Reset Networks: Emergent Topography

by Composition of Convolutional Neural Networks

T. Hannagan[1]

[1] Connecticut Institute for the Brain and Cognitive Sciences, University of Connecticut

Abstract

We introduce Reset networks, which are compositions of several neural networks - typically several levels of CNNs - where outputs at one level are gathered and reshaped into a spatial input for the next level. We demonstrate that Reset networks exhibit emergent topographic organization for numbers, as well as for visual categories taken from CIFAR-100. We outline the implications of this model for theories of the cortex and developmental neuroscience.

Introduction

CNN classifiers are scalable and high-performing deep learning models, that have now been shown beyond any reasonable doubt to predict activity in the visual system. However, a very salient aspect of the latter is the spatial layout of areas selective for faces, houses and other objects found deep in ventral Occipitotemporal cortex [1], and which cannot be explained within the standard CNN classifier framework. This is because categorical areas respond to high level features and yet are spatially extended objects in vOTC, whereas by design, CNN classifiers trade-off spatial dimensions for feature channels as information is fed-forward. In the deepest layers of the network, where features are the most complex, they also have little if any spatial arrangement left.

Another limit of the current CNN-to-visual-cortex mapping endeavor is that most if not all studies attempt to predict cortical responses from a single deep CNN classifier, trained on a single task. Though understandable, these two simplifications nevertheless make the model qualitatively quite different from the visual system, which is shaped by many different tasks other than classification (e.g. visual tracking, naming), and involves different processing streams.

In this article, we take a modeling approach that attempts to factor in the multiplicity of processing streams and versatility of tasks that are characteristic of the visual system. We show that requiring the outputs of many CNNs to serve as input to other CNNs downstream is sufficient for topography to emerge.

Reset networks

Reset networks are compositions of several neural networks - typically several levels of CNNs - whose outputs at one level are reshaped into a spatial input for the next level. They implement a sequence of neural spaces where networks performing similar computations end-up being neighbors, as do units that are selective to the same input.

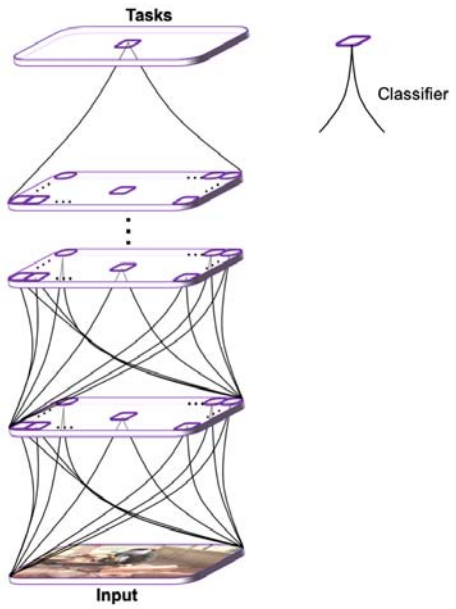

Figure 1. *A Reset network is a differentiable neural system with an arbitrary number of levels, where each level itself consists of a spatial arrangement of deep neural networks.*

The general form of a Reset network is shown in Figure 1. It has an arbitrary depth of levels, each consisting of several networks operating in parallel on the same input. The next three requirements can be relaxed, but will be followed in the remainder of this article. At any level, all networks are independent processors: they do not share any weight parameters and do not project to each other laterally. All networks in a given grid also receive, as a common input, the entirety of the level below. The last level in a reset network is the only output level, where error signals for all k tasks are received.

Reset networks include in particular the family of depth 2 shown in Figure 2, where level 1 is obtained by reshaping and concatenating the outputs of nxn parallel networks into a single map, called "grid" hereafter, which then serves as input for a final network. We refer to such systems as Reset Networks of depth **2** and width **n**, or Reset(n), we will sometimes also write Reset(n, m) to further specify the grid's width **m** in terms of units.

