Category-orthogonal object features guide information processing in recurrent neural networks trained for object categorization

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recurrent neural networks (RNNs) have been shown to perform better than feedforward architectures in visual object categorization, especially in challenging conditions such as cluttered images. However, little is known about the role recurrent information flow plays in these computations. Here we test an RNN trained for object categorization on the hypothesis that recurrence iteratively aids object categorization via the communication of category-orthogonal *auxiliary* variables. Using diagnostic linear readouts, we find that: (a) information about auxiliary variables increases across time in all network layers, (b) this information is indeed present in the recurrent information flow, and (c) its manipulation affects task performance. These observations confirm the hypothesis that category-orthogonal auxiliary variable information is conveyed through recurrent connectivity and is used to optimize category judgements in cluttered environments.

1 Introduction

While feedforward neural networks have been successful at the task of visual object categorization, recurrent neural networks (RNNs) have been shown to outperform them in certain settings (Spoerer et al., 2020; Kubilius et al., 2018), particularly under challenging conditions such as partial object occlusion and clutter (Spoerer et al., 2017; Ernst et al., 2019). Beyond the empirical finding that recurrence can help object categorization, it remains unclear what information is conveyed by recurrent connections and what its functional role is.

In contrast to feedforward networks, unit activations in RNNs are a function of both their input and their prior activations. This enables these networks to process input in a context-dependent, conditional manner. The usefulness of these conditional computations has been observed, for example, in networks processing natural language, in which contextual words, such as negations, can 'steer' a network's trajectory in state space to process subsequent words differently (Maheswaranathan & Sussillo, 2020). In the primate visual cortex, recurrence is believed to underlie computations that can benefit from contextual signals, such as assigning local features to a figure or the background based on global shape consistency (Roelfsema et al., 2007).

In object categorization, contextual modulations would likely take advantage of the structure of real-world data, similar to how recurrence can exploit the part-whole hierarchy of visual shapes in figure-ground segmentation or the branching structure of phrases in natural language. In particular, natural images are a function of both the categories of objects therein and other category-orthogonal *auxiliary* variables, such as the objects' location, orientation or scale. An RNN could select information that is most relevant for object categorization, by first extracting information about the auxiliary variables

Submitted to 3rd Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM 2021) of the Neural Information Processing Systems (NeurIPS) conference.

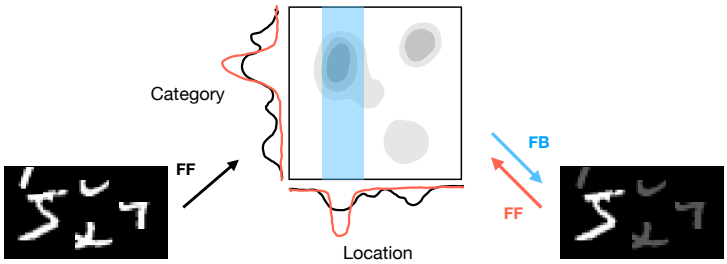

Figure 1: In cluttered images, the feedforward sweep (FF) of a recurrent neural network, trained for view-invariant object recognition, could learn to infer the location of the intact object (category-orthogonal, auxiliary variable) in addition to its category to filter out information from the irrelevant locations in the image (through feedback, FB), to improve the inference of the object's category.

and then conditioning subsequent processing on that information. This mechanism could iteratively focus on the locations and features corresponding to the object in the image and filter out irrelevant noise such as clutter. That is, auxiliary, category-orthogonal information would not be discarded due to it being non-diagnostic for the category identity, but rather recurrent connectivity would focus on this information to guide and improve performance of the main task of object categorization (Fig. 1).

To test this hypothesis, we here train and test an RNN on an object categorization task while presenting target objects in cluttered environments. We use diagnostic readouts across layers and time to characterise the presence of information related to auxiliary variables and perform in-silico causal experiments to further elucidate their computational role in object categorization.

