

# RECURSIVE NEURAL PROGRAMS: VARIATIONAL LEARNING OF IMAGE GRAMMARS AND PART-WHOLE HIERARCHIES

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Human vision involves parsing and representing objects and scenes using structured representations based on part-whole hierarchies. Computer vision and machine learning researchers have recently sought to emulate this capability using capsule networks, object reference frames and active predictive coding, but a generative model formulation has been lacking. We introduce Recursive Neural Programs (RNPs), a neural generative model that addresses the part-whole hierarchy learning problem by modeling images as hierarchical trees of probabilistic sensory-motor programs. These programs recursively reuse learned sensory-motor primitives to model an image within different reference frames, enabling intuitive and explainable composition and allowing for forming recursive image grammars. We express RNPs as structured variational autoencoders (sVAEs) for inference and sampling, and demonstrate parts-based parsing, sampling and one-shot transfer learning for MNIST, Omniglot and ETH-80 datasets. Our results show that RNPs provide an intuitive and explainable way of composing objects and scenes, allowing rich compositionality and intuitive interpretations of objects in terms of part-whole hierarchies.

## 1 INTRODUCTION

Human visual cognition relies heavily on hierarchical relationships between objects and their parts. For example, a human face can be modeled as a hierarchical tree of parts, each part’s relative position specified within a local reference frame: eyes, nose, mouth etc. positioned within the face’s reference frame, the parts of an eye (eyebrow, eyelid, iris, pupil etc.) positioned within the eye’s reference frame, and so on. To emulate such a capability, a computer vision system needs to not only learn what a part looks like (shapes, contours, colors etc. as in current deep convolutional networks) but also the relative transformation of the part within a local reference frame, and do this recursively in order to compose a human face (or a Picasso painting).

Beyond vision, nested structure and hierarchical parts-based decompositions are ubiquitous in human attributes such as natural language (texts, chapters, paragraphs, sentences, words, characters) and complex behaviors (cooking a recipe, driving to work, etc.). Such recursive modeling confers the important property of compositionality Lake et al. (2015): the same building blocks can be hierarchically and recursively composed into an endless variety of possible patterns, allowing an agent to “imagine” novel configurations of parts (e.g., for creating new solutions to problems), and recognize new configurations of known parts for zero-shot generalization. The challenge lies in learning a model of the parts and their transformations that is recursive and composable. Existing approaches for parsing tree-structured data Eslami et al. (2016); Lake et al. (2015); Hinton et al. (2018); Hinton (2021); Mnih et al. (2014); Socher et al. (2011) are either not recursive Eslami et al. (2016); Mnih et al. (2014), not compositional Socher et al. (2011), not generative Hinton et al. (2018); Hinton (2021), or not differentiable Lake et al. (2015). Indeed, the lack of a smooth “program space” has been a challenge in this regard.

We introduce recursive neural programs (RNPs), which address this problem by creating a fully differentiable recursive tree representation of sensory-motor programs. Our model builds on past work on Active Predictive Coding Networks Gklezacos & Rao (2022) in using state and action networks

but is fully generative, recursive, and probabilistic, allowing a structured variational approach to inference and sampling of neural programs. The key differences between our approach and existing approaches are: 1) Our approach can be extended to arbitrary tree depth, creating a "grammar" for images that can be recursively applied 2) our approach provides a sensible way to perform gradient descent in hierarchical "program space," and 3) our model can be made adaptive by letting information flow from children to parents in the tree, e.g., via prediction errors Jiang, Preston L et al. (2021); Gklezakos & Rao (2022).

## 2 RECURSIVE NEURAL PROGRAMS

We describe a 2-level Recursive Neural Program (RNP), though the architecture can be generalized to more levels. Consider the problem of parsing an image of a digit at two levels ( $k = \{1, 2\}$ ) of an abstraction tree (fig. 1), e.g., in terms of larger parts and smaller strokes (henceforth referred to as parts and sub-parts). A top-level program ( $k = 2$ ) generates the digit in terms of parts and a bottom-level program ( $k = 1$ ) generates each large part as a sequence of smaller parts and their transformations within the larger part's reference frame. Each program is expressed as an interaction between two recurrent functions, a state-transition function (or state-based forward model) that predicts the next state  $z_{t+1}^k = f_{state}^k(z_t^k, a_t^k)$ , and an action transition function (policy)  $a_{t+1}^k = f_{policy}^k(z_t^k, a_t^k)$  (fig. 1b, fig. 2, algorithm 1); in this paper, we assume actions correspond to transformations of parts). This is similar to the next-state and policy functions in a partially observable Markov decision process (POMDP Kaelbling et al. (1998)).

