

Learning a Deep Generative Model like a Program: the Free Category Prior

Eli Sennesh¹

¹Khoury College of Computer Science
Northeastern University
440 Huntington Ave
Boston, MA 02115
sennesh.e@northeastern.edu

Humans surpass the cognitive abilities of most other animals in our ability to “chunk” concepts into words, and then combine the words to combine the concepts. In this process, we make “infinite use of finite means”¹, enabling us to learn new concepts quickly (Lake, Salakhutdinov, and Tenenbaum, 2013) and nest concepts within each-other (Lake and Piantadosi, 2019). We construct concepts to model the world around us (Goodman, Tenenbaum, and Gerstenberg, 2014). In contrast, in artificial intelligence we still mostly focus on learning discriminative functions mapping from sensory inputs to decision outputs, without modeling the data-generating process (Zhu et al., 2019).

While program induction and synthesis remain at the heart of foundational theories of artificial intelligence (Solomonoff, 1964; Vitányi and Li, 2000), only recently has the community moved forward in attempting to use program learning as a benchmark task itself (Chollet, 2019). Models of program learning are most often based upon probabilistic context-free grammars (Overlan, Jacobs, and Piantadosi, 2017; Romano et al., 2018), with neural networks and/or reinforcement learning sometimes used to aid inference (Ellis et al., 2018, 2020). However, not every grammatical program in the support of a PCFG yields a meaningful likelihood. Indeed, a recent experiment by Bramley et al. (2018) found that 35% of sampled hypotheses for a logical rule task were either tautologies (true for all individual cases) or contradictions (true for no individual statements).

Here we confront the assumption that a compositional prior over complex programs must necessarily take the form of a (probabilistic) grammar over a symbolic (probabilistic) language of thought (Piantadosi, Tenenbaum, and Goodman, 2016; Frankland and Greene, 2020). While a formal language theorist would recognize strings of symbols, particularly bound and unbound variables, as one way among many to encode a universal programming language, symbol strings are the modality in which most programmers express most programs. The cognitive science community has thus often assumed that *if* the brain has simulation and reasoning capabilities equivalent to a universal computer, *then* it must




Figure 1: The multigraph (or nerve) of a free category generated by deep generative models.

employ a serialized, symbolic representation Marcus (2018). We provide a counterexample in which compositionality is expressed via network structure.

Contributions Here we contribute an alternative generative model for representing programs: free categories. From this structure, we can sample “correct by construction” probabilistic programs, whose support must contain the data, without maintaining a symbol table or performing additional type-checking steps. We show how our formalism allows neural networks to serve as primitives in probabilistic programs, similarly to Valkov et al. (2018). We learn both program structure and model parameters end-to-end.

The free category prior over programs

A category \mathcal{C} consists of a set of *objects* $Ob(\mathcal{C})$, and for every pair of objects $A, B \in Ob(\mathcal{C})$ a set of *morphisms* $\mathcal{C}(A, B)$. Individual morphisms can be denoted $f : A \rightarrow B$, and each object A is equipped with a unique $id_A : A \rightarrow A$. As implied by the notation, in programming terms we can imagine objects to be data-types and morphisms to be functions between them. There are also two major laws which all morphisms must obey to form a category: the associativity of morphism composition, and the identity of unital composition with identity morphisms.

```

def path_between(A, B, g, P):
    loc = A
    f = id_A
    while loc != B:
        probs = [P_{a,B} for a in g.out_arrows(loc)]
        arrow = sample(Categorical(probs))
        f = f >> arrow # Composition of morphisms
        loc = dest(arrow)
    return f

```

Listing 1: The loop for sampling short paths between a source A and a destination B in a free category described by the graph g and biased random walk P .




Figure 2: String diagram drawn from the posterior distribution over deep generative model architectures for Omniglot.

The closure of a directed multigraph under associative composition and self-looping forms the *free* category over that multigraph. When the number of generating morphisms (edges in the multigraph) is finite, we say that the category is a finitely-generated free category, and we can represent it computationally via its underlying multigraph, $G = (V, E)$.

