Predictive Coding, Variational Autoencoders, and Biological Connections

Joseph Marino
Computation & Neural Systems
California Institute of Technology
jmarino@caltech.edu

Abstract

This paper identifies connections between predictive coding, from theoretical neuroscience, and variational autoencoders, from machine learning. These connections imply striking correspondences for biological neural circuits, suggesting that pyramidal dendrites are functionally analogous to non-linear deep networks and lateral inhibition is functionally analogous to normalizing flows. Connecting these areas provides new directions for further investigations.

1 Introduction

1.1 Cybernetics

Machine learning and theoretical neuroscience once overlapped under the field of cybernetics [Wiener, 1948, Ashby, 1956]. Within this field, perception and control, both in biological and non-biological systems, were formulated in terms of negative feedback and feedforward processes. Briefly, negative feedback attempts to minimize error signals by *feed*ing the errors *back* into the system, whereas feedforward processing attempts to preemptively reduce error through prediction. At the heart of its formulation, cybernetics formalized these techniques using probabilistic models, i.e. models that estimate the likelihood of random outcomes, and variational calculus, a mathematical technique for estimating functions, particularly probability distributions [Wiener, 1948]. These techniques resulted in the first computational models of neuron function [McCulloch and Pitts, 1943], a formal definition of information [Wiener, 1942, Shannon, 1948] (with connections to neural systems [Barlow, 1961]), and algorithms for negative feedback perception and control [MacKay, 1956, Kalman, 1960]. However, the field of cybernetics ultimately disbanded, with the new mathematical techniques and ideas surviving in the offshoots of theoretical neuroscience, machine learning, control theory, etc.

1.2 Neuroscience & Machine Learning: Convergence & Divergence

A renewed dialogue between neuroscience and machine learning formed in the 1980s–1990s. Neuroscientists, bolstered by new physiological and functional analyses, began making traction in studying neural systems in probabilistic and information theoretic terms [Laughlin, 1981, Srinivasan et al., 1982, Barlow, 1989, Bialek et al., 1991]. In machine learning, improvements in probabilistic modeling [Pearl, 1986] and artificial neural networks [Rumelhart et al., 1986] combined with ideas from statistical mechanics [Hopfield, 1982, Ackley et al., 1985] to yield new classes of models and training techniques. This convergence of ideas, primarily centered around perception, resulted in new theories of neural processing and improvements in their mathematical underpinnings.

In particular, the notion of *predictive coding* emerged within neuroscience [Srinivasan et al., 1982, Rao and Ballard, 1999]. In its most general form, predictive coding postulates that neural circuits are

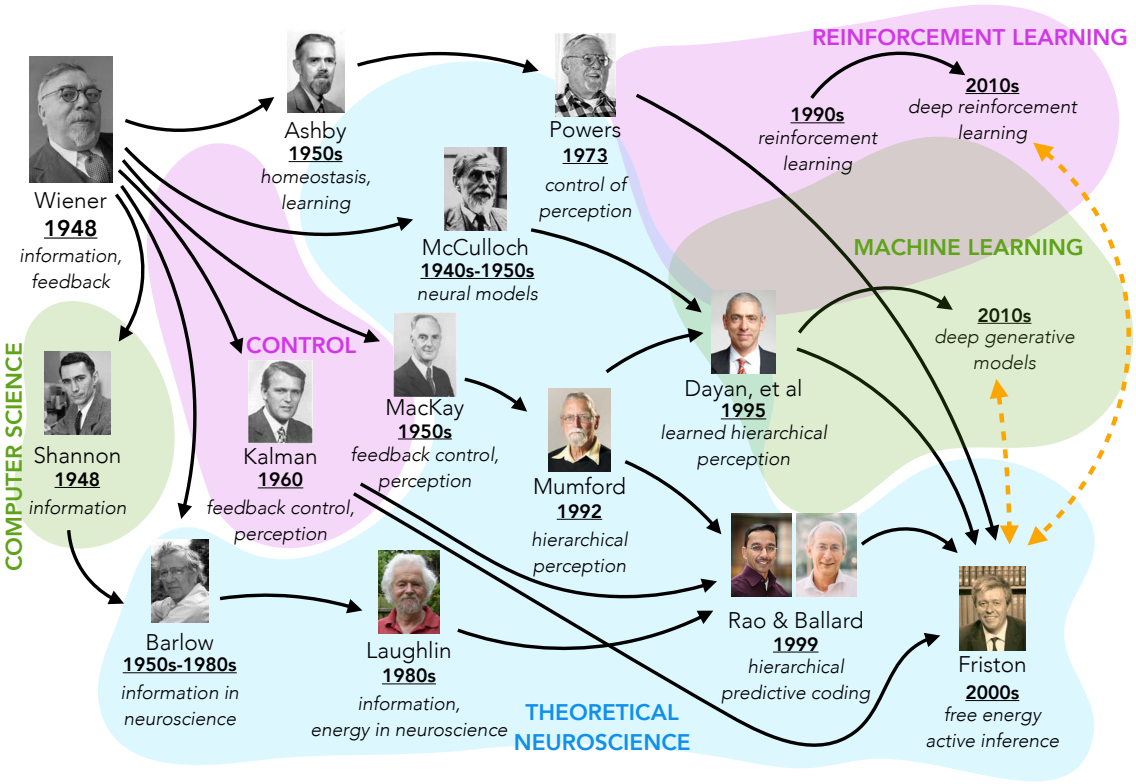

Figure 1: **Conceptual Evolution**. Concepts emphasized by cybernetics (top left) influenced the areas that ultimately became theoretical neuroscience and machine learning. Despite branching into separate areas, there are shared mathematical concepts between these areas (orange dashed lines). This paper explores the connections between predictive coding, from theoretical neuroscience, and variational autoencoders, from machine learning.

fundamentally engaged in estimating probabilistic models of other neural activity and the surrounding environment, with feedback and feedforward processes playing a central role. These models were initially formulated in early sensory areas, e.g. retina [Srinivasan et al., 1982] and thalamus [Dong and Atick, 1995], using feedforward processes to predict future neural activity. Similar notions were also extended to higher-level sensory processing in neocortex. In a series of papers by David Mumford [Mumford, 1991, 1992], top-down neural projections (from higher-level to lower-level sensory areas) were hypothesized to convey hierarchical sensory predictions, whereas bottom-up neural projections were hypothesized to convey prediction errors. Through a negative feedback process, these errors would then update state estimates. These ideas were formalized and analyzed by Rao and Ballard [1999], formulating a simplified artificial neural network model of images, reminiscent of a Kalman filter [Kalman, 1960].

Feedback and feedforward processes also featured prominently in machine learning. Indeed, the primary training algorithm for artificial neural networks, backpropagation [Rumelhart et al., 1986], literally *feeds* (propagates) the output prediction errors *back* through the network, i.e. negative feedback. During this period, the technique of variational inference was rediscovered within machine learning [Hinton and Van Camp, 1993, Neal and Hinton, 1998], recasting probabilistic inference using variational calculus. This technique proved essential in formulating the Helmholtz machine [Dayan et al., 1995, Dayan and Hinton, 1996], a hierarchical probabilistic model parameterized by artificial neural networks. Similar advances were made in autoregressive probabilistic models [Frey et al., 1996,

Bengio and Bengio, 2000], using artificial neural networks to form sequential feedforward predictions, as well as new classes of invertible probabilistic models [Parra et al., 1995, Deco and Brauer, 1995, Bell and Sejnowski, 1997]. Unfortunately, as the chill of an "AI winter" arrived in the late 1990s, progress in these areas slowed.

In neuroscience, these new ideas surrounding variational inference and probabilistic models influenced predictive coding. This effort was led by Karl Friston, utilizing variational inference [Friston, 2005] to formulate hierarchical dynamical models of neocortex [Friston, 2008a]. In line with Mumford's proposal [Mumford, 1992], these models contain multiple levels of variables, with each level attempting to predict its future activity (feedforward) as well as the activity at lower levels, closer to the input data. Through variational inference, prediction errors across levels facilitate updating higher-level estimates (negative feedback). Such models have incorporated many aspects of neuroscience, including local learning rules [Friston, 2005] and attention [Spratling, 2008, Feldman and Friston, 2010, Kanai et al., 2015], and have been coarsely compared with neural circuits [Bastos et al., 2012, Keller and Mrsic-Flogel, 2018, Walsh et al., 2020]. While predictive coding and other forms of Bayesian brain theories have become increasingly popular [Doya et al., 2007, Friston, 2009, Clark, 2013], empirically testing these normative models remains challenging. This is partially due to the difficulty of distinguishing between the large number of specific design choices and the more general theoretical claims of probabilistic learning and inference [Gershman, 2019]. Further, because these models have been limited to simplified implementations, often without learned parameters, it has been difficult to bridge the gap to the complexity of biological neural systems.

The AI winter thawed in the early 2010s, brought on by advances in parallel computing as well as standardized datasets [Deng et al., 2009, Krizhevsky and Hinton, 2009] and environments [Todorov et al., 2012, Bellemare et al., 2013]. In this new era of *deep learning* [LeCun et al., 2015, Schmidhuber, 2015], i.e. artificial neural networks with multiple layers, a flourishing of ideas emerged around probabilistic modeling. Building off of previous work, more expressive classes of deep hierarchical [Gregor et al., 2014, Mnih and Gregor, 2014, Kingma and Welling, 2014, Rezende et al., 2014], autoregressive [Uria et al., 2014, van den Oord et al., 2016b], and invertible [Dinh et al., 2015, 2017] probabilistic models were developed. Of particular importance is a class of models known as *variational autoencoders* (VAEs) [Kingma and Welling, 2014, Rezende et al., 2014], which bear a close resemblance to hierarchical predictive coding models. Unfortunately, despite this similarity, the machine learning community remains largely oblivious to the progress in predictive coding that occurred during the AI winter (and vice versa).

1.3 Predictive Coding Meets Deep Learning

This paper aims to help bridge that divide. While the present work provides unique contributions, many of the insights were inspired by previous works at this intersection. In particular, van den Broeke [2016] outlines the relationship between hierarchical probabilistic models in predictive coding and deep learning. Likewise, Lotter et al. [2017] implement basic predictive coding techniques in deep probabilistic models, later comparing these models with neural phenomena [Lotter et al., 2018].

After establishing several common mathematical concepts (Section 2), we review the basic formulations of the main areas of predictive coding (Section 3) and variational autoencoders (Section 4). We then identify connections between these areas through the techniques of iterative amortization [Marino et al., 2018b] and sequential autoregressive flows [Marino et al., 2020a]. Briefly, iterative amortization facilitates efficiently updating estimates to minimize prediction errors, i.e. negative feedback, whereas sequential autoregressive flows provide a feedforward mechanism for improving predictions across time. By identifying these connections, we arrive at striking new perspectives on the correspondences between machine learning and neuroscience. In particular, these correspondences suggest that

- **the dendrites of pyramidal neurons are functionally analogous to deep networks**, questioning fundamental assumptions on the analogy of biological and artificial neurons [McCulloch and Pitts, 1943] and
- **normalization in neural circuits is functionally analogous to normalizing flows** [Rezende and Mohamed, 2015], providing a general mathematical framework for this canonical neural computation [Carandini and Heeger, 2012].

Like the works of van den Broeke [2016] and Lotter et al. [2017, 2018], we hope that these ideas will inspire future research in exploring this promising direction, bringing together both fields.

2 Background

This section covers the relevant mathematical topics for this paper. It is self-contained, at the expense of added length. Readers familiar with maximum likelihood, probabilistic models, and variational inference may proceed to the following section.

2.1 Maximum Log-Likelihood

Consider a random variable, $X \in \mathbb{R}^M$, with a corresponding distribution, $p_{\text{data}}(X)$, defining the probability of observing each possible observation, $X = \mathbf{x}$. We will use the shorthand notation $p_{\text{data}}(\mathbf{x})$ to denote the probability $p_{\text{data}}(X = \mathbf{x})$. This distribution is the result of an underlying data generation process, e.g. the emission and scattering of photons. While we do not have direct access to p_{data}, we can sample observations, $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$, yielding an empirical distribution, $\widehat{p}_{\text{data}}(\mathbf{x})$.

Often, we wish to model p_{data}, for instance, to predict or compress observations of X. We refer to this model as $p_\theta(\mathbf{x})$, with parameters θ. Estimating the model parameters involves maximizing the log-likelihood of data examples under the model's distribution:

$$\theta^* \leftarrow \arg\max_\theta \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} \left[\log p_\theta(\mathbf{x})\right]. \tag{1}$$

This is the *maximum log-likelihood* objective, which is found throughout machine learning and probabilistic modeling [Murphy, 2012, Goodfellow et al., 2016]. In practice, we do not have access to $p_{\text{data}}(\mathbf{x})$ and must instead approximate the objective using data samples, i.e. using $\widehat{p}_{\text{data}}(\mathbf{x})$.

2.2 Probabilistic Models

2.2.1 Dependency Structure

Formulating a probabilistic model involves considering the dependency structure of the model and the parameterization of these dependencies. The dependency structure of a probabilistic model is the set of conditional dependencies between variables. One common form of dependency structure is found in **autoregressive models** [Frey et al., 1996, Bengio and Bengio, 2000], which utilize the chain rule of probability to model dependencies between variables:

$$p_\theta(\mathbf{x}) = \prod_{j=1}^{M} p_\theta(x_j | x_{<j}). \tag{2}$$

Here, we have induced an arbitrary ordering over the M dimensions of $\mathbf{x}$, allowing us to factor the joint distribution over dimensions, $p_\theta(\mathbf{x})$, into a product of M conditional distributions, each conditioned on the previous dimensions, $x_{<j}$. Slightly abusing notation, at $j = 1$, we have $p_\theta(x_j | x_{<j}) = p_\theta(x_1)$.

