
Geometry of naturalistic object representations in recurrent neural network models of working memory

Xiaoxuan Lei^{1,2}

Takuya Ito³

Pouya Bashivan^{1,2}

¹ Department of Physiology, McGill University, Montreal, Canada

² Mila, Université de Montréal, Montreal, Canada

³ IBM Research, Yorktown Heights, NY, USA

Abstract

Working memory is a central cognitive ability crucial for intelligent decision-making. Recent experimental and computational work studying working memory has primarily used categorical (i.e., one-hot) inputs, rather than ecologically-relevant, multidimensional naturalistic ones. Moreover, studies have primarily investigated working memory during single or few number of cognitive tasks. As a result, an understanding of how naturalistic object information is maintained in working memory in neural networks is still lacking. To bridge this gap, we developed sensory-cognitive models, comprising of a convolutional neural network (CNN) coupled with a recurrent neural network (RNN), and trained them on nine distinct N-back tasks using naturalistic stimuli. By examining the RNN’s latent space, we found that: 1) Multi-task RNNs represent both task-relevant and irrelevant information simultaneously while performing tasks; 2) While the latent subspaces used to maintain specific object properties in vanilla RNNs are largely shared across tasks, they are highly task-specific in gated RNNs such as GRU and LSTM; 3) Surprisingly, RNNs embed objects in new representational spaces in which individual object features are less orthogonalized relative to the perceptual space; 4) Interestingly, the transformation of WM encodings (i.e., embedding of visual inputs in the RNN latent space) into memory was shared across stimuli, yet the transformations governing the retention of a memory in the face of incoming distractor stimuli were distinct across time. Our findings indicate that goal-driven RNNs employ chronological memory subspaces to track information over short time spans, enabling testable predictions with neural data.

1 Introduction

Working memory (WM) – the ability to store and manipulate information over short periods – is a central cognitive capability that enables a wide spectrum of behaviors [Baddeley, 1992]. Over the past few decades, various experimental, computational, and theoretical techniques have been adopted to study WM from both cognitive and neural perspectives. However, several fundamental issues remain unresolved, including the key question of *how high-dimensional sensory information is encoded, maintained, and modulated according to specific task demands*.

A significant body of work has been dedicated to modelling the computations underlying WM. These include many classic models from cognitive science [Meyer and Kieras, 1997a, Ritter et al., 2019,

Baddeley et al., 2021], neuroscience [Miller et al., 1996, Wang, 1999, Emrich et al., 2013], and more recently deep learning [Yang et al., 2019, Ehrlich and Murray, 2022]. Recent approaches have focused on using artificial neural networks, such as recurrent neural networks (RNNs), to understand and model WM due to their ability to learn the complex cognitive tasks that are commonly used to study WM in humans. However, prior studies were limited in three aspects. First, most works used abstract categorical inputs, such as color and location of moving dots that are represented as binary (or one-hot vector) inputs [Panichello and Buschman, 2021, Yang et al., 2019, Piwek et al., 2023]. While training models with such categorical inputs is easier in practice, the resulting models offer limited insights into how real world, naturalistic stimuli (which are embedded in high-dimensional spaces) are processed. Second, most prior work has considered single or a few cognitive tasks [Xie et al., 2022, Mante et al., 2013, Piwek et al., 2023], limiting the generality of these models in explaining the neural computations underlying working memory. Lastly, while some prior work has explored how object features are encoded into population activity [Xie et al., 2022], it remains unclear how information in working memory is sustained across time to support concurrent encoding, retention, and retrieval during a dynamic task—and how these processes might align with classic cognitive theories [Luck and Vogel, 1997, Alvarez and Cavanagh, 2004, Franconeri et al., 2013].

To address these limitations, we investigated how task-optimized RNNs manipulate the multidimensional properties of naturalistic visual inputs during different stages of WM (Figure 1d). Specifically, we examined: 1) How do task-optimized RNNs select task-relevant properties of naturalistic objects during WM? 2) What computational strategies do RNNs employ to dynamically maintain object properties in the face of incoming (distractor) information? To address these questions, we trained *multi-task models on a collection of N-back tasks using naturalistic stimuli*, and analyzed the RNNs’ latent space during concurrent encoding, retention, and retrieval of information. We specifically chose the N-back task, given that it requires the dynamic encoding, retention, and retrieval of information in the face of incoming distractor information. This is in contrast to prior studies, which primarily studied WM tasks such as delayed-match-to-sample tasks, which focus on WM maintenance during stable fixation periods. The nature of the N-back task makes it an excellent testbed to evaluate how naturalistic object features are maintained in a dynamic environment that requires on-the-fly WM updates and decision-making.

Our main contributions are as follows:

- We trained different classes of gated and gateless RNN models on a suite of WM tasks and developed decoder-based analyses to study the geometry of naturalistic object representations during different stages of WM.
- We found that task-relevant and -irrelevant object properties are simultaneously encoded in multi-task RNN models, while only gateless RNNs produced shared and reusable representations across tasks.
- We found that object features are less orthogonalized in the RNNs’ hidden dynamics compared to perceptual representations.
- We found that RNNs solve the N-back task by using chronological memory subspaces to separate object representations presented across time. This finding supports resource-based models of working memory [Alvarez and Cavanagh, 2004, Franconeri et al., 2013] and challenges the classic slot-based model [Luck and Vogel, 1997].

2 Related Works

Models of working memory. Originally rooted in cognitive science, the notion of WM was first formally defined and popularized by Baddeley [1992] who proposed a cognitive system consisting of modality-specific buffers and a shared executive module to control the information flow in and out of the memory buffers. This initial work, along with most early models, portrayed WM as a memory system with three key features: flexibility of information representation, limited capacity, and limited temporal span.

Subsequent models based on this perspective were largely akin to the architecture of the Von-Neumann computers, comprising of input and output channels, volatile memory components, and a central processing system that continuously executed pre-specified computer code according to task goals [Cowan, 1988, Meyer and Kieras, 1997b, Anderson, 2013]. However, experimental work in

neuroscience has showed that the underlying biological circuitry of WM consists of a wide network of brain areas with diverse roles that do not adhere to the clean-cut modules specified in those previous models [Sreenivasan and D’Esposito, 2019].

Compte et al. [2000] was one of the first studies that showed RNNs with closely matched connection parameters to those measured from the brain can reliably store information in the presence of distractors. Later work showed that these networks are not only capable of performing many classic WM tasks but also replicate neural signatures that were previously observed in animals’ brains during these tasks, indicating that similar neural computations may be used by both systems [O’Reilly and Frank, 2006, Mante et al., 2013, Xie et al., 2023, Finkelstein et al., 2021, Masse et al., 2019]. More recent work proposed various neural network architectures that combine linear attention mechanisms with slot-based memory modules and feedback mechanisms as models of working memory [Hwang et al., 2024, Loynd et al., 2020].

Neural network models in neuroscience. Neural network models have been increasingly used in computational neuroscience to model neural computations during different behaviors. These models are valuable tools that incorporate ideas from a diverse range of cognitive architectures, while also making testable behavioral or neural predictions. Various classes of neural network models have now been used to simulate neural activity in sensory brain regions [Schrimpf et al., 2018, Kell et al., 2018, Khaligh-Razavi and Kriegeskorte, 2014, Yamins et al., 2014, Bashivan et al., 2019], language comprehension [Hasson et al., 2020, Schrimpf et al., 2021], and decision making [Mante et al., 2013, Yang et al., 2019]. Furthermore, others have used RNNs to study the dynamical motifs that underlie various behavioral signatures such as memorization, integration and selection of information, and attention. For example, Mante et al. [2013] used goal-directed RNNs to simulate the dynamics of neural populations in the macaque prefrontal cortex during a context-dependent decision making task and found specific dynamical mechanisms for the selection and integration of task-relevant inputs. Yang et al. [2019] identified specialized functional clusters, mixed selectivity, and compositional representation in multi-task RNNs. The shared dynamic motifs across tasks were further extended and described mathematically in Driscoll et al. [2022].

3 Methods

Tasks. We considered N-back tasks ($N \in \{1, 2, 3\}$) based on one of three distinct object properties (i.e. feature; $F \in \{Location, Identity, Category\}$) (denoted as L, I, C), resulting in a total of 9 N-back task variants (Figure 1b). Naturalistic stimuli were generated using 3D object models from the ShapeNet dataset (rendered examples in Figure A1a) [Chang et al., 2015], comprising 4 object categories, each with 2 unique identities rendered from various view angles, and presented at 1 of 4 possible locations. We consider two validation approaches: validating on novel view angles, and validating on novel identities. The training and validation novel angle datasets differed in their viewing angles, necessitating view-invariant processing by the model. In contrast, the validation novel identity includes unseen identity sampled from categories same as the training dataset.