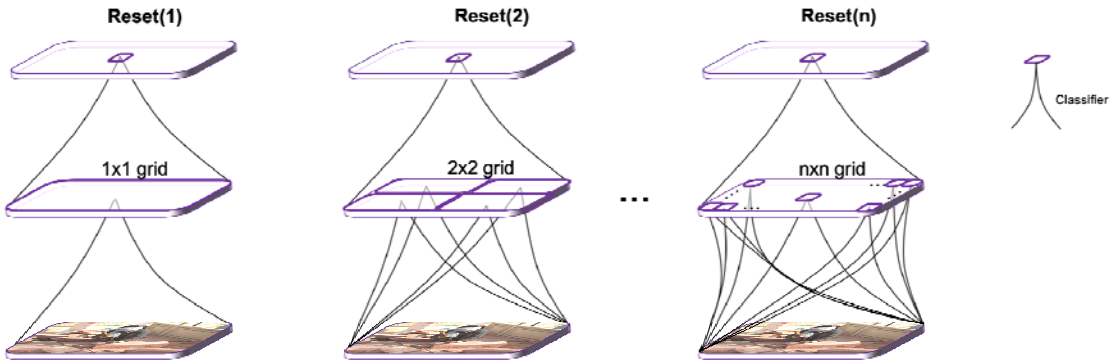

Figure 2. *A family of depth 2 Reset networks, with nxn intermediate grids for increasing n.*

We demonstrate that Reset networks can perform classification and regression at scale while also exhibiting emergent topographic organization. Our code is available on GitHub: https://github.com/THANNAGA/Reset-Networks.

Why are Reset networks relevant to cortical topography?

Cortical topography in the strict sense is the notion that "nearby neurons in the cortex have receptive fields at nearby locations in the world" [2]. When understood as applying also to local fields or voxels as well as to neurons, this is a widespread phenomenon in the brain, imaged throughout the visual cortex as well as in some associative areas.

Despite the aforementioned architectural tension between CNNs and vOTC, recent innovative work has shown that categorical areas can indeed be simulated in Topographic Deep Artificial Neural Networks, or TDANNs [3]. In TDANNs, topography is achieved by invoking a separate entity – dubbed “cortical tissue map” – assigning arbitrary locations on this map to units in the dense layer of the network, before introducing a loss regularizer that penalizes wiring length on the map during training. Since the mechanism realizing this mapping is unspecified, the ontological status of space in the model is problematic. Two different notions of space appear to exist that can contradict each other: the spatial layout of convolutional feature maps in the model, and the spatial dimensions of the cortical tissue map. This issue is not brought to the forefront in extant TDANNs, because cortical tissue maps are restricted to the upper dense layers of the model, where locality is lost. However, there is no reason why cortical tissue maps couldn’t also be invoked for the lower, convolutional levels of the network, with much less interpretability.

By contrast, the way Reset networks achieve topography at scale is conceptually straightforward.

Results

Topography for numbers in parietal cortex

In parietal cortex, voxels selective for similar numbers are more likely to be contiguous, a phenomenon which has been partially explained as a spatial diffusion process of number codes [3]. A Reset Network with a single 8x8 grid can be trained to map images of numbers onto number codes, and succeeds in reproducing topographic organization on the grid.