In Figure 1, we show the multigraph of a free, finitely-generated category. The vertices V correspond to (tuples of) tensor shapes $Ob(\mathcal{C})$, and the edges E correspond to probabilistic decoder networks that act as generating morphisms. There is a privileged object \top , known as the “terminal” or “unit” object; generating morphisms such as $f : \top \rightarrow \mathbb{R}^{64}$ represent prior distributions $p_f(Z)$, $Z \in \mathbb{R}^{64}$ over the individual vector spaces. By designing a probabilistic generative model to sample paths through this graph, we thus obtain a prior over variational autoencoder architectures (Kingma and Welling, 2013; Zhao, Song, and Ermon, 2017), or, equivalently, a domain-specific language for deep generative models.

Syntactically, we have replaced programs sampled from a grammar with string diagrams sampled from a free category. These are conventionally read from the top (leaves of an abstract-syntax tree) to the bottom (root of the tree). An example can be seen in Figure 2.

To sample morphisms from the free category, we conduct a biased random walk from a “source” object i to a “destination” object j . We begin by transforming the directed multi-

graph G into a directed graph $G'_c = (V', E')$, in which each generating morphism $f : A \rightarrow B$ becomes a vertex $v \in V'$ with a single incoming edge (from $A \in V'$) and a single outgoing edge (to $B \in V'$). Denoting the adjacency matrix of G' as $A \in \mathbb{R}^{|V'| \times |V'|}$, we can now construct transition probabilities $P_{i,j}$ for all $(i, j) \in |V'| \times |V'|$,

$$\log P_{i,j} \propto \frac{1}{\beta} (e^A + AW)_{i,j}, \quad (1)$$

$$p(i \rightarrow j | \beta, W) = \text{softmax} \left(\frac{1}{\beta} (e^A + AW) \right)_{i,j}, \quad (2)$$

$$= \frac{\exp \left\{ \frac{1}{\beta} (e^A + AW)_{i,j} \right\}}{\sum_{j=1}^{|V'|} \exp \left\{ \frac{1}{\beta} (e^A + AW)_{i,j} \right\}}. \quad (3)$$

Here, we write e^A for the matrix exponential of A , which describes the “long-run” transition structure of an unbiased random walk, and we also incorporate weights $W \in \mathbb{R}^{|V'| \times |V'|}$ that can encode arbitrary “preferences” for some primitive arrows over others. The first term inside the parentheses thus imposes an inductive bias towards primitive arrows that belong to many multi-step paths through the graph. The second term incorporates a learnable preference for some edges over others. The inverse temperature $\beta > 0$ denotes the relative “confidence” of the soft optimization. Rather than using $-\log P_{i,j}$ as an “intuitive distance” (Baram et al., 2018; Behrens et al., 2018) that compresses shared path segments, we can also hold j constant across a whole path and use $-\log P_{i,j}$ as an intuitive distance to the destination node. This observation inspires the core of our generative model, seen in Listing 1, and can be considered a stochastic generalization of path-based planning techniques (Blum and Furst, 1997; Blum and Langford, 2000).

Since our underlying free category includes Cartesian product objects such as (A, B) , we include “macro” arrows $\top \rightarrow (A, B)$ in the multigraph over which we perform the random walk. When sampled, we “expand” these macros by sampling morphisms $\top \rightarrow A$ and $\top \rightarrow B$ and applying the Cartesian product \times on morphisms. The same type of macro can be used to represent exponential objects B^A (the internal hom-set of morphisms $A \rightarrow B$), by recursively invoking the probabilistic subroutine `path_between()`.

The complete generative model For data $X \in \mathcal{X}$ and latent variables Z sampled by the parameterized morphism f_θ , itself drawn from the free category of the directed multi-graph G , the complete generative model is

$$p_\theta(X, Z, f, \beta, W) = p_\theta(X | Z, f)p_\theta(Z | f) p_G(f | \beta, W)p(\beta)p(W), \quad (4)$$

where the prior over the inverse temperature β and edge weights W are Gamma distributions,

$$\beta \sim \gamma(1, 1), \\ W \sim \gamma(1, 1).$$

Amortized variational inference in the model We implement a data-driven proposal with neural networks,

$$q_\phi(f, \beta, W | X) = p_G(f | \beta, W)q_\phi(\beta | X)q_\phi(W | X).$$

We also assume that each primitive morphism f comes equipped with a trainable stochastic inverse (Stuhlmüller, Taylor, and Goodman, 2013) $f_\phi^\dagger : X \rightarrow Z$ (with density $q_\phi(Z | X; f)$), so we can derive a complete joint proposal density for all latent variables,

$$q_\phi(Z, f, \beta, W | X) = q_\phi(Z | X; f)q_\phi(f, \beta, W | X). \quad (5)$$