Model Architecture. We considered a two-stage model that delineates perceptual and cognitive processes (Figure 1c). At the first stage, the model processes sequences of images, utilizing an ImageNet [Deng et al., 2009] pre-trained ResNet50 [He et al., 2016] model to derive visual embeddings from each image input. All object features including category, identity, and location were highly decodable from these activations (category: 100.00%, identity: 99.57%, location: 100.00%; 2-fold cross-validation). A point-wise convolutional layer reduces the dimensionality of the 2048-channel feature map from ResNet’s penultimate layer (layer 4.2 ReLU) to match the RNN’s hidden size. Next, the vectorized embeddings are concatenated with a task index vector and processed by a fully connected layer (matching the RNN’s latent size) with layer normalization [Ba, 2016], serving as input to the RNN. The RNN’s output is then fed through another fully connected layer and projected to one of three possible responses: match, non-match, or no action.

Each network is trained to perform one (single-task-single-feature) or multiple tasks (multi-feature or multi-task or both). After training, we analyzed activations from the penultimate layer of ResNet50 (i.e. *the perceptual space*), as well as the RNN activations during the stimulus presentation and subsequent timesteps (i.e. *encoding and memory space* respectively, denoted as E and M as shown

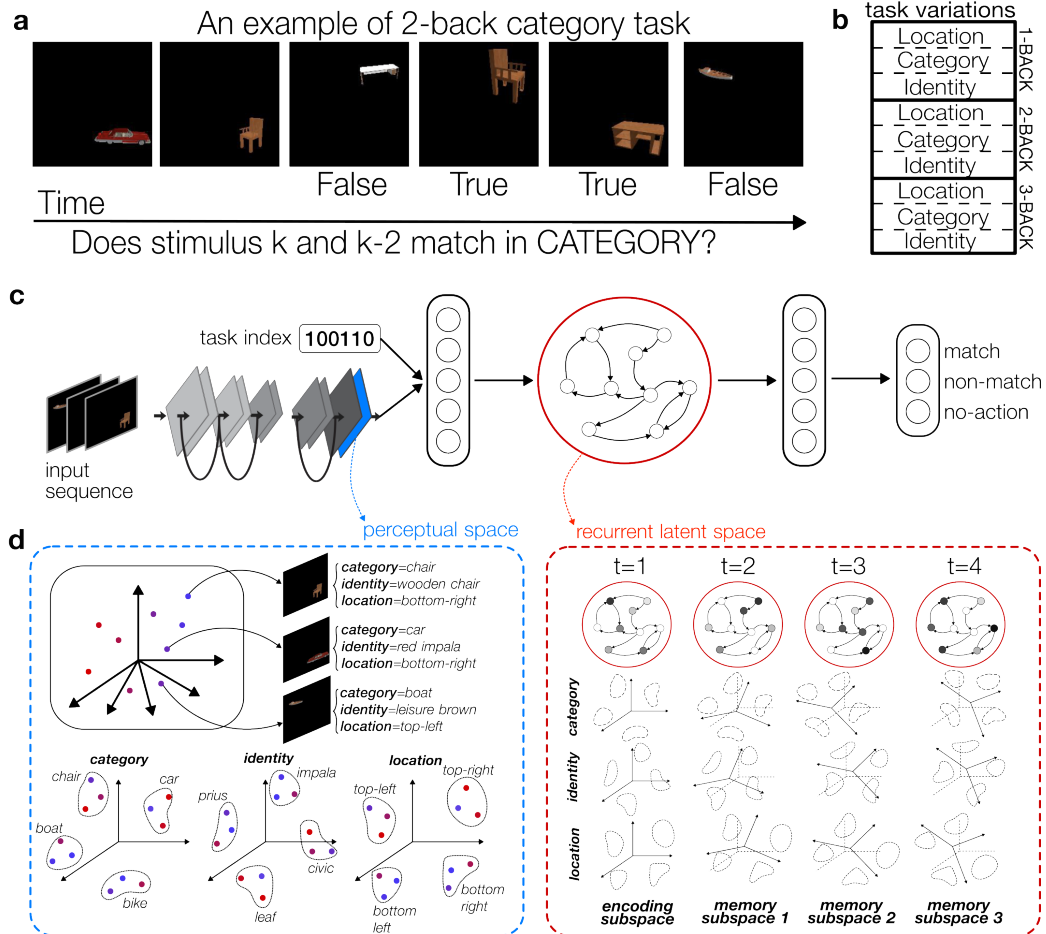


Figure 1: **Tasks and Models:** **a)** Example of a 2-back category task. Each object’s category is compared with the category of the object seen two frames prior. **b)** The suite of n -back tasks considered in the study. **c)** The sensory-cognitive model architecture. **d)** A schematic showing the latent subspaces for category, identity, and locations in the perceptual, encoding, and memory subspaces. Left: Stimuli are encoded in high dimensional latent space of the vision model (CNN). Each object property is encoded in a high dimensional latent subspace of this model; Right: RNN model represents each object property in its encoding latent subspace and retains some or all of the properties within its memory subspaces at later time points.

in Figure 4a). We considered three recurrent architectures for the second stage of the model including the vanilla RNN, GRU [Chung et al., 2014], and LSTM [Hochreiter and Schmidhuber, 1997].

Model Training Each model architecture was trained within three scenarios that differed in their training diet:

- **Single-Task-Single-Feature (STSF):** trained on a single n -back task based on a single object feature (e.g. 1-back location).
- **Single-Task-Multi-Feature (STMF):** trained on a single choice of N for all three feature variations (e.g. 1-back L , or C or I).
- **Multi-Task-Multi-Feature (MTMF):** encompassing all choices of N (1,2,3) and features (L, I, C).

Model parameters were trained using AdamW [Loshchilov and Hutter, 2017] optimizer with an initial learning rate of $3e - 5$ and a Multi-step learning rate decay with $\gamma = 0.1$ for every 100 iterations. Batch size was 256 and all trials were generated on-the-fly using the iWISDM package [Lei et al.,

2024]. Hidden layers of RNN modules are initialized with Kaiming Initialization [He et al., 2015]. For details of decoding analysis and Procrustes analysis, please refer to Appendix .1

4 Experiments

All models reached $> 95\%$ accuracy on train and $> 90\%$ on validation set with novel object angles. Generalization to novel object instances was substantially weaker (Figure A1b, c). Model performance increased with the number of model parameters, and with identical parameter count, vanilla RNNs accuracy was lower compared to their gated counterparts (Figure A1c). The ensuing analyses utilized data collected from models with 512 units for vanilla RNNs, and 256 units for GRUs and LSTMs, to ensure comparable model performance as well as comparable model parameters.

4.1 Encoding of task-relevant and -irrelevant object properties in task-optimized RNNs

We first examined how RNN modules represent various object properties such as location, identity and category in their latent space. In particular, we investigated the following two questions:

1) Do recurrent networks preserve object properties that are not necessary for the task? In order to perform a task, recurrent networks must maintain information about the task-relevant object properties. However, maintaining information about the task-irrelevant factors is not necessary from the perspective of the task objective. Therefore, RNNs may either selectively maintain task-relevant information or maintain full object representations, recalling task-relevant information when prompted.

We trained decoders (i.e. classifiers) to predict each object property from the RNN hidden state activity from the first timestep of each trial (e.g. $F = L_i$ vs. $F = L_{j \neq i}$ decoders, total 4 location decoders). Cross-validated decoding accuracies are shown in Figure 2b for STSF, STMF, and MTMF GRU models. Unsurprisingly, the task-relevant object properties are fully retrievable in all models. Further tests revealed a causal relationship between the subspaces encoding task-relevant information and the network’s generated response (Fig. A7; see section .1 for details.) However, while task-irrelevant object features are not generally well-preserved in STSF models (Figure 2b, left), they are much better preserved in STMF and MTMF models (i.e. decoding accuracy $> 85\%$; Figure 2b middle and right). This finding was consistent across all three RNN architectures (Figure A2a) and across time points (Fig. A5). These results suggested that all RNNs maintained a full representation of objects in their latent spaces regardless of which object properties were required for performing the task.