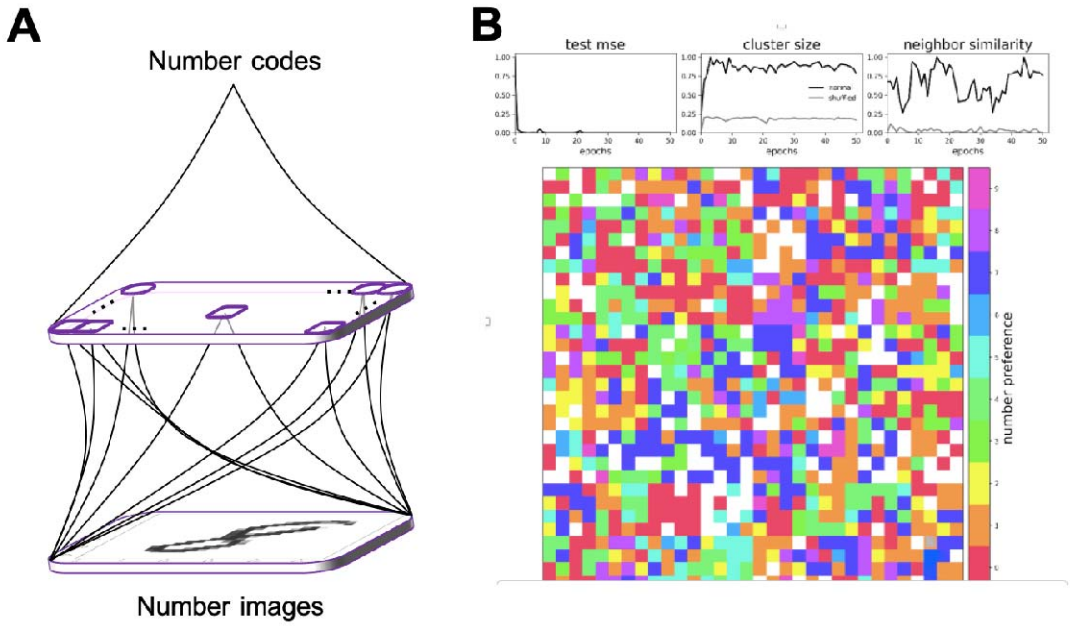

Figure 3. *(A) A Reset network with 8x8 grid and depth 2 is trained to map images of numbers onto number codes. (B) Number preferences on the network grid show topography, which as training progresses and mean squared error on the test set decreases (upper left plot), can be measured by cluster size (upper middle plot) and neighbor similarity (upper right plot). Cluster size and neighborhood index are both significantly above the control values obtained for shuffled maps.*

Figure 3 (left) shows a Reset(8, 32) network, trained for 50 epochs to regress number images onto the number vector codes of [4]. Each subnetwork is a Resnet20 network. The details of the number dataset, training procedure and network analysis can be found in Supplementary material S1 and S2.

Topography is visible on the map of number preferences on Figure 3 (bottom right), and is quantified in the plots above. It can be seen that topography (upper middle plot) and neighborhood similarity (upper middle and right plots) are both significantly above the control levels, which are computed from shuffled selectivity maps.

Topography for numbers can be seen to emerge quickly during training. Also notable is the tendency of subnetworks as a whole to specialize for specific numbers, or for numbers in the same ballpark. This is particularly striking when considering the training histories for Reset(2, 32) and Reset(4, 32) on the same task, available at https://github.com/THANNAGA/Reset-Networks/tree/main/Topography%20for%20numbers.

Topography for CIFAR-100 and categorical areas in vOTC

In ventral occipitotemporal cortex, more than two decades of studies have established the presence of areas selective for various widespread visual categories, in particular faces, bodies, tools, houses, and words. While there is no shortage of computational models able to reproduce many characteristics of the visual system, including some of vOTC, only one [2] arguably achieves both topography and scale at the same time - with topography being problematic as it requires two different notions of space to coexist. By contrast, the way Reset networks achieve topography at scale is conceptually straightforward.