2 Methods

2.1 Network architecture

We analyzed an RNN with two convolutional layers and one fully connected layer, as illustrated in Fig. 2A. The architecture contained both lateral and top-down recurrent connections. Top-down connections were sent from a given layer to all the previous layers, including the input. The RNN was unrolled for 4 timesteps. The activations of the fully-connected layers at the end of each timestep were concatenated and mapped to the classification output with a fully-connected layer.

2.2 Combining the feedforward and recurrent information flow

We operationalize the effect of recurrent connectivity in terms of gain modulation. For this, we summed the incoming lateral and top-down recurrent activations, and multiplied them, element-wise, with the feedforward activations. We refer to the summed recurrent activations fed to a layer as the *recurrent flow* to that layer. For a layer l at timestep t the activation was given by:

$$y_t^l = ReLU(C_{ff}(y_t^{l-1})) \odot (1 + \underbrace{C_{lat}(y_{t-1}^l) + \sum_{m>l} C_{td}^m(y_{t-1}^m)}_{\text{Recurrent flow to layer } l \text{ at timestep } t}) \tag{1}$$

where C_{ff}, C_{lat} and C_{td} correspond to the feedforward, lateral and top-down transformations corresponding to layer l respectively. Activations at the Input layer were clipped between 0 and 1, to ensure the modulated input images always remained in the same range across timesteps.

2.3 Dataset and task

In order to study the effects of recurrent information flow, we trained our RNNs under challenging visual conditions. Specifically, we generated a dataset in which the objects to be categorized were manipulated according to a number of category-orthogonal variables. Target objects were taken from the MNIST (LeCun et al., 1998) and Fashion-MNIST (Xiao et al., 2017) datasets - corresponding to 20 categories in total. The objects were randomly varied in their location, orientation and scale. In

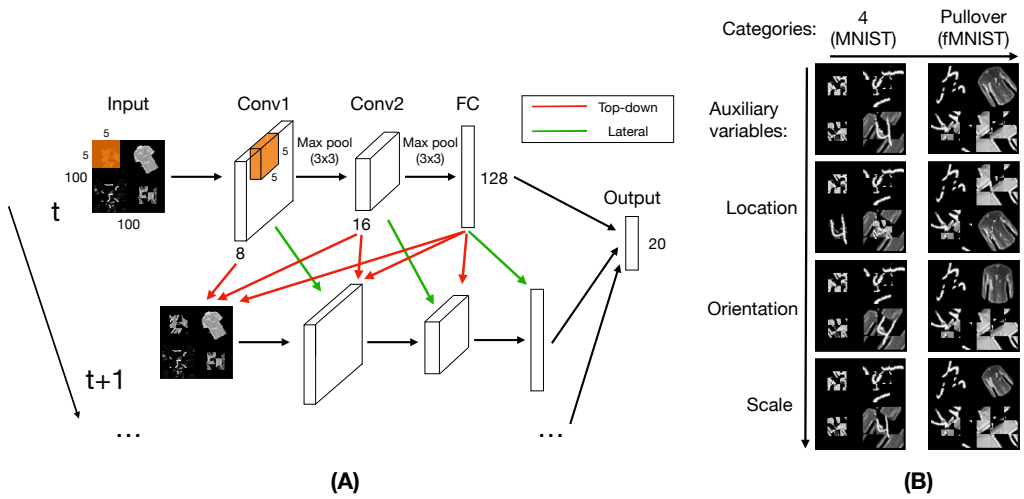

Figure 2: (A) Network architecture. The RNN with multiplicatively interacting feedforward and recurrent flows was unrolled for 4 timesteps. (B) Two example images from the dataset with corresponding transformations along the auxiliary variables.

addition, structured clutter made by randomly sampled fragments of other objects in the dataset was added to the images. Each of the auxiliary variables (horizontal and vertical locations, orientation and scale; Fig. 2B) had two possible values. In the results, the measures for the vertical and horizontal locations are averaged and *location* is treated as one auxiliary variable.

The RNN was trained for 20-way classification. 5 instances of the RNN were trained from random initializations to assess the robustness of our findings (Mehrer et al., 2020).

