Recurrent circuits improve neural response prediction and provide insight into cortical circuits

Harold Rockwell[1 4*], Sicheng Dai[1 3*], Yimeng Zhang[1 2*], Stephen Tsou[1], Ge Huang[1], Yuanyuan Wei[1 3], Tai Sing Lee[1 2 3†]
[1]Neuroscience Institute, [2]Computer Science Department, [3]Biomedical Engineering Department, Carnegie Mellon University, Pittsburgh, PA
[4]Committee on Computational Neuroscience, University of Chicago, Chicago, IL
(* denotes equal contribution) († corresponding author: taislee@andrew.cmu.edu)

Virtual poster QR code

Summary

Inspired by the extensive recurrent connectivity in cortex, we added recurrent structure to standard convolutional neural network models for neural response prediction. We find that recurrent models consistently outperform their parameter- and hyperparameter-matched feedforward models, particularly with limited data, and particularly with wide-field and long stimulus presentations that may evoke more recurrent processing in cortex. To understand this finding, we propose a new formulation of recurrent models as ensembles of weight-tied feedforward processing paths, and show that the "more recurrent" datasets' models are more reliant on longer paths. In addition, we test the models for surround suppression, a response property assumed to be implemented by recurrent connections in cortex, and find that while it is present in both feedforward and recurrent models, the recurrent models learn to decompose it into the recurrent connections. Together, our findings suggest that recurrent network models, combined with stimulus presentation paradigms that evoke recurrent processing, are effective tools to probe the recurrent circuits of visual cortex.

Three Datasets

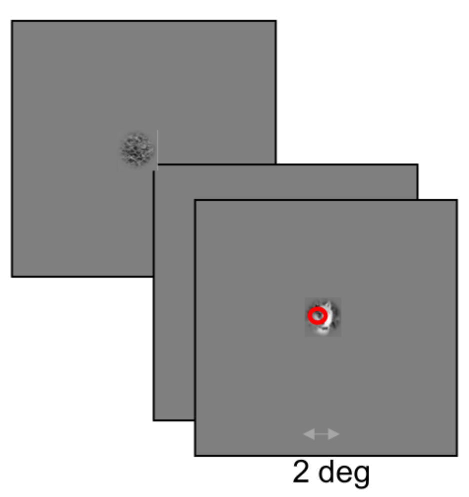

Receptive-field **S**hort presentations **(RF-S)**

- 2 degree aperture
- 60 ms presentation
- 7250 images
- 115 V1 neurons
- Used with permission from the work of Cadena et. al. (2019)