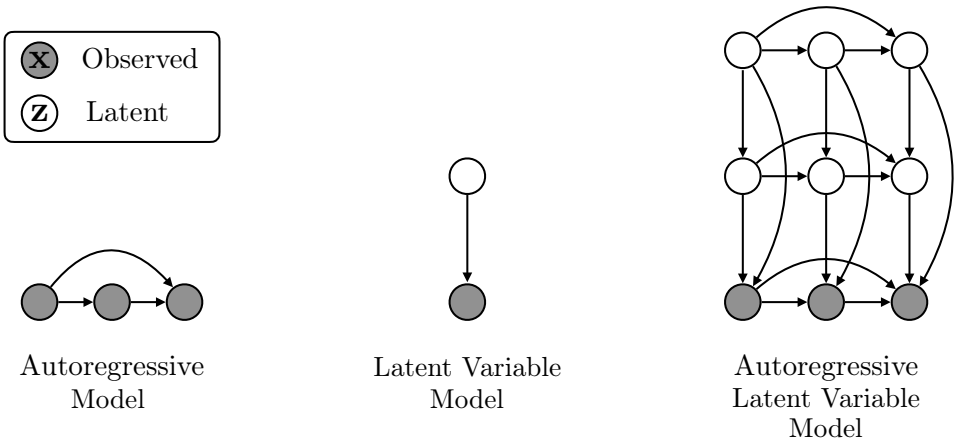

Figure 2: **Dependency Structures**. Each diagram shows a (directed) graphical model, with nodes representing random variables and arrows representing probabilistic dependencies. The two main forms of probabilistic dependency structure are autoregressive (**left**; Eq. 2) and latent variable models (**center**; Eq. 4). These structures can be combined (**right**; Eq. 8) to further improve model capacity.

A natural use-case for this dependency structure arises in modeling sequential data, where time provides an ordering over a sequence of T observed variables, $\mathbf{x}_{1:T}$:

$$p_\theta(\mathbf{x}_{1:T}) = \prod_{t=1}^{T} p_\theta(\mathbf{x}_t|\mathbf{x}_{<t}). \tag{3}$$

Although such models are conventionally formulated in forward temporal order, this is truly a modeling assumption. Likewise, while the chain rule of probability dictates that we must consider *all* previous variables, there may be cases where it is safe to assume conditional independence outside of some window. In the extreme case, in which we only consider pairwise dependencies, we arrive at a Markov chain: $p_\theta(\mathbf{x}_{1:T}) = \prod_{t=1}^{T} p_\theta(\mathbf{x}_t|\mathbf{x}_{t-1})$.

Autoregressive models are also referred to as "fully-visible" models [Frey et al., 1996], as dependencies are only explicitly modeled between observed variables. However, we can also model such dependencies by introducing *latent variables*, denoted as Z. Formally, a **latent variable model** is defined by the joint distribution

$$p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}|\mathbf{z})p_\theta(\mathbf{z}), \tag{4}$$

where $p_\theta(\mathbf{x}|\mathbf{z})$ is the *conditional likelihood* and $p_\theta(\mathbf{z})$ is the *prior*. Again, we have used the shorthand notation $p_\theta(\mathbf{x}, \mathbf{z})$ to denote $p_\theta(X = \mathbf{x}, Z = \mathbf{z})$. Introducing latent variables is one of, if not, *the* primary technique for increasing the flexibility of a probabilistic model. This is because evaluating the probability of an observation now requires marginalizing over the latent variables. If Z is a continuous variable, this involves integration, $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}, \mathbf{z})d\mathbf{z}$, and if Z is discrete, this involves summation, $p_\theta(\mathbf{x}) = \sum_{\mathbf{z}} p_\theta(\mathbf{x}, \mathbf{z})$. In either case, we have

$$p_\theta(\mathbf{x}) = \mathbb{E}_{\mathbf{z}\sim p_\theta(\mathbf{z})}\left[p_\theta(\mathbf{x}|\mathbf{z})\right], \tag{5}$$

which illustrates that $p_\theta(\mathbf{x})$ is a *mixture* distribution, with each mixture component, $p_\theta(\mathbf{x}|\mathbf{z})$, weighted according to $p_\theta(\mathbf{z})$. Thus, even when restricting $p_\theta(\mathbf{x}|\mathbf{z})$ to simple distribution forms, such as Gaussian distributions, $p_\theta(\mathbf{x})$ can take on flexible forms that do not have closed form analytical expressions. In this way, Z can implicitly model dependencies in X, assigning higher probability to particular regions of the observation space.

However, increasing flexibility through latent variables comes with increasing computational overhead. In general, marginalizing over Z is not analytically tractable. This requires us to either 1) adopt approximation techniques, which we discuss in Section 2.3, or 2) restrict the form of the model to ensure computationally tractable evaluation of $p_\theta(\mathbf{x})$. This latter approach is the basis of **flow-based models** [Tabak and Turner, 2013, Rippel and Adams, 2013, Dinh et al., 2015], which define the conditional dependency between X and Z in terms of an invertible transform, $\mathbf{x} = f_\theta(\mathbf{z})$ and $\mathbf{z} = f_\theta^{-1}(\mathbf{x})$. With a prior or *base distribution*, $p_\theta(\mathbf{z})$, we can then express $p_\theta(\mathbf{x})$ using the *change of variables formula*:

$$p_\theta(\mathbf{x}) = p_\theta(\mathbf{z}) \left| \det \left(\frac{\partial \mathbf{x}}{\partial \mathbf{z}} \right) \right|^{-1}, \tag{6}$$

where $\frac{\partial \mathbf{x}}{\partial \mathbf{z}}$ is the Jacobian of the transform and $\det(\cdot)$ denotes matrix determinant. The term $\left|\det \left(\frac{\partial \mathbf{x}}{\partial \mathbf{z}}\right)\right|^{-1}$ can be interpreted as the local scaling of space when moving from Z to X, conserving probability mass in the transform. Flow-based models, also referred to as *normalizing flows* [Rezende and Mohamed, 2015], are the basis of independent components analysis (ICA) [Bell and Sejnowski, 1997, Hyvärinen and Oja, 2000] and non-linear generalizations [Chen and Gopinath, 2001, Laparra et al., 2011]. As such, these models can serve as a general-purpose mechanism for adding and removing statistical dependencies between variables. Although flow-based models avoid the intractability of marginalization, their requirement of invertibility may be overly restrictive or undesirable in some contexts [Cornish et al., 2020]. And while the change of variables formula can also be applied to non-invertible transforms [Cvitkovic and Koliander, 2019], it raises computational intractabilities.

We have presented autoregression and latent variables separately, however, these techniques can, in fact, be combined in numerous ways to model dependencies. For instance, one can create *hierarchical latent variable models* [Dayan et al., 1995], incorporating autoregressive dependencies between latent variables. Considering L levels of latent variables, $Z^{1:L} = \left[Z^1, \ldots, Z^L\right]$, we can express the joint distribution as

$$p_\theta(\mathbf{x}, \mathbf{z}^{1:L}) = p_\theta(\mathbf{x}|\mathbf{z}^{1:L}) \prod_{\ell=1}^{L} p_\theta(\mathbf{z}^\ell|\mathbf{z}^{\ell+1:L}). \tag{7}$$

Hierarchical latent variable models are a repeated application of the latent variables technique in order to create increasingly complex *empirical priors* [Efron and Morris, 1973]. We can also consider incorporating latent variables within sequential (autoregressive) probabilistic models, giving rise to *sequential latent variable models*. Considering a single level of latent variables in a corresponding sequence, $Z_{1:T}$, we have the following joint distribution:

$$p_\theta(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}) = \prod_{t=1}^{T} p_\theta(\mathbf{x}_t|\mathbf{x}_{<t}, \mathbf{z}_{\leq t}) p_\theta(\mathbf{z}_t|\mathbf{x}_{<t}, \mathbf{z}_{<t}), \tag{8}$$

where we have again assumed a forward sequential ordering. This formulation encompasses special cases, such as hidden Markov models or linear Gaussian state-space models [Murphy, 2012]. Beyond hierarchical and sequential latent variable models, there are a variety of other ways to combine autoregression and latent variables [Gulrajani et al., 2017, Razavi et al., 2019].

2.2.2 Parameterizing the Model

The probability distributions that define probabilistic dependencies are ultimately functions. In this section, we discuss forms that these functions may take, restricting our focus here to *parametric* distributions, which are defined by one or more distribution parameters. The canonical example is the Gaussian (or Normal) distribution, $\mathcal{N}(x; \mu, \sigma^2)$, which is defined by a mean, μ, and variance, σ^2.

This can be extended to the multivariate setting, where $\mathbf{x} \in \mathbb{R}^M$ is modeled with a mean vector, $\boldsymbol{\mu}$, and covariance matrix, $\boldsymbol{\Sigma}$, with the probability density written as

$$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{M/2} \det(\boldsymbol{\Sigma})^{1/2}} \exp\left(\frac{-1}{2}(\mathbf{x} - \boldsymbol{\mu})^\intercal \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right). \tag{9}$$

For convenience, we may also consider diagonal covariance matrices, $\boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\sigma}^2)$, simplifying the parameterization and resulting calculations. In particular, the special case where $\boldsymbol{\Sigma} = \mathbf{I}_M$, the $M \times M$ identity matrix, the log-density, up to a constant, becomes the familiar *mean squared error*,

$$\log \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{I}) = -\frac{1}{2}||\mathbf{x} - \boldsymbol{\mu}||_2^2 + \text{const.} \tag{10}$$

With a parametric distribution, conditional dependencies are mediated by the distribution parameters, which are functions of the conditioning variables. For example, we can express an autoregressive Gaussian distribution (Eq. 2) through conditional densities, $p_\theta(x_j|x_{<j}) = \mathcal{N}(x_j; \mu_\theta(x_{<j}), \sigma_\theta^2(x_{<j}))$, where μ_θ and σ_θ^2 are functions taking $x_{<j}$ as input. A similar form applies to autoregressive models on sequences of vector inputs (Eq. 3), with $p_\theta(\mathbf{x}_t|\mathbf{x}_{<t}) = \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_\theta(\mathbf{x}_{<t}), \boldsymbol{\Sigma}_\theta(\mathbf{x}_{<t}))$. Likewise, in a latent variable model (Eq. 4), we can express a Gaussian conditional likelihood as $p_\theta(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_\theta(\mathbf{z}), \boldsymbol{\Sigma}_\theta(\mathbf{z}))$. In the above examples, we have overloaded notation, using a subscript θ for all functions, however, these may be separate functions in practice.

The functions supplying each of the distribution parameters can range in complexity, from constant to highly non-linear. Classical modeling techniques often employ linear functions. For instance, in a latent variable model, we could parameterize the mean as a linear function of $\mathbf{z}$:

$$\boldsymbol{\mu}_\theta(\mathbf{z}) = \mathbf{W}\mathbf{z} + \mathbf{b} \tag{11}$$

where $\mathbf{W}$ is a matrix of weights and $\mathbf{b}$ is a bias vector. Models of this form underlie factor analysis, probabilistic principal components analysis [Tipping and Bishop, 1999], independent components analysis [Bell and Sejnowski, 1997, Hyvärinen and Oja, 2000], and sparse coding [Olshausen and Field, 1996]. Linear autoregressive models are also the basis of many classical time-series models. While linear models are relatively computationally efficient, they are often too limited to accurately model complex data distributions, e.g. those found in natural images or audio.

Recent improvements in deep learning [Goodfellow et al., 2016] have provided probabilistic models with expressive classes of non-linear functions, improving their capacity. In these models, the distribution parameters are parameterized with deep networks, which are then trained by backpropagating [Rumelhart et al., 1986] the gradient of the log-likelihood objective, $\nabla_\theta \mathbb{E}_{\mathbf{x} \sim \widehat{p}_{\text{data}}} [\log p_\theta(\mathbf{x})]$, back through the network. Deep autoregressive models and deep latent variable models have enabled recent advances in speech [Graves, 2013, van den Oord et al., 2016a], natural language [Sutskever et al., 2014, Radford et al., 2019], images [Razavi et al., 2019], video [Kumar et al., 2020], reinforcement learning [Chua et al., 2018, Ha and Schmidhuber, 2018] and many other areas.

We visualize a probabilistic computation graph for a deep autoregressive model in Figure 3. This diagram breaks the variables into their associated distributions and terms in the log-likelihood objective. Here, green circles denote the conditional likelihood at each step, containing a Gaussian mean and standard deviation, which are parameterized by a deep network. The log-likelihood, $\log p_\theta(\mathbf{x}_t|\mathbf{x}_{<t})$, evaluated at the data observation, $\mathbf{x}_t \sim p_{\text{data}}(\mathbf{x}_t|\mathbf{x}_{<t})$ (gray circle), provides the objective (red dot). The gradient of this objective w.r.t. the network parameters is calculated through backpropagation (red dotted line).