3 Results

The network's accuracy (averaged across 5 trained instances of the network) on a test set ($10,000$ images) was 81.2% (chance performance for 20-way classification is 5%), implying that the network successfully learned to classify the images in the dataset. Next, we analyzed the network's activations to assess whether auxiliary variables were extracted or suppressed, and how those variables affected information flow and network performance.

3.1 Category-orthogonal information is expressed incrementally over time in the RNN

To test for the presence of category-orthogonal information in the RNN activation patterns across layers and timesteps, we trained linear diagnostic readouts to predict the true value for each of the auxiliary variables (location, scale, and orientation). Results are shown in Fig. 3A, along with the categorizability of the activations (also determined by linear readout).

These results demonstrate that category information increased with layer depth (related to hierarchical processing) and with timesteps, indicating that the recurrent information flow was carrying information relevant to categorization at each timestep.

In addition, the values of all auxiliary variables could be decoded in all layers and at all timesteps, with increasing performance with increasing timesteps (averaged across layers, network instances, and auxiliary variables, with increasing timesteps: $65.2\%, 73.0\%, 82.5\%, 84.9\%$). In summary, instead of filtering out the category-orthogonal information, the network extracted auxiliary variable information incrementally with each iteration.

Importantly, auxiliary variable information (and also category information) could also be decoded, incrementally through time, from the Input. As the RNNs were set up to feed back information all the way down to the input, the effects can be easily visualised. This analysis revealed that incremental decoding of variables accompanied clutter-reduction over time (Fig. 3B).

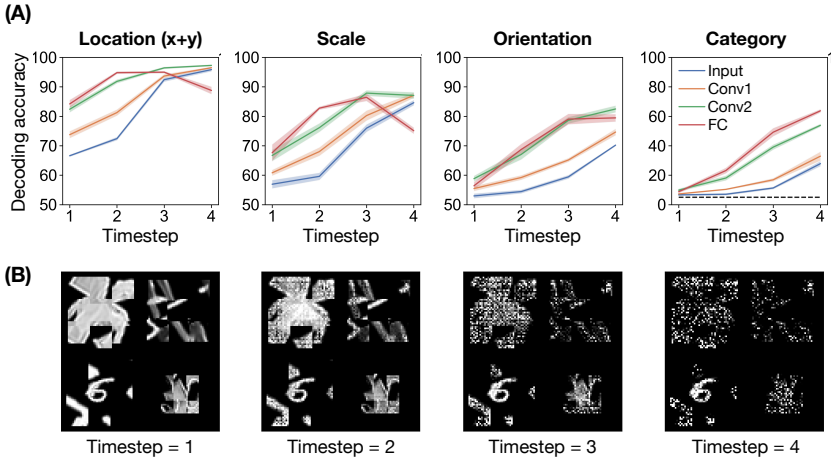

Figure 3: (A) Auxiliary variable information is observed in all the layers of the RNN at all the timesteps. 95% confidence intervals of the average (across 5 RNN instances) accuracies are shown. (B) Clutter-reduction at the Input - an emergent phenomenon.

3.2 Recurrent information flow includes category-orthogonal information that guides subsequent network inference

To determine whether the auxiliary variable information was carried explicitly in the recurrent flow, we decoded the variables from the recurrent flow at each timestep (starting from the second). We found that the information about both auxiliary variables and category increased over time in the recurrent flow to all layers, consistent with the information in the layers' activations reported earlier.

Does successful decoding of auxiliary variables from recurrent information flow imply functional impact, or is it a side-effect of layer activations with no causal role in the networks' categorization performance? To answer this question, we conducted a perturbation analysis in which we exchanged the feedback to a given layer and timepoint with feedback extracted from another, systematically perturbed, image from the dataset (Fig. 4A), i.e. feedback signalling the wrong value of the auxiliary variable. As any manipulation may alter network performance, we included a control condition in which we exchanged the original feedback signal with a randomly rotated version of the perturbation vector (Fig. 4B corresponding to the systematic perturbation, see Appendix for details). This control perturbation signals different, but not entirely misleading, information about the auxiliary variables (as it does not consistently align with changes in any of the auxiliary variables).