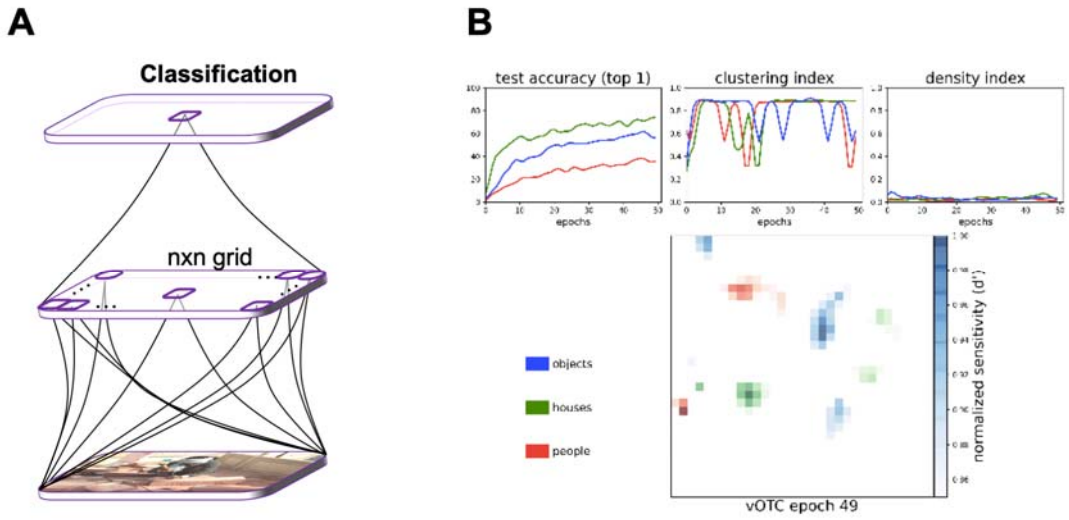

Figure 4. *(A) A depth-2 Reset network is trained on CIFAR-100. (B) Unit preferences on the network grid show clusters for objects, houses and people. These are quantified by means of a clustering index (upper middle plot): proportion of a unit's neighbor which are selective to the same unit), and a density index (upper right plot): proportion of the total activity on the grid that is devoted to each category.*

The left panel in Figure 4 shows a Reset network classifier trained on Cifar-100. The right panel shows category preferences on the grid after training. Only 3 super-categories are considered – objects, houses and people – which were obtained by aggregating the relevant Cifar-100 classes. Clustering is visible in the map, and quantified in the subplots above, although with slightly different measures as before for number topography.

Discussion

The main insight of Reset networks is that local processing at level L exerts a pressure on the grid of networks at L-1 to organize in order to solve the task, distributing work in a way that creates topography. We now discuss some outstanding issues, questions and prospects.

Classification performance

We have shown that Reset networks can classify standard computer vision datasets such as CIFAR-100. However and as the figure below shows, at this stage their performance remains disappointing, only at best matching that of a single Resnet 20, while having many more parameters.

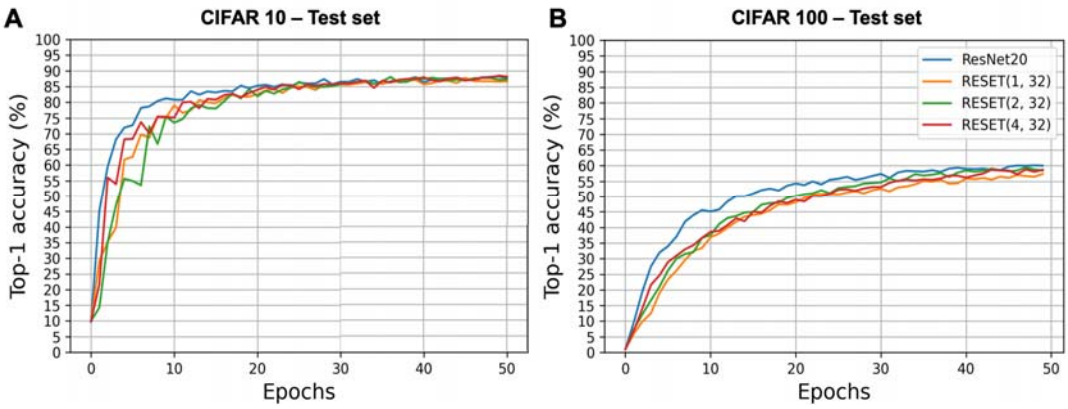

Figure 6. *Top-1 accuracy of Reset Networks on CIFAR-10 (A) and CIFAR-100 (B). Composing several CNN classifiers does not destroy classification abilities, though unsatisfyingly, Reset networks can currently at best only converge to the same performance as a single Resnet20. The notation Reset(n, 32) specifies that the outputs of nxn networks are reshaped into a map of 32x32 units.*

