}, \{u_t^i\}, \{m_t^{j \to i}\}) \right]$$
(11)

and (2) the negative Kullback-Leibler (**KL**) divergence D_{KL} between the approximate posteriors (Eqs. 7–10) and the priors (Eq. 1) over the latent variables.

The expected log-likelihood measures reconstruction accuracy and is estimated by running samples from the approximate posteriors (Eqs. 7-10) through the generative model to evaluate the experiment-dependent output distributions from Eqs. 4–6. The $D_{\rm KL}$ term acts as a regularizing information bottleneck on the latent variables. To control this regularization, we allow rescaling of the $D_{\rm KL}$ term (Higgins et al., 2017; Keshtkaran et al., 2022). Noting that the $D_{\rm KL}$ term decomposes into contributions from the three sets of MR-LFADS latent variables $\{g_0^i\}, \{u_t^i\},$ and $\{m_t^{j \to i}\}$, we weight each contribution differently and treat the weights $(\beta_{q_0}, \beta_u, \beta_m)$ as hyperparameters. To encourage MR-LFADS to infer inputs from unobserved regions only when that information cannot be obtained as communication from an observed region, we propose a structured KL bottleneck with $\beta_u = 10\beta_m$. Other choices of KL regularization structure might be appropriate if a priori knowledge is available about information flow or anatomical connectivity between recorded regions. See Appendix A.1 for further detail on MR-LFADS.

4. Results I: Synthetic Multi-Region Datasets

Here, we evaluate MR-LFADS' ability to recover the ground truth inter-regional communication across a broad range of synthetic multi-region datasets. Each dataset was generated by a unique data-generating network (**DGN**): an ensemble of noisy RNN modules jointly trained to perform a specified cognitive neuroscience task, with each module representing a distinct brain region. Prior to training, we explicitly specified the presence or absence of directed, low-rank communication channels between each region pair.

In Experiments 1 and 2, we manually designed the DGNs to impose specific challenges outlined in Section 1. In Experiment 3, we randomly generated dozens of DGNs, each trained to perform a randomly selected cognitive neuroscience task. Across all experiments, we treated each module's hidden-unit activity as recorded neural activity and retained all external inputs, inter-module connectivity, and communication signals as ground truth. These ground truth quantities are crucial for evaluating communication models but are typically not directly observable in real-data settings.

To benchmark MR-LFADS, we compare it to three established communication modeling techniques—RRR (see Appendix A.2), mp-srSLDS, and MR-SDS (see Section 2). To demonstrate the importance of specific MR-LFADS design features, we also compare against ablated MR-LFADS variants that selectively exclude those features. We focus evaluation on each method's ability to recover a causal model of each DGN—including both the *pathways* and *content* of inter-regional communication—given a multi-region dataset generated by the DGN.

To assess communication *pathways*, we consider recovery of an "effectome" (Pospisil et al., 2024) describing the causal flow of effects along the inter-regional connectome. We represent this effectome as a matrix with each element $(i, j)_{i \neq j}$ indicating the volume of directed communication flow from region *j* to *i*. The effectome reflects both the inter-regional connectivity and the magnitude of communication flow over each directed connection. For each dataset, we compare model-inferred effectomes to the ground truth effectome by computing cosine similarity S_{cos} between the vectorized effectome matrices.

To assess communication *content*, we quantify how well the model-inferred messages capture the information in the ground truth messages. Specifically, we apply linear regression to predict the ground truth messages $m_t^{j \to i}$ from the inferred messages $\mu_{m,t}^{j \to i}$ (Eq. 7), and we report a cross-validated coefficient of determination, denoted $R^2(\mu_m^{j \to i}, m^{j \to i})$. For MR-LFADS, we use an analogous procedure to compare model-inferred inputs $\mu_{u,t}^i$ (Eq. 10) to ground truth external inputs. This comparison is not applicable to RRR, mp-srSLDS, or MR-SDS, as these methods do not infer inputs from unobserved regions. Additional details on experiments, model hyperparameters, and evaluation metrics are provided in Appendix B-D.

4.1. Experiment 1: Inferring Unobserved Inputs

This experiment demonstrates that inferring inputs, rather than manually specifying inputs, enables more accurate identification of inter-regional communication. To evaluate design implications related to input specification, we designed a DGN that implements a dynamical memory function. Each region of this "memory network" receives unique stimulus information from one observed region and one unobserved region, and is tasked with remembering a recent history of those signals (Fig. 2a, left). With each region of the DGN receiving information from both observed and unobserved sources, this setup poses the challenge of disentangling whether each signal arises due to communication or due to external input. A common modeling choice is to provide all known external inputs to all model regions and to let model fitting determine which inputs are needed by each region. However, in this case, such manual input specification can result in a model completely forgoing communication (Fig. 2a, right) because the manually specified inputs contain the information that was actually transmitted as communication in the DGN.

We evaluate MR-LFADS(R), which does not use the stimulus signals during training or evaluation but rather automatically infers external inputs in an unsupervised manner. We also evaluate an ablated MR-LFADS variant that does not automatically infer inputs. Termed MR-LFADS(S), this model receives all stimulus signals as manually specified external inputs to each region's generator GRU^{*i*}_{gen} (replacing u_t^i with $s_t^{1:3}$ in Eq. 2). By comparison, MR-SDS and mp-srSLDS also receive all stimulus signals as manually specified inputs, and RRR neither receives nor infers external inputs.

The MR-LFADS variants (R and S) reconstructed the simulated multi-region activity more accurately than MR-SDS, mp-rSLDS, and RRR (Fig. 2b, Fig. S1a). Critically, MR-LFADS(R) accurately infers the effectome (Fig. 2c, left, Fig. S1b). By contrast, all models that are manually provided stimulus signals (MR-LFADS(S), MR-SDS, mprSLDS) infer less accurate effectomes and demonstrate the failure mode mentioned above, forgoing communication and instead relying on the specified inputs to provide the corresponding signals. RRR also infers a less accurate effectome. These results suggest that manually specifying inputs can discourage models from utilizing—and thus identifying—communication. By automatically inferring inputs, MR-LFADS(R) avoids this failure mode.

Next, we assess the accuracy of inferred inputs and mes-

sages. In the ground truth DGN at time t, region-i receives only s_t^i and $m_t^{j \to i}$. An accurate communication model should therefore infer inputs and messages that encode only these time-t quantities. However, due to structure of the memory task, the DGN's region-i time-t activity contains information about $s_{t-4:t}^{i}$ and $m_{t-4:t}^{j\to i} = s_{t-6:t-2}^{j}$. To re-



Figure 2. Experiment 1. (a) Left: DGN setup. Each region, area A^i , receives a private stimulus s_t^i (red) and communicates a two-step delayed version s_{t-2}^i (blue) to a downstream region. Each region is trained to recall (green) the last five time steps of its private stimulus and its received communication. Right: Potential incorrect communication model capable of accurately reconstructing these synthetic neural data. (b) R^2 scores for data reconstruction. Box plots describe distributions of values across 10 models fit from distinct random initializations (seeds). Boxes represent the interquartile range (IQR), and whiskers extend to the most extreme points within 1.5 IQR from the quartiles. (c) Left: Cosine similarity between inferred effectomes and ground truth. Right: Example fitted models. Color intensity indicates the relative message norm, computed by concatenating all multidimensional messages across trials and time, taking the 2-norm, then normalizing across communication channels. $s^{1:3}$ indicates the ground truth input. (d) R^2 of linear prediction of ground truth stimulus inputs (with time lag $\in \{0,2\}$) from inferred inputs. (e) \mathbb{R}^2 of linear prediction of ground truth messages (with time lag $\in \{2, 4\}$) from inferred messages.

construct the data, a communication model must account for how this time-lagged information becomes represented in region-*i* at time-*t*. MR-LFADS(R) correctly infers the current timestep ground truth stimuli s_t^i as inputs to each region (Fig. 2d, left), while correctly avoiding inferring timelagged versions of those stimuli (Fig. 2d, right), despite their utility for data reconstruction. Similarly, MR-LFADS(R) accurately recovers the current timestep messages $m_t^{j\to i}$ (Fig. 2e, left) and avoids incorrectly inferring time-lagged versions (Fig. 2e, right). By contrast, all other models infer no communication or communication with incorrect temporal lags (Fig. S1c). See Appendix B.2 for further details.