Autoregressive models (without latent variables) have proven useful in many domains, often obtaining better log-likelihoods as compared with latent variable models. However, there are a number of reasons to prefer latent variable models in some contexts. First, autoregressive sampling is inherently sequential, and this linear computational scaling becomes costly in high-dimensional domains. Second, latent variables provide a representational space for downstream tasks, compression,

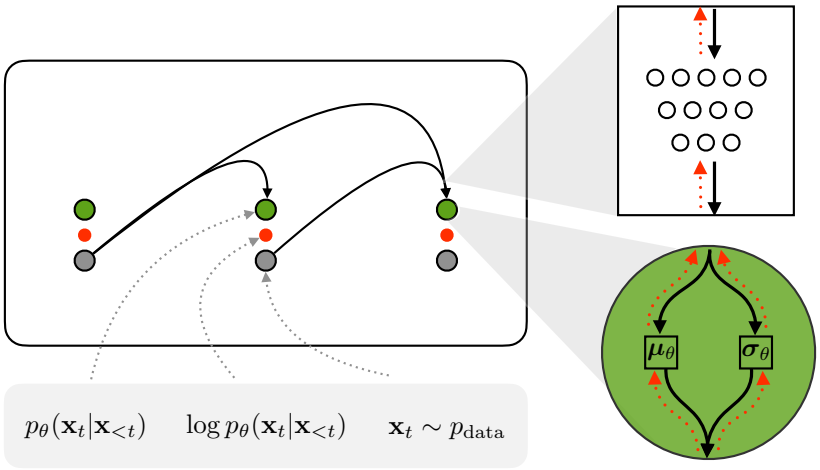

Figure 3: **Model Parameterization & Computation Graph**. The diagram depicts a simplified computation graph for a deep autoregressive Gaussian model. Green circles denote the conditional likelihood distribution at each step, while gray circles again denote the (distribution of) data observations. Smaller red circles denote each of the log-likelihood terms in the objective. Gradients w.r.t. these terms are backpropagated through the networks parameterizing the model's distribution parameters (red dotted lines).

and overall data analysis. Finally, latent variables provide added flexibility, which is particularly useful for modeling continuous random variables with relatively simple, e.g. Gaussian, conditional distributions. For these reasons, we require methods for handling the latent marginalization in Eq. 5. Variational inference is one such method.

2.3 Variational Inference

Training latent variable models through maximum likelihood requires evaluating $\log p_\theta(\mathbf{x})$. However, evaluating $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}, \mathbf{z}) d\mathbf{z}$ is generally computationally intractable. This problem is only exacerbated in deep latent variable models, where computing $p_\theta(\mathbf{z})$ and $p_\theta(\mathbf{x}|\mathbf{z})$ is more expensive. Thus, we require some technique for tractably estimating $\log p_\theta(\mathbf{x})$. Variational inference [Hinton and Van Camp, 1993, Jordan et al., 1998] approaches this problem by introducing an *approximate posterior* distribution, $q(\mathbf{z}|\mathbf{x})$, which provides a tractable lower bound, $\mathcal{L}(\mathbf{x}; q, \theta) \leq \log p_\theta(\mathbf{x})$, on the log-likelihood. This lower bound is referred to as the *evidence lower bound (ELBO)*, variational lower bound, and the negative free energy. By tightening and maximizing the ELBO w.r.t. the model parameters, θ, we can approximate maximum likelihood training while avoiding marginalization.

Variational inference converts probabilistic inference into an optimization problem. Given a family of distributions, $\mathcal{Q}$, e.g. Gaussian, non-parametric, etc., variational inference attempts to find the distribution, $q \in \mathcal{Q}$, that minimizes $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x}))$:

$$q(\mathbf{z}|\mathbf{x}) \leftarrow \arg\min_q D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})), \tag{12}$$

where $p_\theta(\mathbf{z}|\mathbf{x})$ is the *posterior* distribution, $p_\theta(\mathbf{z}|\mathbf{x}) = \frac{p_\theta(\mathbf{x},\mathbf{z})}{p_\theta(\mathbf{x})}$. Because $p_\theta(\mathbf{z}|\mathbf{x})$ includes the intractable

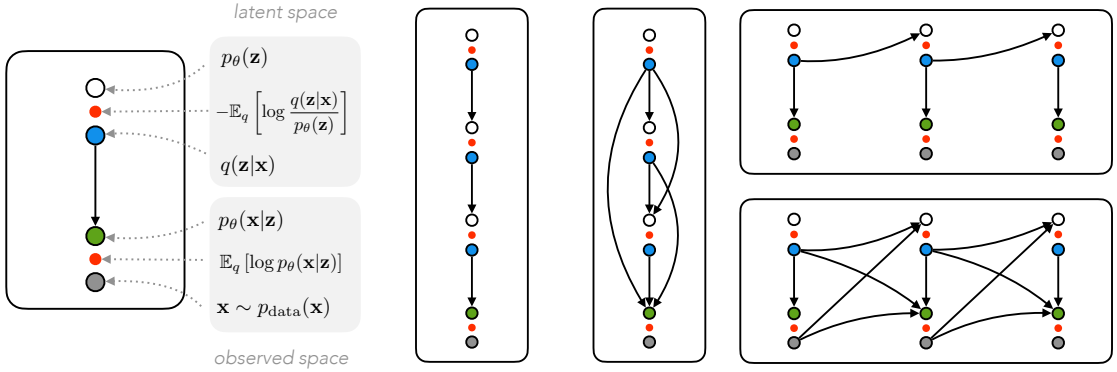

(a) Computation Graph (b) Hierarchical (c) Sequential

Figure 4: **ELBO Computation Graphs**. **(a)** Basic computation graph for variational inference. Outlined circles denote distributions. Smaller red circles denote terms in the ELBO objective. Arrows, again, denote conditional dependencies. This notation can be used to express **(b)** hierarchical and **(c)** sequential models with various model dependencies.

$p_\theta(\mathbf{x})$, we cannot minimize the KL divergence in Eq. 12 directly. Instead, we can rewrite this as

$$
\begin{aligned}
D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})) &= \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x})}\left[\log q(\mathbf{z}|\mathbf{x}) - \log p_\theta(\mathbf{z}|\mathbf{x})\right] && (13)\\
&= \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x})}\left[\log q(\mathbf{z}|\mathbf{x}) - \log\left(\frac{p_\theta(\mathbf{x},\mathbf{z})}{p_\theta(\mathbf{x})}\right)\right] && (14)\\
&= \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x})}\left[\log q(\mathbf{z}|\mathbf{x}) - \log p_\theta(\mathbf{x},\mathbf{z})\right] + \log p_\theta(\mathbf{x}) && (15)\\
&= -\mathcal{L}(\mathbf{x};q,\theta) + \log p_\theta(\mathbf{x}). && (16)
\end{aligned}
$$

In Eq. 16, we have defined $\mathcal{L}(\mathbf{x};q,\theta)$, as

$$
\boxed{
\begin{aligned}
\mathcal{L}(\mathbf{x};q,\theta) &\equiv \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x})}\left[\log p_\theta(\mathbf{x},\mathbf{z}) - \log q(\mathbf{z}|\mathbf{x})\right] && (17)\\
&= \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{x})}\left[\log p_\theta(\mathbf{x}|\mathbf{z})\right] - D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z})). && (18)
\end{aligned}
}
$$

Rearranging terms in Eq. 16, we have

$$\log p_\theta(\mathbf{x}) = \mathcal{L}(\mathbf{x};q,\theta) + D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})). \tag{19}$$

Because KL divergence is non-negative, $\mathcal{L}(\mathbf{x};q,\theta) \leq \log p_\theta(\mathbf{x})$, with equality when $q(\mathbf{z}|\mathbf{x}) = p_\theta(\mathbf{z}|\mathbf{x})$. As the LHS of Eq. 19 does not depend on $q(\mathbf{z}|\mathbf{x})$, maximizing $\mathcal{L}(\mathbf{x};q,\theta)$ w.r.t. q implicitly minimizes $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x}))$ w.r.t. q. Together, these statements imply that maximizing $\mathcal{L}(\mathbf{x};q,\theta)$ w.r.t. q tightens the lower bound on $\log p_\theta(\mathbf{x})$. With this tightened lower bound, we can then maximize $\mathcal{L}(\mathbf{x};q,\theta)$ w.r.t. θ. This alternating optimization process is referred to as the variational expectation maximization (EM) algorithm [Dempster et al., 1977, Neal and Hinton, 1998], consisting of approximate inference (E-step) and learning (M-step).

As with autoregressive models, we can represent latent variable models and the ELBO objective as a computation graph. In Figure 4, we illustrate examples of these graphs. Each variable contains a red circle, denoting a term in the ELBO objective. In comparison with the fully-observed autoregressive model, we now have an additional objective term for the latent variable, corresponding to the KL divergence. This graphical representation also allows us to visualize the variational objective for more complex hierarchical and sequential models (Figures 4b & 4c).

3 Predictive Coding

Predictive coding, as described within neuroscience, can be divided into two distinct settings, spatiotemporal and hierarchical, roughly corresponding to the two main forms of structured probabilistic dependencies. In this section, we review each of these settings, discussing previously hypothesized correspondences with neural anatomy. Finally, we outline the current empirical support for predictive coding in neural systems, highlighting the need for large-scale, testable models.

3.1 Spatiotemporal Predictive Coding

Spatiotemporal predictive coding [Srinivasan et al., 1982], as the name implies, involves forming predictions across spatial dimensions and temporal sequences. These predictions then produce the resulting "code" as the prediction error. Concretely, in the temporal setting, we can consider a Gaussian autoregressive model, p_θ, defined over observation sequences, $\mathbf{x}_{1:T}$. The conditional probability at time t can be written as

$$p_\theta(\mathbf{x}_t|\mathbf{x}_{<t}) = \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_\theta(\mathbf{x}_{<t}), \text{diag}(\boldsymbol{\sigma}^2_\theta(\mathbf{x}_{<t}))).$$

Introducing auxiliary variables, $\mathbf{y}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, we can use the reparameterization trick to express $\mathbf{x}_t = \boldsymbol{\mu}_\theta(\mathbf{x}_{<t}) + \boldsymbol{\sigma}_\theta(\mathbf{x}_{<t}) \odot \mathbf{y}_t$, where $\odot$ denotes element-wise multiplication. Conversely, we can express the inverse, normalization or *whitening* transform as

$$\mathbf{y}_t = \frac{\mathbf{x}_t - \boldsymbol{\mu}_\theta(\mathbf{x}_{<t})}{\boldsymbol{\sigma}_\theta(\mathbf{x}_{<t})}. \tag{20}$$

An example of temporal normalization with video, adapted from Marino et al. [2020a], is shown in Figure 5b. Note that one special case of this transform involves setting $\boldsymbol{\mu}_\theta(\mathbf{x}_{<t}) \equiv \mathbf{x}_{t-1}$ and $\boldsymbol{\sigma}_\theta(\mathbf{x}_{<t}) \equiv \mathbf{1}$, in which case, $\mathbf{y}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$, i.e. temporal differences. For sequences that change slowly relative to the temporal step-size, this is a reasonable assumption. This inverse transform can remove temporal redundancy in the input sequence. Thus, by Shannon's source coding theorem [Shannon, 1948], we can encode or compress $\mathbf{y}_{1:T}$ more efficiently than $\mathbf{x}_{1:T}$. The benefit of this sequential predictive coding scheme was recognized in the early days of information theory [Harrison, 1952, Oliver, 1952], forming the basis of modern video [Wiegand et al., 2003] and audio [Atal and Schroeder, 1979] compression.

A similar process can also be applied within $\mathbf{x}_t$ to remove spatial dependencies. For instance, we could also apply an autoregressive affine transform over spatial dimensions, predicting the i^{th} dimension, $x_{i,t}$, as a function of previous spatial dimensions, $\mathbf{x}_{1:i,t}$. With linear functions, this corresponds to Cholesky whitening [Pourahmadi, 2011, Kingma et al., 2016]. However, this requires imposing an arbitrary ordering over spatial dimensions. Perhaps a more reasonable approach in the spatial setting is to learn a set of *symmetric* dependencies between dimensions. Here, the linear case corresponds to ZCA whitening [Kessy et al., 2018], shown in Figure 5a. In the natural image domain, both of these whitening schemes generally result in center-surround spatial filters, extracting edges from the input.