A program at tree depth  $k$ , represented by the state vector  $z_t^k$ , generates a fixed-length sequence of  $\tau^k$  lower level states  $S^{k-1} = \{z_1^{k-1}, \dots, z_{\tau^k}^{k-1}\}$  and their transformations  $T^{k-1} = \{a_1^{k-1}, \dots, a_{\tau^k}^{k-1}\}$ . The state  $z_t^k$  can be decoded into an image patch  $\hat{x}_t^k$  that corresponds to a stroke or other image feature, then transformed according to  $g(\hat{x}_t^k, a_t^k)$  to place it on a "canvas" (here  $a$  refers to parameters of an affine transform on a grid, where  $g$  is the bilinear interpolation function Jaderberg et al. (2015)). The transformed images are added together at each time step, such that each step increasingly approximates the target image represented by  $z^k$  (fig. 2b). This method allows us to reuse the same strokes with different transformations. For example, if  $z_t^k$  represents a 4,  $S^{k-1}$  can represent three straight lines, and  $T^{k-1}$  are the transformations that orient and place them in the configuration of a 4 (fig. 1a).

The above model can be made recursive, with generation performed in a depth-first manner: each  $z_t^k$  generates the program for a sequence  $\{z_1^{k-1}, \dots, z_{\tau^k}^{k-1}\}$ .  $z_{t+1}^k$  begins after  $z_t^k$  terminates. Here we use the decoded patches  $\{\hat{x}_1^{k-1}, \dots, \hat{x}_{\tau^k}^{k-1}\}$  as accumulated evidence to update  $z_t^k$  (similar to other predictive coding models Jiang, Preston L et al. (2021); Gklezakos & Rao (2022)).

### 2.1 MODEL ARCHITECTURE

In a two-level RNP (fig. 1, fig. 2), the top-level program  $z^2$  parameterizes two recurrent neural networks (RNNs)  $f_{state}^2$  and  $f_{policy}^2$  via hypernetworks (state hypernetwork  $H_{state}$  and policy hypernetwork  $H_{policy}$ ) Ha et al. (2016) (fig. 2b,c). Each hypernetwork generates parameters for the level-specific networks responsible for state or policy: a single-layer encoder  $\hat{e} = Enc_{state/policy}^k(\hat{z}^k, a_t^k)$ , where  $\hat{e}$  is the input to the  $f_{state/policy}^k$  network;  $f_{state/policy}^k$ , an RNN with hidden state  $|z|$ , and their initial hidden states; and a decoder  $D^k(z^k)$  that generates an image patch  $\hat{x}_{t+1}^k$  (state network) or affine transform parameters  $a_{t+1}^k$  (policy network; scaling, translation, rotation and shear) that transform the image patch  $\hat{x}_{t+1}^k$ . The hypernetworks provide initialization values  $\hat{x}_0^k, a_0^k$  to initialize the sequence generation. More implementation and training details can be found in the Appendix.

We train the model described above by exploiting the end-to-end differentiability of the architecture, minimizing the reconstruction loss between all transformed sub-parts and the target image  $\hat{x}$ , regularized by the reconstruction at the level of parts:

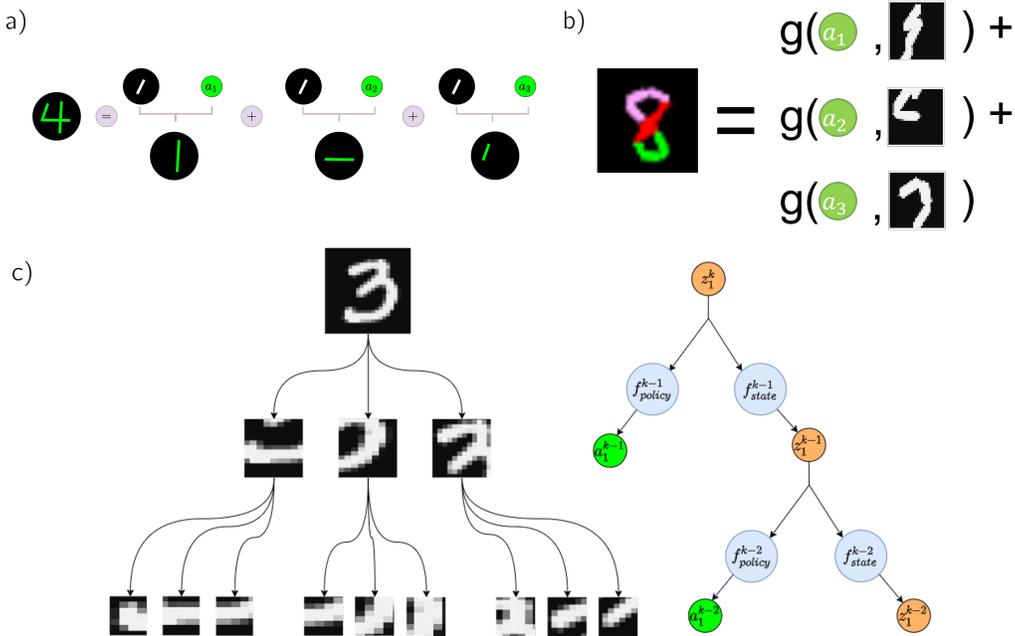


Figure 1: **Image Parsing as a Sequence of Transformed Primitives.** (a) A 4 can be constructed by generating three identical straight lines (black circles) and transforming them according to parameters  $a$  to place them in the appropriate location. (b) As in (a), but a digit (here an eight, left) is generated by a two-level RNP that uses only transformations of one straight line to generate parts (right) that are then transformed into a full digit (placement order: red  $\rightarrow$  green  $\rightarrow$  purple) (c) Left: An MNIST digit decomposed into an abstraction tree of parts, each of which is further decomposed into smaller sub-parts. Right: Schematic representation of a parsing tree produced by a recursive neural program. The digit is described as a “program,” represented by the vector  $z^k$ , which generates functions  $f_{state}^k$  and  $f_{policy}^k$  to construct the digit by generating parts and transforming them according to the action vector  $a^k$  (position, scale, ...) within the digit’s reference frame. Each part is in turn described by a program  $z^{k-1}$ , which generates smaller parts transformed according to  $a^{k-1}$  within that part’s reference frame.

$$\mathcal{L} = \left\| \sum_{t_2=1}^{\tau^2} g\left(\sum_{t_1=1}^{\tau^1} g(\hat{x}_{t_1}^1, a_{t_1}^1), a_{t_2}^2\right) - x \right\|_2^2 + \frac{1}{\tau^2} \left\| \sum_{t_2=1}^{\tau^2} g(\hat{x}_{t_2}^2, a_{t_2}^2) - x_{t_2}^{patch} \right\|_2^2 \quad (1)$$

where  $\tau^2$  and  $\tau^1$  are the number of level-2 and level-1 time steps respectively,  $x$  is the target image and  $x_{t_2}^{patch}$  is the image patch generated by transforming  $x$  with  $g^{-1}(a_{t_2}^2)$  (i.e. zooming in instead of scaling down). We note that RNPs can be trained one depth at a time to decrease training time and resources.

To allow probabilistic sampling of programs, we can express an RNP as a structured variational auto-encoder Kingma & Welling (2014) (VAE) to learn an approximate posterior  $q(z^K|x) \approx p(z^K|x)$  of an image  $x$  given prior  $p(z^K) \sim \mathcal{N}(0, 1)$ , where  $z^K$  is the highest level state vector. We therefore use an encoder network to parameterize the approximate posterior  $q(z^K|x)$  and regularize eq. (1) with the  $KL(q||p)$  term.

### 3 RESULTS

We first demonstrate how our RNPs can recursively parse input images of MNIST digits LeCun et al., Yann (1998), Omniglot characters Lake et al. (2015) and ETH-80 objects Leibe & Schiele (2003)