Taken together, these results indicate that MR-LFADS(R) learned region-specific population dynamics consistent with those that implement the memory functions in the ground truth DGN. Moreover, these results demonstrate the unique ability of MR-LFADS(R) to disentangle region-specific external inputs, communication between recorded regions, and local population dynamics, all in an unsupervised manner that mitigates biases associated with manual specification of external inputs. See Appendix E.1 for further analyses linking this disentangling (Miller et al., 2024) to MR-LFADS(R)'s structured information bottlenecks.

4.2. Experiment 2: Data-Constrained Communication

This experiment demonstrates the implications of messageinference design choices. We designed MR-LFADS(R) to infer messages as affine functions of the source-region predicted firing rates r_t^i (Eqs. 7–8), which are tied to the observed source-region neural activity x_t^i through the data reconstruction term in the ELBO (Eq. 11). This dataconstrained communication architecture contrasts with that of MR-SDS and mp-srSLDS, which infer communication as a function of less-constrained, source-region latent dynamical states (see Table S3). To directly evaluate the implications of this design choice, all within the MR-LFADS framework, we designed model variants with less-constrained factor-based and generator-based communication, termed MR-LFADS(F) and MR-LFADS(G), respectively.

To highlight the significance of this design choice, we evaluate models on data from a two-region "pass-decision" DGN that computes perceptual decisions based on time-varying sensory evidence (Fig. 3a, left). An upstream region, area A^P , receives a white noise stimulus s_t and is trained to *pass* that stimulus through to a readout, effectively learning an identity function routing input to output. A downstream *decision* region, area A^D , receives this routed stimulus as communication $m_t^{P \to D}$, integrates that stimulus over time into a decision variable d_t , and reports ± 1 choices indicating the sign of that decision variable (Mante et al., 2013).

This setup again challenges models to disentangle external inputs, communication, and local dynamics—and in partic-

ular, accurately identifying and localizing the ground truth pass-through and integration computations. Though the integration dynamics are localized to A^D in the DGN, an overly flexible model might instead learn integration in A^P , yielding accurate data reconstruction but mislocalizing the computation, e.g., if A^P integrates s_t into d_t and communicates d_t as $m^{P \to D}$ to A^D (Fig. 3a, right).

On this pass-decision dataset, all MR-LFADS models achieved comparably high-fidelity data reconstruction, slightly outperforming MR-SDS and mp-srSLDS (Fig. 3b, Fig. S2a). RRR reconstructed the data poorly, likely due to its inability to capture the timescale of the integration computation, i.e., d_t cannot be predicted from any single-timestep value $s_{t'}$.

MR-LFADS(R) inferred the most accurate effectome (Fig. 3c, Fig. S2b), identifying $A^P \rightarrow A^D$ communication but not $A^D \rightarrow A^P$. By contrast, all other models, including MR-LFADS(F) and MR-LFADS(G), identified spurious



Figure 3. Experiment 2. (a) Left: DGN setup, with stimulus (red), communication (blue), trained readouts (green), and computations (I: identity; \int : integration). Right: Potential failure mode for a learned model. (b) R^2 scores for data reconstruction. (c) Left: Cosine similarity between inferred effectomes and ground truth. Right: Example fitted models. Color intensity indicates relative message norm. (d) R^2 of linear prediction of ground truth input s to region P, from inferred inputs. (e) Left: R^2 when predicting ground truth messages $m^{P \to D} = s$ from inferred messages, $\mu_m^{P \to D}$, Right: R^2 when predicting the decision variable d from $\mu_m^{P \to D}$, indicating mislocalization of integration.

$A^D \rightarrow A^P$ communication.

All MR-LFADS variants correctly inferred inputs to A^P that encoded s_t (Fig. 3d). Because MR-SDS and mp-srSLDS do not infer unsupervised inputs, we estimated their effective inputs by passing the manually specified inputs through their corresponding trained input mappings. In MR-SDS, these effective inputs to A^P carried markedly less information about s_t relative to MR-LFADS and mp-srSLDS.

Accurate identification of communication requires inferred messages $m_t^{P \to D}$ to encode the stimulus s_t . By contrast, $m_t^{P \to D}$ instead encoding the decision variable d_t would imply mislocalization of the integration dynamic (Fig. 3a, right). Only MR-LFADS(R), mp-srSLDS, and RRR correctly encoded s_t in $m_t^{P \to D}$ (Fig. 3e, left) without incorrectly encoding d_t (Fig. 3e, right).

Taken together, these Experiment 2 results again demonstrate MR-LFADS(R) outperforming existing methods at disentangling external inputs, communication, and local population dynamics, without sacrificing data reconstruction. The comparisons to MR-LFADS(F) and MR-LFADS(G) specifically highlight the importance of MR-LFADS(R)'s data-constrained communication, which improves identification of communication by reducing ambiguity inherent to overly flexible models. See Appendix E.2 for further analyses into this excessive flexibility.

4.3. Experiment 3: Generalization Across Random Multi-Region Networks

The previous experiments utilized datasets synthesized by DGNs designed to highlight specific failure modes of communication models. However, real brain-wide networks exhibit a broad range of architectures and computations. To assess generalization across such a broad range of settings, here we evaluate communication models on datasets generated by a wide variety of randomly configured multi-region DGNs (Fig. 4a, top), each trained to perform a randomly selected cognitive neuroscience task (Fig. 4a, bottom). We generated 35 multi-region datasets, each from a unique tasktrained DGN consisting of three or four regions and randomized inter-regional connectivity. Tasks were drawn from the set described by Yang et al. (2019), spanning multiple variants of decision-making, working memory, categorization, and inhibitory control (see Appendix D).



Figure 4. Experiment 3. (a) Top: Each DGN is configured with a random number of areas (3 or 4), random inter-regional connectivity, and is trained on a randomly selected task. Bottom: Example task setup. On each trial, the DGNs received a noisy fixation stimulus $s_{\text{fix},t}$ and two noisy task stimuli, represented in polar coordinates as $s_t^k = (a_t^{(k)}, \theta_t^{(k)})$ for $k \in \{1, 2\}$. DGNs were trained to process these inputs into task-dependent outputs: a response angle θ^{resp} and a task-dependent saccadic eye movement. Here, multisensory decision-making is depicted as an example task. When the fixation cue disappears, the output area A^N must saccade and report the $\theta^{(i)}$ value corresponding to the stimulus with larger $a^{(i)}$. (b) R^2 scores for data reconstruction. Box plots describe distributions of values across 35 DGN fits. (c) Cosine similarity (S_{cos}) of inferred effectomes relative to ground truth, compared to that of MR-LFADS(R), with $\Delta S_{cos} = S_{cos}^{model} - S_{cos}^{MR-LFADS(R)}$. One-tailed t-test p-values: p = 0.00016, 0.12, 0.006, 0.0, 0.01. NS: not significant. (d) Left: R^2 scores, relative to MR-LFADS(R), for linear prediction of ground truth messages from inferred messages. One-tailed t-test p-values: p = 0.0, 0.0, 0.001, 0.0, 0.14. Right: R^2 scores, relative to MR-LFADS(R), for linear prediction of ground truth messages from models fit to datasets from a three-region DGN (top) and a four-region DGN (bottom). We chose the datasets on which MR-LFADS(R) achieved its median S_{cos} scores (across the 35 generated datasets). Color intensity indicates relative message norm.

We fit MR-LFADS(R), (F), and (G), along with MR-SDS, mp-srSLDS, and RRR, to each of these datasets. Aggregating results across all datasets, the MR-LFADS models achieved the best data reconstruction, which was indistinguishable across model variants (Fig. 4b, Fig. S3). MR-LFADS(R) and MR-LFADS(G) inferred the most accurate effectomes, with statistically indistinguishable S_{cos} distributions (Fig. 4c). To evaluate the accuracy of inferred message content, we attempted to linearly decode the ground truth messages from the inferred messages and quantified accuracy using the R^2 scores of these predictions (Fig. 4d, left). We also performed the reverse, predicting the inferred messages from the ground truth and interpreted lower R^2 scores as an indication that inferred messages contained additional information beyond that present in the ground truth (Fig. 4d, right). MR-LFADS(R) was the only model that performed best across both of these metrics. Inferred effectomes from two example datasets are shown in Fig. 4e. Taken together, these results demonstrate that MR-LFADS(R) outperforms existing communication models across a broad range of neuroscience-relevant synthetic multi-region datasets.