Srinivasan et al. [1982] investigated the principles of spatiotemporal predictive coding in the retina, where compression is essential for transmission through the optic nerve. Estimating the autocorrelation function of input sensory signals, i.e. a linear prediction, they showed that spatiotemporal predictive coding provides a reasonable fit to retinal ganglion cell recordings from flies, allowing the retina's output neurons to more fully utilize their dynamic range. It is now generally accepted that retina, in part, performs stages of spatial and temporal normalization through center-surround receptive fields and on-off responses [Hosoya et al., 2005, Graham et al., 2006, Pitkow and Meister, 2012, Palmer et al., 2015]. Dong and Atick [1995] applied similar predictive coding ideas to the thalamus, proposing an additional stage of temporal normalization. Likewise, Friston's use of

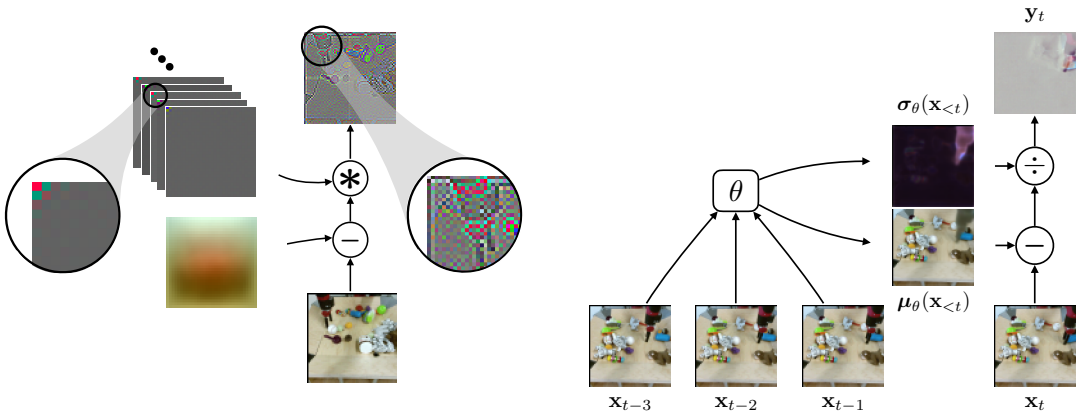

(a) Spatial (b) Temporal

Figure 5: **Spatiotemporal Predictive Coding**. **(a)** Spatial predictive coding models remove spatial dependencies. In the domain of natural images, one version of linear predictive coding is ZCA whitening, which yields center-surround filters (left). As a result, the whitened image contains highlighted edges (right). **(b)** Temporal predictive coding models remove temporal dependencies. In the domain of natural video, this tends to remove static backgrounds. Images are from the BAIR Robot Pushing dataset [Ebert et al., 2017].

generalized coordinates [Friston, 2008a], i.e. modeling multiple orders of temporal derivatives, can be approximated using finite temporal differences through repeated application of predictive coding. That is, $\frac{d\mathbf{x}}{dt} \approx \Delta\mathbf{x}_t \equiv \mathbf{x}_t - \mathbf{x}_{t-1}$. Thus, spatiotemporal predictive coding may be utilized at multiple stages of sensory processing to remove redundancy [Huang and Rao, 2011].

In neural circuits, spatiotemporal normalization often involves inhibitory interneurons [Carandini and Heeger, 2012], carrying out operations similar to those in Eq. 20 (though other mechanisms are also possible). For instance, retinal inhibitory interactions take place between photoreceptors, via horizontal cells, and between bipolar cells, via amacrine cells. This enables unpredicted motion to be computed, e.g. an object moving relative to the background, using inhibitory interactions from amacrine cells [Ölveczky et al., 2003, Baccus et al., 2008]. Similar inhibitory interactions are present in the lateral geniculate nucleus (LGN) in thalamus, with interneurons inhibiting relay cells originating from retina [Sherman and Guillery, 2002]. As mentioned above, this is thought to implement a form of temporal normalization [Dong and Atick, 1995], removing, at least, linear dependencies [Dan et al., 1996]. Inhibition via lateral inhibitory interactions is also a prominent feature of neocortex, with distinct classes of local interneurons playing a significant role in shaping the responses of principal pyramidal neurons [Isaacson and Scanziani, 2011]. While these distinct classes may serve separate computational roles, part of this purpose appears to be for spatiotemporal normalization [Carandini and Heeger, 2012]. Finally, while we have focused largely on early stages of sensory processing, inhibitory interneurons are also prevalent in other areas of neocortex, as well as in central pattern generator (CPG) circuits [Marder and Bucher, 2001], found in the spinal cord. These circuits are repsonsible for the rhythmic generation of movement, such as locomotion. Thus, just as inhibitory interactions *remove* spatiotemporal dependencies in early sensory areas, similar computational operations can *add* spatiotemporal dependencies in motor activation.

3.2 Hierarchical Predictive Coding

The other main form of predictive coding, mathematically formulated by Rao and Ballard [1999], Friston [2005], involves hierarchies of latent variables and, as such, has been postulated as a model

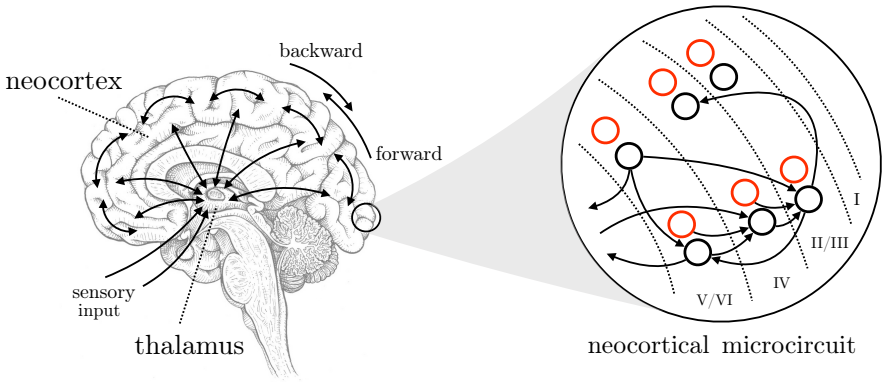

Figure 6: **Brain Anatomy & Cortical Circuitry**. **Left**: Sensory inputs enter first-order relays in thalamus from sensory organs. Thalamus forms reciprocal connections with neocortex. Neocortex consists of hierarchies of cortical areas, with both forward and backward connections. **Right**: Neocortex is composed of six layers (I–VI), with specific neuron classes and connections at each layer. The simplified schematic depicts two cortical columns. Black and red circles represent excitatory and inhibitory neurons respectively, with arrows denoting major connections. This basic circuit motif is repeated with slight variations throughout neocortex.

of hierarchical cortical processing. The neocortex (Figure 6) is a sheet-like structure involved in many aspects of sensory and motor processing. It is composed of six layers (I–VI), containing particular classes of neurons and connections. Across layers, neurons are arranged into columns, which are engaged in related computations [Mountcastle et al., 1955]. Columns interact locally via inhibitory interactions from interneurons while also forming processing hierarchies through longer-range excitatory interactions from pyramidal neurons. Particularly in earlier sensory areas, longer-range connections are generally grouped into forward (up the hierarchy) and backward (down the hierarchy) directions. Forward connections are traditionally thought to be driving (evoking neural activity) [Girard and Bullier, 1989, Girard et al., 1991]. Backward connections are traditionally thought to be modulatory, however, they have also been shown be driving [Covic and Sherman, 2011, De Pasquale and Sherman, 2011], which can be inverted through inhibition [Meyer et al., 2011]. These sets of connections, repeated with slight variations throughout neocortex, constitute a canonical *neocortical microcircuit* [Douglas et al., 1989], which could suggest a single processing algorithm [Hawkins and Blakeslee, 2004], capable of adapting to a variety of inputs [Sharma et al., 2000].

In formulating a theory of neocortex, Mumford [1992] proposed that thalamus acts as an 'active blackboard,' with the cortical hierarchy attempting to reconstruct or predict the thalamic input and activity in areas throughout the hierarchy. Backward (top-down) projections would convey predictions, while forward (bottom-up) projections would use prediction errors to update the estimates throughout the hierarchy. Through a dynamic process of activation, the entire system would settle to a consistent pattern of activity, minimizing prediction error. Over longer periods of time, the model parameters would be adjusted to yield improved predictions. In this way, cortex would use negative feedback, both in inference and learning, to use and construct a generative model of its inputs. This notion of generative state estimation dates back (at least) to Helmholtz [Von Helmholtz, 1867], and the notion of correcting predictions based on prediction errors is inline with concepts from cybernetics [Wiener, 1948, MacKay, 1956], which influenced techniques like Kalman filtering [Kalman, 1960], a ubiquitous Bayesian filtering algorithm.

A more complete mathematical formulation of this hierarchical predictive coding model, with many similarities to Kalman filtering (see Rao [1998]), was provided by Rao and Ballard [1999], with the generalization to variational inference provided by Friston [2005]. To illustrate this setup, consider

a simple model consisting of a single level of continuous latent variables, $\mathbf{z}$, modeling continuous data observations, $\mathbf{x}$. We will use Gaussian densities for each distribution and assume we have

$$p_\theta(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; f(\mathbf{W}\mathbf{z}), \text{diag}(\boldsymbol{\sigma}_\mathbf{x}^2)), \tag{21}$$

$$p_\theta(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_\mathbf{z}, \text{diag}(\boldsymbol{\sigma}_\mathbf{z}^2)), \tag{22}$$

where f is an element-wise function (e.g. logistic sigmoid, tanh, or the identity), $\mathbf{W}$ is a weight matrix, $\boldsymbol{\mu}_\mathbf{z}$ is the constant prior mean, and $\boldsymbol{\sigma}_\mathbf{x}^2$ and $\boldsymbol{\sigma}_\mathbf{z}^2$ are constant vectors of variances.

In the simplest approach to inference, we can find the maximum-a-posteriori (MAP) estimate, i.e. estimate the $\mathbf{z}^*$ which maximizes $p_\theta(\mathbf{z}|\mathbf{x})$. While we cannot tractably evaluate $p_\theta(\mathbf{z}|\mathbf{x})$ directly, we can use Bayes' rule to write

$$\begin{aligned}
\mathbf{z}^* &= \arg\max_\mathbf{z} p_\theta(\mathbf{z}|\mathbf{x}) \\
&= \arg\max_\mathbf{z} \frac{p_\theta(\mathbf{x}, \mathbf{z})}{p_\theta(\mathbf{x})} \\
&= \arg\max_\mathbf{z} p_\theta(\mathbf{x}, \mathbf{z}).
\end{aligned}$$

Thus, rather than evaluating the posterior distribution, $p_\theta(\mathbf{z}|\mathbf{x})$, we can perform this maximization using the joint distribution, $p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}|\mathbf{z})p_\theta(\mathbf{z})$, which we can tractably evaluate. We can also replace the optimization over the probability distribution with an optimization over the log probability, since $\log(\cdot)$ is a monotonically increasing function and will not affect the optimization. We then have

$$\begin{aligned}
\mathbf{z}^* &= \arg\max_\mathbf{z} \left[\log p_\theta(\mathbf{x}|\mathbf{z}) + \log p_\theta(\mathbf{z})\right]. \\
&= \arg\max_\mathbf{z} \left[\log \mathcal{N}(\mathbf{x}; f(\mathbf{W}\mathbf{z}), \text{diag}(\boldsymbol{\sigma}_\mathbf{x}^2)) + \log \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_\mathbf{z}, \text{diag}(\boldsymbol{\sigma}_\mathbf{z}^2))\right].
\end{aligned}$$

Each of the terms in this objective is a weighted squared error. For instance, the first term is the weighted squared error in reconstructing the data observation:

$$\log \mathcal{N}(\mathbf{x}; f(\mathbf{W}\mathbf{z}), \text{diag}(\boldsymbol{\sigma}_\mathbf{x}^2)) = \frac{-M}{2}\log(2\pi) - \frac{1}{2}\log\det\left(\text{diag}(\boldsymbol{\sigma}_\mathbf{x}^2)\right) - \frac{1}{2}\left|\left|\frac{\mathbf{x} - f(\mathbf{W}\mathbf{z})}{\boldsymbol{\sigma}_\mathbf{x}}\right|\right|_2^2,$$

where M is the dimensionality of $\mathbf{x}$ and $||\cdot||_2^2$ denotes the squared L2 norm. Plugging these terms into the objective and dropping terms that do not depend on $\mathbf{z}$ yields

$$\begin{aligned}
\mathbf{z}^* &= \arg\max_\mathbf{z} \left[\frac{-1}{2}\left|\left|\frac{\mathbf{x} - f(\mathbf{W}\mathbf{z})}{\boldsymbol{\sigma}_\mathbf{x}}\right|\right|_2^2 - \frac{1}{2}\left|\left|\frac{\mathbf{z} - \boldsymbol{\mu}_\mathbf{z}}{\boldsymbol{\sigma}_\mathbf{z}}\right|\right|_2^2\right], \\
&= \arg\max_\mathbf{z} \mathscr{L}(\mathbf{z}; \theta),
\end{aligned} \tag{23}$$

where we have defined the objective as $\mathscr{L}(\mathbf{z}; \theta)$. For purposes of illustration, let us assume that $f(\cdot)$ is the identity function, i.e. $f(\mathbf{W}\mathbf{z}) = \mathbf{W}\mathbf{z}$. We can then evaluate the gradient of $\mathscr{L}(\mathbf{z}; \theta)$ w.r.t. $\mathbf{z}$, yielding

$$\nabla_\mathbf{z} \mathscr{L}(\mathbf{z}; \theta) = \mathbf{W}^\intercal \left(\frac{\mathbf{x} - \mathbf{W}\mathbf{z}}{\boldsymbol{\sigma}_\mathbf{x}}\right) - \frac{\mathbf{z} - \boldsymbol{\mu}_\mathbf{z}}{\boldsymbol{\sigma}_\mathbf{z}}.$$

The transposed weight matrix, $\mathbf{W}^\intercal$, comes from differentiating $\mathbf{W}\mathbf{z}$, and translates the error in reconstruction into an update in $\mathbf{z}$. If we define the following terms as weighted errors:

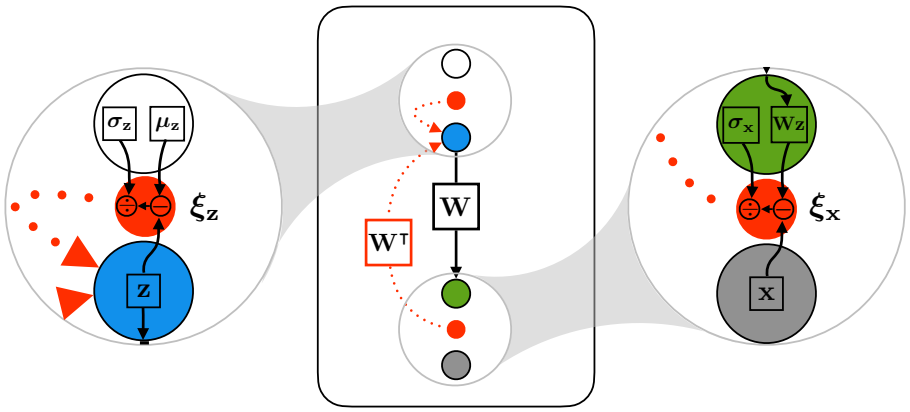

Figure 7: **Hierarchical Predictive Coding**. The diagram shows the basic computation graph for a Gaussian latent variable model with MAP inference. The insets show the weighted error calculation for the latent (left) and observed (right) variables.

$$\boldsymbol{\xi}_{\mathbf{x}} \equiv \frac{\mathbf{x} - \mathbf{W}\mathbf{z}}{\boldsymbol{\sigma}_{\mathbf{x}}}, \qquad \boldsymbol{\xi}_{\mathbf{z}} \equiv \frac{\mathbf{z} - \boldsymbol{\mu}_{\mathbf{z}}}{\boldsymbol{\sigma}_{\mathbf{z}}},$$

then we can re-write the gradient using these terms:

$$\nabla_{\mathbf{z}} \mathscr{L}(\mathbf{z}; \theta) = \mathbf{W}^\intercal \boldsymbol{\xi}_{\mathbf{x}} - \boldsymbol{\xi}_{\mathbf{z}}. \tag{24}$$

Thus, if we want to perform inference using gradient-based optimization, e.g. $\mathbf{z} \leftarrow \mathbf{z} + \alpha \nabla_{\mathbf{z}} \mathscr{L}(\mathbf{z}; \theta)$, we need 1) the weighted errors, $\boldsymbol{\xi}_{\mathbf{x}}$ and $\boldsymbol{\xi}_{\mathbf{z}}$, and 2) the transposed weights, $\mathbf{W}^\intercal$, or more generally, the Jacobian of the conditional likelihood mean. This overall scheme is depicted in Figure 7.

To learn the weight parameters, we can differentiate $\mathscr{L}(\mathbf{z}; \theta)$ (Eq. 23) w.r.t. $\mathbf{W}$:

$$\begin{aligned} \nabla_{\mathbf{W}} \mathscr{L}(\mathbf{z}; \theta) &= \left(\frac{\mathbf{x} - \mathbf{W}\mathbf{z}}{\boldsymbol{\sigma}_{\mathbf{x}}} \right) \mathbf{z}^\intercal \\ &= \boldsymbol{\xi}_{\mathbf{x}} \mathbf{z}^\intercal. \end{aligned}$$

This gradient is the product of a local error term, $\boldsymbol{\xi}_{\mathbf{x}}$, and the latent variable, $\mathbf{z}$, possibly suggesting a biologically-plausible learning rule [Whittington and Bogacz, 2017].

Predictive coding identifies the conditional likelihood (Eq. 21) with backward (or top-down) cortical projections, whereas inference updating (Eq. 24) is identified with forward (or bottom-up) cortical projections [Friston, 2005]. Such connections are thought to be mediated by pyramidal neurons. Scaling this model up in size and structure, each cortical column could contain the necessary computational elements involved in predicting and estimating a latent variable. Interneurons within columns could be involved in error calculation ($\boldsymbol{\xi}_{\mathbf{x}}$ and $\boldsymbol{\xi}_{\mathbf{z}}$). Although we only discussed diagonal covariance matrices ($\boldsymbol{\sigma}_{\mathbf{x}}^2$ and $\boldsymbol{\sigma}_{\mathbf{z}}^2$), interneurons involved in local lateral inhibition could parameterize (the inverse of) full covariance matrices, i.e. $\boldsymbol{\Sigma}_{\mathbf{x}}$ and $\boldsymbol{\Sigma}_{\mathbf{z}}$. This is an instance of *spatial* predictive coding, which we discussed in the previous section. Note that these factors, which weight $\boldsymbol{\xi}_{\mathbf{x}}$ and $\boldsymbol{\xi}_{\mathbf{z}}$, effectively modulate the *gain* of each error term, possibly providing a form of "attention" [Feldman and Friston, 2010]. Possible neural correspondences are summarized in Table 1.

Here, we have discussed a simplified model of hierarchical predictive coding, with a single latent level and no dynamics. However, a full theory of hierarchical predictive coding would include these

Table 1: Proposed Neural Correspondences of Hierarchical Predictive Coding.

Neuroscience	**Predictive Coding**
Top-Down Cortical Projections	Generative Model Conditional Mapping
Bottom-Up Cortical Projections	Inference Updating
Lateral Inhibition	Covariance Matrices
(Pyramidal) Neuron Activity	Latent Variable Estimates & Errors
Cortical Column	Corresponding Estimate & Error

additional aspects and others. Karl Friston has explored various design choices throughout multiple papers [Friston et al., 2007, Friston, 2008a,b], yet the core aspects of probabilistic generative modeling and variational inference remain largely the same. Further elaborating and comparing these design choices will be essential for empirically validating the details of hierarchical predictive coding.

3.3 Empirical Support

Empirically validating predictive coding in neural circuits is an active area of research. This remains challenging, as it is difficult to disentangle the theory itself from the wide array of possible design choices, e.g. distributions, parameterizations, etc. [Gershman, 2019]. Nevertheless, many of the core aspects of predictive coding do appear to have *some* empirical support. We briefly outline some of these studies here, but we refer the reader to the multiple review papers on the topic [Huang and Rao, 2011, Bastos et al., 2012, Clark, 2013, Keller and Mrsic-Flogel, 2018, Walsh et al., 2020].

Spatiotemporal As discussed above, various works have investigated spatiotemporal predictive coding in early sensory areas, primarily retina [Srinivasan et al., 1982, Atick and Redlich, 1992]. This typically involves fitting retinal ganglion cell responses to a spatial whitening (or decorrelation) process [Graham et al., 2006, Pitkow and Meister, 2012], which is dynamically adjusted based on lighting conditions [Hosoya et al., 2005]. Similar analyses suggest that retina employs temporal predictive coding as well [Srinivasan et al., 1982, Palmer et al., 2015]. While the exact mathematical details of these neural computations have not been fully characterized, the corresponding models contain stages of linear decorrelating filters (e.g. center-surround) followed by non-linearities. Importantly, non-linearities have been shown to be an essential aspect in explaining retinal ganglion cell responses [Pitkow and Meister, 2012], possibly inducing an added degree of sparsity [Graham et al., 2006]. As previously noted, similar spatiotemporal predictive coding computations may be found in thalamus [Dong and Atick, 1995] and cortex. While Dan et al. [1996] provide some supporting evidence, such investigations are complicated by the presence of backward and modulatory interactions.

Hierarchical Early work toward empirically validating hierarchical predictive coding came from explaining extra-classical receptive field effects [Rao and Ballard, 1999, Rao and Sejnowski, 2002], whereby top-down processing in cortex can alter classical visual receptive fields, suggesting that top-down influences play an important role in sensory processing [Gilbert and Sigman, 2007]. Likewise, temporal influences have been demonstrated in the form of repetition suppression [Summerfield et al., 2006], in which cortical activity diminishes in response to repeated, i.e. predictable, stimuli. This effect may reflect the suppression of errors through improved predictions. Predictive coding has also been postulated as an explanation of biphasic responses in LGN [Jehee and Ballard, 2009], in which reversing the visual input with an anti-correlated image results in a large neural response, presumably due to prediction errors. Predictive signals have been documented in auditory [Wacongne et al., 2011] and visual [Meyer and Olson, 2011] processing. Activity seemingly corresponding to prediction errors has also been observed in a variety of areas and contexts, including visual flow

in primary visual cortex in mice [Keller et al., 2012, Zmarz and Keller, 2016], auditory cortex in monkeys [Eliades and Wang, 2008] and rodents [Parras et al., 2017], and visual cortex in humans [Murray et al., 2002, Alink et al., 2010, Egner et al., 2010]. While further studies are needed, it appears that sensory cortex is engaged in some form of hierarchical and temporal prediction, with prediction error signals playing a key role in driving the perceptual process.

The empirical evidence for spatiotemporal and hierarchical predictive coding is suggestive, but given the complexity of neural systems, some aspects of the theory are undoubtedly incorrect, incomplete, or under-specified. In particular, the complexity of biological systems makes it is difficult to isolate and assess detailed aspects of predictive coding. For instance, it appears that cortex calculates some form of prediction error, but without access to fine-grained recordings of all relevant signals, e.g. dendritic currents, neuromodulators, etc., it is difficult to determine the exact computational form of the circuit. Thus, while general aspects of predictive coding appear supported, we are unable to probe into the details of such models, making predictive coding a largely *normative* theory. The purpose of this paper is to establish connections between predictive coding and machine learning. Ideally, by building larger-scale models and training them on similar sensory data, we can form more fine-grained empirical predictions for biological neural systems. Building off of the example of Rao and Ballard [1999], Lotter et al. [2018] provided another step in this direction, comparing the responses of neural systems and their hierarchical predictive coding model. In the current thesis, we have attempted to help further build the foundation for this collaborative effort.

4 Variational Autoencoders

Variational autoencoders (VAEs) [Kingma and Welling, 2014, Rezende et al., 2014] are latent variable models parameterized by deep networks. As in hierarchical predictive coding, these models typically contain Gaussian latent variables and are trained using variational inference. However, rather than performing inference optimization directly, VAEs *amortize* inference [Gershman and Goodman, 2014].

4.1 Amortized Variational Inference

Amortization refers to spreading out costs. In amortized inference, these "costs" are the computational costs of performing inference optimization. Thus, with $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_q, \text{diag}(\boldsymbol{\sigma}_q^2))$ and $\boldsymbol{\lambda} \equiv [\boldsymbol{\mu}_q, \boldsymbol{\sigma}_q]$, rather than separately optimizing $\boldsymbol{\lambda}$ for each data example, we amortize this optimization cost using a learned optimizer, i.e. an *inference model*. By using this meta-optimization procedure, we can perform inference optimization far more efficiently. The concept of inference models is deeply embedded with deep latent variable models, popularized by the Helmholtz Machine [Dayan et al., 1995], which was formulated as an *autoencoder* [Ballard, 1987]. Formally, in such setups, the inference model is a direct mapping from $\mathbf{x}$ to $\boldsymbol{\lambda}$:

$$\boldsymbol{\lambda} \leftarrow f_\phi(\mathbf{x}), \tag{25}$$

where f_ϕ is a model (deep network) with parameters ϕ. Conventionally, we denote the approximate posterior as $q_\phi(\mathbf{z}|\mathbf{x})$ to denote the parameterization by ϕ. Now, rather than optimizing $\boldsymbol{\lambda}$ using gradient-based techniques, we periodically update ϕ using $\nabla_\phi \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \boldsymbol{\lambda}} \frac{\partial \boldsymbol{\lambda}}{\partial \phi}$, thereby letting f_ϕ learn to optimize $\boldsymbol{\lambda}$. This procedure is incredibly simple, as we only need to tune the learning rate for ϕ, and efficient, as we have an estimate of $\boldsymbol{\lambda}$ after only one forward pass through f_ϕ. Amortization is also widely applicable: if we can estimate $\nabla_{\boldsymbol{\lambda}} \mathcal{L}$ using stochastic gradient estimation (see above), we can continue differentiating through the chain $\phi \rightarrow \boldsymbol{\lambda} \rightarrow \mathbf{z} \rightarrow \mathcal{L}$.

To differentiate through $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x})$, we can use the pathwise derivative estimator, sometimes referred to as the reparameterization estimator [Kingma and Welling, 2014]. This is accomplished by reparameterizing $\mathbf{z}$ in terms of an auxiliary random variable. The most common example is reparameterizing $\mathbf{z} \sim \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_q, \text{diag}(\boldsymbol{\sigma}_q^2))$ as $\mathbf{z} = \boldsymbol{\mu}_q + \boldsymbol{\epsilon} \odot \boldsymbol{\sigma}_q$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\epsilon}; \mathbf{0}, \mathbf{I})$ and $\odot$ denotes

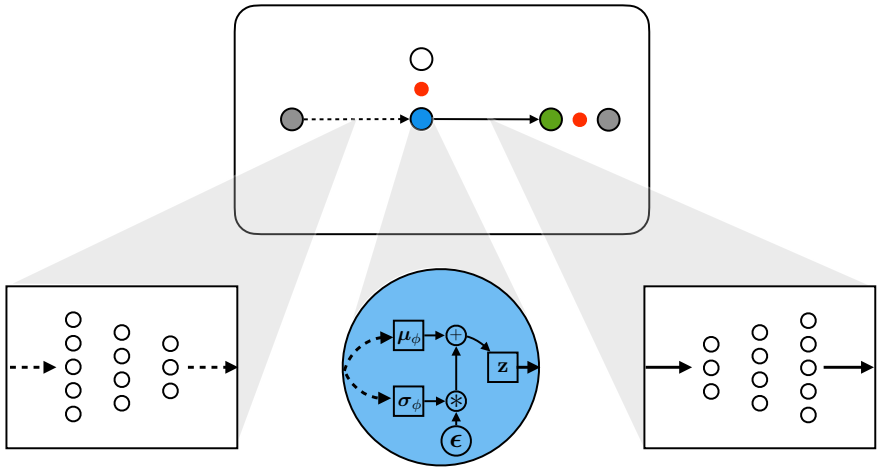

Figure 8: **Variational Autoencoder (VAE)**. VAEs combine direct amortization (Eq. 25) and the pathwise derivative estimator with Gaussian approximate posteriors to train deep latent variable models. In the model diagram (top), the amortized inference model (dashed arrow) acts as an *encoder*, with the conditional likelihood (solid arrow) acting as a *decoder*. Each are parameterized by deep networks.

element-wise multiplication. We can then estimate $\nabla_{\boldsymbol{\mu}_q}\mathcal{L}$ and $\nabla_{\boldsymbol{\sigma}_q}\mathcal{L}$, allowing us to calculate the inference model gradients, $\nabla_\phi\mathcal{L}$.