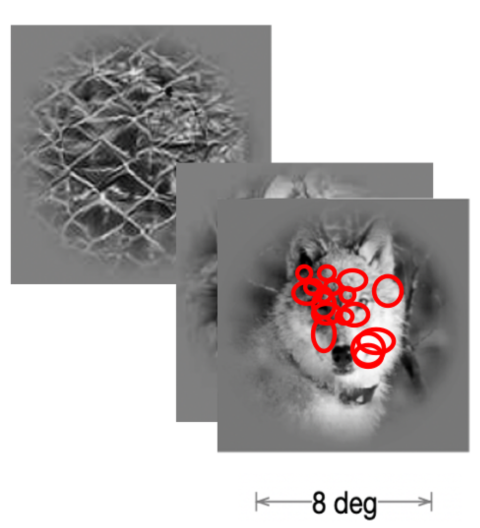

Wide-field **S**hort presentations **(WF-S)**

- 8 degree aperture
- 47 ms presentation
- 8000 images
- 79 V1 and V2 neurons

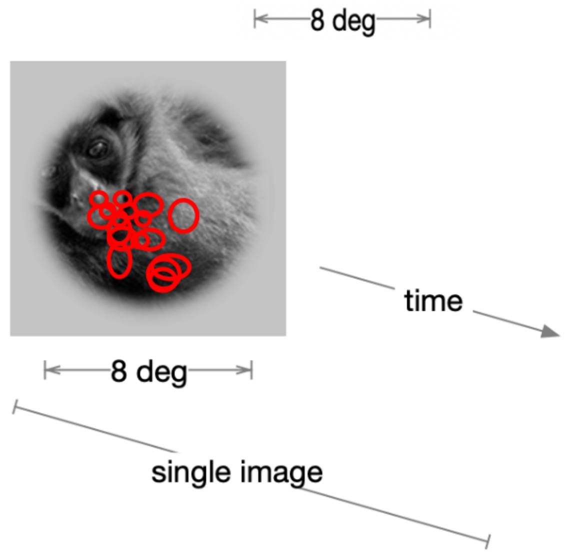

Wide-field **L**ong presentations **(WF-L)**

- 8 degree aperture
- 500 ms presentation
- 2250 images
- 34 V1 and V2 neurons

Monkeys passively viewed stimuli while fixating for a juice reward. Spike counts during presentation were averaged over 4-10 trials per stimulus and used as targets for the neural network models.

References

Cadena et. al., 2019. PLOS Comp. Biol.
Klindt et. al., 2017. NeurIPS.
Schoppe et. al., 2016. Frontiers in Comp. Neuro.
Spoerer et. al., 2017. Frontiers in Psychology.
Zhang et. al., 2022. arxiv.org/abs/2110.00825