5. Results II: Multi-Region Electrophysiology

Here, we apply MR-LFADS(R) to large-scale electrophysiological data from multiple simultaneously recorded Neuropixel probes in mice performing a decision-making task (Fig. 5a) (Chen et al., 2024). We test whether MR-LFADS can predict the effects of causal circuit perturbations that were held out from training. We also assess the reliability of inferred inter-regional communication across random initializations, comparing MR-LFADS(R) to MR-LFADS(G).

We trained a 5-region MR-LFADS(R) model on simultaneously recorded population activity from anterior lateral motor cortex (**ALM**), midbrain reticular nucleus (**MRN**), superior colliculus (**SC**), thalamic regions with strong, reciprocal connections to ALM (**Thal(A**)), and other thalamic regions (**Thal(O**)). In a subset of trials excluded from training, ALM was transiently photoinhibited (Fig. 5a). MR-LFADS was trained only on unperturbed ("control") trials, with photoinhibition trials held out for validation. See Appendix F for further detail.

We first verified that MR-LFADS accurately reconstructed multi-regional neural activity in held-out control trials (Fig. 5b). We then adapted MR-LFADS(R) to predict the effects of ALM photoinhibition on the remaining recorded regions (Fig. 5c). To mimic ALM photoinhibition in silico, we ablated MR-LFADS communication from ALM to the other regions by zeroing all outgoing messages from ALM. We summarized the temporal influence of ALM photoinhibition by computing the differences between conditionaveraged population activity in photoinhibition (\bar{r}_t^{\dagger}) and control (\bar{r}_t^c) trials. MR-LFADS predicted these photoinhibition effects (Fig. 5d), despite never seeing photoinhibition trials during training. Namely, MR-LFADS predicted that Thal(A) would be most affected by ALM photoinhibition, MRN least affected, and SC and Thal(O) intermediately affected. These results demonstrate that MR-LFADS learned



Figure 5. MR-LFADS(R) applied to multi-region, high-density electrophysiology. (a) Mice receive a high- or low-tone auditory stimulus ("sample") and respond by licking left or right ("response"). (b) MR-LFADS single-trial predicted firing rates in held-out control trials (blue), recorded spike times (black vertical ticks), and smoothed, binned spike counts (black; causal, exponential filter) for example neurons in each modeled brain region. (c) Condition-averaged smoothed spike counts of example neurons in control and photoinhibition trials (left-hit condition). Shaded regions indicate the standard deviation across trials. (d) Photoinhibition-related changes in population recordings, as observed experimentally (left) and as predicted by MR-LFADS (right). (e) Cosine similarity of inferred effectomes across models with different random initializations (seeds). (f) Message norms inferred by MR-LFADS(R) for all connections across example seeds k. (g) Correlation of inferred message norms across pairs of seeds (k, l). Box plots describe distributions of values across 5 models fit from distinct random initializations.

a model of inter-regional communication that is accurate enough to predict multi-regional effects of causal circuit perturbations.

Finally, we evaluated the consistency of MR-LFADS across random initializations of the model parameters. We compared MR-LFADS(R) and MR-LFADS(G), each trained using five different random seeds, on simultaneous population recordings from ALM, thalamus, and MRN (see Appendix F2). Both models inferred consistent effectomes (Fig. 5e), but message content was more consistent across seeds in MR-LFADS(R) (Fig. 5f, g), further highlighting the benefits of data-constrained communication—particularly for improving the reproducibility of scientific conclusions derived from the model.

6. Discussion

Understanding how brain regions interact to support distributed computation requires communication models that can disentangle inter-regional communication from local population dynamics, accounting for region-specific inputs from unrecorded regions. In this work, we identified critical failure modes that limit existing communication modelsincluding misidentification due to manually specified external inputs and mislocalization of neural dynamics due to overly flexible communication architectures. We introduced MR-LFADS, a communication model specifically designed to mitigate such failures through three key design features: (1) automatic inference of region-specific inputs from unobserved sources, (2) data-constrained communication inferred from reconstructed firing rates, and (3) structured regularization that promotes disentangling and prevents inferring inputs from unobserved sources when the same information can be obtained via communication from observed regions. These features discourage the model from learning spurious solutions that explain the data but misrepresent inter-regional interactions. While MR-LFADS is one concrete implementation, the three design principles are more general, and alternative architectures-for example, different formulations of the local-region dynamical systems-may also succeed if they incorporate these same principles.

Using synthetic datasets designed to rigorously test communication models, we demonstrated that MR-LFADS outperforms existing approaches in accurately recovering both the structure and content of inter-regional communication. Crucially, ablated model variants indicated that these performance gains stem directly from MR-LFADS design features. Applying MR-LFADS to real multi-region electrophysiological recordings further validated its utility. MR-LFADS inferred inter-regional interactions that accurately predicted brain-wide effects of causal perturbations, despite these perturbations being absent during training. In this setting, MR-LFADS models were also more reproducible across random model initializations compared to a model variant that removed data-constraints on inferred communication.

Despite its advantages, MR-LFADS has potential limitations. MR-LFADS can be sensitive to hyperparameter (HP) settings, mirroring a known limitation of SR-LFADS (Keshtkaran et al., 2022). Thus, significant computational resources might be required to adequately optimize HPs. As in SR-LFADS, HPs specifying the prior distributions (Eq. 1) can shape the representations of inferred latent variables, which can in turn shape the inferred population dynamics. Although inferred inputs have been shown to reflect the presence, identity, and timing of external inputs (Pandarinath et al., 2018), future work is needed to interpret the representations of those signals (Sedler et al., 2023; Versteeg et al., 2024). The data-constrained communication we propose in this work might mitigate this representational sensitivity in the context of inferred communication.

Another set of MR-LFADS HPs (the β 's from Section 3) control the model's preference to infer communication rather than inputs from unobserved regions whenever possible. While the settings we chose were effective in this study, different choices may be needed in other scenarios—especially when an unrecorded region directly drives multiple recorded regions. This classical correlation–causation confound cannot be resolved from passively observed neural activity alone. While MR-LFADS recovers more causal structure than existing methods, its greatest value may be in generating hypotheses about communication and motivating targeted causal circuit perturbations to test them.

Acknowledgments

This work was supported by NIH award T32-MH132518 (JS), the Paul G. Allen Foundation (MDG), and NIH award R00-MH121533 (MDG). We thank O. Karniol-Tambour for valuable discussion and support with MR-SDS (Karniol-Tambour et al., 2024). We are grateful to S. Chen, N. Steinmetz, E. Shea-Brown, and J.N. Kutz for insightful input on the project, and to H. Gurnani, T. Kim, and J. Pemberton for thoughtful feedback on the manuscript. We are also grateful to D. Sussillo, K. Shenoy, W. Newsome, and C. Pandarinath for many years of guidance and inspiration relevant to this project.

Impact Statement

This work aims to advance the fields of machine learning and neuroscience, with potential societal impacts including the development of neuroengineering technologies and treatments for neurological injuries, diseases, and neuropsychiatric conditions.

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- Allen, W. E., Kauvar, I. V., Chen, M. Z., Richman, E. B., Yang, S. J., Chan, K., Gradinaru, V., Deverman, B. E., Luo, L., and Deisseroth, K. Global representations of goal-directed behavior in distinct cell types of mouse neocortex. *Neuron*, 94(4):891–907, 2017.
- Allen, W. E., Chen, M. Z., Pichamoorthy, N., Tien, R. H., Pachitariu, M., Luo, L., and Deisseroth, K. Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*, 364(6437):eaav3932, 2019.
- Bennett, C., Ouellette, B., Ramirez, T. K., Cahoon, A., Cabasco, H., Browning, Y., Lakunina, A., Lynch, G. F., McBride, E. G., Belski, H., et al. Shield: Skull-shaped hemispheric implants enabling large-scale electrophysiology datasets in the mouse brain. *Neuron*, 112(17): 2869–2885, 2024.
- Biswas, T., Bishop, W. E., and Fitzgerald, J. E. Theoretical principles for illuminating sensorimotor processing with brain-wide neuronal recordings. *Current Opinion in Neurobiology*, 65:138–145, 2020.
- Chen, S., Liu, Y., Wang, Z. A., Colonell, J., Liu, L. D., Hou, H., Tien, N.-W., Wang, T., Harris, T., Druckmann, S., et al. Brain-wide neural activity underlying memoryguided movement. *Cell*, 187(3):676–691, 2024.
- Dai, B., Wang, Y., Aston, J., Hua, G., and Wipf, D. Connections with robust pca and the role of emergent sparsity in variational autoencoder models. *Journal of Machine Learning Research*, 19(41):1–42, 2018.
- Duncker, L. and Sahani, M. Dynamics on the manifold: Identifying computational dynamical activity from neural population recordings. *Current Opinion in Neurobiology*, 70:163–170, 2021.
- Durstewitz, D., Koppe, G., and Thurm, M. I. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24(11):693–710, 2023.
- Gilad, A., Gallero-Salas, Y., Groos, D., and Helmchen, F. Behavioral strategy determines frontal or posterior location of short-term memory in neocortex. *Neuron*, 99 (4):814–828, 2018.

- Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L., and Linderman, S. Recurrent switching dynamical systems models for multiple interacting neural populations. *Advances in Neural Information Processing Systems*, 33: 14867–14878, 2020.
- Gokcen, E., Jasper, A. I., Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., and Yu, B. M. Disentangling the flow of signals between populations of neurons. *Nature Computational Science*, 2(8):512–525, 2022.
- Gokcen, E., Jasper, A., Xu, A., Kohn, A., Machens, C. K., and Yu, B. M. Uncovering motifs of concurrent signaling across multiple neuronal populations. *Advances in Neural Information Processing Systems*, 36, 2024.
- Goris, R. L., Movshon, J. A., and Simoncelli, E. P. Partitioning neuronal variability. *Nature Neuroscience*, 17(6): 858–865, 2014.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. betavae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- IBL, Benson, B., Benson, J., Birman, D., Bonacchi, N., Bougrova, K., Bruijns, S. A., Carandini, M., Catarino, J. A., Chapuis, G. A., et al. A brain-wide map of neural activity during complex behaviour. *bioRxiv preprint*, 2023.
- Jia, X., Siegle, J. H., Durand, S., Heller, G., Ramirez, T. K., Koch, C., and Olsen, S. R. Multi-regional module-based signal transmission in mouse visual cortex. *Neuron*, 110 (9):1585–1598, 2022.
- Kang, B. and Druckmann, S. Approaches to inferring multi-regional interactions from simultaneous population recordings. *Current Opinion in Neurobiology*, 65: 108–119, 2020.
- Karniol-Tambour, O., Zoltowski, D. M., Diamanti, E. M., Pinto, L., Brody, C. D., Tank, D. W., and Pillow, J. W. Modeling state-dependent communication between brain regions with switching nonlinear dynamical systems. In *The Twelfth International Conference on Learning Representations*, 2024.
- Kass, R. E., Bong, H., Olarinre, M., Xin, Q., and Urban, K. N. Identification of interacting neural populations: methods and statistical considerations. *Journal of Neurophysiology*, 130(3):475–496, 2023.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I., and Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17(3): 440–448, 2014.

- Keeley, S. L., Zoltowski, D. M., Aoi, M. C., and Pillow, J. W. Modeling statistical dependencies in multi-region spike train data. *Current Opinion in Neurobiology*, 65: 194–202, 2020.
- Keshtkaran, M. R., Sedler, A. R., Chowdhury, R. H., Tandon, R., Basrai, D., Nguyen, S. L., Sohn, H., Jazayeri, M., Miller, L. E., and Pandarinath, C. A large-scale neural network training framework for generalized estimation of single-trial population dynamics. *Nature Methods*, 19 (12):1572–1577, 2022.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. arXiv preprint, 1312.6114, 2013.
- Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J. E., and Stoica, I. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.
- Linderman, S. W., Miller, A. C., Adams, R. P., Blei, D. M., Paninski, L., and Johnson, M. J. Recurrent switching linear dynamical systems. *arXiv preprint*, 1610.08466, 2016.
- MacDowell, C. J., Libby, A., Jahn, C. I., Tafazoli, S., Ardalan, A., and Buschman, T. J. Multiplexed subspaces route neural activity across brain-wide networks. *Nature Communications*, 16(1):3359, 2025.
- Makino, H., Ren, C., Liu, H., Kim, A. N., Kondapaneni, N., Liu, X., Kuzum, D., and Komiyama, T. Transformation of cortex-wide emergent properties during motor learning. *Neuron*, 94(4):880–890, 2017.
- Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013.
- Miller, K., Eckstein, M., Botvinick, M., and Kurth-Nelson, Z. Cognitive model discovery via disentangled rnns. Advances in Neural Information Processing Systems, 36, 2024.
- Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., and Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, 22(10):1677–1686, 2019.
- Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., et al. Inferring singletrial neural population dynamics using sequential autoencoders. *Nature Methods*, 15(10):805–815, 2018.
- Perich, M. G. and Rajan, K. Rethinking brain-wide interactions through multi-region 'network of networks' models. *Current Opinion in Neurobiology*, 65:146–151, 2020.

- Perich, M. G., Gallego, J. A., and Miller, L. E. A neural population mechanism for rapid learning. *Neuron*, 100 (4):964–976, 2018.
- Perich, M. G., Arlt, C., Soares, S., Young, M. E., Mosher, C. P., Minxha, J., Carter, E., Rutishauser, U., Rudebeck, P. H., Harvey, C. D., et al. Inferring brain-wide interactions using data-constrained recurrent neural network models. *bioRxiv preprint*, 2020.
- Pospisil, D. A., Aragon, M. J., Dorkenwald, S., Matsliah, A., Sterling, A. R., Schlegel, P., Yu, S.-c., McKellar, C. E., Costa, M., Eichler, K., et al. The fly connectome reveals a path to the effectome. *Nature*, 634(8032):201–209, 2024.
- Raposo, D., Kaufman, M. T., and Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nature neuroscience*, 17(12):1784–1792, 2014.
- Ruff, D. A. and Cohen, M. R. Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nature Neuroscience*, 22(10):1669–1676, 2019.
- Sedler, A. R., Versteeg, C., and Pandarinath, C. Expressive architectures enhance interpretability of dynamics-based neural population models. *Neurons, Behavior, Data Analysis, and Theory*, 2023, 2023.
- Semedo, J. D., Zandvakili, A., Machens, C. K., Byron, M. Y., and Kohn, A. Cortical areas interact through a communication subspace. *Neuron*, 102(1):249–259, 2019.
- Semedo, J. D., Gokcen, E., Machens, C. K., Kohn, A., and Byron, M. Y. Statistical methods for dissecting interactions between brain areas. *Current Opinion in Neurobiol*ogy, 65:59–69, 2020.
- Shenoy, K. V., Sahani, M., and Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annual Review of Neuroscience*, 36(1):337–359, 2013.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, 592(7852):86–92, 2021.
- Sofroniew, N. J., Flickinger, D., King, J., and Svoboda, K. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *eLife*, 5:e14472, 2016.

- Song, A., Charles, A. S., Koay, S. A., Gauthier, J. L., Thiberge, S. Y., Pillow, J. W., and Tank, D. W. Volumetric two-photon imaging of neurons using stereoscopy (vTwINS). *Nature Methods*, 14(4):420–426, 2017.
- Steinmetz, N. A., Zatka-Haas, P., Carandini, M., and Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–273, 2019.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364 (6437):eaav7893, 2019.
- Versteeg, C., Sedler, A. R., McCart, J. D., and Pandarinath, C. Expressive dynamics models with nonlinear injective readouts enable reliable recovery of latent features from neural activity. In *Proceedings of the 2nd NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, volume 228 of *Proceedings of Machine Learning Research*, pp. 255–278. PMLR, 2024.
- Veuthey, T., Derosier, K., Kondapavulur, S., and Ganguly, K. Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, 11(1):4057, 2020.
- Vyas, S., Golub, M. D., Sussillo, D., and Shenoy, K. V. Computation through neural population dynamics. *Annual Review of Neuroscience*, 43(1):249–275, 2020.
- Watanabe, S. Tree-structured parzen estimator: Understanding its algorithm components and their roles for better empirical performance. *arXiv preprint arXiv:2304.11127*, 2023.
- Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., and Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, 22(2):297–306, 2019.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S., Shenoy, K. V., and Sahani, M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in Neural Information Processing Systems*, 21, 2008.
- Zhu, F., Grier, H. A., Tandon, R., Cai, C., Agarwal, A., Giovannucci, A., Kaufman, M. T., and Pandarinath, C. A deep learning framework for inference of single-trial neural population dynamics from calcium imaging with subframe temporal resolution. *Nature Neuroscience*, 25 (12):1724–1734, 2022.