2) Are object properties encoded within a subspace that is shared across different tasks or distinct ones within each task? Having observed that both task-relevant and -irrelevant information are retained by multi-feature RNNs, we next asked whether RNN encoding of object properties is task-dependent or -independent. To probe this, we trained decoders to predict object properties from the RNNs’ activations during one task, and tested the decoder on RNN activations when performing another task (i.e., cross-task decoding). We quantified the generalization performance of MTMF models across all three architectures (Figure 2a). We found that gated RNNs (GRU and LSTM, Figure A2b) utilized highly task-specific subspaces for encoding object properties, while vanilla RNN encoded object properties within a subspace that was shared across all task-variations (Figure 2c). This suggests that gated RNNs tend to learn task-specific representations that do not generalize across tasks, potentially impacting their ability to generalize to new tasks.

4.2 Representational orthogonalization in task-optimized RNNs.

To improve their performance, RNN weights might form structured and separable representations for each task-relevant feature. For this to be true, the RNN latent space may orthogonalize feature representations beyond their perceptual representation (see schematics in Figure 3a). To quantify orthogonalization, we calculated the angles between all pairs of decision hyperplanes using cosine similarity (bootstrapped 10 times). We then summarized these angles into a single orthogonality measure by computing the Frobenius norm of the difference between this matrix and the identity matrix (which represents complete orthogonalization). We defined this measure as the orthogonalization

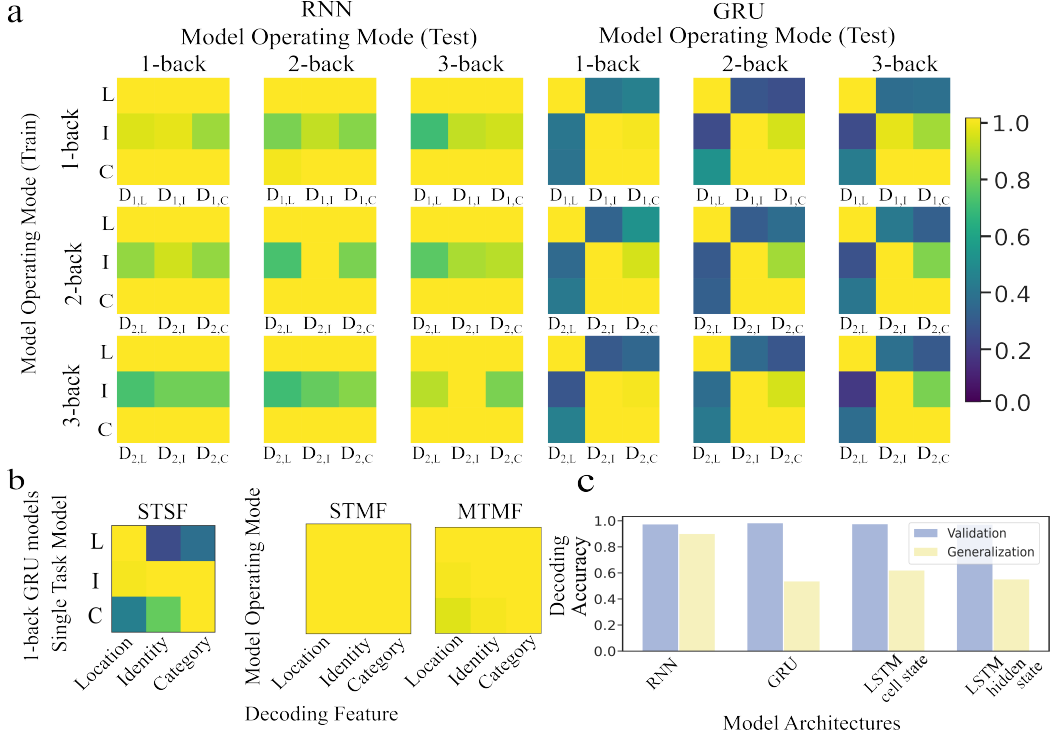


Figure 2: **Representation of task-relevant/-irrelevant object properties:** (a) Decoding generalization accuracy for each object property is displayed across tasks and operating modes for vanilla RNN and GRU. Rows and columns of 3×3 matrices correspond to the N -back task on which the decoders are fitted and tested on respectively. Matrix columns correspond to particular decoders denoted by $D_{k,F}$ ($k \in \{1, 2, 3\}$, $F \in \{L, I, C\}$) (indicating which task and decoding feature the decoder was fitted on), while matrix rows correspond to the object property of the task the decoder was tested on. (b) Validation accuracy of decoders trained on RNN latent space activations from the first time step of each trial to predict different object properties. Each column represents the object property the decoder was trained on, while each row corresponds to a model. (c) Quantification of the validation accuracy (within the same task, indicated in purple) and generalization accuracy (across tasks with different task-relevant features, indicated in yellow) across all model architectures.

index (O):

$$\tilde{W}_{ij} = 1 - \text{abs}(\cos(W_i, W_j))$$

$$O = \mathbb{E}(\text{triu}(\tilde{W})) \quad (1)$$

where W_i is the normal vector of the decision hyperplane that separates points assigned with feature value i from the rest, $\cos(W_i, W_j)$ is the cosine similarity between the two normal vectors. We take the absolute value since the relative direction does not matter. $\text{triu}(\cdot)$ is the upper triangle operator.

We evaluated the degree of representation orthogonalization within the perceptual and encoding spaces. Contrary to our hypothesis, we observed that relative to the perceptual space, the RNN latent space slightly de-orthogonalize the axes along which distinct object features are represented (Figure 3b). Similar results were found when PCA was used to equalize dimensionality between the two spaces (Figure A3). Although more orthogonalized representations generally facilitate structured and enhanced separation of task-relevant features, our behavioral results still indicate optimal performance. One possible interpretation of these results is that the reduced orthogonalization in the RNN latent space produces a more efficient (lower dimensional) representation. (In contrast, increased orthogonalization in the representational space would increase the overall dimensionality of the object representations.) In practice, only a subset of dimensions need to contain orthogonalized representations for successful task performance. In turn, this would make it easier to train the read out weights of the RNN to produce the correct task responses.

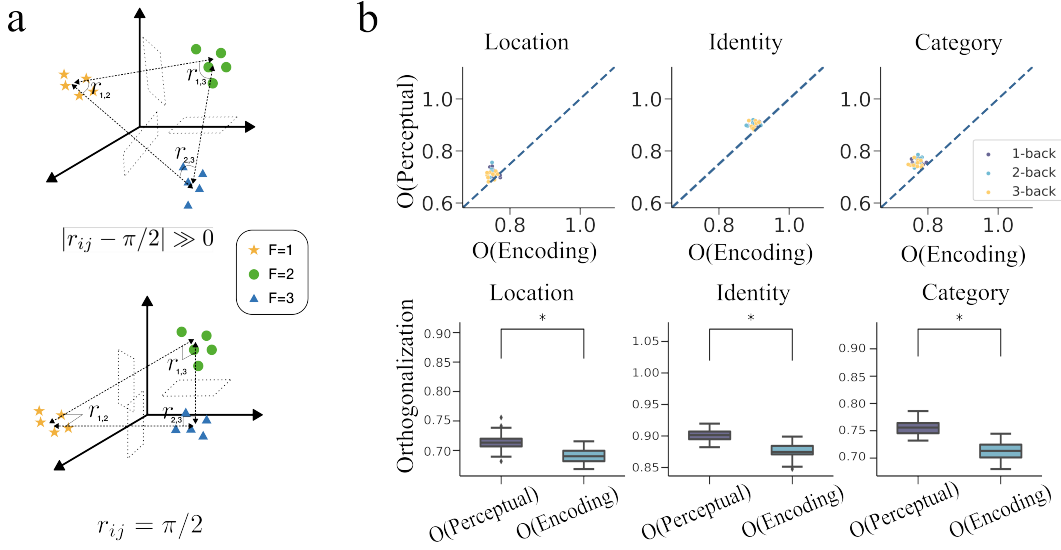


Figure 3: **Orthogonalization:** a) A schematic of two hypothetical object spaces in 3D. $r_{i,j}$ represents the angle formed by the decision hyperplanes that separate feature value i and j from each other. Top: non-orthogonalized representation; Bottom: orthogonalized representation. b) Upper panel: Normalized orthogonalization index, for both perceptual and encoding spaces respectively (denoted as $O(\text{Perceptual})$ and $O(\text{Encoding})$). In most models, a less orthogonalized representation of feature values emerges in the RNN encoding space compared to the perceptual space (CNN output). Lower panel: Statistical comparison of the relative orthogonalization levels between the perceptual and encoding spaces. A two-sample t-test was performed to assess differences between the distributions of orthogonalization indices in the perceptual space and the encoding space.

