A Metalearned Neural Circuit for Nonparametric Bayesian Inference

Jake C. Snell[1] Gianluca M. Bencomo[1] Thomas L. Griffiths[1,2]

[1]Department of Computer Science, Princeton University [2]Department of Psychology, Princeton University

Overview

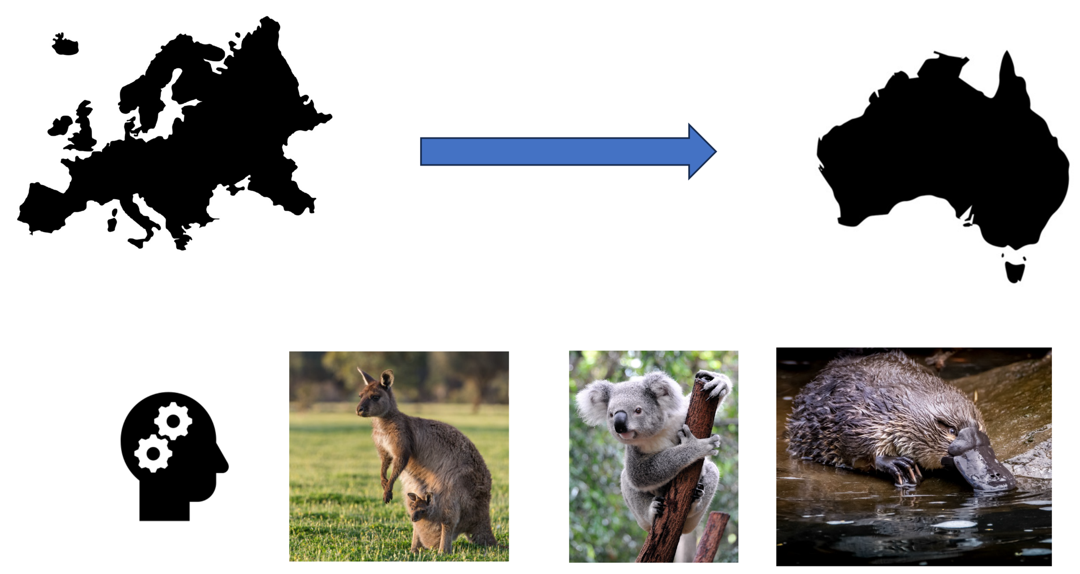

We want to build flexible deep neural networks (DNNs) in order to solve challenging tasks like open-set classification.

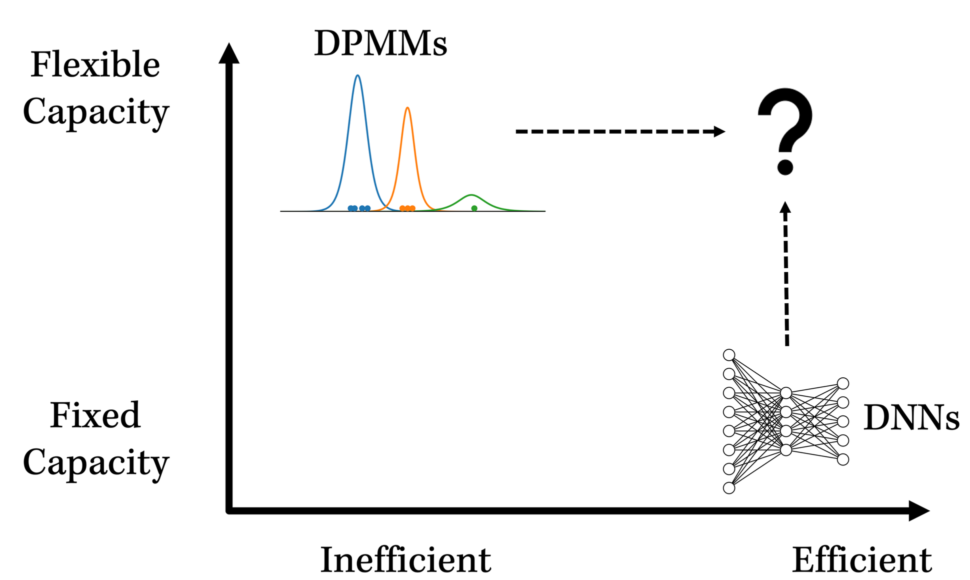

Nonparametric Bayesian (NPB) methods such as Dirichlet process mixture models (DPMMs) capture this flexibility but are computationally inefficient. Can we transfer the flexible inductive bias from DPMMs to DNNs?