The functional importance of an auxiliary variable in the recurrent flow was quantified as the decrease in performance after the injection of the systematically perturbed recurrent flow, corresponding to a change in the image along that auxiliary variable, as compared to the performance after the injection of the control recurrent flow, all normalized with the performance of the original network (Fig. 4C).

Location information in the recurrent flow to the Input was functionally important at timesteps 3 and 4 (Fig. 4D), which corresponds to the major part of the clutter-reduction observed in Fig. 3B. In addition to location, information about the object features - orientation, scale, and category - was functionally relevant in the recurrent flow to the convolutional layers. Neither the auxiliary variables nor category information in the lateral information flow to FC were functionally important (data not shown). In line with this observation, ablating that lateral connection from the trained RNN led to no performance reduction. In summary, auxiliary variable information in the recurrent flows seems to causally guide the information processing in the RNN to support object categorization.

4 Discussion

We investigated the role of recurrent information flow in RNNs trained for object categorization in cluttered environments. We hypothesized that, to improve categorization performance, category-orthogonal variables are extracted, rather than filtered out, and used by the RNN to constrain subsequent information processing. Consistent with this hypothesis, we found that information about

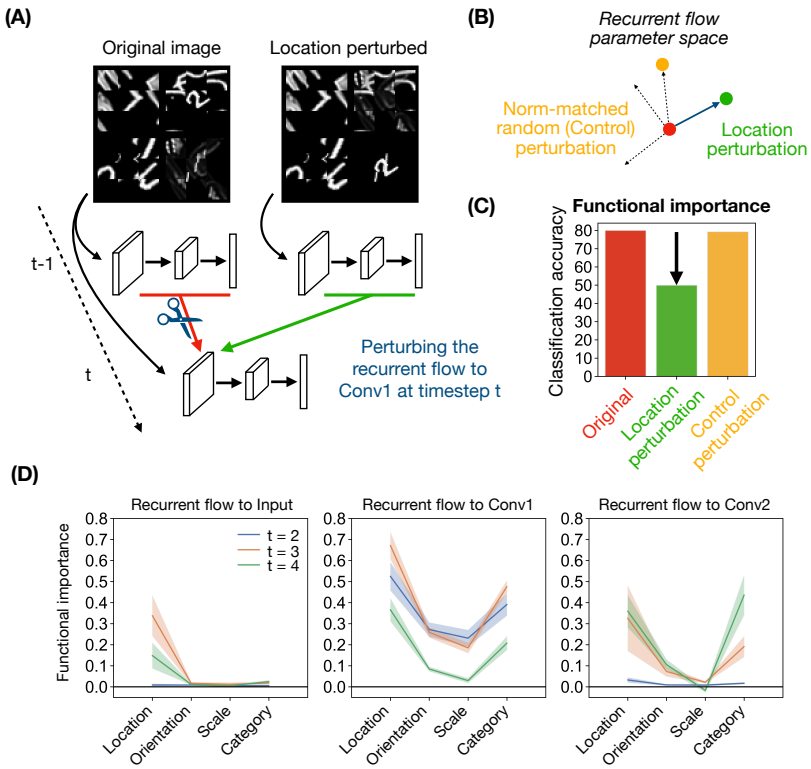

Figure 4: (A) An example perturbation: perturbing the object's location and injecting the resulting recurrent flow into the network, at Conv1, processing the original image at timestep t. (B) Illustration of a control perturbation matched in magnitude to the systematic perturbation corresponding to a specific variable (location). (C) Functional importance = (Performance of the control perturbation - Performance of the systematic perturbation) / Original performance. (D) Functional importance of perturbing the auxiliary variables at different layers and timesteps. Location information is functionally important in the recurrent flow to the Input at later stages of processing. In addition to location, orientation, scale, and category are functionally important in the recurrent flow to the convolutional layers. 95% confidence intervals of the averages (across 5 RNN instances) are shown.

all auxiliary variables was present at all network layers, that this information became more prominent across time, and that perturbing this information in the recurrent flow reduced network performance.