4.3 Neural mechanisms of concurrent encoding, maintenance, and retrieval in RNN models of WM

Having examined the encoding of objects in the RNN’s latent space, we next investigated how RNN dynamics enable simultaneous encoding, maintenance, and retrieval of information. Performing our N-back task suite required the RNN to keep track of prior objects’ properties while simultaneously encoding incoming stimuli with minimal interference.

We reasoned that the RNN may implement one of three possible mechanisms to perform the N-back working memory task suite (Figure 4e):

- **H1: Slot-based memory subspaces [Luck and Vogel, 1997].** Where the RNN latent space is divided into separate subspaces that are indexed in time within the sequence. Each object is encoded into its corresponding subspace (i.e. slot) and is maintained there until retrieved. By definition, the subspace assigned to each memory slot is distinct and “sustained” in time.
- **H2: Relative chronological memory subspaces.** Where the RNN latent space is divided into separate subspaces that each maintains object information according to their age (i.e. how long ago they were encoded). Such a mechanism will require a dynamic process for updating the content of each memory space at each time step during the task.
- **H3: Stimulus-specific relative chronological memory subspaces.** This is similar to the relative chronological memory hypothesis but with independent subspaces assigned to each object. Each observation in the sequence is thus encoded into a distinct subspace and encoding of each stimulus is in turn distinctly transformed into associated memory representations.

To identify the hypothesis that best matches the computations performed by the RNN, we analyzed how the RNN latent subspace encodes and transforms each object property ($E_{(S,T)}$) across time into memory ($M_{(S,T)}$) (Fig. 1d).

We first tested whether object information is maintained in a temporally stable subspace (i.e. sustained working memory representation) which aligns with **H1** prediction (i.e. $E_{(S=i,T=1)} \stackrel{?}{=} M_{(S=i,T=k)}$, $k \in \{2, 3, 4, \dots\}$; Fig. 4a). For this, we trained decoders to predict the value of each object property using the RNN unit activity during the encoding phase (i.e. Encoding Space) and evaluated its generalization performance in consecutive steps (i.e. Memory Spaces). We reasoned that if the object information is encoded in a subspace that is stable across time (as in a memory slot), the decoders' generalization performance should be high and comparable to its performance during the encoding phase. Contrary to H1's prediction, we found that the decoders do not generalize well (Figure 4b), suggesting that the object information is not stably encoded in a temporally-fixed RNN latent space.

However, we observed that in STMF and MTMF models, the cross-time decoding accuracy is consistently higher during recall (Figure 4b and c). Interestingly, this suggests that the object representation is partially realigned with its original encoding representation when it is retrieved (Fig. 4b).

Next, we examined whether the object encoding space is shared between incoming stimuli, irrespective of the specific object or time (**H2** vs. **H3**; i.e. $E_{(S=i,T=i)} \stackrel{?}{=} E_{(S=j,T=j)}$, for $i \neq j$). We thus fitted classifiers to decode each object property using the hidden activity from the encoding space of each stimulus within the sequence (i.e. decoding $S = i$ from $E_{(S=i,T=i)}$), testing it on the stimuli appearing at other time steps (i.e. decoding $S = j$ from $E_{(S=j,T=j)}$). We performed this analysis for all object properties and for all models and all tasks. The validation and generalization accuracies were almost identical (Figure 4d), suggesting a stable encoding representation ($E_{(S=i,T=i)} = E_{(S=j,T=j)}$) consistent with **H2**. In other words, each object in the sequence was encoded according to its chronological age (i.e., when it was placed into memory), regardless of the object property.

Having examined how the RNN latent space allows concurrent encoding, retention, and retrieval of information, we next investigated what transformations underlie the conversion of information from one subspace to another. Specifically, we inquired whether the transformation of feature subspaces across timesteps is stable with respect to the same encoded stimulus (i.e. $\mathcal{T}_i = \mathcal{T}_{i+1}$; see Fig. 4e). As detailed in Appendix .1, we adopted the orthogonal Procrustes analysis to obtain rotation matrix $R_{S,T}$ to characterize the transformation. The orthogonal Procrustes analysis is a statistical shape analysis which discovers simple rigid transformations that superimpose a set of vectors/points onto another. We used this analysis to inquire whether each set of object feature decoders can be rotated to align with the set of decoders for the same object properties, but across time. Thus, the above test can be reformulated as

$$R_{(S=i,T=j)} \stackrel{?}{=} R_{(S=i,T=j+1)} \quad (2)$$

Additionally, we also tested if the transformation is consistent across stimuli within the sequence:

$$R_{(S=i,T=j)} \stackrel{?}{=} R_{(S=i+k,T=j+k)} \quad (3)$$

We first checked whether the feature representation subspace transformation across timesteps are structured (Figure 4f-left). In other words, we evaluated how well the Procrustes analysis can capture the rotation transformation of feature representation subspaces. The reconstructed decision hyperplanes resulting from rotating the original decoders by the Procrustes transformation matrix R were highly accurate (Fig. 4f-right), indicating that a rotation was able to capture the transformation performed by the RNN across time (also see Figure A4).

Next, we tested Eqs. 2 and 3 by swapping the rotation matrix at $R_{(S=i,T=j)}$ with $R_{(S=i,T=j+1)}$ or $R_{(S=i+k,T=j+k)}$ respectively, and plotted the accuracy of reconstructed decision hyperplane for MTMF models in Figure 4g. We reasoned that if these rotation operations were shared across time steps and stimuli, swapping them would not significantly affect the decoding accuracy. Across all model architectures and tasks, we found that replacing $R_{(S=i+k,T=j+k)}$ consistently yields good accuracy, whereas replacing $R_{(S=i,T=j+1)}$ does not. Similar results were obtained in models trained on larger N (N=1-4), stimuli with naturalistic texture backgrounds, and more number of categories and identities (Fig. A6). These results suggest that while the transformation remains consistent across different stimuli, it is not stable over time.