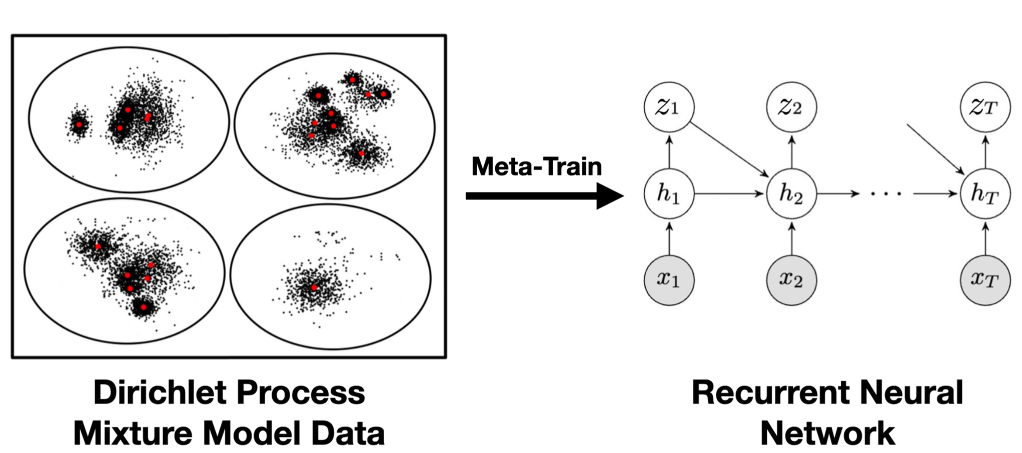

By simulating data from a DPMM, we show how to metalearn a sequence model called a "neural circuit" that performs inference over an unlimited set of classes.

Our experiments show that the metalearned neural circuit achieves comparable or better performance than particle filter-based methods that explicitly use NPB inference.

Dirichlet Process Mixture Models

The DPMM follows a Markov process on class labels $z_1, \ldots, z_T$ and assumes class-conditional independence for $\mathbf{x}_1, \ldots, \mathbf{x}_T$.

$$p(z_{1:T}, \mathbf{x}_{1:T}) = p(z_1)p(\mathbf{x}_1|z_1)\prod_{t=2}^{T} p(z_t|z_{1:t-1})p(\mathbf{x}_t|z_t)$$

The conditional distribution on class memberships follows a Chinese Restaurant Process (CRP):

$$p(z_t = k \mid z_{1:t-1}) \propto \begin{cases} n_k & k \text{ previously observed} \\ \alpha & k \text{ is a new class} \end{cases}$$

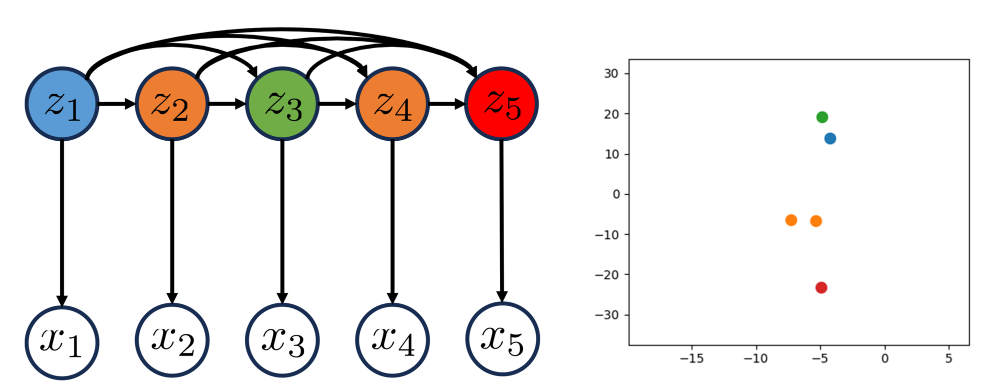

Metalearning a Neural Circuit

In metalearning, we sample tasks from a distribution and take gradient descent steps to improve the performance on the task. Our key insight is to use a NPB model to define the task distribution.