While object categorization has traditionally been cast as the task of extracting representations that are invariant to all category-orthogonal variables, our results indicate a functional role for such variables in selecting information for categorization. Previous research (DiCarlo & Cox, 2007; Patel et al., 2016) has proposed that optimally separating categories might require explicitly extracting auxiliary variables, and such explicit representations have been found empirically in the later layers of feedforward networks as well as in the primate visual system (Hong et al., 2016). Understanding how auxiliary variable information constrains information processing is more complicated in feedforward networks as, unlike recurrent networks, they don't possess separate architectural pathways in which auxiliary variables can interact with the category extraction in a modulatory manner. Recurrent connections might provide an inductive bias to iteratively converge to the best joint configuration of variables characterizing the target object - category and auxiliary.

While we focused our analysis on multiplicatively interacting recurrent and feedforward information, we observed similar results with additive interactions in terms of the decodability of auxiliary variables across time and layers (see Appendix). However, at the Input we could not observe similar clutter reduction, suggesting that the solution found by the multiplicative interactions for combining recurrent and feedforward information was different, and more closely resembled our intuitive notions of information selection. Multiplicative interactions have been proposed to be a useful inductive bias for conditional computations (Jayakumar et al., 2020), so clarifying their role in information selection for object categorization in more detail presents a promising avenue for future research.

References

James J DiCarlo and David D Cox. Untangling invariant object recognition. *Trends in cognitive sciences*, 11(8):333–341, 2007.

Markus Roland Ernst, Jochen Triesch, and Thomas Burwick. Recurrent connections aid occluded object recognition by discounting occluders. In *International Conference on Artificial Neural Networks*, pp. 294–305. Springer, 2019.

Ha Hong, Daniel LK Yamins, Najib J Majaj, and James J DiCarlo. Explicit information for category-orthogonal object properties increases along the ventral stream. *Nature neuroscience*, 19(4):613, 2016.

Siddhant M Jayakumar, Wojciech M Czarnecki, Jacob Menick, Jonathan Schwarz, Jack Rae, Simon Osindero, Yee Whye Teh, Tim Harley, and Razvan Pascanu. Multiplicative interactions and where to find them. *International Conference on Learning Representations*, 2020.

Jonas Kubilius, Kohitij Kar, and James DiCarlo. Task-driven convolutional recurrent models of the visual system. *Neural Information Processing Systems*, 2018.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Niru Maheswaranathan and David Sussillo. How recurrent networks implement contextual processing in sentiment analysis. *International Conference on Machine Learning*, 2020.

Johannes Mehrer, Courtney J Spoerer, Nikolaus Kriegeskorte, and Tim C Kietzmann. Individual differences among deep neural network models. *Nature communications*, 11(1):1–12, 2020.

Ankit B Patel, Tan Nguyen, and Richard G Baraniuk. A probabilistic theory of deep learning. *Neural Information Processing Systems*, 2016.

Pieter R Roelfsema, Michiel Tolboom, and Paul S Khayat. Different processing phases for features, figures, and selective attention in the primary visual cortex. *Neuron*, 56(5):785–792, 2007.

Courtney J Spoerer, Patrick McClure, and Nikolaus Kriegeskorte. Recurrent convolutional neural networks: a better model of biological object recognition. *Frontiers in psychology*, 8:1551, 2017.

Courtney J Spoerer, Tim C Kietzmann, Johannes Mehrer, Ian Charest, and Nikolaus Kriegeskorte. Recurrent neural networks can explain flexible trading of speed and accuracy in biological vision. *PLoS computational biology*, 16(10):e1008215, 2020.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No] The work features a set of exploratory analyses of how a recurrent neural network functions. No claims about the exhaustiveness of the analyses are made.
   (c) Did you discuss any potential negative societal impacts of your work? [No] To our knowledge there are none.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] Code and data will be uploaded upon acceptance.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] We used an internal high-performance cluster for training and testing the networks.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No] Both the MNIST and fashion-MNIST datasets are freely available for re-use with citation.
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Appendix