One reason for this could be that in our simulations, spatial resets between levels were always done by reshaping the subnetworks' outputs, which constitute an information bottleneck. Reshaping prior to the subnetwork's output, e.g. the dense layer or before, might be a more astute choice. We also observe that the full resources of the Reset network don't seem to be used: some subnetwork units are more active than others. This can be alleviated to some extent by using dropout, or another kind of regularization on the grid.

Regularization by auto-encoding

In the course of our investigations (see Supplementary Material 3), we have observed that Reset networks performed much better when the second level had 2 subnetworks: one that classified the input, and another that tried to reconstruct the input from the grid. Auto-encoding in this situation appears to act as an efficient regularizer for classification, forcing activation to be distributed across the whole grid rather than to be drawn by one, or just a few subnetworks. Such regularization effects of auto-encoding have been reported before for standard classifiers [5]. The novelty in Reset networks is that input reconstruction must be accomplished using the information from the whole grid: this suggests that in visual cortex, some feedback connections between distal cortical areas actually function as regularizers of cortical spaces.

Topography

Reset networks constitute a novel mechanism for topography to emerge in deep learning. They also provide a way to implement the mapping between foveal/peripheral input and lateral/medial in visual cortex, a gradient which is not easily captured within the standard assumptions of CNNs.
We have presented firm evidence that Reset networks can reproduce at least two examples of topographic organization: in parietal cortex, when regressing number images into number codes, and in ventral Occipitotemporal cortex, when classifying CIFAR images into classes. The fact that topography appeared for different tasks and inputs suggests that it is specific to neither of those, but inherent to the model's architecture. Consequently, we would also expect topography for numbers, even if a Reset network was trained to classify natural images into number classes. It would also be interesting to find out whether topography for numbers could emerge in an auto-encoding Reset network, in absence of any teaching signal related to number.

Adding networks when necessary: the width and depth of Reset networks

The proposed approach aligns well with a view of neural development in which, as an alternative to recycling neural material, new resources are recruited if needed. Learning a new task could require only to widen the system by adding a network at the current level, with different networks possibly trained on different tasks. If expertise from previously learned tasks is required, the system could be made deeper by reshaping network outputs at the current level and creating a new level. In order to really contribute to continual learning theory, it is now necessary to better specify the mechanisms of network growth within the Reset network approach, that would prevent interference of old and new functions.

VOTC topography and the Visual Word Form Area

A closely related topic is that of the so-called Visual Word Form Area, which, with the benefit of insight and despite its discoverers' best intent upon naming it, is neither visual (congenital blind subjects have it too), word-level specific (it is also active for individual letters), nor actually a single area (it appears to be organized in patches). But names have great inertia, and this one does convey well the idea of a localized region selective for stimuli related to words. While some efforts have gone into modeling the VWFA [6], currently no model can account for its specific place within the topography of vOTC. The network in Figure 6 describes what a Reset network of vOTC and the VWFA could look like.