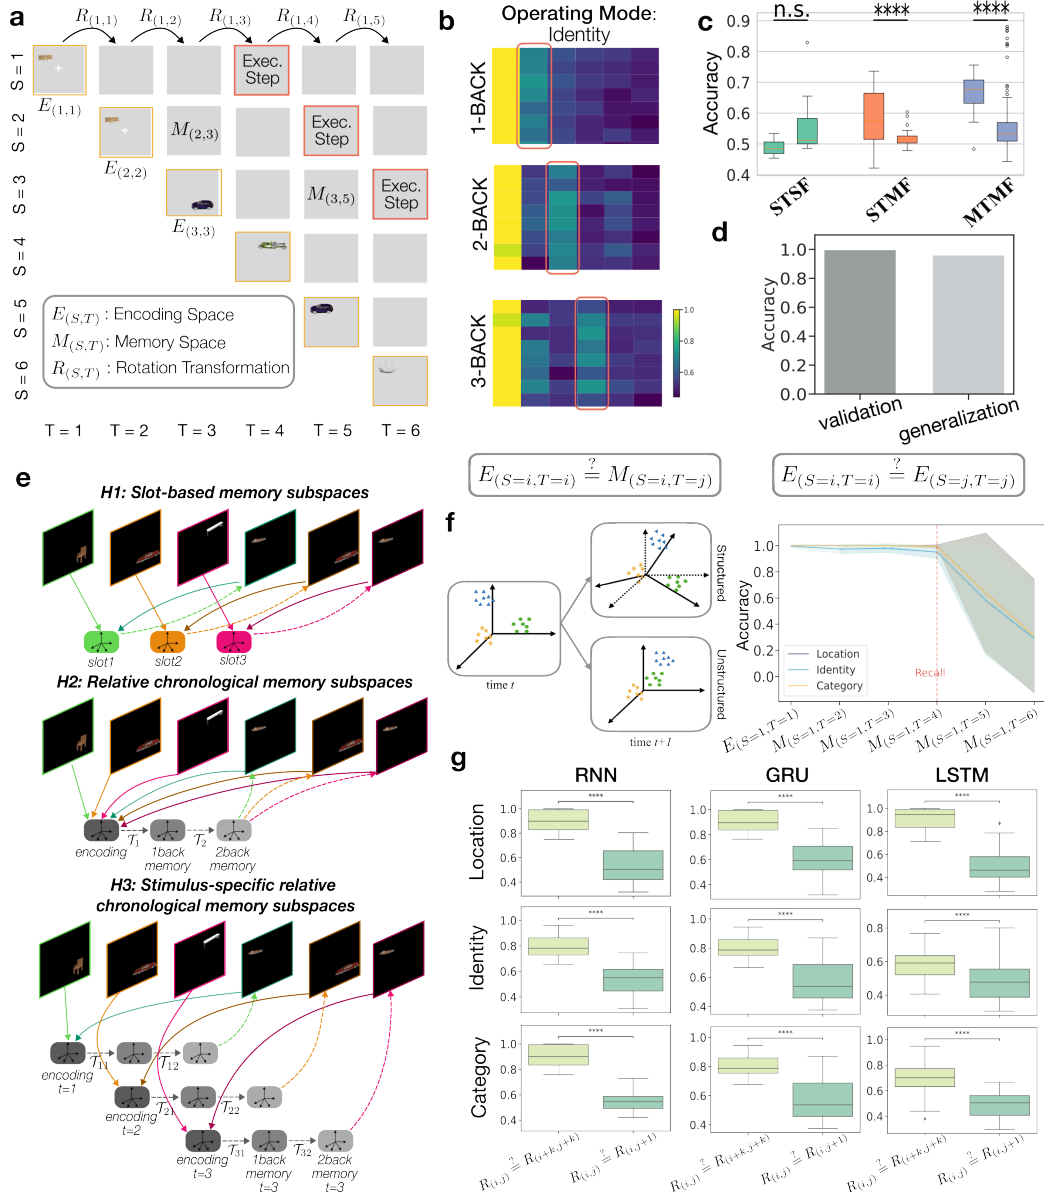


Figure 4: RNN dynamics during n-back task a) schematic of the 3-back task for a trial of 6 inputs. Model encodes each observed object in its respective Encoding Space denoted as $E_{(i,j)}$ (diagonal frames with yellow borders). For each stimulus, various object properties are retained over time in their respective Memory Space denoted as $M_{(i,j)}$. On executive steps (frames with red borders) model produces a response according to the memory of the stimulus and the newly observed stimulus at that time. b) Decoding accuracy for predicting object identity at different time steps where the decoder is fit to data from the encoding step of a MTMF GRU during 1/2/3-back identity tasks. Red box indicates the executive steps. c) For each model type, we measured the generalization accuracy on executive (left boxplot) and non-executive (right boxplot) steps. d) Decoding accuracy for decoders trained and tested on the same $E_{i,i}$ space (validation) or tested on other $E_{j,j}, j \neq i$ spaces. e) Schematic of the three hypotheses. f) Left: Schematic of the two latent space transformations. Structured transformation preserves the topology (i.e. the transformation can be captured solely by a common scaling factor and a rotation matrix). Unstructured transformation: does not preserve the topology. Right: Decoding accuracy for fitted decoders (solid line) and reconstructed decoders (dotted line) using the rotation matrix $R_{(i,i)}$ from the Procrustes analysis. The small accuracy gap between fitted and reconstructed decoders suggests a structured transformation. g) Decoding accuracy of the reconstructed decoder when the original rotation matrix is substituted with another (indicated by the x-axis labels). Rows and columns corresponds to object properties and MTMF network architectures respectively.

5 Discussion and Conclusions

We investigated how naturalistic objects are represented in recurrent models during a dynamic and difficult WM task. In contrast to most prior work that build computational models of WM using abstract categorical stimuli, which typically study WM during a stable fixation/delay period, we trained a range of sensory-cognitive models to perform N-back tasks using naturalistic stimuli. We found that models trained on multiple features and multiple tasks retained object information regardless of their task-relevance. While prior studies investigating WM tasks using RNNs have identified shared representations and dynamical motifs across related tasks [Yang et al., 2019, Driscoll et al., 2022], we found that representations were largely task-specific in gated RNNs such as GRUs and LSTMs. The representational differences between gated and gateless RNNs provides an opportunity for future work to adjudicate between these recurrent architectures and their cognitive significance (i.e., gated vs. gateless) with empirical neural data.

While increased orthogonalization could in theory enable better task performance through the formation of increased separation of object representations in the recurrent module, we found the opposite. Despite reduced orthogonalization of task-relevant features in the RNN, this did not appear to influence task generalization behavior. One possible interpretation of this is that the reduced orthogonalization in the RNN latent space produces a more efficient (lower dimensional) representation. In practice, only a subset of latent need to contain orthogonalized task-relevant representations for successful task performance. Having fewer latent dimensions that orthogonalize task-relevant representations could make it easier to decode/read out the correct response information from this latent subspace. Clarifying this distinction further will be important for future work.

Lastly, we found that RNNs solve the N-back task by leveraging chronological memory subspaces to maintain information about different objects distinct. This finding is consistent with the “resource” model of working memory [Alvarez and Cavanagh, 2004, Franconeri et al., 2013], which proposes that memory resources are flexibly distributed across all items. This is contrast to the “slot-based” model [Luck and Vogel, 1997], which suggests that memory is composed of discrete, independent slots for each item. Furthermore, the observed dynamics align with the conceptual framework in Whittington et al. [2023], where different memory slots can transfer information between each other during working memory. Our experiments reveal that similar slot-like subspaces naturally emerge in task-optimized RNNs, with information transfer between them guided by RNN transformations based on task demands. However, as these subspaces are carved out of a shared neural space defined by the RNN units and dependent on chronology, they are not necessarily non-overlapping (unlike that assumed in Whittington et al. [2023]). Such possible overlaps may account for previous findings on the influence of memory load on working memory performance [Ma et al., 2014]. Altogether, these results provide testable predictions to evaluate the neural basis of WM in humans in future work.

6 Limitations

Our study has several limitations:

1. **Task-Specific Findings:** Our results are specific to the N-back task structure, and it remains uncertain whether similar computational strategies would emerge in other working memory tasks.
2. **Analysis for Novel Objects:** While we observed reduced performance with novel objects, we did not analyze the representational geometry associated with these stimuli or investigate the reasons for diminished performance. Our use of naturalistic stimuli was intended to avoid imposing the representational geometry typical of abstract inputs, as seen in previous studies (e.g., Piwek et al. [2023], Yang et al. [2019]).
3. **Architectural Constraints:** Our findings are restricted to commonly used RNN architectures, such as vanilla RNNs and LSTMs. Therefore, we cannot make definitive claims about how different neural network architectures might affect representational geometry. Further research could explore whether other architectures yield similar or distinct patterns.
4. **Impact of Network Scaling:** We did not explore how scaling the network size might influence the results. It is possible that increasing the network size could alter the strategies employed by models to perform dynamic WM tasks.

Acknowledgments

This research was supported by the Healthy-Brains-Healthy-Lives startup supplement grant, the NSERC Discovery grant RGPIN-2021-03035, and CIHR Project Grant PJT-191957. P.B. was supported by FRQ-S Research Scholars Junior 1 grant 310924, and the William Dawson Scholar award. All analyses were executed using resources provided by the Digital Research Alliance of Canada (Compute Canada) and funding from Canada Foundation for Innovation project number 42730.

References

- G.A. Alvarez and P. Cavanagh. The Capacity of Visual Short-Term Memory is Set Both by Visual Information Load and by Number of Objects. *Psychological Science*, 15(2):106–111, February 2004. ISSN 0956-7976. doi: 10.1111/j.0963-7214.2004.01502006.x. URL <https://doi.org/10.1111/j.0963-7214.2004.01502006.x>. Publisher: SAGE Publications Inc.
- John R Anderson. *The architecture of cognition*. Psychology Press, 2013.
- Jimmy Lei Ba. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Alan Baddeley. Working memory. *Science*, 255(5044):556–559, 1992.
- Alan D Baddeley, GJ Hitch, and Richard Allen. A multicomponent model of working memory. *Working memory: State of the science*, pages 10–43, 2021.
- Pouya Bashivan, Kohitij Kar, and James J. DiCarlo. Neural population control via deep image synthesis. *Science*, 364(6439):eaav9436, 2019. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aav9436. URL <https://www.science.org/doi/10.1126/science.aav9436>.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Albert Compte, Nicolas Brunel, Patricia S Goldman-Rakic, and Xiao-Jing Wang. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral cortex*, 10(9):910–923, 2000.
- Nelson Cowan. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin*, 104(2):163, 1988. ISSN 1939-1455. doi: 10.1037/0033-2909.104.2.163. URL <https://psycnet.apa.org/fulltext/1989-00336-001.pdf>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- Laura Driscoll, Krishna Shenoy, and David Sussillo. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *bioRxiv*, pages 2022–08, 2022.
- Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 571–576. 2023.
- Daniel B Ehrlich and John D Murray. Geometry of neural computation unifies working memory and planning. *Proceedings of the National Academy of Sciences*, 119(37):e2115610119, 2022.
- Stephen M Emrich, Adam C Riggall, Joshua J LaRocque, and Bradley R Postle. Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *Journal of Neuroscience*, 33(15):6516–6523, 2013.