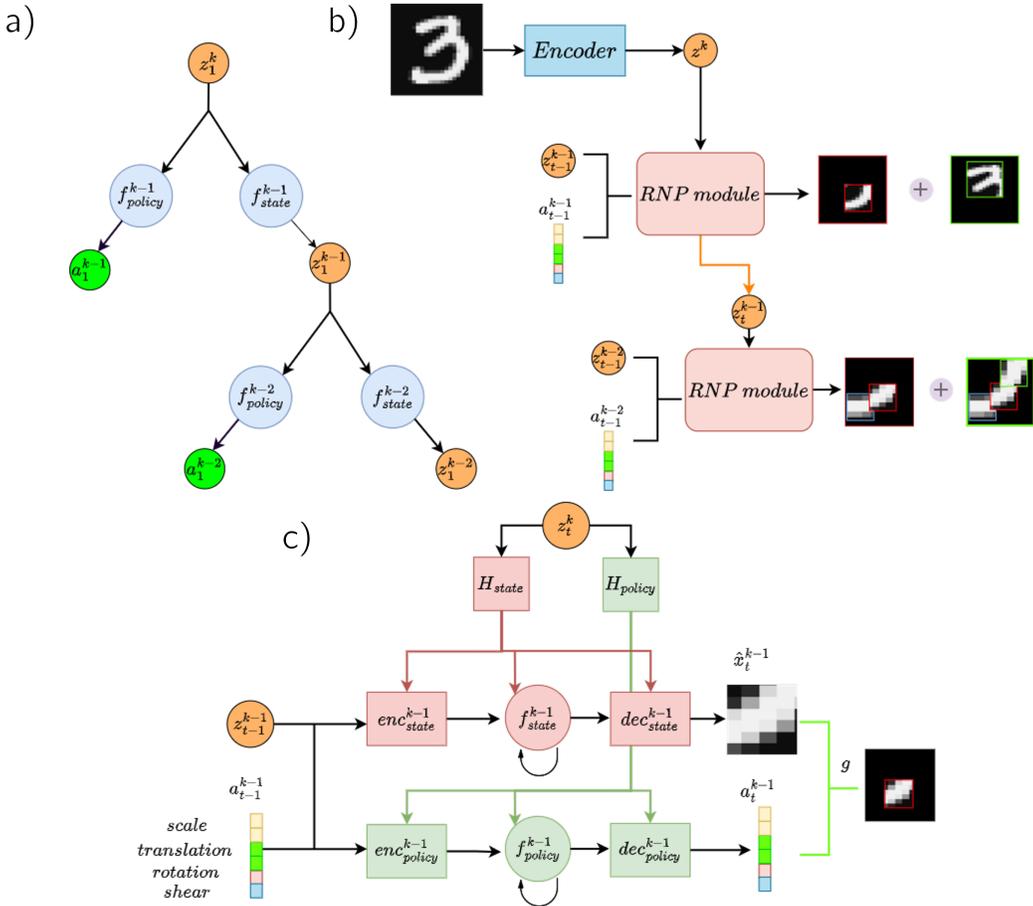


Figure 2: **Recursive Neural Program for Parsing Images.** (a) Parsing tree for a digit, same as fig. 1b. (b) Representation of parsing tree in (a) with a neural network, where  $z^k$  is generated by an encoder network. At time  $t$ , the network receives the most recent estimate of the part/sub-part and  $a_{t-1}^k$  as input, and generates a prediction of the next transformed part or sub-part. Transformed parts and sub-parts are summed at their respective levels. Purple lines indicate recurrence. In our implementation, the top level  $k$  receives the output of the program  $k - 1$  as input (dotted purple line) as opposed to its own output (dashed purple line). (c) The RNP module consists of two hypernets  $H_{state}$  and  $H_{policy}$ , which parameterize  $f_{state}^k$ ,  $f_{policy}^k$  and auxiliary networks to perform the computation shown in (b).

into parts and sub-parts. We then characterize the embedding space of state vectors at two levels and show how learned representations at various tree-depths can be composed to generate previously unseen image types.

### 3.1 IMAGE PARSING INTO PARTS AND SUB-PARTS

We trained RNPs to reconstruct MNIST digits and Omniglot characters as two-level generative programs. An encoder network was trained to map the input image to the top-level program (embedding vector)  $z^2$ . As described above,  $z^2$  parameterizes  $f_{state}^1$  and  $f_{policy}^1$  via the hypernetworks  $H_{state}$  and  $H_{policy}$ , and  $z^1$  is the latent code corresponding to the parts (larger patches, 6x6px - 12x12px; fig. 1b).  $z^1$  is then passed through the same hypernetworks to synthesize sub-parts (smaller patches, 1.5x1.5px - 4x4px; fig. 2b). We force the network to learn a part-wise representation by constraining each part to be smaller than its parent, therefore requiring a sequence of steps to reconstruct it. Figure 3 shows examples of MNIST digits (fig. 3a), Omniglot characters (fig. 3b) and ETH-80 objects

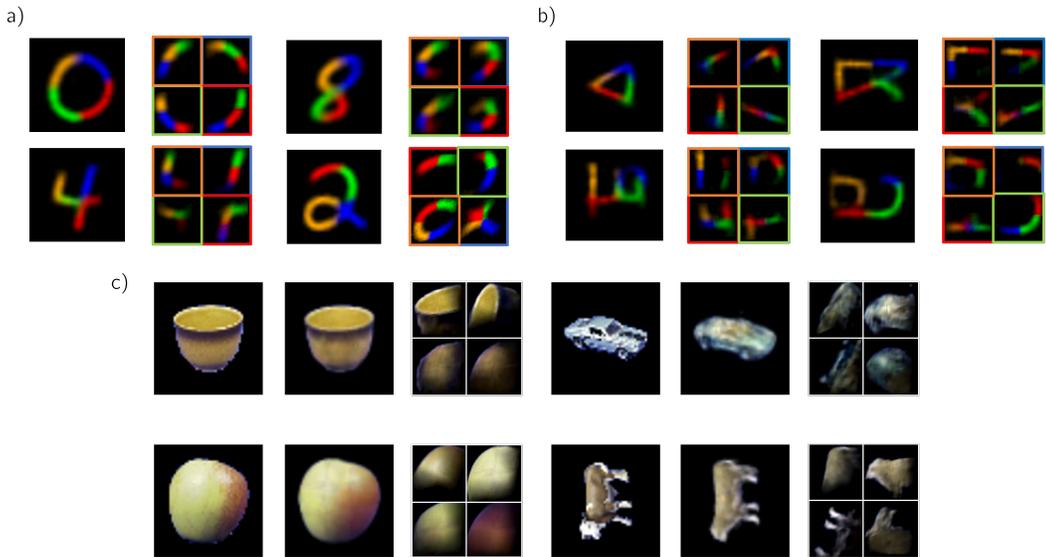


Figure 3: **Hierarchical Parts-Based Decomposition by the Model after Learning:** Parsing of (a) four MNIST digits, (b) four Omniglot characters and (c) four objects in the ETH-80 dataset by the model. (a,b) Two levels of the hierarchical representation are shown, parts generated by  $z^2$  (left; each part is denoted by a different color) and sub-parts generated by  $z^1$  (right, bordered boxes, each sub-part is denoted by a different color). Order: blue  $\rightarrow$  red  $\rightarrow$  green  $\rightarrow$  orange. Each bordered box shows the output of a program generated by  $z^1$  to construct a part as a combination of sub-parts transformed and placed within the reference frame of the part. (c) Leftmost column: original image; middle column: reconstruction; right column: individual parts (sum of sub-parts).