When direct amortization is combined with the pathwise derivative estimator in deep latent variable models, the resulting setup is referred to as a *variational autoencoder* [Kingma and Welling, 2014, Rezende et al., 2014]. In this autoencoder interpretation, $q_\phi(\mathbf{z}|\mathbf{x})$ is an *encoder*, $\mathbf{z}$ is the latent *code*, and $p_\theta(\mathbf{x}|\mathbf{z})$ is a *decoder*. A computation graph is shown in Figure 8. This direct encoding scheme seems intuitively obvious: in the same way that $p_\theta(\mathbf{x}|\mathbf{z})$ directly maps $\mathbf{z}$ to a distribution over $\mathbf{x}$, $q_\phi(\mathbf{z}|\mathbf{x})$ directly maps $\mathbf{x}$ to a distribution over $\mathbf{z}$. Indeed, with perfect knowledge of $p_\theta(\mathbf{x}, \mathbf{z})$, f_ϕ could act as a lookup table, precisely mapping each $\mathbf{x}$ to the corresponding optimal $\boldsymbol{\lambda}$. However, in practice, direct amortization of this form tends to result in suboptimal estimates of $\boldsymbol{\lambda}$ [Cremer et al., 2018], motivating the need for more powerful amortized inference techniques.

4.2 Iterative Amortized Inference

One method for improving direct amortization involves incorporating iterative updates [Hjelm et al., 2016, Krishnan et al., 2018, Kim et al., 2018, Marino et al., 2018b], replacing a one-step inference procedure with a more accurate multi-step procedure. Iterative amortized inference [Marino et al., 2018b] maintains an inference model, but uses it to perform iterative updates on the approximate posterior estimate. Following the notation from the previous section, the basic form of an iterative amortized inference model is given as:

$$\boldsymbol{\lambda} \leftarrow f_\phi(\boldsymbol{\lambda}, \nabla_{\boldsymbol{\lambda}}\mathcal{L}). \tag{26}$$

Iterative inference models take in the current estimate of $\boldsymbol{\lambda}$, as well as the gradient, $\nabla_{\boldsymbol{\lambda}}\mathcal{L}$, and output an updated estimate of $\boldsymbol{\lambda}$. As with direct inference models, iterative inference model parameters are updated using estimates of $\nabla_\phi\mathcal{L}$. Note that Eq. 26 generalizes stochastic gradient-based optimization. For instance, a special case is $\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} + \alpha\nabla_{\boldsymbol{\lambda}}\mathcal{L}$, where α is a step-size parameter, however, Eq. 26 also contains more general non-linear updates [Andrychowicz et al., 2016].

As we saw in Section 3.2, in latent Gaussian models, $\nabla_{\boldsymbol{\lambda}}$ is defined by the weighted errors, $\boldsymbol{\xi}_{\mathbf{x}}$ and $\boldsymbol{\xi}_{\mathbf{z}}$, and the Jacobian of the conditional likelihood, $\mathbf{J}$ (this was $\mathbf{W}$ in the linear model). Because the conditional likelihood does not change across data examples, in latent Gaussian models, we can consider iterative amortized inference models of the special form:

$$\boldsymbol{\lambda} \leftarrow f_{\phi}(\boldsymbol{\lambda}, \boldsymbol{\xi}_{\mathbf{x}}, \boldsymbol{\xi}_{\mathbf{z}}). \tag{27}$$

This is a learned, non-linear mapping from errors to updated estimates of the approximate inference distribution, i.e. learned negative feedback. The distinction between direct and iterative amortization is shown in Figures 10b & 10c. Iterative amortization can be readily extended to sequential models [Marino et al., 2018a], resulting in a general predict-update inference scheme based on prediction errors, highly reminiscent of Kalman filtering [Kalman, 1960]. Marino et al. [2020b] also applied iterative amortized inference to perform policy optimization in reinforcement learning.

4.3 Normalizing Flows

Another development within VAEs is the use of normalizing flows [Rezende and Mohamed, 2015]. As discussed in Section 2.2.1, flow-based distributions are *invertible* latent variable models, enabling exact log-probability calculation. That is, through a sequence of invertible transforms, a flow-based distribution attempts to map a variable with independent (or less dependent) dimensions, typically from a standard Normal, to a sample from a more structured distribution. Conversely, using the inverse transforms, we can attempt to remove these structured dependencies, i.e. *normalization*. While such models can operate as standalone generative models [Dinh et al., 2015, 2017, Papamakarios et al., 2017, Kingma and Dhariwal, 2018], they can also define the distributions within VAEs.

One can use normalizing flows to define the approximate posterior [Rezende and Mohamed, 2015, Kingma et al., 2016], prior [Huang et al., 2017], and conditional likelihood [Agrawal and Dukkipati, 2016] in VAEs. In each case, a deep network defines the parameters (e.g., mean and variance) of a *base distribution* over the less structured, normalized variable. Separate deep networks then parameterize the invertible transforms, which add dependencies between dimensions of the less structured variable, or, conversely, remove structure from the more structured, unnormalized variable.

Example: As a generic example, consider a less structured variable, $\mathbf{u}$, defined by the distribution $p_{\theta}(\mathbf{u}|\cdot) = \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_{\theta}(\cdot), \mathrm{diag}(\boldsymbol{\sigma}^2_{\theta}(\cdot)))$, where $\boldsymbol{\mu}_{\theta}$ and $\boldsymbol{\sigma}_{\theta}$ are output by deep networks, with $\cdot$ denoting conditioning (input) variables. We then consider an *affine* invertible transform [Dinh et al., 2017], defined by a shift vector, $\boldsymbol{\alpha}_{\theta}$, and a scale matrix, $\boldsymbol{B}_{\theta}$, each of which may be functions of $\mathbf{u}$ (again, parameterized by deep networks). This affine transform defines a more structured variable, $\mathbf{v}$:

$$\mathbf{v} = \boldsymbol{\alpha}_{\theta} + \boldsymbol{B}_{\theta}\mathbf{u}, \tag{28}$$

which can now contain affine dependencies between dimensions. Note that for Eq. 28 to be invertible, we require that $\boldsymbol{B}_{\theta}$ itself is invertible, i.e., non-zero determinant. Thus, $\boldsymbol{B}_{\theta}$ is a square matrix, and $\mathbf{u}$ and $\mathbf{v}$ are the same dimensionality. If, instead, we are given an observation of $\mathbf{v}$ and want to calculate its log-probability, we can apply the normalizing inverse transform to get $\mathbf{u}$:

$$\mathbf{u} = \boldsymbol{B}_{\theta}^{-1}(\mathbf{v} - \boldsymbol{\alpha}_{\theta}), \tag{29}$$

then use the change of variables formula (Eq. 6). This allows us to convert the log-probability calculation from the more structured (complex) space of $\mathbf{v}$ to the less structured (simpler) space of $\mathbf{u}$. However, we need to account for the local scaling of space induced by $\boldsymbol{B}_{\theta}$ through its determinant. Note that the multivariate Gaussian density (Eq. 9) is a special case of this transform, taking a standard Gaussian variable, e.g., $\mathbf{u} \sim \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{I})$, and adding linear dependencies to yield a multivariate Gaussian variable, $\mathbf{v} \sim \mathcal{N}(\mathbf{v}; \boldsymbol{\alpha}_{\theta}, \boldsymbol{B}_{\theta}^{\intercal}\boldsymbol{B}_{\theta})$. In this case, where $\boldsymbol{\alpha}_{\theta}$ and $\boldsymbol{B}_{\theta}$ are constant, applying the inverse transform removes linear dependencies between dimensions in $\mathbf{v}$.

Normalizing flows provide a general method for improving the flexibility of the distributions within VAEs. This allows us to parameterize more flexible approximate posteriors for variational inference [Rezende and Mohamed, 2015, Kingma et al., 2016, Tomczak and Welling, 2016, van den Berg et al., 2018]. However, we can also apply normalizing flows on the model's distributions, i.e. the prior [Huang et al., 2017] and conditional likelihood [Agrawal and Dukkipati, 2016], allowing us to model variable estimates and observations in a normalized space. While many prior works apply normalizing flows across spatial dimensions, we can also add and remove temporal dependencies [van den Oord et al., 2018]. In particular, building off of the Gaussian special case discussed above, Gaussian autoregressive models (e.g. Figure 3) are *equivalent* to affine autoregressive normalizing flows [Kingma et al., 2016]. In the linear case, this is also equivalent to Cholesky whitening [Pourahmadi, 2011].

Using this idea, Marino et al. [2020a] parameterized the conditional likelihood in sequential latent variable models with autoregressive normalizing flows across time, i.e. *sequential autoregressive flows*. At the observation level, the autoregressive flow acts as a moving frame of reference, as shown in Figure 5b, removing simple temporal dependencies and allowing the latent variable model to focus on modeling more complex spatiotemporal structure. In practice, this involves extracting a low-level temporal prediction error (as in Eq. 20), then modeling this error using the latent variable model. This technique has been successfully applied to video compression [Yang et al., 2020], as well as parameterizing temporal dependencies in reinforcement learning policies [Guerra and Marino, 2020], serving as a form of feedforward control.

5 Connections

We now connect two recent machine learning techniques, iterative amortized inference [Marino et al., 2018b] and sequential autoregressive flows [Marino et al., 2020a], to their origins in predictive coding.

5.1 Iterative Amortization

Iterative amortized inference [Marino et al., 2018b] (Figures 9a & 10c) was inspired by the inference scheme proposed by Rao and Ballard [1999] and Friston [2005]. In these early formulations of hierarchical predictive coding, approximate inference is performed using gradient-based optimization of a point estimate of the latent variables. These works made it clear that prediction (or reconstruction) errors drive both inference and learning optimization, and this procedure can be readily extended to sequential settings [Friston, 2008b]. However, such procedures typically assume that the inference gradients, supplied by forward connections, can be easily calculated, but the weights of these forward connections are, in fact, the Jacobian of the backward connections [Rao and Ballard, 1999] (Section 3). This is an example of the *weight transport* problem [Grossberg, 1987], i.e. the weights of one set of connections (forward) depends on the weights from another set of connections (backward). This is generally regarded as not being biologically-plausible.

Amortization [Dayan et al., 1995] provides a simple solution to this problem: learn to perform inference optimization. That is, rather than transporting the generative weights to the inference connections, amortization learns a separate set of inference weights, potentially using similar local learning rules [Bengio, 2014, Lee et al., 2015]. Thus, despite criticism from Friston [2018], amortization may offer a more biologically-plausible account of inference. Further, as demonstrated in Marino et al. [2018b,a, 2020b], by using non-linear functions, amortization is capable of automatically adjusting update step sizes, yielding accurate estimates with exceedingly few inference iterations. These substantial benefits in computational efficiency provide another argument for amortization over the gradient-based schemes often employed in predictive coding.

Iterative amortization is an example of the more general concept of negative feedback. As noted earlier, negative feedback was the core concept of cybernetics, which went on to inspire predictive

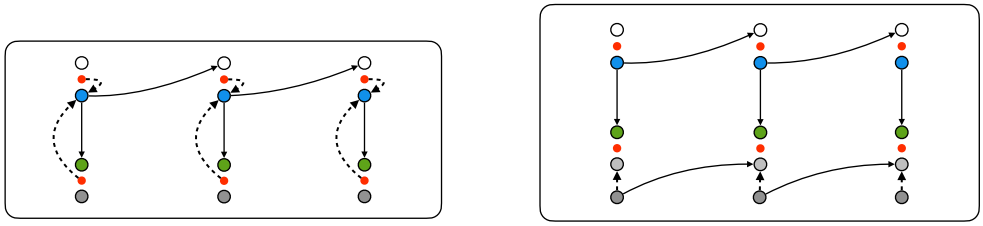

(a) Iterative Amortized Inference (b) Sequential Autoregressive Flows

Figure 9: **Connections**. Computation graphs illustrating **(a)** iterative amortized inference (applied across time steps [Marino et al., 2018a]) and **(b)** sequential autoregressive flows. Iterative amortized inference learns to update approximate posterior estimates using prediction errors or gradients. Sequential autoregressive flows learn to perform temporal normalization, simplifying dynamics modeling. Here, dashed arrows between gray circles denote the "inference" (normalizing) direction of the normalizing flow. Iterative amortization and sequential autoregressive flows were inspired by hierarchical and spatiotemporal predictive coding, respectively.

coding. While hierarchical predictive coding has largely focused on perceptual inference in cortex, the principles of negative feedback appear to apply more broadly to neural systems. Indeed, even at the outset of cybernetics, it was clear that cerebellum plays a central role in negative feedback control [Wiener, 1948]. From more recent studies of cerebellum and other cerebellum-like structures [Ito, 1998, Bell, 2001, Kennedy et al., 2014], we are beginning to understand how such circuits correct sensorimotor prediction errors. One prominent example is given by the Purkinje cells of the cerebellum, which appear to take in error signals as inputs and output motor corrections. This follows the general paradigm of iterative amortization, mapping errors to updates. Casting these neural circuits in terms of amortization, i.e. learned negative feedback, may provide insights into how such error-correcting mechanisms are learned from experience.