This proposal can then be made to approximate the true Bayesian posterior by optimizing the Evidence Lower Bound (ELBO),

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_q \left[\frac{p_\theta(X, Z, f, \beta, W)}{q_\phi(Z, f, \beta, W | X)} \right]. \quad (6)$$

Implementation and optimization We have implemented this generative model in Pyro (Bingham et al., 2019) atop the DisCoPy (de Felice, Toumi, and Coecke, 2020) library for applied category theory in Python and the NetworkX (Hagberg, Schult, and Swart, 2008) library for graph analysis. NVIL (Mnih and Gregor, 2014) and graph-based (Schulman et al., 2015) gradient estimation techniques were used to perform Stochastic Variational Inference (Hoffman et al., 2013).

The library implementing the fundamental operations described here is called `Discopyro` and is available online, along with the [code for reproducing our experiments](#). Please note that we regard the exponentiated adjacency matrix e^A as a constant for optimization purposes, calculating it once upon construction of the multigraph without allowing gradients to flow through it. When we have attempted to use differentiable approximate implementations of the matrix exponential, we have found that the approximate posterior distributions collapse down to a single preferred morphism per minibatch of input data due to the extreme effects of small changes to the real-valued adjacency matrix.

Application: variational autoencoder architecture search

As an example domain, we consider structured variational autoencoder architectures for 28×28 grayscale image data. The structures we consider include standard variational

| | VAE | VLAE | SPAIR |
|---------------|-----|------|-------|
| MNIST | ✓ | | |
| Fashion MNIST | | ✓ | |
| Omniglot | | ✓ | ✓ |

Table 1: Dataset-architecture combinations

| | Epochs | \mathcal{L} (validation) | \mathcal{L} (test) |
|---------------|--------|----------------------------|----------------------|
| MNIST | 500 | 8.36×10^5 | n/a |
| Fashion MNIST | 500 | 4.87×10^5 | n/a |
| Omniglot | 1200 | n/a | 1.46×10^5 |

Table 2: Evidence lower bounds on validation and test data for our three example datasets

autoencoders (Kingma and Welling, 2014) with Continuous Bernoulli likelihoods (Loaiza-Ganem and Cunningham, 2019), variational ladder autoencoders (Zhao, Song, and Ermon, 2017) with continuous Bernoulli likelihoods, and the spatially-invariant Attend-Infer-Repeat variant (Crawford and Pineau, 2019) with Gaussian likelihoods. All proposals and priors over vector latent variables are Gaussian.

We list the combinations of network architecture and dataset tested in Table 1. We include MNIST (LeCun, Cortes, and Burges, 2010) for a standard of comparison, Fashion MNIST (Xiao, Rasul, and Vollgraf, 2017) for a more up-to-date benchmark, and Omniglot (Lake, Salakhutdinov, and Tenenbaum, 2015) as a challenge dataset for learning from few examples per class. The former two datasets were split 90/10 into training and validation data, while Omniglot provided a full test-set of handwritten characters not present in the training dataset. All evaluations are performed on the validation and test sets, respectively. Minibatching was used to subsample independent, identically distributed samples from each class within each dataset, treating alphabets as classes in Omniglot.

Learning a latent-variable model over 28×28 images amounts to sampling a morphism $f : \mathbb{T} \rightarrow \mathbb{R}^{784}$ whose path-length in the underlying multigraph is at least two (which we enforce via extra looping and edge-filtering conditions on the code in Listing 1).

Results

Table 2 shows training times and validation losses.

MNIST Figure 3 shows a VAE architecture sampled from the amortized approximate posterior over the MNIST validation data. Figure 4 compares the original test images and their reconstructions. Reconstructions are satisfactory.