(fig. 3c) generated by RNPs, with reconstructions at the level of parts (untiled-) and sub-parts (tiled images).

### 3.2 TOPOGRAPHY OF NEURAL PROGRAMS

A notable challenge in optimizing and representing probabilistic programs has been the absence of a continuous program space that can be interpretably manipulated. As we use the same hypernetworks to generate programs at all levels, we should expect that programs at different tree depths inhabit different areas of  $|z|$ -dimensional space, i.e. programs representing digits cluster separately from programs representing parts. Analyzing the embedding space of  $z^2$  and  $z^1$  vectors that represent the trained data (MNIST digits or Omniglot characters) reveals that  $z^2$  and  $z^1$  “neural program” vectors do cluster separately (fig. 4a,b).

To test the expressiveness of our model, we investigated the space between learned  $z^2$  and  $z^1$  program clusters by linearly interpolating in the latent “neural program” space occupied by the  $z^2$  and  $z^1$  vectors. Sampling from regions between clusters produced programs that generated novel images (fig. 5), showing that the model can exploit the latent structure of the program embedding space to synthesize previously unseen patterns by combining the learned parts.

### 3.3 COMPOSITIONALITY AND TRANSFER LEARNING

Compositionality is a main goal of our architecture. With a generative model over programs, we are able to sample program space in regions outside those representing the trained data (fig. 4). This can be demonstrated by interpolating between clusters in  $\{z^2, z^1\}$  (fig. 4c), or sampling randomly from  $z^2 \sim \mathcal{N}(0, I)$  (fig. 5). Figure 5 shows that the model can generate novel characters by synthesizing learned primitives in different, often novel, combinations of parts.

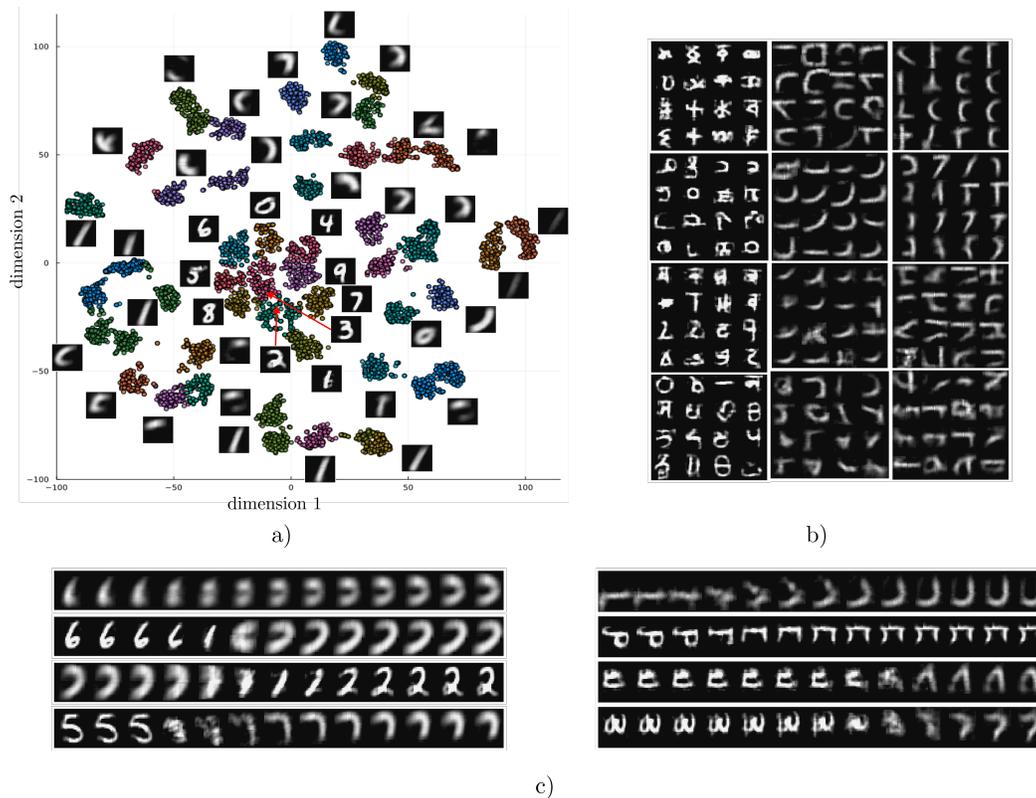


Figure 4: **Topography of Neural Program Space.** (a) t-SNE clustering of  $z^2$  and  $z^1$  vectors in a model trained on MNIST. A representative image is shown for each cluster. (b) Example clusters of sampled images from  $z^2$  (left column) and  $z^1$  (remaining columns) for a model trained on Omniglot. (c) Example linear interpolations in  $z$  space from the center of one cluster (leftmost image) to the center of another cluster (rightmost image) show novel generated images from neural programs in the intermediate space. Left: MNIST, right: Omniglot.