Feedforward and Recurrent Convolutional Network Models

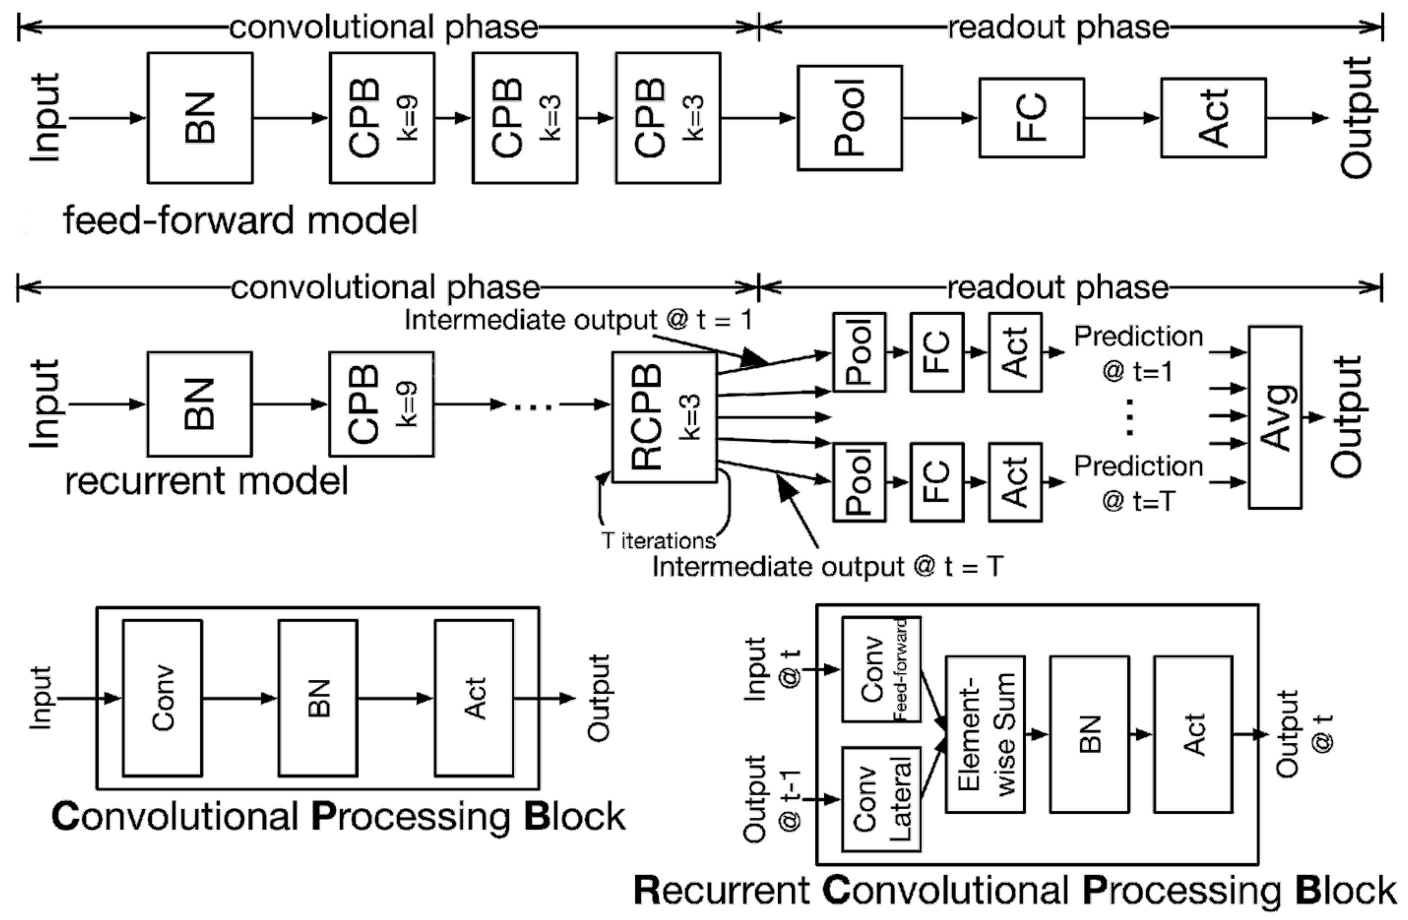

We trained convolutional neural network models via gradient descent to predict the spike counts of recorded neurons from the pixels of the presented stimuli. Models were trained on each dataset to predict the activity of all neurons in that dataset with a shared computational core. Each feedforward model had a parameter-matched recurrent model with one RCPB to match two CPBs in the feedforward model. We trained at least 288 models on each dataset, exhaustively exploring hyperparameter combinations to ensure results that were robust to a wide variety of choices, including activation function, loss function, number of channels, and others.

Results: Consistent Performance Improvements

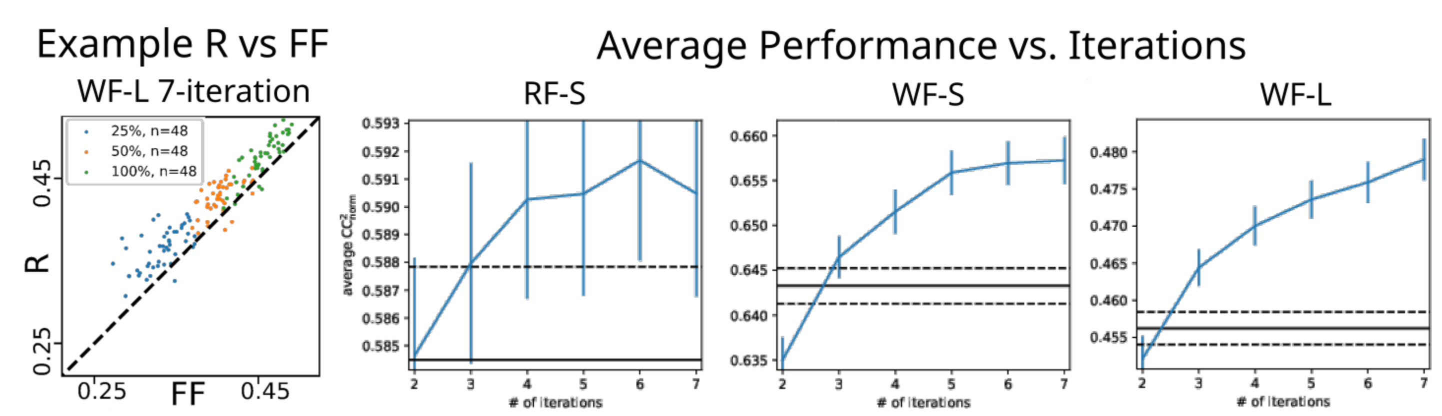

With the full amount of data, recurrent models improve over feedforward ones, averaged over the hundreds of hyperparameter combinations, for all datasets. These improvements increase with the number of iterations trained, and for the the WF-S and WF-L data, increase further with lower amounts of training data. In addition, the benefit of recurrence increases as the stimulation paradigm changes to evoke more recurrent processing, with WF-L improving more than WF-S, and WF-S improving more than RF-S (even when only V1 cells are included).