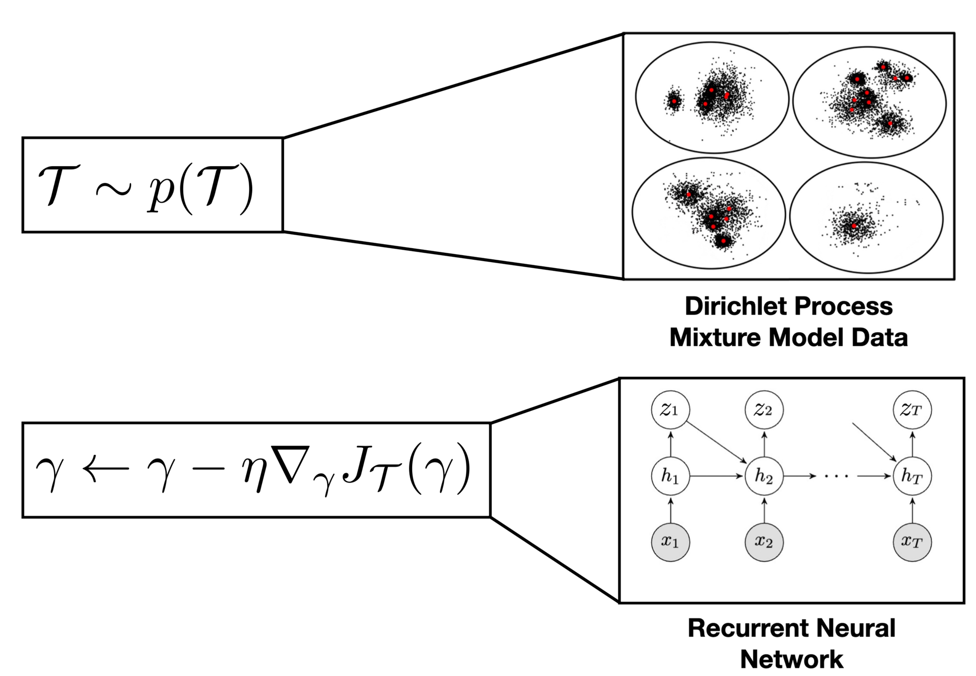

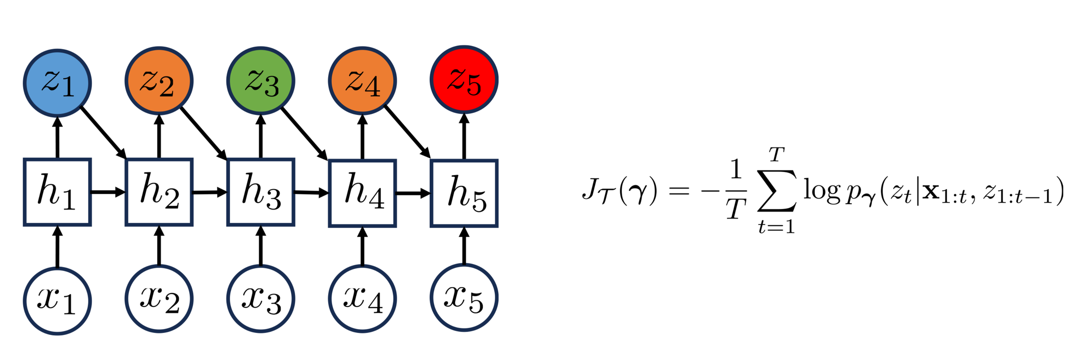

$$p_\gamma(z_t|\mathbf{x}_{1:t}, z_{1:t-1}) = \textsc{Softmax}(\mathbf{a}_t + \mathbf{m}_t)$$
$$\mathbf{a}_t = \mathbf{W}\mathbf{u}_t + \mathbf{b}$$
$$m_{tk} = \begin{cases} 0 & k \le 1 + \max z_{1:t-1} \\ -\infty & \text{otherwise,} \end{cases}$$
$$\mathbf{u}_t, \mathbf{h}_t \leftarrow \mathrm{RNN}_{\boldsymbol{\theta}}([\mathbf{x}_t, \textsc{OneHot}(z_{t-1})], \mathbf{h}_{t-1})$$

Modeling Synthetic DPMM Data

We first compare performance on synthetic DPMM data. We use a normal-inverse-gamma prior and Gaussian class conditional distributions with unknown mean and variance:

$$\sigma_d^2 \sim \Gamma^{-1}(a, b) \quad \mu_d \sim \mathcal{N}(m, \sigma_d^2/\lambda) \quad x_d \sim \mathcal{N}(\mu_d, \sigma_d^2)$$