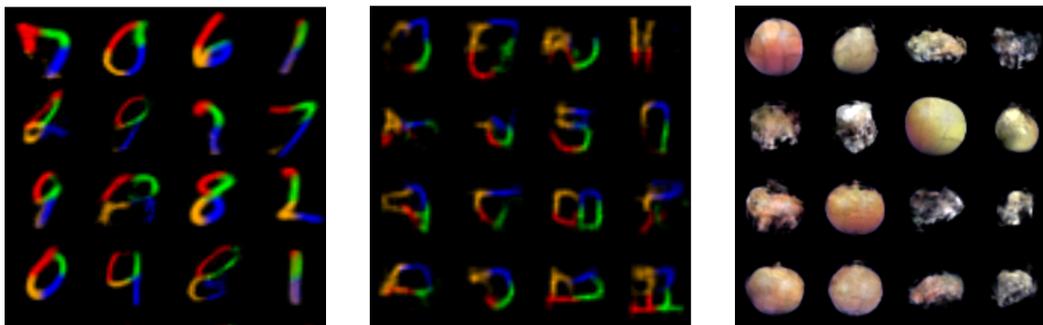


Figure 5: **Sampling from the Prior.** Sampling of  $z^2$  from  $\mathcal{N}(0, 1)$  for a model trained on MNIST digits (left), Omniglot characters (middle) and ETH-80 objects (right). As in (fig. 3), part order is blue  $\rightarrow$  red  $\rightarrow$  green  $\rightarrow$  orange.

We further tested the compositional ability of our model in two transfer learning tasks. Firstly, we trained RNPs on all MNIST classes but one (7 or 8), or on the Omniglot training dataset. By adjusting the weights of the encoder network (but not the decoder hypernetworks  $H_{state}, H_{policy}$ ), RNPs were able to synthesize parts for the unseen class (fig. 6a,b). Secondly, we adapted a model

trained on the ETH-80 dataset to reconstruct colored Fashion-MNIST Xiao et al. (2017) objects by keeping the  $H_{state}$  network and training a new Encoder and  $H_{policy}$  network (fig. 6c).

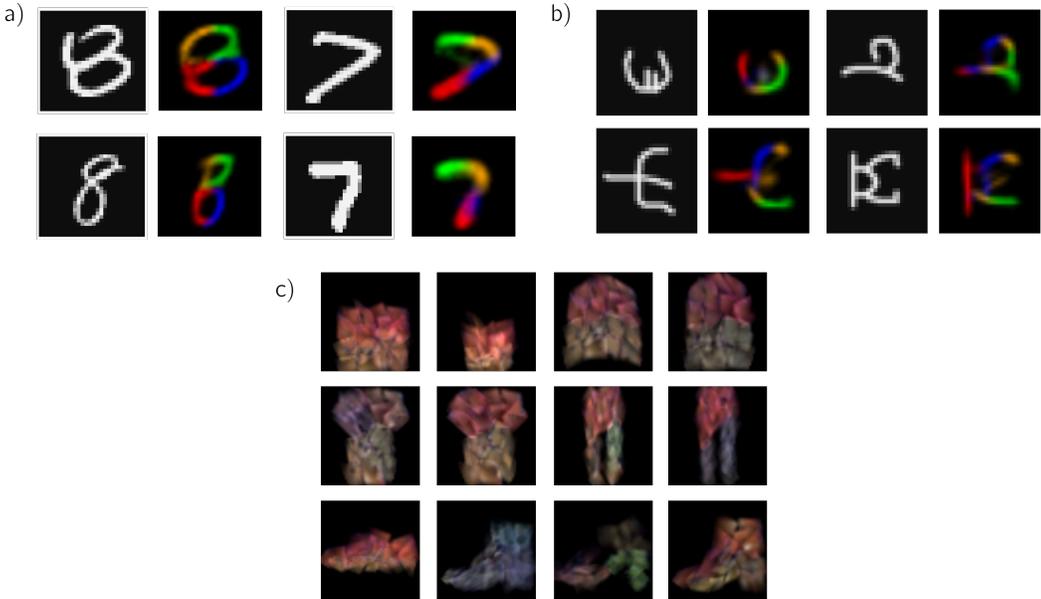


Figure 6: **Transfer Learning.** RNPs trained on a training set of classes (see text) are able to explain novel examples from unseen classes and synthesize the parts for MNIST digits (a) and Omniglot characters (b). (c) Example Fashion-MNIST objects generated by an RNP with  $H_{state}$  trained on the ETH-80 dataset.

### 3.4 EXTENDING RNPs BEYOND TWO RECURSION LEVELS

To demonstrate that RNPs can be trained at a recursion depth greater than two, we trained a three-level RNP on images containing pairs of MNIST digits (fig. 7a). We then show that the three-level model constitutes a two-level RNP on MNIST digits, i.e. we reduce the recursion depth to two and train a new Encoder to generate centered MNIST digits (fig. 7b).