5.2 Sequential Autoregressive Flows

The technique of sequential autoregressive flows [Marino et al., 2020a] (Figure 9b) was, in part, inspired by the temporal normalization schemes from Srinivasan et al. [1982] and Dong and Atick [1995], which are thought to occur in retina and first-order relays of thalamus. Unlike these earlier works, which were limited to linear functions of previous inputs, sequential autoregressive flows can utilize non-linear functions to parameterize the normalizing affine transform. Likewise, because these feedforward transforms are learned using the prediction errors on the normalized variables, they can adapt to meet the demands or limitations of higher-level models.

A related technique is that of generalized coordinates [Friston, 2008a], decomposing a sequence into its temporal derivatives. Friston has suggested that this may be a general modeling technique employed by neural circuits. As we have seen, a simplified version of sequential autoregressive flows, using the previous variable as the affine shift, extracts an approximation of temporal derivatives. Thus, given a short enough time step, $\Delta t = t_1 - t_0$, sequential autoregressive flows provide a technique for automatically learning an approximation of generalized coordinates.

The *same* technique can be applied to control [Guerra and Marino, 2020], serving as a low-level dynamical policy, i.e. a dynamical motor basis for control. While early sensory processing and motor processing are often considered separately, we see that spatiotemporal dependencies are central to both areas. Normalization (and its inverse) plays a singular role in both cases, simplifying estimation for upstream models. If normalization is truly a canonical neural computation [Carandini and Heeger, 2012], then similar temporal normalization operations may parameterize dynamics estimation throughout cortex, operating in conjunction with spatial normalization. This was demonstrated

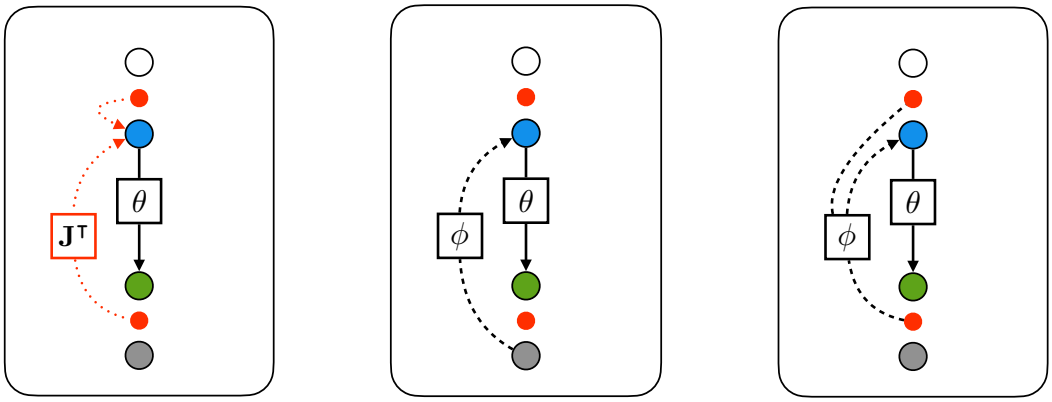

(a) Predictive Coding (b) VAE (Direct) (c) VAE (Iterative)

Figure 10: **Hierarchical Predictive Coding & VAEs**. Computation diagrams for **(a)** hierarchical predictive coding [Rao and Ballard, 1999], **(b)** VAE with direct amortized inference [Kingma and Welling, 2014, Rezende et al., 2014], and **(c)** VAE with iterative amortized inference [Marino et al., 2018b]. $\mathbf{J}^\intercal$ denotes the transposed Jacobian matrix of the generative model's conditional likelihood. Red dotted lines denote gradients, and black dashed lines denote amortized inference. Hierarchical predictive coding and VAEs are highly similar in both their model formulation and inference approach.

in Marino et al. [2020a], where a modified VideoFlow model [Kumar et al., 2020] utilized *spatial* normalization within time steps and *temporal* normalization across time steps, demonstrating the utility of these complementary procedures.

Marino et al. [2020a] exclusively investigated normalizing flows with affine transforms [Dinh et al., 2017], due to their simplicity and their similarity to proposed neural computations. Given their connection with Gaussian densities and Gaussianization [Chen and Gopinath, 2001] (*normal*ization is, ultimately, the process of converting a data distribution into a standard *Normal* (Gaussian) density), transforms of this form are prominent in the statistics literature [Friedman, 1987, Kessy et al., 2018]. However, the change of variables formula readily applies to all invertible transforms, including non-affine transforms. This more general perspective, afforded by normalizing flows, offers a method toward improving spatiotemporal normalization in current models. Similarly, it suggests that neural circuits may implement multiple computational forms of normalization transforms, ranging from simple, affine or constant transforms to more complex, non-affine transforms [Durkan et al., 2019]. These may rely on computational mechanisms within individual neurons as well as spatiotemporal interactions between neurons.

6 Correspondences

Having drawn connections between recent machine learning approaches and their inspirations in predictive coding, we can now traverse this bridge from machine learning, through predictive coding, to neuroscience. In this section, we identify correspondences implied by this bridge. In particular, we explore the consequences of two implied correspondences: **1**) pyramidal neurons and deep networks and **2**) lateral inhibition and normalizing flows. These correspondences should be interpreted at a *functional* level, potentially shedding new light on the computational and learning mechanisms employed in biological neural circuits. In some aspects, these correspondences offer a departure from the current paradigm, providing an alternative approach toward connecting these areas.

Hierarchical predictive coding and deep latent variable models, particularly variational autoencoders (VAEs) [Kingma and Welling, 2014, Rezende et al., 2014], are highly related in both their

model formulations and inference approaches (Figure 10). Specifically,

- **Model Formulation**: Both areas focus on hierarchical latent Gaussian models with non-linear dependencies between latent levels, as well as dependencies within levels via covariance matrices (predictive coding) or normalizing flows (VAEs). Note that a covariance matrix is computationally equivalent to an affine normalizing flow with linear dependencies [Kingma et al., 2016].
- **Inference**: Both areas use variational inference, often with Gaussian approximate posteriors. While predictive coding employs gradient-based optimization and VAEs employ amortized optimization, these are just different design choices in solving the same inference optimization problem.

While previous works within predictive coding and deep latent variable models have explored distinct design choices, e.g. in parameterizing dynamics, the two areas share a common mathematical foundation, inherited from cybernetics and descendant areas.

With this connection explicitly established, we can now identify the biological correspondences implied by the connection from VAEs to predictive coding to neuroscience. Looking at Table 1, we see that top-down and bottom-up cortical projections, each mediated by pyramidal neurons, respectively parameterize the generative model and inference updates. Mapping this onto VAEs implies that **deep (artificial) neural networks are in correspondence with pyramidal neuron dendrites** (Figure 11). This analogy is not perfect, as each output dimension of a deep network shares parameters with the other outputs through previous layers. In this regard, the analogy to pyramidal dendrites in cortex would specifically imply a separate deep network *per variable* in a VAE. Or, conversely, a deep network corresponds to a collection of pyramidal dendrites operating in parallel. **Lateral interneurons, which parameterize the inverse covariance matrices at each latent level in predictive coding, map onto normalizing flows**. As mentioned above, normalizing flows are a non-linear generalization of linear covariance matrices, suggesting the possibility of non-linear normalization computations in cortex and elsewhere. We note that normalizing flows are also implemented using deep networks, however, the effect that they have on the core probabilistic computation tends to be restricted and simple (e.g. affine) to ensure tractability. These correspondences are obviously quite coarse-grained, and many details are left to be filled-in. However, they may provide a useful starting point for shifting the current analogies between machine learning and neuroscience. Below, we explore some of the consequences of these correspondences.

6.1 Pyramidal Neurons & Deep Networks

Non-linear Dendritic Computation Placing deep networks in correspondence with pyramidal dendrites departs from the traditional one-to-one correspondence of biological and artificial neurons [McCulloch and Pitts, 1943]. This suggests that (some) individual biological neurons may be better computationally described as non-linear functions. Evidence from neuroscience supports this analogy. Early work in simulations proposed that individual pyramidal neurons, through dendritic processing, could operate as multi-layer artificial networks [Zador et al., 1992, Mel, 1992]. This was later supported by empirical findings that hippocampal pyramidal dendrites act as computational 'subunits,' yielding the equivalent of a two-layer artificial network [Poirazi et al., 2003, Polsky et al., 2004]. More recently, Gidon et al. [2020] demonstrated that individual L2/3 pyramidal neurons are capable of computing the XOR operation, known for requiring non-linear processing [Minsky and Papert, 1969]. Likewise, Jones and Kording [2020] recently demonstrated a non-linear dendritic model capable of performing classification tasks, and Beniaguev et al. [2020] demonstrated that multi-layer spatiotemporal convolutional networks are required to accurately model L5 pyramidal neurons. This potential functional correspondence between deep networks and pyramidal dendrites posits a substantial role for dendritic computation [London and Häusser, 2005], moving beyond

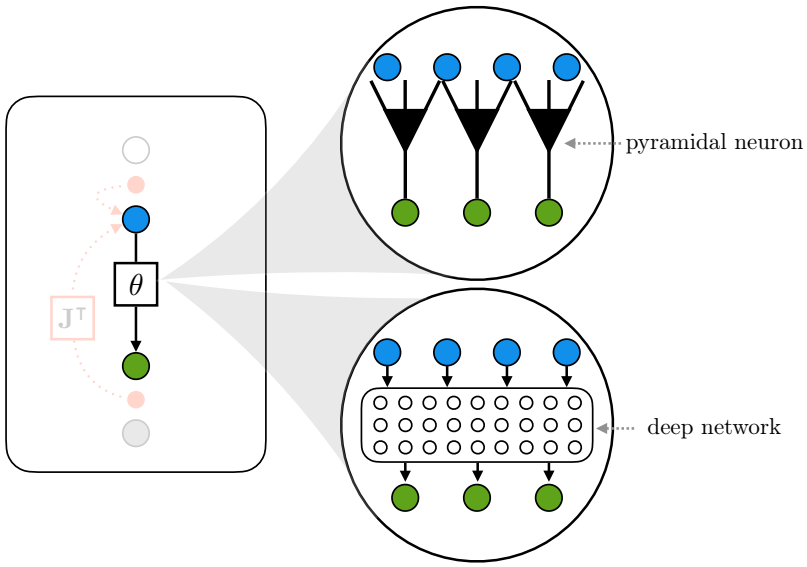

Figure 11: **Pyramidal Neurons & Deep Networks**. Connecting deep latent variable models with predictive coding places deep networks (bottom) in correspondence with the dendrites of pyramidal neurons (top). This is in contrast with conventional one-to-one analogies of biological and artificial neurons, suggesting a larger role for non-linear dendritic computation and alternative correspondences for backpropagation.

the overly simplistic comparison of biological and artificial neurons. Further, rather than assuming every neuron is equivalent, i.e. linear summation with non-linearity, separate classes of neurons would represent distinct function classes, likely derived from their morphology. This places a greater emphasis on understanding the particular intricacies of neural circuits, rather than assuming a uniform network of identical computational elements.

Amortization It is feasible that similar hardware and mechanisms underlying generative predictions could also perform inference updating, i.e. amortization. This is the insight of deep latent variable models: deep networks can parameterize conditional probabilities in *both* directions. The computational components are identical, with different inputs and output targets. Building off of the correspondence of pyramidal dendrites and deep networks, in cortex, we see particular classes of pyramidal neurons with separate apical and basal dendrites. These segregated dendritic compartments selectively take inputs from top-down and bottom-up pathways respectively [Bekkers, 2011, Guergiuev et al., 2016, Richards], thought to perform separate computations. These pyramidal neurons could implement a form of iterative amortized inference model [Marino et al., 2018b], separately processing top-down and bottom-up error signals to update inference estimates (Figure 12). This agrees with the conjecture from predictive coding that separate neurons in the forward pathway perform inference updating. Amortization also resolves the weight-transport issue from predictive coding, as separate inference weights are *learned*. While some empirical evidence appears to support amortization [Yildirim et al., 2015, Dasgupta et al., 2018], we note that Friston [2018] remains skeptical of its biological-plausibility. Further experiments, particularly at the cortical circuit level, are needed to resolve this question. For instance, amortization would require some form of stochastic gradient estimation, e.g. reparameterization gradients [Kingma and Welling, 2014]. Although all of the necessary error signals are local to the amortized optimizer, the details remain unclear.