A.1 Dataset generation

Each image in our dataset was 100×100 pixels in size. Each image contained one intact target object (originally 28×28 px), rotated and scaled. 7 other *scrambled* objects comprised the clutter, each of which were also rotated and scaled, then divided into 9 square blocks, and permuted randomly. The image was divided into four quadrants, corresponding to the two horizontal and two vertical locations. Each quadrant contained two objects - one overlaid on top of the other. One of the quadrants contained the intact object which was always overlaid on the scrambled object, to ensure it wouldn't be completely occluded while still being challenging to parse from the background. Both target and scrambled objects could be drawn with equal probability from either the MNIST or Fashion-MNIST dataset.

The horizontal and vertical location of the target object was chosen randomly for each image, meaning that the object would be located at the center of one of the four quadrants, jittered both horizontally and vertically between -2.5 and 2.5 pixels (uniform distribution). The orientation of the objects (both target and scrambled) was randomly chosen to be either 30 degrees clockwise or counterclockwise, with a uniformly distributed jitter of ± 5 pixels. Scaling was either 0.9 or 1.5 times the original object size, with a jitter of ± 0.1.

The split between train, validation and test sets was done by randomly drawing images only from the corresponding splits of the original datasets. The datasets provided the $110,000$ possible source images for the training set (MNIST and Fashion-MNIST combined), $10,000$ for the validation set and $20,000$ for the test set. Training images were generated on the fly for each training batch. The test set contained $10,000$ images.