A. Models

A.1. Multi-Region LFADS (MR-LFADS)

Reconstruction validation. During training, we held out 10% of the neurons in each real or synthetic brain region for validation. After training MR-LFADS, we fit a separate linear decoder *post hoc* on the training trials, regressing the inferred factors onto the held-out neurons' activity. Since the decoder is not trained end-to-end with MR-LFADS, held-out neurons do not influence model fitting. Goodness of fit is reported as the R^2 between predicted and observed held-out-neuron activity on validation trials.

KL penalty. A critical hyperparameter during training is the scale of the KL penalty. The KL penalty coefficient for the inferred inputs, β_u , is always set higher than that for communication, β_m , an implicit assumption that encourages the network to prioritize learning information via communication channels whenever possible. The timing of when KL penalties are introduced also impacts results, though to a lesser extent. The schedule we found to work well begins with an initial stage where no KL penalty is applied, allowing the model to overfit to the data. Next, the penalty for inferred inputs is introduced, discouraging the model's reliance on these inputs. Finally, the penalty for communication is added, limiting the model from learning excessive information through communication channels.

Weight regularization. We apply light L_2 regularization to all GRU network recurrent weights.

Hyperparameters. Table S1 summarizes key hyperparameters used in the MR-LFADS models. Overall, we find that KL coefficients have the most impact on held-out neuron loss, S_{cos} , and R^2 scores for messages compared to other hyperparameters. SR-LFADS was originally described with a factor layer that is potentially lower dimensional than the number of generator units or modeled neurons. Here, we remove the rank constraint from the generator hidden states to rates by setting $N_{fac} = N_{neu}$. Additionally, the inferred input and message channel dimensions only need to exceed the estimated true dimensionality of these quantities, as KL penalties naturally suppress redundant channels by driving their activity to zero, as discussed in Appendix E.1.

Hyperparameter	value	Description		
learning rate	$\in [10^{-5}, 0.004]$	Scheduled by PyTorch's ReduceLROnPlateau; initial value: 0.004		
Т	190	Total time steps used for inferring inferred inputs		
τ	10	Total time steps used for inferring the initial condition		
total epoch	350	Total number of epochs		
β_u		KL penalty coefficient for <i>u</i> ; performs search for this hyperparameter		
β_u start epoch	50	Epoch at which β_u starts increasing from 0		
β_u increase epoch	200	Number of epochs for β_u to reach the maximum value		
β_m		KL penalty coefficient for m; performs search for this hyperparameter		
β_m start epoch	150	Epoch at which β_m starts increasing from 0		
β_m increase epoch	100	Number of epochs for β_m to reach the maximum value		
β_{g_0}	β_u	KL penalty coefficient for g_0		
α	10^{4}	L2 penalty coefficient		
α start epoch	0	Epoch at which α starts increasing from 0		
α increase epoch	80	Number of epochs for α to reach the maximum value		
N ⁱ _{neu}		Number of neurons, (64, 16, 64) for exp 1, 2 and 3 respectively		
$N_{\rm gen}^i$	$2N_{ m Neu}^i$	Generator size		
$N_{\rm fac}^i$	$N^i_{ m Neu}$	Factor size		
$N_{\rm inp}^i$		Inferred input dimension, (4, 8, 6, 8) for exp 1, 2, 3 and 4 respectively		
N ⁱ _{msg}		Inferred message dimension, (4, 16, 10, 8) for exp 1, 2, 3 and 4 respectively		

Table S1. Key hyperparameters for MR-LFA	OS models. Experiment 4 refers to applications	s to multi-region electrophysiology data.

Unidirectional Encoder and Controller. To ensure that inferred inputs reflect only causal information—so that messages, in turn, are causal and thereby allow a more mechanistic interpretation in MR-LFADS—we modify the original LFADS model so that both the encoder and controller used for input inference are entirely unidirectional:

$$e_{t}^{i} = \text{GRU}_{\text{enc},u}^{i}(e_{t-1}^{i}, x_{t}^{i})$$

$$e_{t}^{i} = \text{GRU}_{\text{con}}^{i}(e_{t-1}^{i}, [e_{t}^{i}; f_{t-1}^{i}])$$
(12)

The inferred inputs are then given by:

$$q(u_{t}^{i} \mid x_{1:t}^{i}) = q(u_{t}^{i} \mid c_{t}^{i}) = \mathcal{N}(\mu_{u,t}^{i}, \Sigma_{u,t}^{i})$$

$$\mu_{u,t}^{i} = W_{\mu_{u}}^{i}(c_{t}^{i}) \quad \Sigma_{u,t}^{i} = \text{diag}\Big(\exp\left(W_{\sigma_{u}}^{i}(c_{t}^{i})\right)\Big)$$
(13)

By contrast, the encoder for the initial condition can remain bidirectional, since it operates on data preceding t = 1 and thus does not violate causality:

$$e_t^{i,-} = \text{GRU}_{\text{enc},g}^{i,-}(e_{t+1}^{i,-}, x_t^i)$$

$$e_t^{i,+} = \text{GRU}_{\text{enc},g}^{i,+}(e_{t-1}^{i,+}, x_t^i)$$
(14)

$$q(g_{0}^{i} \mid x_{-\tau:0}^{i}) = q(g_{0}^{i} \mid [e_{-\tau}^{i,-}; e_{0}^{i,+}]) = \mathcal{N}(\mu_{g_{0}}^{i}, \Sigma_{g_{0}}^{i})$$

$$\mu_{g_{0}}^{i} = W_{\mu_{g_{0}}}^{i}([e_{-\tau}^{i,-}; e_{0}^{i,+}])$$

$$\Sigma_{g_{0}}^{i} = \operatorname{diag}\left(\exp\left(W_{\sigma_{g_{0}}}^{i}([e_{-\tau}^{i,-}; e_{0}^{i,+}])\right)\right)$$
(15)

A.2. Reduced-Rank Regression

The RRR model used in this study is based on the inter-regional communication subspace model of MacDowell et al. (2025), which integrates reduced-rank regression with ridge regression. The key difference in our implementation is that we apply the rank constraint separately to each source brain region, allowing us to disentangle the contributions of individual areas. Specifically, rather than concatenating activity from all regions into a single input matrix governed by a shared rank-constrained weight matrix—which conflates signals across regions—we assign a dedicated weight submatrix to each source region A^j , each with its own rank constraint r^j .

Therefore, for the communication subspace model, we have:

$$W^{\text{rrr}} = \underset{\text{rank}(W)=r}{\arg\min} \|Y - XW\|_{F}^{2} + \alpha \|W\|_{F}^{2}$$
which is equivalent to:
$$W^{\text{ridge}} = \arg\min_{W} \|Y - XW\|_{F}^{2} + \alpha \|W\|_{F}^{2}$$

$$W^{\text{rrr}} = \arg\min_{W} \|Y - XW^{\text{ridge}}\|_{F}^{2}$$

$$+ \|XW^{\text{ridge}} - XW\|_{F}^{2}$$
(16)

which is then equivalent to:

$$W^{\text{rrr}} = W^{\text{ridge}} V_r V_r^T$$
, where $U \Sigma V^T = X W^{\text{ridge}}$

where $X \in \mathbb{R}^{T \times N_{sre}}$ represents the activity of all source regions concatenated together, and $Y \in \mathbb{R}^{T \times N_{tar}}$ represents the activity of the target region. W^{ridge} is the weight matrix obtained after applying ridge regression, and α is the ridge regularization parameter. W^{rrr} is the reduced-rank regression matrix obtained after ridge regression is applied. In the final step, singular value decomposition (SVD) is applied to XW^{ridge} , where $U\Sigma V^T$ is the decomposition, and V_r corresponds to the top r components of V.