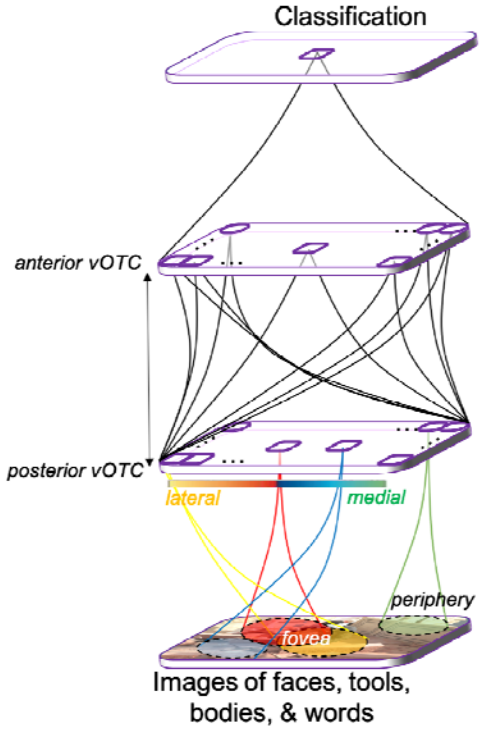

Figure 5. *Design of a Reset network of depth 3 for modeling vOTC and the VWFA. The network has 2 intermediate grids, standing for posterior and anterior vOTC. Networks in the posterior grid see different parts of the input depending on where they are: left-located ("lateral") networks on the grid receive input from the center of the image, whereas right-located ("medial") networks receive input from its periphery.*

First, this Reset network would have 2 intermediate grids, P and A, standing for the posterior and anterior axis in vOTC. This is not an innovation, but now Reset networks allow for something interesting to happen. In addition to the posterior-to-anterior gradient, we can capture a lateral-to-medial gradient by ensuring that networks in the P grid see different parts of the input depending on where they are: left-located (lateral) networks on the P grid would receive input from the center of the image, whereas right-located (medial) networks would receive input from the periphery. In other words, we build into the model a lateral-to-medial gradient in vOTC by exploiting its well-documented correspondence with center/periphery processing [7]. Such a relation cannot easily be built into a CNN, because of location invariance.

Conclusion

Reset networks show that topography must emerge in deep CNN classifiers, when composed with one another. In this view, the cortex should not be modeled as a single classifier, however deep and richly organized, but as a sequence of levels of neural network classifiers. This in turn rests on the idea that the cortex has the ability to compose networks when necessary, and it predicts that the outputs, or the late

computational stages, of cortical classifiers are either spatially organized, or somehow reshaped spatially during the course of composition.

Acknowledgments

We thank Thibault Fouqueray for early and stimulating discussions where this idea was seeded, and Florence Bouhali for ongoing discussions, original ideas and encouragements on Reset networks. We also thank Hyodong Lee for exchanges that provided useful information on Topographic Deep Artificial Neural Networks.

References

[1] Grill-Spector K, Weiner KS. (2014) The functional architecture of the ventral temporal cortex and its role in categorization. Nat. Rev. Neurosci. 15, 536–548.

[2] Patel GH, Kaplan DM, Snyder LH. Topographic organization in the brain: searching for general principles. Trends Cogn Sci. 2014;18(7):351-363. doi:10.1016/j.tics.2014.03.008.

[3] Lee H, Margalit E, Jozwik KM, Cohen MA, Kanwisher N, Yamins DL, DiCarlo JJ. Topographic deep artificial neural networks reproduce the hallmarks of the primate inferior temporal cortex face processing network. 2020 bioRxiv.

[4] Hannagan T, Nieder A, Viswanathan P, Dehaene S. A random-matrix theory of the number sense. Phil. Trans. R. Soc. B. 2018;373:20170253. doi:10.1098/rstb.2017.0253

[5] Le L, Patterson A, White M. Supervised autoencoders: Improving generalization performance with unsupervised regularizers. In Advances in Neural Information Processing Systems. 2018. 107–117.

[6] Hannagan T, Agrawal A., Cohen L, Dehaene S. Emergence of a compositional neural code for written words: Recycling of a convolutional neural network for reading. Proceedings of the National Academy of Sciences Nov 2021, 118 (46) e2104779118; doi: 10.1073/pnas.2104779118.