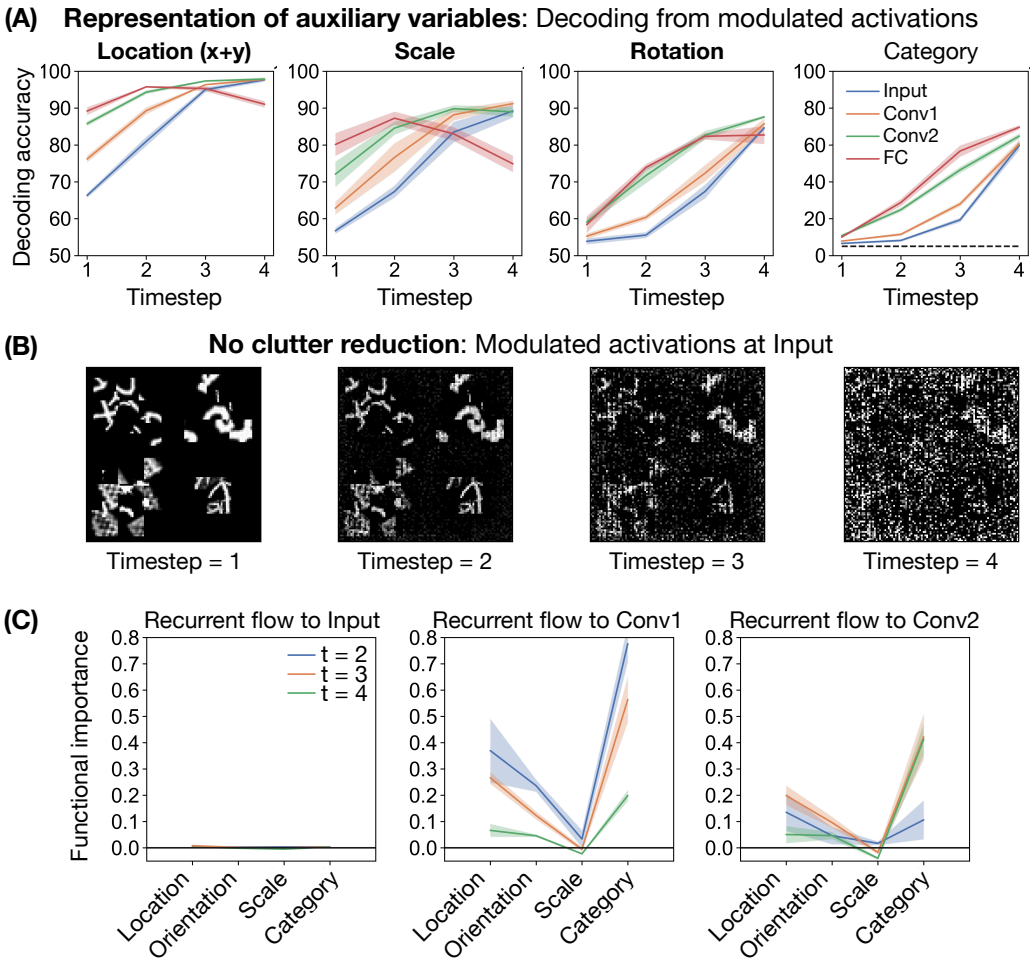

Figure S1: (A) Auxiliary variable information in all layers of the network with additive recurrent interactions, at all timesteps. 95% confidence intervals of the average (across 5 RNN instances) accuracies are shown. (B) No clutter reduction is observed at the Input. (C) Functional importance of different variables at each layer and timestep. Consistent with the lack of clutter reduction, no variable was found to be functionally important in the recurrent flow to the Input.

A.2 Architecture details

The RNN contained two convolutional layers, both with 5×5 kernels, stride of 1, no padding, and with 8 and 16 channels respectively, and one fully-connected layer with 128 units. 3×3 max pooling was applied after each convolutional layer. The lateral recurrent connections from each convolutional layer to itself were convolutional layers with the same number of input and output channels (8 and 16), kernel size 5×5, stride 1 and padding 2. The lateral connection for the fully connected layer was a fully connected layer with 128 units. The top-down connections from convolutional layers to the input and other convolutional layers were transposed convolutional layers, upsampling to the size of the target feature map: Conv1 to Input had kernel size 7×7, stride 3, and no padding; Conv2 to Input had kernel size 20×20 stride 10 and no padding; Conv2 to Conv1 had kernel size 16×16, stride 10 and no padding. Top-down connections from the fully-connected layer were fully-connected layers whose outputs were restructured to match the convolutional layers.

A.3 Training the RNN

The network was trained for 20-way classification using a cross-entropy loss between the network's output and the correct object category. The training images were generated on the fly. We used the Adam optimizer for training, with a batch size of 32, and learning rate of 10^{-4} (after manual tuning