Table 1: average % improvement over matched feedforward model

(data %age)	RF-S	WF-S	WF-S (V1-only)	WF-L
25%	-4.3±1.0	5.2±0.6	5.2±0.6	10.1±0.1
50%	1.5±0.6	3.5±0.2	3.1±0.3	5.3±0.5
100%	1.0±0.2	2.1±0.2	1.7±0.3	5.0±0.4

Acknowledgements

This work was supported by NSF grant CISE RI 1816568 and NIH RO1 EY030226-01A1.

Multi-Path Reformulation

Recurrent models are well-approximated by weight-tied feedforward networks with multiple, parallel computational "paths" for each iteration. The multi-path models' simplified computational graph allows explicit decomposition of the contributions of paths of different lengths.

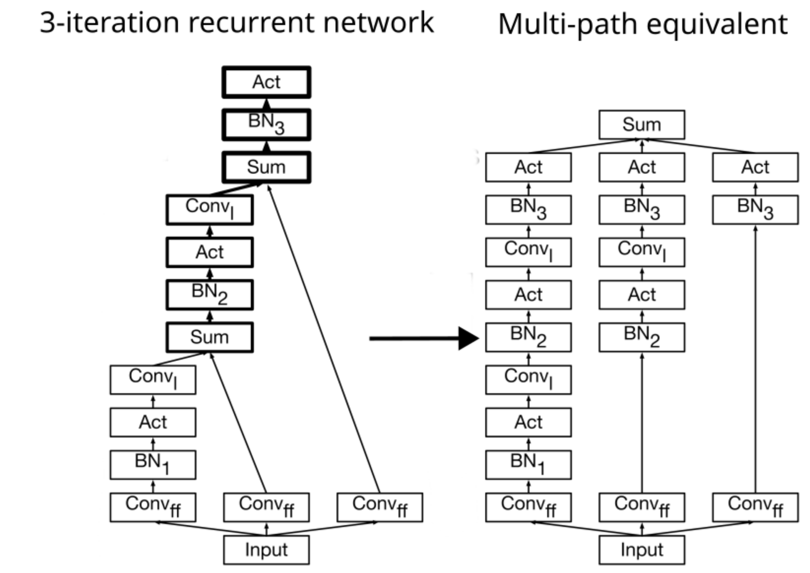

Results: Cortex-Like Recurrent Circuits

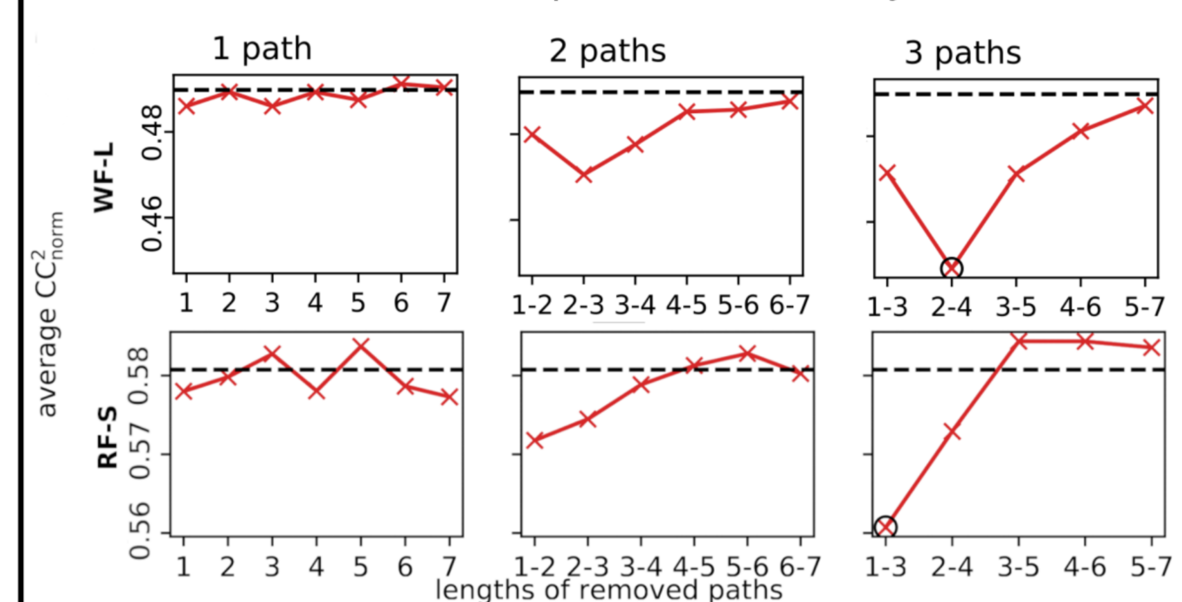

Multi-path models retrained with some paths missing show varying effects on performance. Removing any one path has little effect, as paths of similar length may compensate. Removing 3+ shorter paths has the greatest deleterious effect. Performance on the RF-S data is more dependent on the shortest paths than the WF-L data is.