Fashion MNIST Figure 5 shows a learned VLAE architecture sampled from the amortized approximate posterior over the Fashion MNIST validation data. Figure 6 compares the original test images and their reconstructions. Reconstructions capture only high-level shape and texture.




Figure 3: String diagram with vanilla VAE architecture sampled from MNIST’s approximate posterior.




Figure 4: **Above:** original MNIST digits. **Below:** Reconstructions from the learned model.

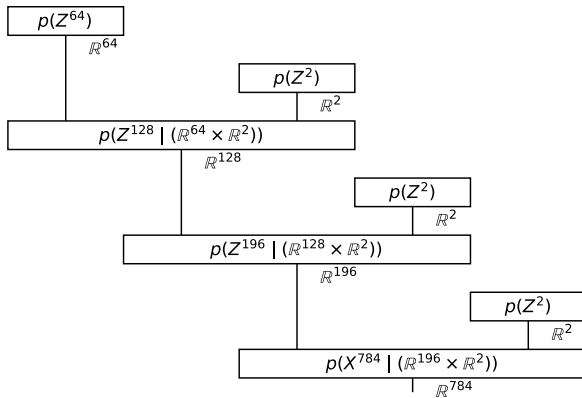


Figure 5: String diagram with VLAE architecture sampled from Fashion MNIST’s approximate posterior.

Omniglot Figure 7 shows a learned architecture sampled from the amortized approximate posterior over the Omniglot evaluation data. The architecture makes use of a spatial attention mechanism (Jaderberg et al., 2015) to improve learning and reconstruction, and comes from the same approximate posterior distribution as the VLAE architecture above. Figure 9 diagrams the associated dagger morphism. Figure 8 compares the original test images and their reconstructions. Reconstructions appear to capture the relevant details of each character with high quality, most likely thanks to spatial attention being used to encode glimpses of whole character




Figure 6: **Above:** original Fashion MNIST images. **Below:** Reconstructions from the learned model.

components in the inference model.

Discussion

Future work We hope that future work can extend what we have presented here in several ways. The most clearly necessary extension, if the intended application is amortized/autoencoding inference, is to cover nonstochastic morphisms, which do not sample their outputs as random variables. This would amount to the separation of a type system (in which morphisms map from a source to a destination object) from an effect system (in which morphisms add random variables to a trace), with a more complex form of the stochastic planning-graph techniques presented here.

An applied extension of this work would consider would consider an challenging concept-learning benchmark designed, such as ARC (Chollet, 2019), CURI (Vedantam et al., 2020), or Bongard-LOGO (Nie et al., 2020). Any program-learning problem with a well-defined, differentiable likelihood could be considered within our framework, including those in which intermediate latent variables are computed deterministically rather than sampled.

This work could also be extended by passing from a strictly finite multigraph underlying the free category prior to an infinite multigraph. This infinite multigraph would include structure created by (recursively) applying endofunctors (such as universal constructions or polymorphic algebraic data types) and natural transformations to the base category, and stochastic path-planning would take account of this structure. This would extend the free category prior to a Bayesian nonparametric model, and end-to-end training would then require truncated or Russian Roulette (Xu, Srivastava, and Sutton, 2019) ELBO estimation.

Conclusions We have presented the free category prior, a neuro-categorical model of compositional program learning in custom domain-specific languages. Our generative model allows for sampling programs which are well-typed, and therefore correct-by-construction in the sense of assigning a nonzero, non-unity likelihood to the data. We emphasize that our generative model does not involve any rejection-sampling steps, nor other complex deterministic heuristics.

We have applied the free category prior to learning modular variational autoencoder architectures for three common datasets, and have shown that the primitives in a free category’s domain-specific language can be parameterized and learned end-to-end by variational Bayes.