- Arseny Finkelstein, Lorenzo Fontolan, Michael N Economo, Nuo Li, Sandro Romani, and Karel Svoboda. Attractor dynamics gate cortical information flow during decision-making. *Nature neuroscience*, 24(6):843–850, 2021.
- Steven L. Franconeri, George A. Alvarez, and Patrick Cavanagh. Flexible cognitive resources: competitive content maps for attention and memory. *Trends in Cognitive Sciences*, 17(3):134–141, March 2013. ISSN 1364-6613, 1879-307X. doi: 10.1016/j.tics.2013.01.010. URL [https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613\(13\)00024-7](https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(13)00024-7). Publisher: Elsevier.
- John C Gower and Garnt B Dijkstra. *Procrustes problems*, volume 30. OUP Oxford, 2004.
- Uri Hasson, Samuel A Nastase, and Ariel Goldstein. Direct fit to nature: an evolutionary perspective on biological and artificial neural networks. *Neuron*, 105(3):416–434, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- Dongseong Hwang, Weiran Wang, Zhuoyuan Huo, Khe Chai Sim, and Pedro Moreno Mengibar. TransformerFAM: Feedback attention is working memory, 2024. URL <http://arxiv.org/abs/2404.09173>.
- Alexander JE Kell, Daniel LK Yamins, Erica N Shook, Sam V Norman-Haignere, and Josh H McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.
- Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.
- Xiaoxuan Lei, Lucas Gomez, Hao Yuan Bai, and Pouya Bashivan. iwisdm: Assessing instruction following in multimodal models at scale. *arXiv preprint arXiv:2406.14343*, 2024.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Ricky Loynd, Roland Fernandez, Asli Celikyilmaz, Adith Swaminathan, and Matthew Hausknecht. Working memory graphs. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6404–6414. PMLR, 2020. URL <https://proceedings.mlr.press/v119/loynd20a.html>. ISSN: 2640-3498.
- Steven J Luck and Edward K Vogel. The capacity of visual working memory for features and conjunctions. *Nature*, 193(1996):1996–1998, 1997. ISSN 0028-0836. doi: 10.1038/36846. Number: 1996.
- Wei Ji Ma, Masud Husain, and Paul M. Bays. Changing concepts of working memory. *Nature Neuroscience*, 17(3):347–356, March 2014. ISSN 1546-1726. doi: 10.1038/nn.3655. URL <https://www.nature.com/articles/nn.3655>. Number: 3 Publisher: Nature Publishing Group.
- Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474):78–84, 2013.
- Nicolas Y Masse, Guangyu R Yang, H Francis Song, Xiao-Jing Wang, and David J Freedman. Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nature neuroscience*, 22(7):1159–1167, 2019.

- David E Meyer and David E Kieras. A computational theory of executive cognitive processes and multiple-task performance: Part i. basic mechanisms. *Psychological review*, 104(1):3, 1997a.
- David E. Meyer and David E. Kieras. A computational theory of executive cognitive processes and multiple-task performance: Part 2. accounts of psychological refractory-period phenomena. 104: 749–791, 1997b. ISSN 1939-1471. doi: 10.1037/0033-295X.104.4.749. Place: US Publisher: American Psychological Association.
- Earl K Miller, Cynthia A Erickson, and Robert Desimone. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of neuroscience*, 16(16):5154–5167, 1996.
- Randall C. O’Reilly and Michael J. Frank. Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, 18(2):283–328, 2006. ISSN 08997667. doi: 10.1162/089976606775093909. Number: 2.
- Matthew F Panichello and Timothy J Buschman. Shared mechanisms underlie the control of working memory and attention. *Nature*, 592(7855):601–605, 2021.
- Emilia P Piwek, Mark G Stokes, and Christopher Summerfield. A recurrent neural network model of prefrontal brain activity during a working memory task. *PLoS Computational Biology*, 19(10): e1011555, 2023.
- Frank E Ritter, Farnaz Tehranchi, and Jacob D Oury. Act-r: A cognitive architecture for modeling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3):e1488, 2019.
- Amy Ross. Procrustes analysis. *Course report, Department of Computer Science and Engineering, University of South Carolina*, 26:1–8, 2004.
- Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J Majaj, Rishi Rajalingham, Elias B Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, et al. Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, page 407007, 2018.
- Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118, 2021.
- Kartik K. Sreenivasan and Mark D’Esposito. The what, where and how of delay activity. *Nature Reviews Neuroscience*, 20(8):10–14, 2019. ISSN 1471-003X. doi: 10.1038/s41583-019-0176-7. URL <http://www.nature.com/articles/s41583-019-0176-7>. Number: 8 Publisher: Springer US.
- Xiao-Jing Wang. Synaptic basis of cortical persistent activity: the importance of nmda receptors to working memory. *Journal of Neuroscience*, 19(21):9587–9603, 1999.
- James CR Whittington, William Dorrell, Timothy EJ Behrens, Surya Ganguli, and Mohamady El-Gaby. On prefrontal working memory and hippocampal episodic memory: Unifying memories stored in weights and activity slots. *bioRxiv*, pages 2023–11, 2023.
- Yang Xie, Peiyao Hu, Junru Li, Jingwen Chen, Weibin Song, Xiao-Jing Wang, Tianming Yang, Stanislas Dehaene, Shiming Tang, Bin Min, et al. Geometry of sequence working memory in macaque prefrontal cortex. *Science*, 375(6581):632–639, 2022.
- Yudi Xie, Yu Duan, Aohua Cheng, Pengcen Jiang, Christopher J Cueva, and Guangyu Robert Yang. Natural constraints explain working memory capacity limitations in sensory-cognitive models. *BioRxiv*, 2023.
- Daniel L K Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23):8619–24, 2014. ISSN 1091-6490. doi: 10.1073/pnas.1403112111. URL <http://www.ncbi.nlm.nih.gov/pubmed/24812127> <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4060707>. Number: 23.

Guangyu Robert Yang, Madhura R Joglekar, H Francis Song, William T Newsome, and Xiao-Jing Wang. Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2):297–306, 2019.

Appendix

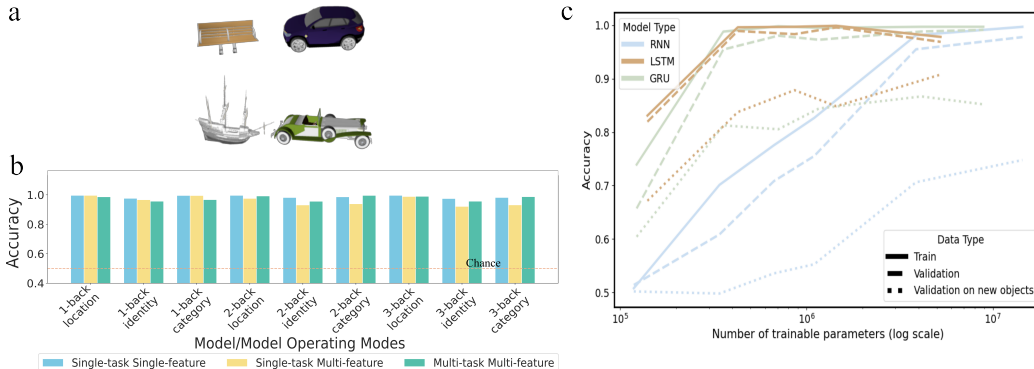


Figure A1: **Stimuli and Model performance** a) rendered stimuli examples from Shapenet. b) 9 task variations of N-back constructed from different choices of task-relevant features (L, I, C) and N (1,2,3) index. c) Model performance on train, validation novel angle and validation novel object datasets. Three architectures are tested with various number of hidden size, with the number of trainable parameters indicated on x-axis.

.1 Methods

Model training We used cross-entropy loss for training, and the identity of the task was encoded in a 6-digit binary format: the first 3 digits represented the one-hot encoding of the feature (e.g., stimulus location, category, or identity), and the second 3 digits represented the one-hot encoding of the n-back choice of n. For the single-task single-feature model, we used the same task identity vector as in the multi-task models. The multi-task multi-feature model typically takes around 4-8k iterations with a batch size of 256, and we cut off training at 14k iterations. The sequence length is fixed at 6 for both the training and validation sets.