## 4 CONCLUSION

In this paper, we introduced Recursive Neural Programs (RNPs), a new model for differentially learning tree-structured data as sensory-motor sequences in a way that allows flexible composition of learned primitives using a recursive “grammar.” We demonstrated our model’s ability to generate images using a hierarchy of parts and their transformations. Our architecture can also be applied to learning in arbitrary domains, such as audio, video and other dynamical processes such as motor behavior.

There are several potential directions for future research. Using the same hypernetworks at different levels allows natural recursion, but limits the expressive power of the model. This can be addressed by learning different hypernetworks for different levels, or introducing a level-specific input. Hypernetworks describing different data modalities (e.g. audio, visual, etc.) could be combined to generate richer multi-modal neural programs, provided constraints on the size of the primary network are taken into account Galanti & Wolf (2020). Training deep RNPs across levels and across time steps can be challenging. This could be addressed by training RNPs at different depths in parallel. Another potential area for improvement is replacing bilinear interpolation (which is not smoothly differentiable) used for transformation of image primitives, with smoother functions to sample images (e.g. Klocek et al. (2019)).

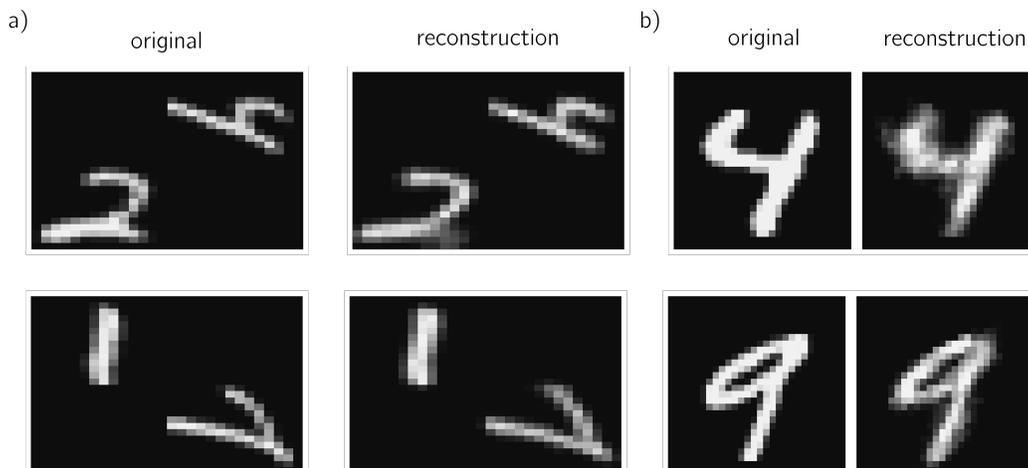


Figure 7: **Three-level RNP** RNPs trained on pairs of MNIST digits (a): Left column: digit pair; right column: model reconstruction. (b): MNIST digits retrieved from two-level truncation of the three-level model by training an encoder on centered (non-transformed) MNIST digits.

Finally, message passing between nodes at different tree depths could allow for bidirectional information flow: predictions from parents to children, and belief updates from children to parents (using, e.g., prediction errors). We intend to explore such predictive coding-based architectures for RNPs in future work.

#### ACKNOWLEDGMENTS

We thank Dimitrios Gklezakos for his help with hypernetworks, and Preston Jiang for feedback on probabilistic aspects of the model. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0021, a Weill Neurohub Investigator award, a grant from the Templeton World Charity Foundation, and a Cherng Jia & Elizabeth Yun Hwang Professorship to RPNR. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the funders.

#### REFERENCES

- S. M. Ali Eslami et al. Attend, Infer, Repeat: Fast Scene Understanding with Generative Models. In *Advances in Neural Information Processing Systems*, 2016. URL <https://proceedings.neurips.cc/paper/2016/hash/52947e0ade57a09e4a1386d08f17b656-Abstract.html>.
- Tomer Galanti and Lior Wolf. On the Modularity of Hypernetworks. *arXiv:2002.10006 [cs, stat]*, November 2020. URL <http://arxiv.org/abs/2002.10006>. arXiv: 2002.10006.
- Dimitrios C. Gklezakos and Rajesh P. N. Rao. Active Predictive Coding Networks: A Neural Solution to the Problem of Learning Reference Frames and Part-Whole Hierarchies. preprint, Neuroscience, January 2022.
- David Ha, Andrew Dai, and Quoc V. Le. HyperNetworks. *arXiv:1609.09106 [cs]*, December 2016. URL <http://arxiv.org/abs/1609.09106>. arXiv: 1609.09106.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. pp. 770–778, 2016. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2016/html/He\\_Deep\\_Residual\\_Learning\\_CVPR\\_2016\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html).
- Geoffrey Hinton. How to represent part-whole hierarchies in a neural network. *arXiv:2102.12627*, February 2021. doi: 10.48550/arXiv.2102.12627. URL <https://arxiv.org/abs/2102.12627v1>.