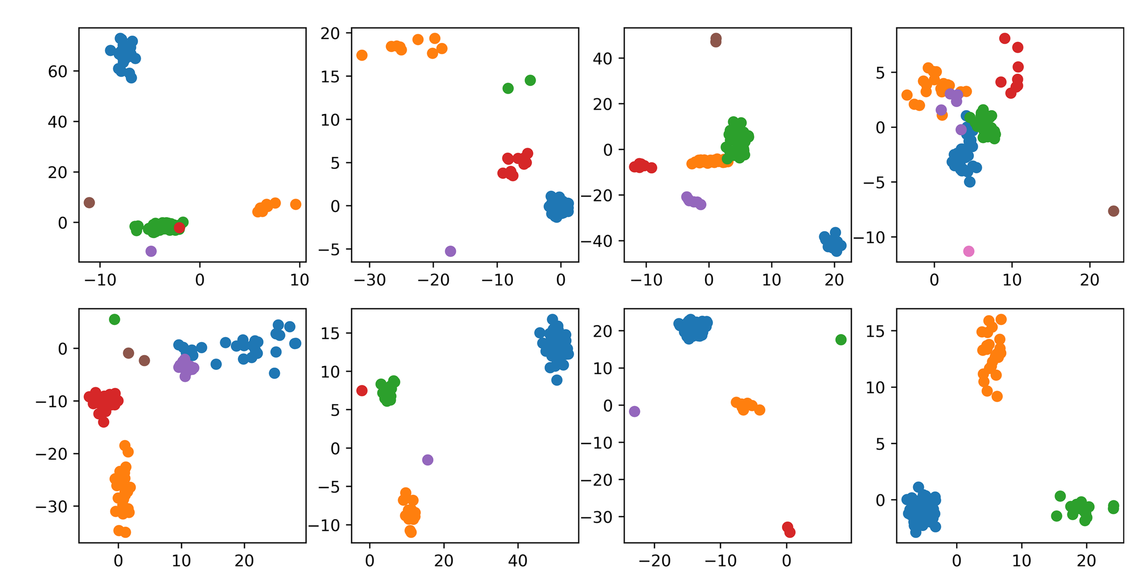

After metalearning on this synthetic data, the neural circuit learns to produce sensible predictive distributions that flexibly allow for the possibility of a previously unseen class.