In this version, the W^{ridge} matrix is divided into chunks corresponding to different regions A^{j} :

$$W^{\text{ridge}} = [W^{\text{ridge},1}; W^{\text{ridge},2}; \dots; W^{\text{ridge},N}], \tag{17}$$

where each $W^{\text{ridge},j}$ corresponds to the contribution of source region A^j , and $[W^1; ...; W^N]$ represents a vertical stack of the matrices. SVD is then applied to each individual $W^{\text{ridge},j}$ matrix:

$$U^{j}\Sigma^{j}(V^{j})^{T} = W^{\operatorname{ridge},j}.$$
(18)

The reduced-rank version of $W^{ridge,j}$ is computed as:

$$W^{\mathrm{rrr},j} = W^{\mathrm{ridge},j} V^j_{ri} (V^j_{rj})^T.$$
⁽¹⁹⁾

Finally, all $W^{\operatorname{rrr},j}$ matrices are concatenated to form the complete W^{rrr} matrix. This ensures that the rank reduction applied to each $W^{\operatorname{ridge},j}$ only compresses the information within that specific region's contribution, preserving the interpretability of the communication pathway from A^j to the target region.

The hyperparameters for this model include α and a matrix $R \in \mathbb{R}^{N \times N}$, where each element r^{ij} represents the rank associated with the communication from source region A^j to target region A^i . Additionally, since time delays may exist between regions—and such delays are explicitly configured in the synthetic datasets—a delay parameter d^{ij} is introduced for each source-target communication channel.

For fitting the memory and pass-decision networks, to tune the model, we perform an iterative search based on cross-validation performance, optimizing the hyperparameters in the following order: $D = \{d^{ij} : i, j = 1, ..., N, i \neq j\}$, α , and $R = \{r^{ij} : i, j = 1, ..., N, i \neq j\}$. This process is repeated until the hyperparameter values converge. The final values used are provided in Table S2. While the rank values R for both networks did not converge exactly to the true ranks of the messages, they were close. Notably, providing the true number of latents did not necessarily lead to better results. The delay values D were accurately learned for both models.

For fitting the networks in Experiment 3, the true delay is directly provided, and other hyperparameters are iterated in the same order for 10 epochs.

To increase the robustness of the RRR model fit, we implemented a bagging approach. For each model, 10 trials were bootstrapped from the training set, with each trial containing 200 time steps. A total of 87 fitted models were averaged to obtain the matrix W^{rrr} . This specific number of models was chosen to ensure that the total number of trials used during training remained consistent with other models.

the set of				
	Memory	Pass-Decision		
α	0.055	0.01		
R	$ \begin{pmatrix} 12 & 24 \\ 12 & 32 \\ 24 & 18 \end{pmatrix} $	$\begin{pmatrix} & 1 \\ 6 & \end{pmatrix}$		
$d^{ij}, i \neq j$	2	0		

Table S2. Hyperparameter search results for RRR models.

A.3. Multi-Region Switching Dynamical Systems

We consider two variants of multi-region switching dynamical system models. The first is mp-srSLDS (Glaser et al., 2020), which consists of linear transitions, dynamics, and emissions. The relevant hyperparameters are the number of latent states per region, the number of discrete switching states, amount of L_1 and L_2 regularization on the weights, and the learning rate. Additionally, we consider MR-SDS (Karniol-Tambour et al., 2024), which is an extension that uses nonlinear transitions, dynamics, and emissions. It consists of two components: an inference network and a latent state-space model. The inference network is a transformer that performs the amortized inference of latent variables given observed neural activity. The latent state-space model is composed of a number of networks and functions as a structured prior on the latent variables. Specifically, we consider additive communication and input terms to the latent dynamics of the state-space model. That is, messages from other regions and external inputs affected the latent dynamics via additive terms. Relevant hyperparameters include the number of latent states per region, the number of discrete switching states, and the sizes of each sub-network.

For both models, we did extensive hyperparameter tuning to find the best model for each of the synthetic datasets and then computed all metrics on a held-out test set. We used the Tree-structured Parzen Estimator (TPE) algorithm (Watanabe, 2023)

with the Optuna backend (Akiba et al., 2019) in Ray Tune (Liaw et al., 2018). The algorithm fits two Gaussian mixture models (GMMs), one to the set of parameter values associated with the best objective and another to the remaining ones. It chooses new parameters to explore by maximizing the ratio of the likelihood between these two GMMs. As such, it is a search strategy which uses results from prior tested hyperparameters to inform the next choice of hyperparameters to test. We used the TPE algorithm to search over all relevant hyperparameters above. Additionally, for MR-SDS, we used dropout for regularization with the default settings in the provided implementation. We also manually picked a good learning rate and number of epochs for training. Finally, we also made use of co-smoothing for evaluation. This holds out a set of neurons from the inference network, and computes the fit on the reconstruction of these held-out neurons.

A.4. Model Comparisons

We outline the key design features of MR-LFADS variants and existing communication models in Table S3. Models are compared across four criteria: (1) region-specific dynamics, (2) unsupervised inferred inputs, (3) data-constrained communication, and (4) structured information bottlenecks. Only MR-LFADS(R) incorporates all four features.

MR-LFADS(S) ablates the controller, removing inferred inputs and instead using manually specified external inputs for each region. MR-LFADS(F) and MR-LFADS(G) both communicate via latent variables not directly grounded in observed data—using factors and generator states, respectively.

MR-SDS and mp-srSLDS include region-specific dynamics but rely on external inputs and latent-variable-based messaging, without any regularization on inputs or communication. RRR infers communication from observable quantities but lacks dynamics and inputs altogether. While it enforces a rank constraint on the communication subspace, this is not equivalent to explicit regularization on messages.

	U		U	
	Region-Specific	Unsupervised	Data-Constrained	Structured Information
	Dynamics	Inferred Inputs	Communication	Bottlenecks
MR-LFADS(R)	\checkmark	\checkmark	\checkmark	\checkmark
MR-LFADS(S)	✓	×	\checkmark	×
MR-LFADS(F)	✓	✓	×	\checkmark
MR-LFADS(G)	✓	✓	×	\checkmark
MR-SDS	✓	×	×	×
mp-srSLDS	✓	×	×	×
RRR	X	X	\checkmark	×

Table S3. Design features of MR-LFADS variants and existing communication models.

B. Evaluation Metrics

B.1. Quantifying Effectome Similarity

In a trained MR-LFADS model, we define the inferred effectome to be a matrix of pairwise message norms, M, with element $M_{i,j}$ as the average value of $||\mu_{m,t}^{j\to i}||_2$ (Eq. 7) across all trials and timesteps. In Experiments 1-2, we compared model-inferred effectomes to the corresponding ground truth connectivity matrix M_{true} , consisting of ones and zeros to indicate the presence or absence of a communication channel in the DGN, respectively. In contrast, Experiment 3 features networks in which not all connections are actively used; in this case, we define M_{true} analogously to M, but computed using ground truth messages $m^{j\to i}$ instead of inferred messages $\mu_{m,t}^{j\to i}$. To assess similarity between M and M_{true} , we flatten these matrices into vectors— \vec{m} and \vec{m}_{true} —and compute their cosine similarity:

$$S_{\rm cos} = \frac{\langle \vec{m}, \vec{m}_{\rm true} \rangle}{\|\vec{m}\|_2 \cdot \|\vec{m}_{\rm true}\|_2} \in [0, 1],\tag{20}$$

where perfect alignment is indicated when $S_{cos} = 1$.

To visualize an inferred effectome (e.g., Fig. 2c, right), we plot arrows whose color intensity reflects the relative message norm, computed by concatenating all multi-dimensional messages across trial and time, taking the 2-norm, and then

normalizing across all channels, with inferred communication and inputs normalized separately. The arrows in Fig. 4e, bottom are further scaled by a sigmoid function for visual contrast. All reported values elsewhere are computed without applying any thresholds or scaling, and heatmap visualizations of the effectome are also provided in Fig. S1b, Fig. S2b.

B.2. Evaluating Information Encoded in Learned Messages

In Experiment 1, we tested whether the inferred inputs and communications encoded information about the past ground truth values, as the network's hidden units activity contains information about past inputs. A correct model should only learn the ground truth inputs or communications. For the results shown in Fig. 2d-e, right, the r-squared values were calculated with a time lag d as $R^2(\mu_{m,t}^{j\to i}, m_{t-d}^{j\to i})$.