[7] Op de Beeck HP, Pillet I, Ritchie JB. Factors Determining Where Category-Selective Areas Emerge in Visual Cortex. Trends Cogn Sci. 2019 Sep;23(9):784-797. doi: 10.1016/j.tics.2019.06.006.

Supplementary material

S1. Modeling Number Topography

Training dataset
The custom-made dataset comprised 6000 exemplars of number images between 0 and 9, paired with as many number codes from [4].

Number images were black on white 32x32 pixel images, in 9 possible fonts (arial, lato, openSans, ostrich, oswald, PTN57F, raleway, roboto and tahoma), 6 x-locations and 24 y-locations.

The 10 number codes onto which those images were regressed were 100 dimensional vectors taken from [3], obtained by power iteration of a randomly and locally connected matrix.

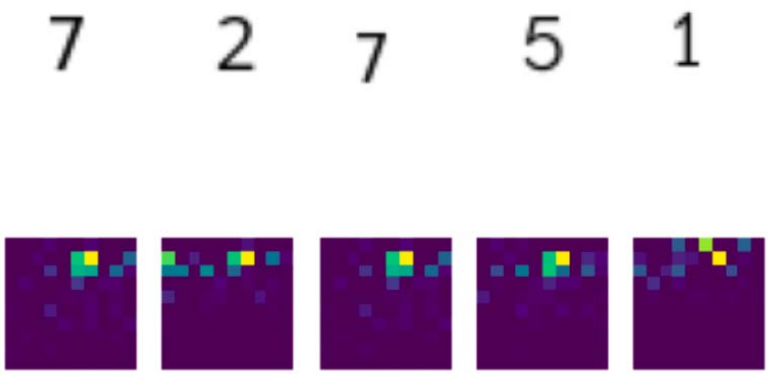

Figure S1. *Input and labels of the number dataset used for our number topography simulations.*

Models and training
The models had the following number of parameters:
Resnet20: 275572
Reset1: 610918
Reset2: 1418134
Reset4: 4646998
Reset8: 17562454
The models were trained for 50 epochs using a Binary Cross-Entropy loss and the Adam optimizer.

S2. Quantifying clustering

Clustering is surprisingly difficult to quantify in a rigorous way. Our procedure, shown in Figure S2, is the following:

1. Collect average activations for categories A and B on the grid.
2. Possibly smooth this activation using a 2d Gaussian kernel (this step was skipped in our analyses).
3. Compute the d-prime of A over B for each unit on the grid.
4. Threshold the resulting map of d-primes.
5. The clustering index is the average number of units in the connected components of this map.

We also perform exactly the same computations for shuffled activations, in order to obtain a reference clustering index. A network is said to exhibit clustering if its index is larger than the reference index obtained for shuffled activations.