Causal test To establish the causal relevance between the decoder-defined subspace and the network’s behavioral performance, we perturbed the network’s representations by shifting them along the direction of the normal vector to a given decision hyperplane. By passing the resulting hidden states through consecutive timesteps, we computed the probabilities of the three possible actions. We subsampled matched trials and perturbed the hidden states at various magnitudes in the direction of the corresponding decision hyperplane. As shown in Figure A7, the probability of obtaining a match action dropped significantly as the hidden states traversed the hyperplane, while the probability of no action increased. The probability of a non-match action remained largely unaffected, except for an increase in variance as the hidden states crossed the boundary. These results support the causal relationship, indicating that the subspace defined by the decoding analysis is actively utilized by the network in solving the task.

Decoding analyses We consider the latent space of the RNN as a D-dimensional space \mathbb{R}^D where D is the number of units in the RNN model. We used support vector classifiers (SVC) to perform all the decoding analysis. The decoder should be able to classify one feature value from the rest, for example, location bottom left to all the other locations. Each decoder was fitted with activations from either perceptual, encoding, or memory spaces and labels from one of the three possible features (L, I, C) of the stimuli from current or previous time steps¹. We adopted 10-fold cross-validation as well as grid search to find the best regularization values from $\{0.001, 0.01, 1, 10\}$. All the classifiers reached cross-validated accuracy of $\geq 85\%$.

Each classifier is then described by its normal vector w to its decision hyperplane d and a bias term b .

¹We used Scikit-Learn’s implementation of SVC for fitting the classifier parameters.

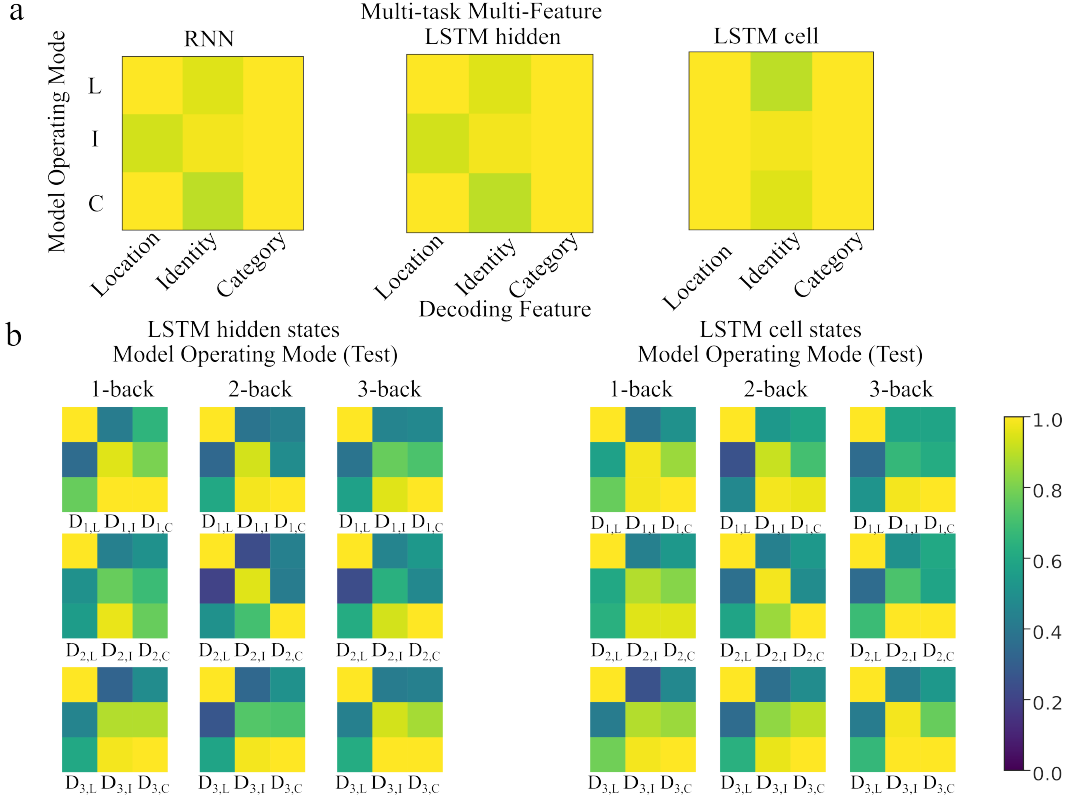


Figure A2: **Object representation efficiency:** a) Similar to Figure 2b, right panel, but for RNN and LSTM trained on MFMT. b) Similar to Figure 2a, but for LSTM hidden and cell states (model trained on MFMT)




Figure A3: **Orthogonalization:** Similar to Figure 3b, but with PCA to reduce the dimensionality to the same level of the perceptual and encoding space.

$$d = \text{data} \cdot w + b \quad (4)$$

In the Procrustes analyses (see the Procrustes analysis section below), the vector w is derived by rotating the original decoder using the Procrustes transformation and the bias term is obtained from the original decoder. When $d \geq 0$, the corresponding data point is assigned to one class, and to the alternative class when $d \leq 0$.

The set of N vectors $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N\}$ in D -dimensional space \mathbb{R}^D , represent N decoders in that space. These vectors span a subspace of \mathbb{R}^D that is the set of all possible linear combinations of these vectors. Mathematically, the subspace S spanned by these vectors is defined as: $S = \text{span}\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N\} = \left\{ \sum_{i=1}^N \alpha_i \mathbf{w}_i \mid \alpha_i \in \mathbb{R} \right\}$.

Procrustes analysis Procrustes analysis [Ross, 2004] is a powerful method for identifying shape correspondence that relies on the orthogonality of the rotation matrix. Here, we consider a shape spanned by vertices specified by the normal vectors of decision hyperplanes obtained by classifying

one feature value from the rest. For example, for four possible locations, the shape is defined by four vertices each corresponding to a classifier that discriminates one location from the rest. We calculate such descriptions of shapes at each time step, for each feature, and across all possible tasks for each model. The goal of our analysis is to transform decision hyperplanes obtained under one condition to target decision hyperplanes obtained under another condition. In other words, we want to align source shape with the target shape. To do so, we take the following steps:

- Train the decoders, obtain the normal vectors of the source and target shape, denoted as w_{source} and w_{target}
- Standardise w_{source} and w_{target} :

$$w'_{source} = \frac{w_{source} - \overline{w_{source}}}{\|w_{source} - \overline{w_{source}}\|_2} \quad (5)$$

$$w'_{target} = \frac{w_{target} - \overline{w_{target}}}{\|w_{target} - \overline{w_{target}}\|_2} \quad (6)$$

We denote $\overline{w_{target}}$ as \mathbf{b} and $\|w_{target} - \overline{w_{target}}\|_2$ as \mathbf{S} .

- Perform Orthogonal Procrustes Analysis [Gower and Dijksterhuis, 2004] to align w'_{source} with w'_{target} , which returns a rotation matrix ($\mathbf{R}_{source \rightarrow target}$) and a global scaling factor (s)
- Transform w'_{source} to w'_{target} by

$$w'_{reconstructed} = (w'_{source} \cdot \mathbf{R}_{source \rightarrow target})s \quad (7)$$

- Apply the inverted standardization to get the final reconstructed weights:

$$w_{reconstructed} = w'_{reconstructed} \cdot \mathbf{S} + \mathbf{B} \quad (8)$$

The obtained $w_{reconstructed}$ can be used to reconstruct SVC, and the decoding accuracy of the reconstructed SVC can be used as a measure of the alignment quality. In future analysis, we swap $\mathbf{R}_{source \rightarrow target}$ by rotation matrix obtained by aligning other pairs of shapes, evaluate the resulting reconstructed SVC's performance and use it to quantify the similarity between two rotation matrices.