Physiological response properties

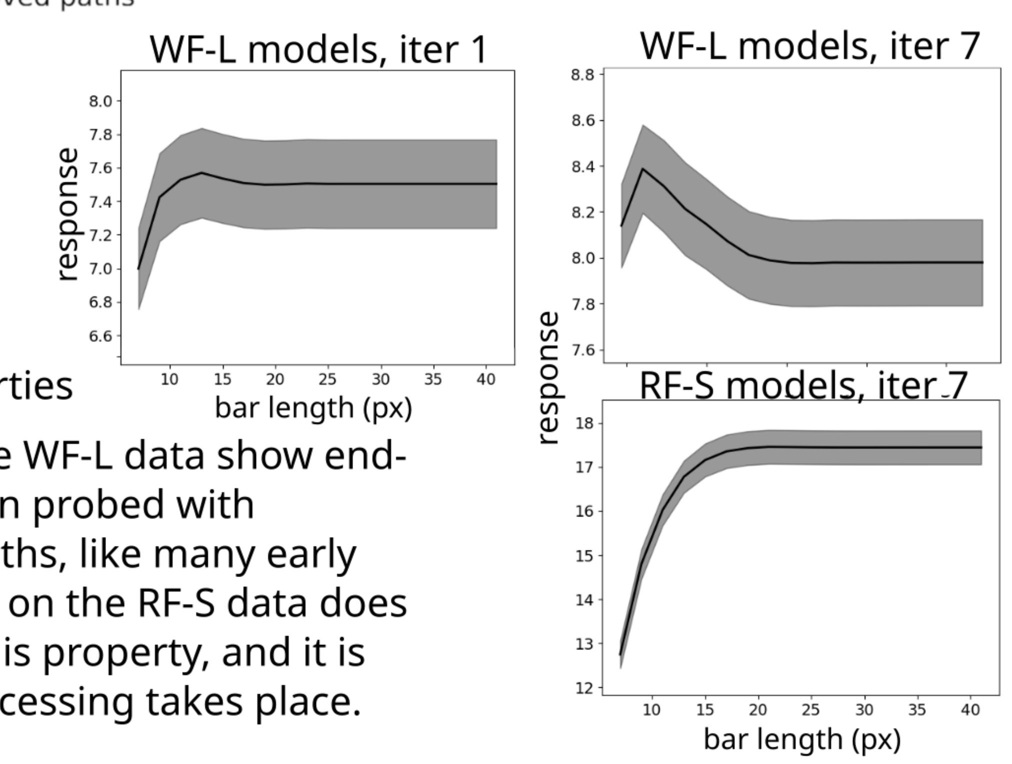

Recurrent models trained on the WF-L data show end-stopping tuning properties when probed with oriented bars of increasing lengths, like many early visual cortical neurons. Training on the RF-S data does not result in model units with this property, and it is only present after recurrent processing takes place.

Conclusions

- Recurrent models provide better neural response prediction (consistently across an extensive set of hyperparameters) than feedforward ones, particularly in the low data regime and with longer, wider stimulus presentations.
- By re-expressing the computational graph of the recurrent models as a collection of parallel feedforward paths, we can show that longer processing paths are more important for models trained on data from longer, wider stimulus presentations, and that the ensemble of diverse paths in the computational graph provided by recurrence may allow for more flexible computation.
- These models also display the well-known physiological effect of end-stopping, unlike the ones trained on short, receptive-field-sized stimulus presentations.
- Our results suggest that recurrent networks can be effective tools for characterizing the response properties of early visual cortical neurons, and particularly the response properties resulting from recurrent activity, if an appropriate stimulus presentation paradigm is used.