C. Synthetic Datasets for System Identification Issues

Synthetic datasets for networks from Experiments 1-3 have 1024, 1024 and 820 total trials respectively, of which 85% is used for training, and 15% for validation. The length of each trial is 200 time steps.

C.1. Memory Network

In this synthetic network, each region A^i is modeled as an GRU network with 64 units that receives a private stimulus $s_t^i \sim \mathcal{N}(0,\mathbb{I})$ with dimensions r^i . To simulate communication channels carrying different amounts of independent information, the dimensions are set as $(r^1, r^2, r^3) = (2, 3, 4)$. Each region has a linear readout W_{out} , and the outputs are required to encode information about the history of all inputs (i.e., stimuli and communication) for up to 5 time steps, enforced using a mean squared error loss. Similarly, the messages transmitted between regions are trained to match s_{t-2}^i and are also optimized via a mean squared error loss. Additionally, each region is subjected to dynamic noise $\xi \sim \mathcal{N}(0, 0.01 \, \mathbb{I})$, introduced as perturbations to the RNN activity at each time step.

After training the synthetic network, for MR-LFADS, we performed a hyperparameter sweep over the KL penalty coefficients for inferred inputs ($\beta_u \in \{0.01, 0.1, 1, 10\}$) and communication ($\beta_m \in \{0.001, 0.01\}$). The coefficient pair that resulted in the lowest held-out neuron loss, (β_u, β_m) = (0.1, 0.01), was selected. Using these optimized coefficients, we ran the MR-LFADS fit across 10 different seeds, which randomizes model initialization and subsequent sampling of inferred quantities during training, but does not change the allocation of training versus validation data.

Comparing Experiments 1 and 2, it is shown that mp-srSLDS and MR-SDS activity reconstruction performance underperforms compared to the MR-LFADS variants in Experiment 1 (Fig. 2b), but not in Experiment 2 (Fig. 3b). One possible explanation for this discrepancy is the amount of information that the latent variables must encode at each time step. For example, in area A^1 , the private stimulus is $s_t^1 \in \mathbb{R}^2$, and the incoming message from area A^3 is $m_t^{3\to 1} = s_{t-2}^3 \in \mathbb{R}^4$. As a result, the latent representation at time t must capture information spanning 5 time steps and 6 variables in total. Under a standard hyperparameter tuning scheme (Section A.3), latent variable models like mp-srSLDS and MR-SDS may struggle to represent all this information accurately.

C.2. Pass-Decision Network

In this synthetic network, each area is an GRU that contains 16 units. The stimulus s_t is two-dimensional and independently sampled from an exponential distribution with rate parameter of 3 time steps. Using a non-Gaussian distribution tests whether the model can learn the correct solution despite structural mismatches between the data and the model, as both the prior and posterior of inferred inputs and messages are Gaussian.

Each region has a linear readout W_{out} , whose output is required to match its corresponding latent variables (s_t for area A^P and d_t for A^D). The message sent from A^P to A^D is trained to represent s_t . Additionally, A^D must encode whether d_t is greater or less than 0 at all times, mimicking a binary decision-making process.

For the pass-decision network, a low KL penalty for inferred inputs ($\beta_u = 0.0075$) was necessary to achieve good held-out neuron loss, while the KL penalty for communication ($\beta_m = 0.001$) was set to be approximately one order of magnitude smaller. Using these coefficients, we ran the MR-LFADS fit across 10 different seeds.





Figure S1. Memory network. (a) Example traces of data reconstruction. *Row:* example neuron from an area, *column:* different trials. (b) Inferred effectomes visualized as heatmaps. *Top:* Inferred inputs, *bottom:* communication. Color intensity represents message norms normalized by the largest input $(\max_i ||\mu_u^i||_2)$ or message norm $(\max_{i,j} ||\mu_m^{j\to i}||_2)$ within each model. These largest inferred inputs are, from left to right: 76, 175, 60, 188, 0. The largest messages are: 195, 22, 95, 1431, 5. (c). R^2 of linear prediction of ground truth messages (with 2 time step lag on the left, 4 time step lag on the right) via inferred messages. MR-LFADS(R) results (left) are replicated from Fig. 2e.



Figure S2. Pass-decision network. (a) Example traces of data reconstruction. *Row:* example neuron from an area, *column:* different trials. (b) Inferred effectomes visualized as circuit diagrams and heatmaps. *Top:* circuit diagram, *middle:* inferred inputs, *bottom:* communication. Color intensity represents message norms normalized by the largest input or message norm within each model.

D. Randomly Generated Multi-Region Networks

We generated a distribution of networks designed to perform computational tasks inspired by Yang et al. (2019). Each network consists of either 3 or 4 regions. A connection probability $p \in \{0.5, 0.6, 0.7\}$ is specified, and the connectome is randomly drawn. To be considered valid, the connectome must meet two criteria: (1) each region must have at least one input connection and one output connection to ensure no region is redundant, and (2) all regions must be within a maximum distance of 2 steps from the output region. Once a valid connectome is generated, the network is trained on one of the computational tasks. During training, dynamic noise $\xi \sim \mathcal{N}(0, 0.01 \mathbb{I})$ is applied to all regions, and the only loss is based on whether the output region produces the correct response. Networks that meet the performance thresholds (train accuracy > 0.8 and validation accuracy > 0.6) are selected as synthetic datasets.

For this case, since we aim to collect a distribution of results, we do not perform a hyperparameter sweep over the KL penalties. Instead, we fit one instance of each communication model to each synthetic dataset using a single random seed. The computational tasks inspired by Yang et al. (2019) are described below.

All trials are 200 time steps in length. Each task receives a fixation input $s_{\text{fix},t}$, stimuli from two channels $s_t^{(1)} = (a^{(1)}, \theta^{(1)})$ and $s_t^{(2)} = (a^{(2)}, \theta^{(2)})$, and requires a saccade response r_t^{sacc} and an additional response $r_{\text{resp},t} = (1, \theta_{\text{resp}})$, which differ based on the specified task. For different tasks, stimuli may come from one or both of the channels. Both the stimuli and responses are expressed in polar coordinates, with a resolution of 10 degrees per angle.

The tasks are described in terms of 3 different families: the go task family, context-dependent decision-making family, and matching family. For all tasks, $t_{\text{start}} \in [30, 50)$ denotes the onset of the first (and sometimes only) stimulus (Table S4).

Task	Δ_{offset}	Δ_{delay}	t _{sacc}	$ heta^{\mathbf{resp}}$
Go	N/A	∞	$t_{ m start} + \Delta_{ m dur}$	$ heta^{(i)}$
Anti-Go	N/A	∞	$t_{ m start}+\Delta_{ m dur}$	$\pi + \theta^{(i)}$
Delay-Go	N/A	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m delay}$	$ heta^{(i)}$
Delay-Anti-Go	N/A	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m delay}$	$\pi + \theta^{(i)}$
DM1	0	[30, 50)	$t_{ m start} + \Delta_{ m dur}$	$ heta^{(1)}$
DM2	0	[30, 50)	$t_{ m start} + \Delta_{ m dur}$	$ heta^{(2)}$
MultSen DM	0	[30, 50)	$t_{ m start} + \Delta_{ m dur}$	$\theta^{(i)}, \ i = \arg \max_i r^{(i)}$
Delay-DM1	[10, 20)	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m offset}$	$ heta^{(1)}$
Delay-DM2	[10, 20)	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m offset}$	$ heta^{(2)}$
Delay MultSen DM	[10, 20)	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m offset}$	$\theta^{(i)}, \ i = \arg \max_i r^{(i)}$
Angle	[10, 20)	[30, 50)	$t_{\text{start}} + \Delta_{\text{dur}} + \Delta_{\text{offset}} \text{ if } \theta^{(1)} = \theta^{(2)}$	$ heta^{(2)}$
Anti-Angle	[10, 20)	[30, 50)	$t_{ m start} + \Delta_{ m dur} + \Delta_{ m offset} ext{ if } heta^{(1)} = heta^{(2)}$	$\pi + \theta^{(2)}$
Category	[10, 20)	[30, 50)	$t_{\text{start}} + \Delta_{\text{dur}} + \Delta_{\text{offset}} \text{ if } \operatorname{sign}(\theta^{(1)}) = \operatorname{sign}(\theta^{(2)})$	$ heta^{(2)}$
Anti-Category	[10, 20)	[30, 50)	$t_{\text{start}} + \Delta_{\text{dur}} + \Delta_{\text{offset}} \text{ if } \operatorname{sign}(\theta^{(1)}) = \operatorname{sign}(\theta^{(2)})$	$\pi + \theta^{(2)}$

Table S4. Task parameters for all families in multi-region data generating networks.