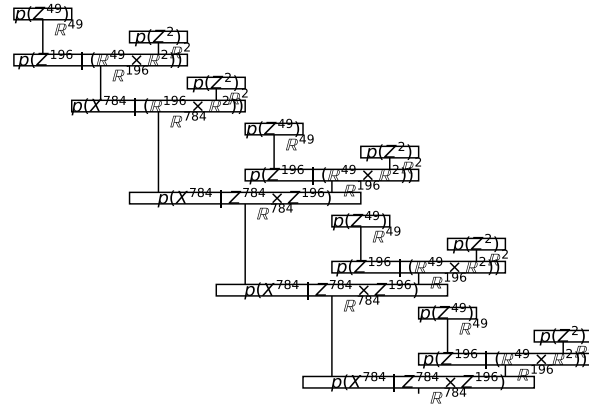


Figure 7: String diagram sampled from Omniglot’s approximate posterior.




Figure 8: **Above:** original Omniglot characters. **Below:** Reconstructions from the learned SPAIR model.

References

- Baram, A. B.; Muller, T. H.; Whittington, J. C.; and Behrens, T. E. 2018. Intuitive planning: global navigation through cognitive maps based on grid-like codes. *bioRxiv* 421461.
- Behrens, T. E.; Muller, T. H.; Whittington, J. C.; Mark, S.; Baram, A. B.; Stachenfeld, K. L.; and Kurth-Nelson, Z. 2018. What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* 100(2):490–509.
- Bingham, E.; Chen, J. P.; Jankowiak, M.; Obermeyer, F.; Pradhan, N.; Karaletsos, T.; Singh, R.; Szerlip, P.; Horsfall, P.; and Goodman, N. D. 2019. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research* 20(1):973–978.
- Blum, A. L., and Furst, M. L. 1997. Fast planning through planning graph analysis. *Artificial intelligence* 90(1-2):281–300.
- Blum, A. L., and Langford, J. C. 2000. Probabilistic planning in the graphplan framework. *Lecture Notes in Artificial Intelligence (Subseries of Lecture Notes in Computer Science)* 1809:319–332.
- Bramley, N. R.; Rothe, A.; Tenenbaum, J. B.; Xu, F.; and Gureckis, T. M. 2018. Grounding Compositional Hypothesis Generation in Specific Instances. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society* 1:1392–1397.
- Chollet, F. 2019. On the Measure of Intelligence. 1–64.
- Crawford, E., and Pineau, J. 2019. Spatially invariant unsupervised object detection with convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3412–3420.
- de Felice, G.; Toumi, A.; and Coecke, B. 2020. DisCoPy: Monoidal Categories in Python. In *Applied Category Theory Conference*, 1–20.
- Ellis, K.; Morales, L.; Sablé-meyer, M.; Paris-saclay, E. N. S.; and Solar-lezama, A. 2018. Library Learning for Neurally-Guided Bayesian Program Induction. In *Neural Information Processing Systems*, 1–11.
- Ellis, K.; Wong, C.; Nye, M.; Sable-Meyer, M.; Cary, L.; Morales, L.; Hewitt, L.; Solar-Lezama, A.; and Tenenbaum, J. B. 2020. Dreamcoder: Growing generalizable, interpretable knowledge with wake-sleep bayesian program learning. *arXiv preprint arXiv:2006.08381*.
- Frankland, S. M., and Greene, J. D. 2020. Concepts and Compositionality: In Search of the Brain’s Language of Thought. *Annual Review of Psychology* 71(1):273–303.
- Goodman, N. D.; Tenenbaum, J. B.; and Gerstenberg, T. 2014. Concepts in a Probabilistic Language of Thought. In *Concepts: New Directions*, number 010. 1–25.
- Hagberg, A. A.; Schult, D. A.; and Swart, P. J. 2008. Exploring network structure, dynamics, and function using networkx. In Varoquaux, G.; Vaught, T.; and Millman, J., eds., *Proceedings of the 7th Python in Science Conference*, 11 – 15.
- Hoffman, M. D.; Blei, D. M.; Wang, C.; and Paisley, J. 2013. Stochastic variational inference. *The Journal of Machine Learning Research* 14(1):1303–1347.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; and Kavukcuoglu, K. 2015. Spatial Transformer Networks. In *Advances in Neural Information Processing Systems*, volume 2.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kingma, D. P., and Welling, M. 2014. Auto-Encoding Vari-