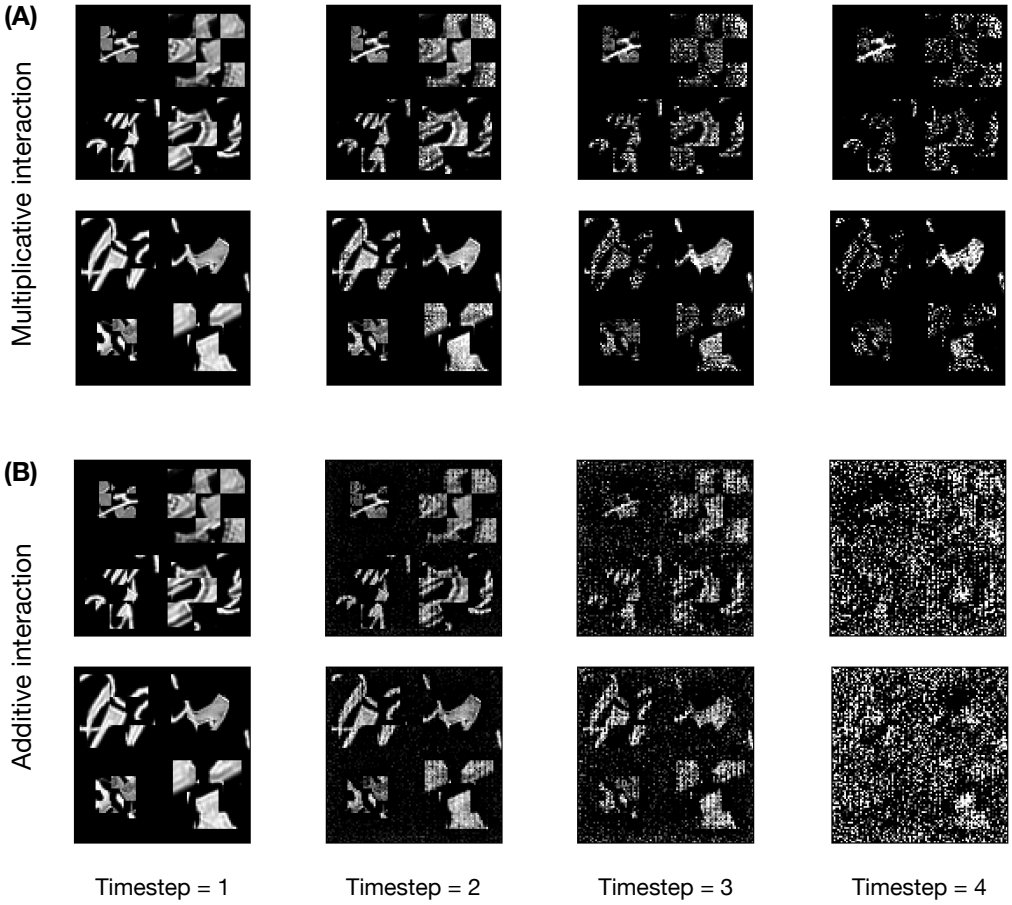

Figure S2: (A) More examples of input-level modulation for the RNN with multiplicative interactions. (B) Examples for the RNN with additive interactions. Both the RNNs with multiplicative or additive interactions classified both these images correctly, but with different strategies.

to get the best validation accuracy). The network was trained for $300,000$ iterations. Before training, the weights were initialized using Xavier initialization and the biases were initialized as zeros.

A.4 Decoding approach

To measure the amount of explicit information about a given variable present in the activations of each layer at each timestep, we trained and tested a linear classifier (a support vector machine, SVM) to classify the correct value of that variable. We used the Python scikit-learn function `svm.LinearSVC()`, with default parameters, which is able to handle binary as well as multi-class classification, which was needed to decode category. Each classifier was trained with 800 images and tested on 200 images (all randomly drawn from the test set). This procedure was repeated 10 times (accuracies averaged), for each of the 5 RNN instances.

A.5 Perturbation analysis

To generate the *perturbed* feedback to the network at a given timestep t and layer l, we ran two copies of the network. One was fed the original input image, and the other the same image, altered in one particular variable (auxiliary or category). For example, vertical location could be changed from top to bottom (Fig. 4A). From this network we extracted the perturbed recurrent flow to l at t, and then fed this recurrent flow, at timestep t, to layer l of the network which had received the original image, thus combining the activations for the original image with the recurrent flow resulting

from the perturbed image. We then compared this network's systematically perturbed classification accuracy to the network's accuracy for the same image without any perturbations.

To control for the unique role of auxiliary variables in determining network performance, we also measured the accuracy of a network that received the same perturbed recurrent flow, but with the elements of the recurrent flow vector randomly permuted (control recurrent flow). This corresponds to a recurrent flow that differs from the original one by the same amount (vector length) as the relevant perturbed recurrent flow, but along a direction that, on average, does not align with any of the auxiliary variables (Fig. 4B).

The functional importance score ($\frac{control - systematic\ perturbation}{original}$, Fig. 4C) for each variable manipulation was computed as the average across 1000 images (randomly drawn from the test set). This procedure was repeated 5 times (scores averaged), for each of the 5 RNN instances.

A.6 Results for additive recurrent interactions

The profile of results for the network with additive recurrent interactions had both similarities and differences to the multiplicative ones. The decoding of variables at all layers and timesteps showed a similar profile (Fig. S1A), but in visualizing the modulation at the input image level, no clutter reduction comparable to that in the multiplicative network was visible (Fig. S1B). Consistent with this observation, no variable was found to be functionally important in the recurrent flow to the input, while other layers showed a pattern of variable importance similar to that of the network with multiplicative recurrent interactions (Fig. S1C).