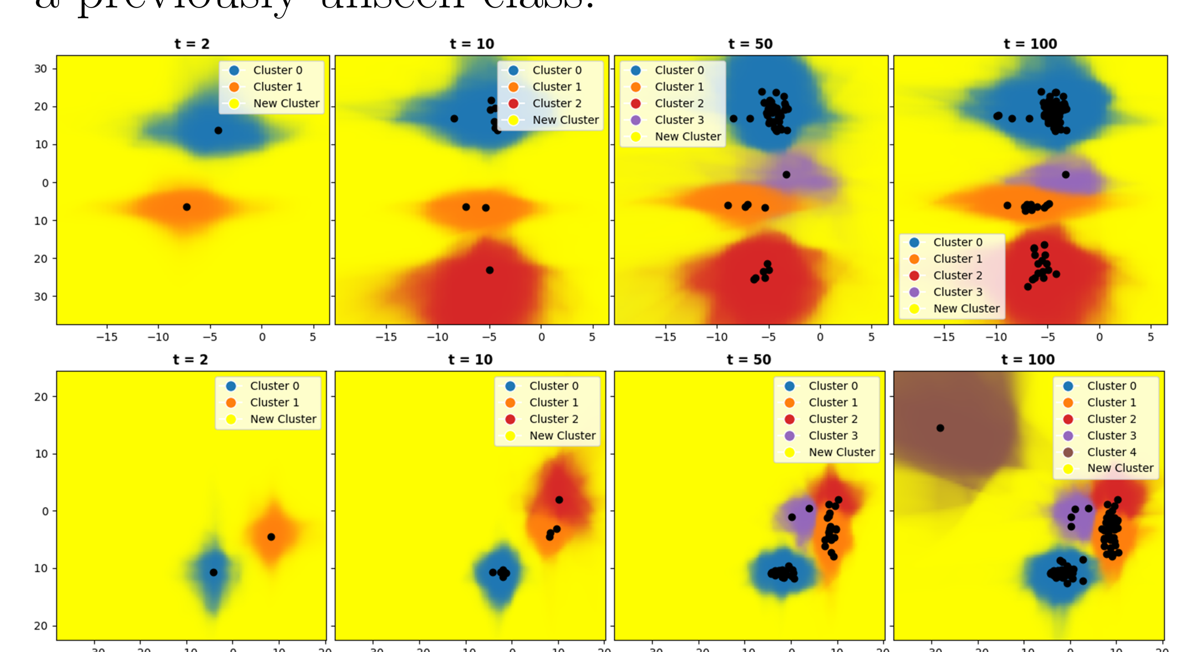

We find that the neural circuit achieves comparable performance to a particle filter [1] but is significantly more efficient.

Method	NLL (↓)	ARI (↑)	Time (ms)
CRP	1.006	0.010	**0.019**
Softmax + Energy	0.929	0.388	1679.716
Particle Filter	**0.048**	0.769	1.617
Neural Circuit	0.076	**0.921**	0.059

ImageNet-CRP Experiments

We also metalearn neural circuits on natural images. We sample a sequence of labels from the CRP as usual, but sample images by randomly choosing an image from the dataset with the corresponding label. We call this dataset ImageNet-CRP.

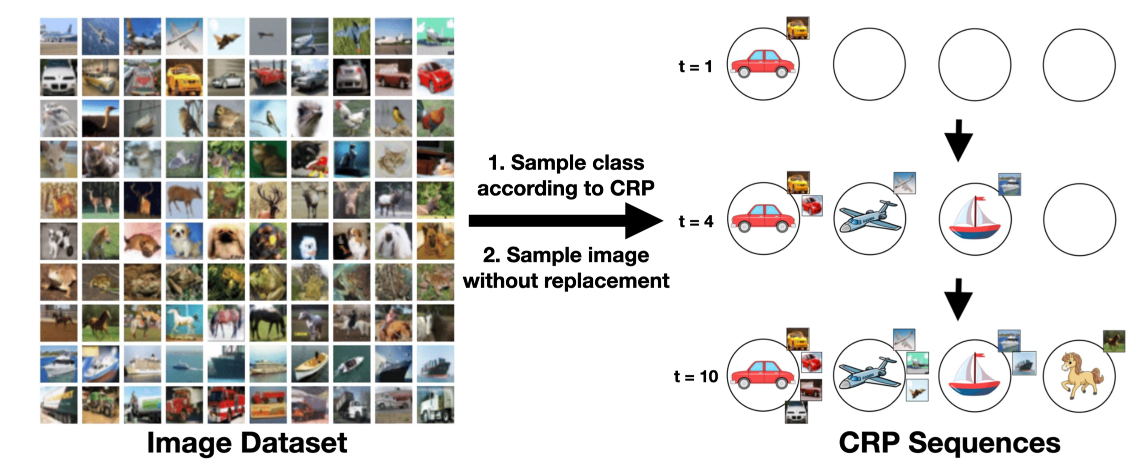

	NLL (↓)		
Method	Meta-train	Meta-test	Time (ms)
CRP	1.005	1.003	**0.019**
SM + Energy	3.196	3.471	1883.066
Particle Filter	0.848	0.933	2.407
Neural Circuit	**0.255**	**0.680**	0.067

iNaturalist 2018 Experiments

We transfer neural circuits trained using ImageNet-CRP to the long-tailed dataset iNaturalist 2018, which is labeled at 7 levels. We freeze the neural circuit and learn only the weights of an affine transform + ReLU.

	CRP		Neural Circuit	
Taxonomy	α	NLL	α	NLL
Kingdom	0.6	0.70	1	**0.42**
Phylum	2.1	1.30	2	**0.84**
Class	5.3	1.82	5	**1.31**
Order	33.1	**2.24**	20	2.28
Family	144.5	**1.56**	100	1.62
Genus	758.6	**0.53**	200	0.59
Species	1584.9	**0.28**	200	0.39

References & Acknowledgments

[1] Fearnhead, P. (2004) Stats. & Computing.

Acknowledgment: This work was supported by grant N00014-23-1-2510 from the Office of Naval Research.