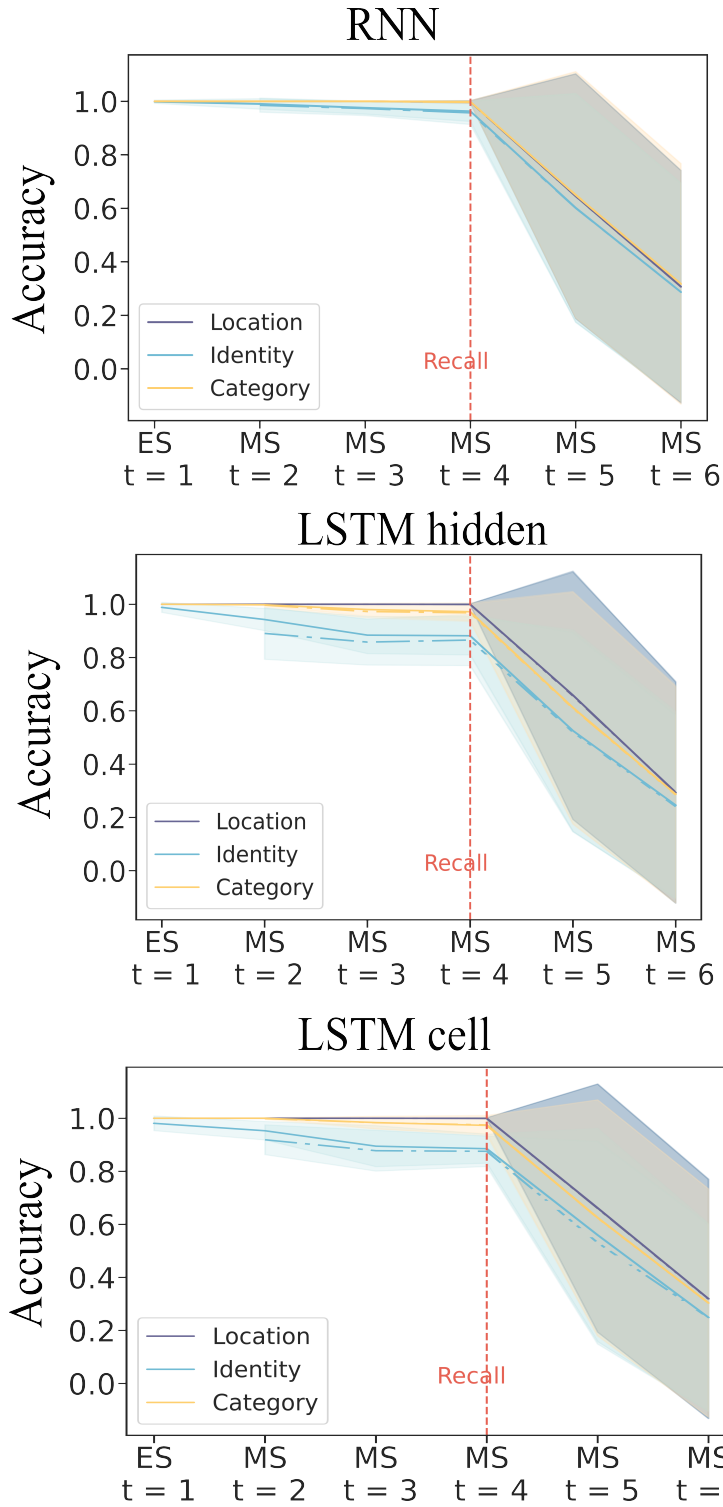


Figure A4: **Validation of Procrustes Analysis:** Similar to Figure 4 g, but for RNN and LSTM trained on MFMT




Figure A5: **Within-timestep Decoding Analysis:** At each timestep, we trained SVMs on activations from the recurrent module for task-relevant features (left) and task-irrelevant features (right), plotting the validation accuracies averaged across different feature values. The results shown are for an example GRU model trained on a multi-task, multi-feature task set. As expected, both task-relevant and task-irrelevant features were well represented at their corresponding encoding times. In addition, task-relevant features were more robustly encoded and distinctly separated compared to task-irrelevant ones.

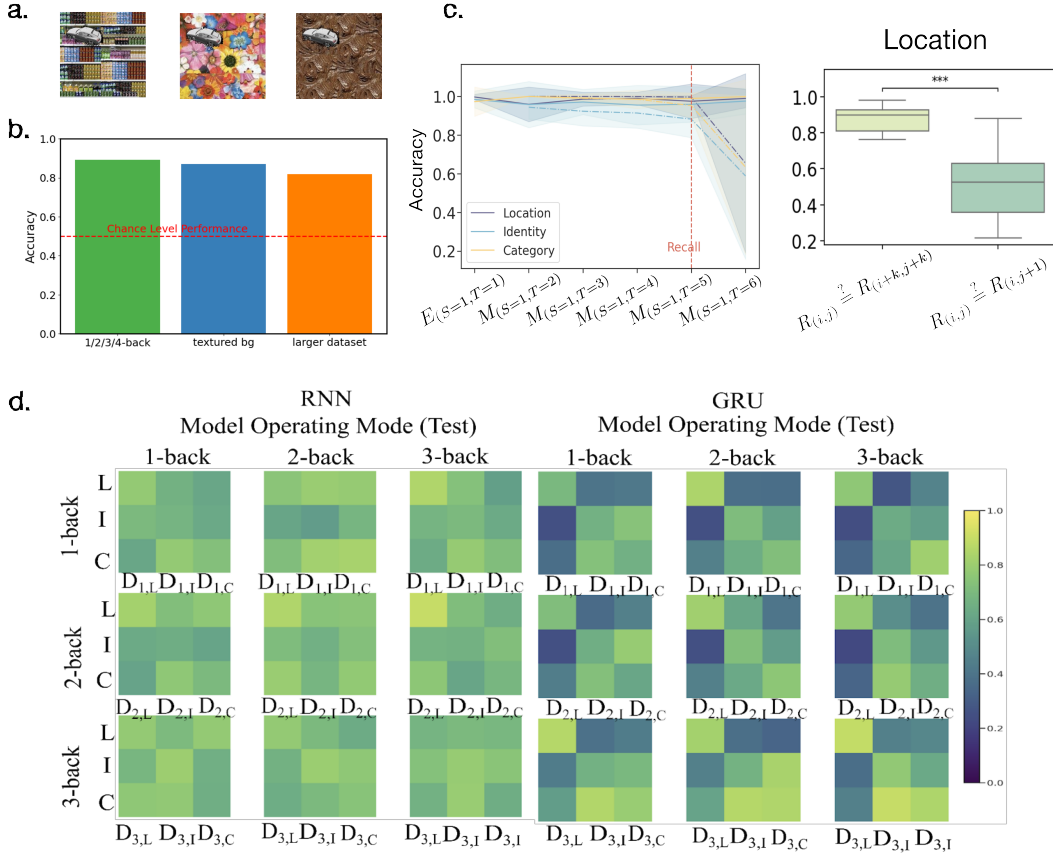


Figure A6: **a) Frames with different natural backgrounds:** We overlaid our original 3D object stimuli on synthesized natural background textures [Efros and Freeman, 2023](<https://github.com/Devashi-Choudhary/Texture-Synthesis>). Three examples of the resulting frames are shown. **b) Task performance for models trained on different datasets:** Accuracies were calculated on the validation set with novel object angles. Multi-task, multi-feature models were trained on more N-back tasks (1/2/3/4-back), visual frames with naturalistic texture backgrounds (textured bg), or tasks generated with a larger stimulus dataset (8 categories with 4 identities each; larger dataset). **c) Rotation transformation analysis:** We identified consistent patterns in the model’s response to rotation transformations across various tasks, similar to the results in Figures 4f-g. Left side plot shows an example for the MTMF model performing 4-back tasks. **d) Cross-task decoding analysis:** Consistent decoding results, with higher generalization accuracy for vanilla RNN models compared to gated models.




Figure A7: **Causality test:** We subsampled trials with matched actions and trained feature-based two-way decoders. We then perturbed the hidden states, obtained during the presentation of the first stimulus, along the direction of the corresponding decision hyperplane. The distance moved is proportional to the normalized norm vector (x-axis, distance moved in units). The perturbed hidden states were passed through the recurrent module along with the paired stimulus to compute the probabilities of the three possible actions.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: Our main claims are consistent with experimental findings.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: We point out the limitations of our model and analysis methods in the discussion.

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: [NA]

Justification: We do not have theoretical results

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: [Yes]

Justification: Although the code base is not available, all the model training details as well as analysis methods are clearly described.

Guidelines:

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

5. Open access to data and code

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

Answer: [No]

Justification: We plan to release the code base at a later phase.

Guidelines:

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

6. Experimental Setting/Details

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

Answer: [Yes]

Justification: We have provided full details to reproduce the results.

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [Yes]

Justification: We provided statistical testing results whenever necessary.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

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

Answer: [No]

Justification: Although we did not provide explicitly information about required computational resources, this is not a computational heavy project.

Guidelines:

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

9. Code Of Ethics

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

Answer: [Yes]

Justification: We comply with NeurIPS Code of Ethics.

Guidelines:

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

10. Broader Impacts

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

Answer: [NA]

Justification: This question is not applicable to our work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: This questions is not applicable to our work.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [NA]

Justification: This question is not applicable to our work.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

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

Answer: [NA]

Justification: This question is not applicable to our work

Guidelines:

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

14. **Crowdsourcing and Research with Human Subjects**

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

Answer: [NA]

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

Guidelines:

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

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

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

Answer: [NA]

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

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.