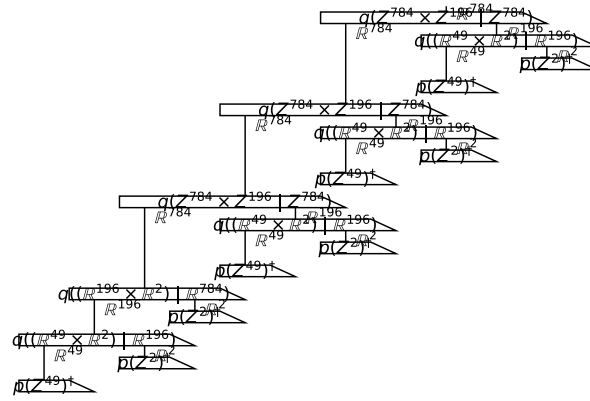


Figure 9: Dagger diagram learned as stochastic inverse of Figure 8.

- ational Bayes. In *International Conference on Learning Representations*, 1–14.
- Lake, B. M., and Piantadosi, S. T. 2019. People infer recursive visual concepts from just a few examples. *Computational Brain & Behavior*.
- Lake, B. M.; Salakhutdinov, R. R.; and Tenenbaum, J. 2013. One-shot learning by inverting a compositional causal process. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems (NIPS)*, 2526–2534. Curran Associates, Inc.
- Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science* 350(6266):1332–1338.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist> 2.
- Loaiza-Ganem, G., and Cunningham, J. P. 2019. The continuous bernoulli: fixing a pervasive error in variational autoencoders. In *Advances in Neural Information Processing Systems*, 13287–13297.
- Marcus, G. F. 2018. *The algebraic mind: Integrating connectionism and cognitive science*. MIT press.
- Mnih, A., and Gregor, K. 2014. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, 1791–1799.
- Nie, W.; Yu, Z.; Mao, L.; Patel, A. B.; Zhu, Y.; and Anandkumar, A. 2020. Bongard-logo: A new benchmark for human-level concept learning and reasoning. *Advances in Neural Information Processing Systems* 33.
- Overlan, M. C.; Jacobs, R. A.; and Piantadosi, S. T. 2017. Learning abstract visual concepts via probabilistic program induction in a Language of Thought. *Cognition* 168:320–334.
- Piantadosi, S. T.; Tenenbaum, J. B.; and Goodman, N. D. 2016. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological Review* 123(4):392–424.
- Romano, S.; Salles, A.; Amalric, M.; Dehaene, S.; Sigman, M.; and Figueira, S. 2018. Bayesian validation of grammar productions for the language of thought. *PLoS ONE* 13(7):1–20.
- Schulman, J.; Heess, N.; Weber, T.; and Abbeel, P. 2015. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, 3528–3536.
- Solomonoff, R. J. 1964. A formal theory of inductive inference. Part I. *Information and Control* 7(1):1–22.
- Stuhlmüller, A.; Taylor, J.; and Goodman, N. D. 2013. Learning stochastic inverses. *Advances in Neural Information Processing Systems* 1–9.
- Valkov, L.; Chaudhari, D.; Sutton, C.; Srivastava, A.; and Chaudhuri, S. 2018. Houdini: Lifelong learning as program synthesis. In *Advances in Neural Information Processing Systems*, volume 2018-Decem, 8687–8698.
- Vedantam, R.; Szlam, A.; Nickel, M.; Morcos, A.; and Lake, B. 2020. Curi: A benchmark for productive concept learning under uncertainty. *arXiv preprint arXiv:2010.02855*.
- Vitányi, P. M., and Li, M. 2000. Minimum description length induction, Bayesianism, and Kolmogorov complexity. *IEEE Transactions on Information Theory* 46(2):446–464.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.
- Xu, K.; Srivastava, A.; and Sutton, C. 2019. Variational russian roulette for deep bayesian nonparametrics. In *International Conference on Machine Learning*, 6963–6972.
- Zhao, S.; Song, J.; and Ermon, S. 2017. Learning Hierarchical Features from Generative Models. In *International Conference on Machine Learning*.
- Zhu, Y.; Gao, T.; Fan, L.; Huang, S.; Edmonds, M.; Liu, H.; Gao, F.; Zhang, C.; and Qi, S. 2019. Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Human-like Commonsense. *Engineering*.