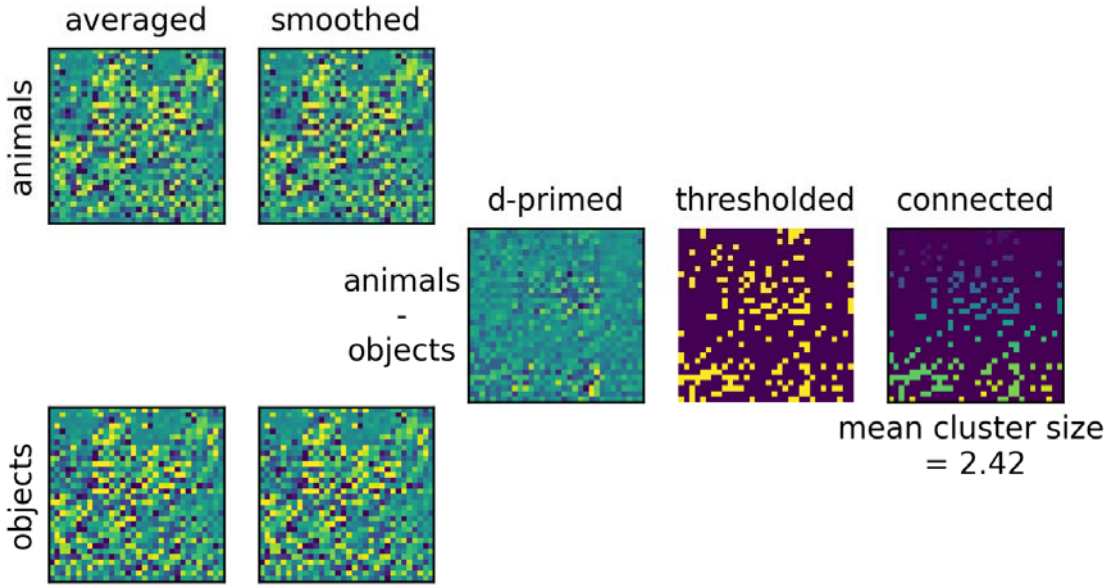

Figure S2. *Quantifying clustering using d-primed maps of activations. After thresholding of the map of d-primes, we compute its clusters as connected components (color coded in the rightmost plot). Our clustering index is the average number of units in the clusters.*

Quantifying topography

Consider a unit x and its neighborhood V(x). We quantify topography as the average over x of the proportion of units in V(x) that have the same preference as x.

S3. Auto-encoding as regularization for Reset networks

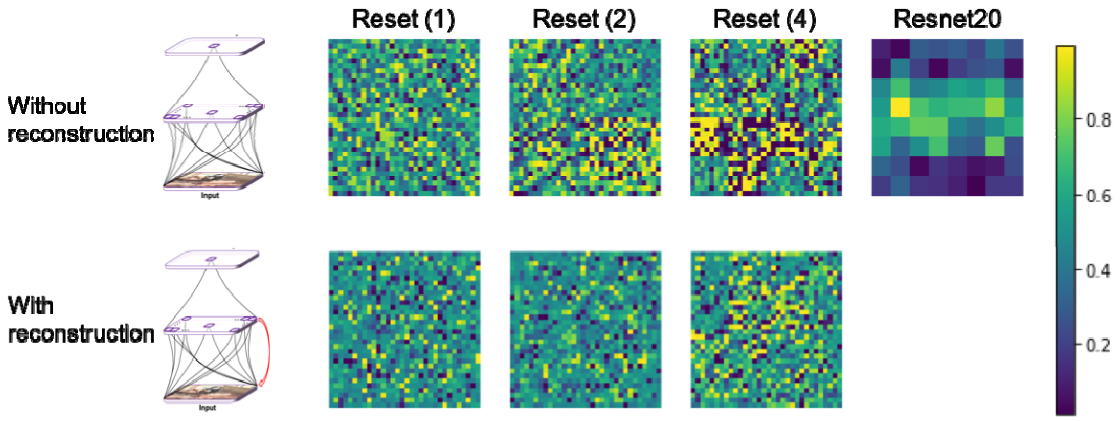

Figure S3. *Normalized activation on the grid of Reset networks of width 1, 2 and 4, trained to classify Cifar-100, with and without input reconstruction from the grid. Reconstruction forces the networks to use grid resources more evenly. Activity in the (2d-reshaped) dense layer of Resnet20 is shown for comparison.*

We trained Reset networks of widths 1, 2 and 4 to classify Cifar-100, with and without adding a subnetwork to reconstruct the input. We then presented 1000 images from the test set to each of the 6 networks, and collected their activations on the grids.

Figure S3 (upper line) shows that activation is not evenly distributed on a Reset network's grid. Units often become polarized, in the sense that many units are rarely activated if ever (dark purple), while many others are very often activated (yellow). This polarization effect becomes stronger within some subnetworks, as network width grows.

Figure S3 (lower line) also shows that polarization can be alleviated by introducing a level 2 subnetwork whose task is to reconstruct the input (red arrow on the model's figure).