- Geoffrey Hinton, Sara Sabour, and Nicholas Frosst. Matrix Capsules with EM Routing. *ICLR*, pp. 15, 2018.
- Max Jaderberg et al. Spatial Transformer Networks. In *Advances in Neural Information Processing Systems*, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/33ceb07bf4eeb3da587e268d663aba1a-Abstract.html>.
- Jiang, Preston L, Gklezakos, Dimitrios, and Rajesh P. N. Rao. Dynamic Predictive Coding with Hypernetworks | bioRxiv, 2021. URL <https://www.biorxiv.org/content/10.1101/2021.02.22.432194v2.abstract>.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, May 1998. ISSN 00043702. doi: 10.1016/S0004-3702(98)00023-X. URL <https://linkinghub.elsevier.com/retrieve/pii/S000437029800023X>.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. Technical Report arXiv:1412.6980, arXiv, January 2017. URL <http://arxiv.org/abs/1412.6980>. arXiv:1412.6980 [cs] type: article.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv:1312.6114 [cs, stat]*, May 2014. URL <http://arxiv.org/abs/1312.6114>. arXiv: 1312.6114.
- Sylwester Klocek et al. Hypernetwork Functional Image Representation. In *Artificial Neural Networks and Machine Learning – ICANN 2019: Workshop and Special Sessions*, 2019. ISBN 978-3-030-30493-5. doi: 10.1007/978-3-030-30493-5\_48.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, December 2015. doi: 10.1126/science.aab3050. URL <https://www.science.org/doi/full/10.1126/science.aab3050>.
- LeCun et al., Yann. Gradient-based learning applied to document recognition | IEEE Journals & Magazine | IEEE Xplore, 1998. URL <https://ieeexplore.ieee.org/document/726791>.
- B. Leibe and B. Schiele. Analyzing appearance and contour based methods for object categorization. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pp. II–409, June 2003. doi: 10.1109/CVPR.2003.1211497. ISSN: 1063-6919.
- Volodymyr Mnih et al. Recurrent Models of Visual Attention. In *Advances in Neural Information Processing Systems*, 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/09c6c3783b4a70054da74f2538ed47c6-Abstract.html>.
- Richard Socher et al. Parsing Natural Scenes and Natural Language with Recursive Neural Networks. January 2011. URL <https://openreview.net/forum?id=SyEunWObH>.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. Technical Report arXiv:1708.07747, arXiv, September 2017. URL <http://arxiv.org/abs/1708.07747>. arXiv:1708.07747 [cs, stat] type: article.

## A APPENDIX

### MODEL AND TRAINING DETAILS

#### PARAMETERS

RNPs consist of two hypernets  $H_{state}$  and  $H_{policy}$  that generates the parameters of an autoregressive network for a level  $k$ . All hypernetworks used in this study are 6-layer MLP’s (64 units, elu activations), with seven heads, parameterizing networks of the same size (except to reflect different

sizes of  $|z|$ ). All networks for a given  $k$  consisted of fully connected layers of 64 units, except the RNNs  $f_{state}^k$  and  $f_{policy}^k$ , which retained the dimensionality of  $|z|$ .

The encoder network consisted of five ResNet blocks (32 channels) He et al. (2016) and four fully connected layers (64 units).

## TRAINING

We trained all models using the ADAM optimizer Kingma & Ba (2017) with a learning rate of  $4e-5$ , which reliably showed convergence. We trained our models for 200 epochs, except on the Omniglot dataset where we trained for 400 epochs. We used  $|z| = 32$  on MNIST and Fashion-MNIST, and  $|z| = 96$  for Omniglot. Models were trained on a single GPU (Nvidia Quadro RTX 6000).

### A.1 RNP ALGORITHM

---

#### Algorithm 1 Recursive image generation

---

```

1: procedure RNPDECODER( $z, level$ )
2:    $enc_{state}^{k-1}, f_{state}^{k-1}, dec_{state}^{k-1}, z_0^{k-1} \leftarrow H_{state}(z^k)$ 
3:    $enc_{policy}^{k-1}, f_{policy}^{k-1}, dec_{policy}^{k-1}, a_0^{k-1} \leftarrow H_{policy}(z^k)$ 
4:    $p_t^{k-1} = \mathbf{0}$ 
5:   for  $t = 1 : \tau_{k-1}$  do
6:      $z_t^{k-1} = f_{state}^{k-1}(a_{t-1}^{k-1}, z_{t-1}^{k-1})$ 
7:      $a_t^{k-1} = f_{policy}^{k-1}(a_{t-1}^{k-1}, z_{t-1}^{k-1})$ 
8:      $\hat{x}_t^{k-1} = dec_{state}^{k-1}(z_t^{k-1})$ 
9:     if  $level > 0$  then
10:       $p_t^{k-1} \leftarrow p_{t-1}^{k-1} + RNPdecoder(z^{k-1}, level - 1)$ 
11:     else
12:       return
13:   return  $p_t^{k-1}$ 
14: procedure RNP( $x, levels = k$ )
15:    $\mu^k, logvar^k = Encoder(x)$ 
16:    $z^k = \mu^k + \exp(logvar) * rand(\mathcal{N}(0, 1))$ 
17:   return RNPdecoder( $z^k, k$ )

```

---