The duration of all stimulus pulses in a trial is represented by $\Delta_{dur} \in [30, 50)$, and Δ_{offset} specifies the offset between the two stimuli, if applicable. The time between the last stimulus offset and the fixation cue offset is given by Δ_{delay} , where $\Delta_{delay} = \infty$ indicates that the fixation cue never disappears. The onset of the saccade is denoted as t_{sacc} . Each parameter— t_{start} , Δ_{dur} , Δ_{offset} , and Δ_{delay} —is drawn independently and uniformly from its specified half-open interval $[t_1, t_2)$.

D.1. Go Task Family

The common characteristic of tasks in this family is that only one of the stimulus channels contains the signal, which varies between trials. Depending on the specific task, the network must saccade dependently or independently of the fixation cue. The response is required to be either in the direction of the signal pulse, $\theta^{(i)}$, or in the opposite direction, $\pi + \theta^{(i)}$. The individual tasks are summarized in Table S4.

D.2. Context-Dependent Decision-Making Family

For this family of tasks, stimulus pulses occur in both channels, and the network must report either $\theta^{(1)}$ or $\theta^{(2)}$, depending on the specific task type. In some tasks, the pulses occur at different times, requiring the network to maintain memory of the stimuli. The task parameters for this family are summarized in Table S4.

D.3. Matching Family

In the matching tasks, the network determines whether to saccade based on whether the two stimulus angles "match." In the "Angle" tasks, the network saccades only if $\theta^{(1)} = \theta^{(2)}$ under the given resolution (10 degrees per angle). In the "Category" tasks, the network saccades if sign($\theta^{(1)}$) = sign($\theta^{(2)}$), meaning both angles are either positive or negative. Task details are provided in Table S4. Regardless of whether the angles match, the response is always set to report the angle $\theta^{(2)}$ (or the opposite of it).



Figure S3. Randomly generated multi-region network: example traces of data reconstruction for networks with median MR-LFADS(R) performance for 3 regions (a) and 4 regions (b).

E. Implications of Constrained Architectural Choices in Restricting Message Content

E.1. Input and Message Inference with KL Penalties

KL penalties with standard Gaussian priors in variational autoencoders are known to reduce latent space dimensionality by pruning unnecessary dimensions (Dai et al., 2018; Miller et al., 2024). Consequently, with sufficiently high KL penalty coefficients (β_u , β_m), MR-LFADS is incentivized to use only the communication channels essential for data reconstruction. This effect is evident when comparing the most active channel (i.e., the one with the highest input or message norm across trials and time) to the least active ones (Fig. S4a, b). Examining input and message norms across all channels further confirms that some channels are effectively silenced (Fig. S4c).



Figure S4. The effect of KL penalties on restricting information across inferred inputs and communication channels in Experiment 1. (a) Inferred input norm over time, $||\mu_{u,t}^{[k]}||_2$, for the most (dark red) and least (light red) active channels, where k denotes channel number. (b) Inferred message norm over time, $||\mu_{m,t}^{j\to i,[k]}||_2$, for the most (dark blue) and least (light blue) active channels. (c) Inferred input (top) and message (bottom) norm across channels for area A^1 in descending order of its scalar norm across trials and time.

E.2. Message Inference via Rates versus Factors

Since MR-LFADS factors and reconstructed rates share the same dimensionality and are related by a linear transformation, it may not be immediately clear how message inference from these variables leads to substantial differences. To investigate this, we examine MR-LFADS(F) trained on the pass-decision synthetic data. We perform SVD on projection matrices (Fig. S5a) from factors to rates, W_r (Eq. 5), and from factors to communication, W_{μ_m} (Eq. 8 with r_t^j replaced by f_t^j). Our analysis shows that W_r has small singular values, indicating that certain factor dimensions contribute minimally to the reconstructed rates (Fig. S5b). In contrast, W_{μ_m} exhibits relatively uniform singular values (within the same order of magnitude), suggesting that all factor dimensions are utilized in communication (Fig. S5c). This implies that some factor dimensions play a role in message inference while being largely detached from rate reconstruction (Fig. S5a).

To further investigate, we re-express the factor space of area P using the left singular vectors of W_r , denoted U. In Experiment 2, MR-LFADS(F) is shown to encode both the stimulus s and decision variable d in its $m^{P \to D}$ messages (Fig. 3e). To examine how these variables are distributed across the subspaces of U and whether this aligns with the under-constrained dimensions identified earlier, we project the factors f^P onto the subspace spanned by the top k singular vectors of U, denoted $U^{[1:k]}$, varying k from 1 to N_{msg}^D . We then compute the R^2 values for predicting s and d from the projected values f' (Fig. S5d, e, cyan lines). We repeat this process in the reverse direction, projecting onto the bottom k singular vectors instead (Fig. S5d, e, orange lines). The results reveal a clear separation: decoding accuracy for s improves more when projecting onto the top singular vectors, whereas decoding accuracy for d increases more rapidly when projecting onto the bottom singular vectors. This suggests that information not used for rate reconstruction—such as d—is preferentially encoded in the under-constrained dimensions of the latent space, reinforcing the idea that message inference utilizes latent dimensions beyond those needed for rate prediction.



Figure S5. SVD Analysis on MR-LFADS(F). (a) Illustration of factors containing unconstrained information. Rates, factors, and communication are all from area P, i.e. $r \equiv r^p$, $f \equiv f^p$, $\mu_m \equiv \mu_m^{P \to D}$. k is the effective rank of W_r . (b) Ranked singular values for W_r . (c) Ranked singular values for W_{μ_m} . Shaded regions indicate standard deviation across different random initializations (seeds). (d) R^2 value for decoding s from f^P projected onto subspaces spanned by top / bottom left singular vectors of W_r of an example seed. (e) Same as (d), but for decoding d.

F. Application to Multi-Region Electrophysiological Data

Data analyzed are from mouse ID:440959. For the photoinhibition experiment, we analyzed a session with the following recorded regions:

- Anterior Lateral Motor cortex (ALM): MOs2/3, MOs5, and MOs6 (layer 6)
- Thalamus (ALM, A): VM and VAL
- Thalamus (Other, O): Anterior Ventral (AV) and Lateral Dorsal (LD)

- Midbrain Reticular Nucleus (MRN): MRN
- Superior Colliculus (SC): intermediate gray (SCig), intermediate white (SCiw), optic (SCop), superficial gray (SCsg), and zonal layer (SCzo)

For the model consistency experiment, we analyzed a session with the these recorded regions:

- ALM: secondary motor cortex, layers 2/3 and 5 (MOs2/3, MOs5)
- Thalamus: Ventral Medial (VM) and Ventral Anterior-Lateral (VAL)
- MRN: MRN

Trials were filtered to include only those with durations between 4.5 and 5.5 seconds. Each trial was binned into 500 time steps, with each bin corresponding to a 10 ms interval.

F.1. Comparison of Photoinhibition Effects

We selected one session of the data involving five brain regions previously implicated in a decision-making task, as identified in Chen et al. (2024). For both control and photoinhibition trials, we only used trials from the same condition per comparison—either left-hit or right-hit—where "hit" indicates making a correct choice, and "left" or "right" refers to the correct choice. For photoinhibition trials, we focused on those with perturbations within the delay period, aligning all such trials to the photoinhibition onset. Firing rate \bar{r}_t is smoothed from raw spike counts using a causal exponential filter with rate parameter of 7.1. For each region, we computed the absolute difference in trial-averaged activity between photoinhibited and control trials, averaged over neurons, to estimate the influence of photoinhibition.

F.2. Consistency of Model Inference Across Random Seeds

To ensure a fair comparison between models, we evaluated each using the same hyperparameters and same number of random seeds. We then computed all pairwise similarities between inferred effectomes and messages across seeds to assess the consistency of model inference.