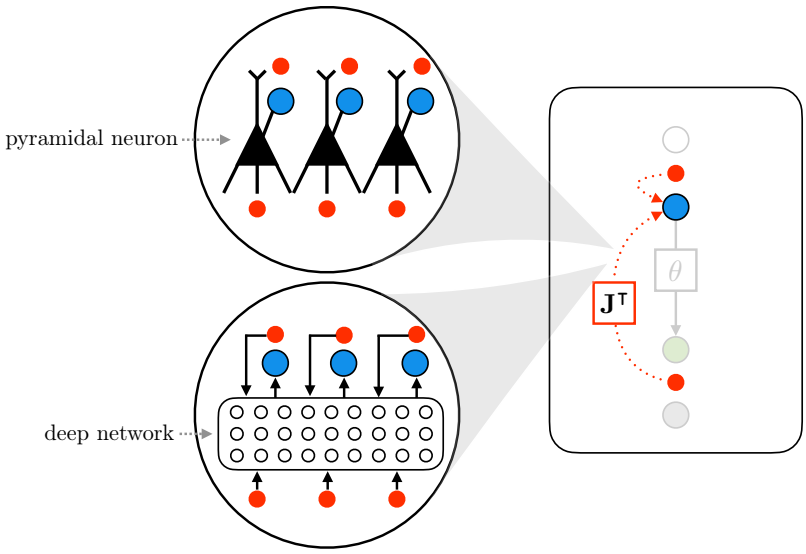

Figure 12: **Pyramidal Neurons & Amortization**. In predictive coding, inference updating is implemented using forward pyramidal neurons in cortex, taking prediction errors as input. In deep latent variable models, iterative amortized inference plays a similar role, continuing the analogy of pyramidal neurons and deep networks. Interestingly, this suggests a separation of processing in apical (upper) and basal (lower) dendrites, incorporating errors from the current and lower level.

Neural Oscillations Oscillations are a common feature of neural circuits, giving rise to various frequency bands in the local field potential (LFP). These frequencies are thought to arise from the synchronous activity of populations of neurons, such as recurrent activity in hippocampus resulting in the theta frequency band (4–10 Hz). Bastos et al. [2015] identify distinct frequency bands associated with forward (gamma, 30–80 Hz) and backward (beta, 10–30 Hz) activity in neocortex. Likewise, Walsh et al. [2020], reviewing the literature, note that violations of expectations, i.e. large prediction errors, are associated with increased gamma amplitude, whereas beta amplitude tends to increase in preparation for a predicted stimulus [Fujioka et al., 2009]. In agreement with hierarchical predictive coding, this supports the conjecture that backward projections convey predictions, while forward projections perform inference updating using prediction errors. If we consider the gamma and beta frequency bands as distinct network-wide "clock rates" associated with inference and dynamics estimation, respectively, this suggests that cortex performs anywhere from 1–8 inference iterations per time step. Each time step would constitute 30–100 ms. As seen with amortized variational filtering [Marino et al., 2018a], as well iterative amortized policy optimization [Marino et al., 2020b], even with a *single* inference iteration per time step, amortization can yield reasonably accurate inference estimates. Note that direct amortization [Dayan et al., 1995] would suggest that these frequencies should be identical, in disagreement with empirical observations. Gradient-based optimization, in contrast, as suggested and employed by Rao and Ballard [1999] and Friston [2005], may not be sufficient to provide accurate inference estimates with so few iterations. This may provide further support for the claim that forward pyramidal neurons in cortex implement a form of iterative amortized inference.

Backpropagation Training deep networks at scale appears to require gradient-based parameter optimization, i.e. backpropagation [Werbos, 1974, Rumelhart et al., 1986]. However, the biological plausibility of backpropagation remains an open question [Lillicrap et al., 2020]. Critics argue that backpropagation requires non-local learning signals and other techniques [Grossberg, 1987, Crick,

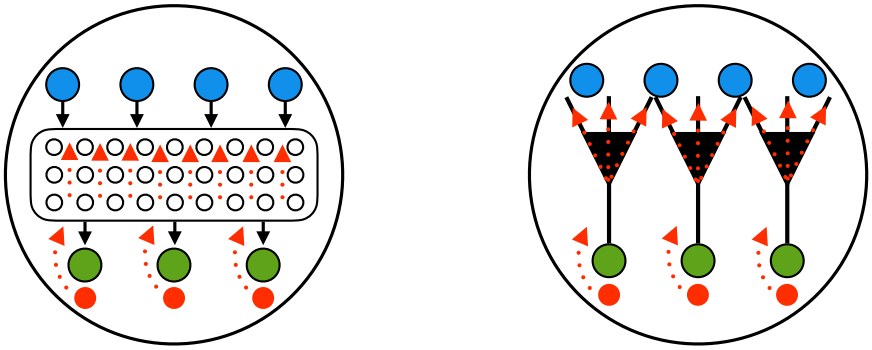

Figure 13: **Backpropagation**. Placing deep networks in correspondence with pyramidal neuron dendrites suggests an alternative perspective on the biological-plausibility of backpropagation. In deep latent variable models, backpropagation is only performed across variables that are directly connected through a conditional probability (left). From the perspective presented here, this corresponds to learning *within* pyramidal neurons. One possible implementation may be through backpropagating action potentials, perhaps combined with other neuromodulatory inputs (right).

1989], whereas the brain relies largely on local learning rules [Hebb, 1949, Markram et al., 1997, Bi and Poo, 1998]. A number of "biologically-plausible" formulations of backpropagation have been proposed [Stork, 1989, Körding and König, 2001, Xie and Seung, 2003, Hinton, 2007, Lillicrap et al., 2016], attempting to reconcile this disparity and others. Yet, consensus is still lacking regarding the biological implementations of these proposed techniques. From another perspective, the apparent biological implausibility of backpropagation may instead be the result of incorrectly assuming a one-to-one correspondence between biological and artificial neurons.

If deep networks are in correspondence with pyramidal dendrites, this suggests a different perspective on the biological-plausibility of backpropagation. In hierarchical latent variable models, prediction errors at each level of the latent hierarchy provide a local learning signal [Friston, 2005, Bengio, 2014, Lee et al., 2015, Whittington and Bogacz, 2017]. Thus, the global objective is decomposed into local errors, with learning within each latent level performed through gradient-based optimization. This is exemplified by deep latent variable models, which utilize backpropagation *within* each latent level, but not (necessarily) *across* latent levels. Again, considering the correspondence of pyramidal dendrites and deep networks, this suggests that learning *within* pyramidal neurons may be more analogous to backpropagation (Figure 13). Not surprisingly, one possible candidate is backpropagating action potentials [Stuart and Sakmann, 1994, Williams and Stuart, 2000]. These occur in the dendrites of pyramidal neurons, actively propagating a signal of neural activity back to synaptic inputs [Stuart et al., 1997, Brunner and Szabadics, 2016]. This results in the location-dependent influx of calcium, leading to a variety of synaptic changes throughout the dendritic tree [Johenning et al., 2015]. Indeed, Schiess et al. [2016] recently investigated a computational model of gradient backpropagation within dendritic trees. While many details remain unclear, this overall perspective of backpropagation *within* neurons, rather than across networks of neurons, offers a more biologically-plausible alternative; all signals are local to the pyramidal neurons/dendrites within the cortical circuit. Given the proposed theoretical role and empirical observations, this possible correspondence between backpropagation and backpropagating action potentials warrants further investigation.

6.2 Lateral Inhibition & Normalizing Flows

Sensory Input Normalization Sensory stimuli are highly redundant in both space and time. Examples include luminance at neighboring photoreceptors or pressure on adjacent mechanoreceptors, each of which tend to persist over time intervals. One of the key computational roles of early sensory areas, e.g. retina, appears to be in reducing these redundancies through normalization. In retina, such normalization operations are carried out through lateral inhibition via horizontal and amacrine cells. As a result, the transmitted output signals are less correlated [Graham et al., 2006, Pitkow and Meister, 2012]. As discussed in this paper and elsewhere, normalization and prediction are inseparable, i.e. one must form a prediction in order to normalize. Accordingly, previous works have framed early sensory processing in terms of (spatiotemporal) predictive coding [Srinivasan et al., 1982, Hosoya et al., 2005, Palmer et al., 2015]. This is often motivated in terms of increased sensitivity or efficiency [Srinivasan et al., 1982, Atick and Redlich, 1990] due to redundancy reduction [Barlow et al., 1961, 1989], i.e. compression.

If we consider cortex as a hierarchical latent variable model of sensory inputs, then early sensory areas are implicated in parameterizing the conditional likelihood. The ubiquity of normalization operations in early sensory areas is suggestive of normalization in a flow-based model. That is, early sensory areas may implement the "inference" direction of a flow-based conditional likelihood [Agrawal and Dukkipati, 2016, Winkler et al., 2019]. This would create a learned, normalized space in which cortex makes predictions. In addition to the sensitivity and efficiency arguments above, this could simplify downstream generative modeling and improve generalization, as demonstrated in Marino et al. [2020a]. Interestingly, Rao and Ballard [1999] employ a similar whitening scheme on image inputs, presumably to imitate retina. Normalizing flows offers a generalization of this idea to non-affine parameterizations and multiple stages of normalization. Further, framing early sensory areas in terms of normalizing flows connects these computations conceptually with the rest of the cortical generative model, i.e. truly evaluating the data-level predictions of cortex would require inverting the normalization of early sensory areas. While normalizing flows may not provide a perfect description of early sensory processing (e.g. these operations may not be completely invertible), this framework may help to unify many disparate input normalization circuits.

Normalization in Thalamus & Cortex Normalization is also thought to be a key aspect of first-order relays in thalamus, such as the lateral geniculate nucleus (LGN). Dong and Atick [1995] proposed that inhibition across time in LGN could provide a mechanism for temporally decorrelating the input from retina, with some supporting evidence provided by Dan et al. [1996]. Again, this can be considered as a form of temporal predictive coding [Srinivasan et al., 1982], removing easily predictable temporal information. Following the interpretation above, normalization in thalamus may provide a second stage of normalizing flow, further removing redundancy. In Marino et al. [2020a], we implemented this general technique using autoregressive flows in deep latent variable models. There, we saw that low-level temporal normalization removes static backgrounds, improving modeling and generalization. Further work is needed to assess the functional form (affine/non-affine) and types of dependencies (linear/non-linear) implemented in first-order thalamic relays, though suggestions are given by Dong and Atick [1995].

Normalization, via local lateral inhibition, is also found throughout cortex [King et al., 2013]. Friston [2005] suggested that lateral inhibition plays the computational role of inverse covariance (precision) matrices, modeling dependencies between dimensions within the same latent level of the hierarchy. This corresponds to parameterizing approximate posteriors [Rezende and Mohamed, 2015, Kingma et al., 2016] and/or conditional priors [Huang et al., 2017] with affine normalizing flows with linear dependencies.[1] Again, note that normalizing flows offers a more general mathematical framework for describing these normalization computations. Further, predictive coding often assumes

[1]Specifically, Friston's proposal corresponds to ZCA whitening [Friston, 2005], whereas those within machine learning have explored Cholesky whitening [Kingma et al., 2016].

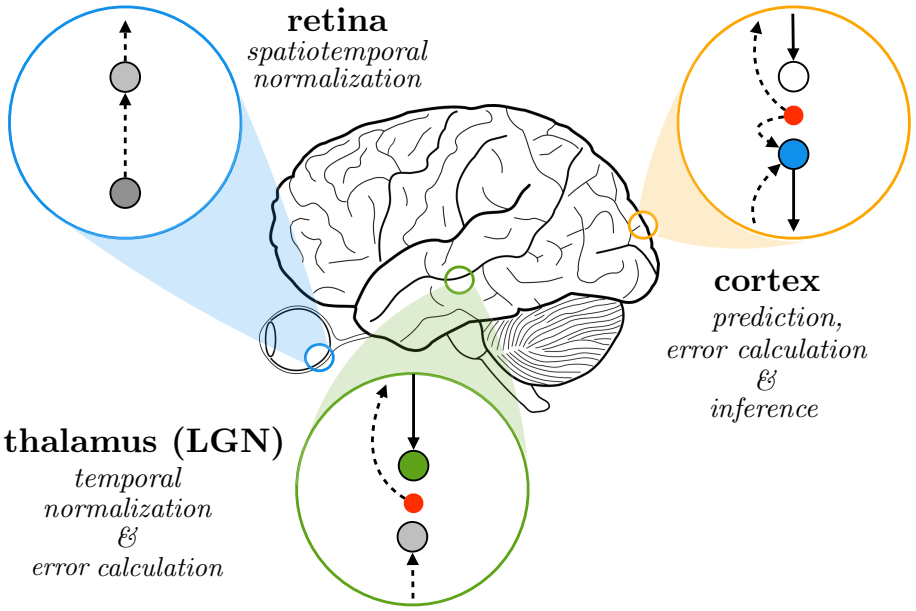

Figure 14: **Computational Schematic of the Visual Pathway**. Interpreting the early visual pathway from the perspective of a latent variable model, we can assign computational functions to retina, LGN, and cortex. Retina and LGN are interpreted as implementing normalizing flows, i.e. spatiotemporal predictive coding, reducing spatial and temporal redundancy in the visual input. This is shown as dashed arrows between gray circles. LGN also serves as the lowest level for hierarchical predictions, which are computed through backward connections from cortex. Using prediction errors throughout the hierarchy, forward cortical connections update latent estimates.

that these dependencies are modeled using *symmetric* lateral weights [Friston, 2005]. In contrast, normalizing flows also permits *non-symmetric* schemes, e.g. using spatially autoregressive models [Kingma et al., 2016] or ensembles of such models [Uria et al., 2014]. Interestingly, Vahdat and Kautz [2020] recently applied normalizing flows in large-scale hierarchical latent variable models, but restricted dependencies to local spatial patches for efficiency, perhaps more in-line with local lateral inhibition in cortex. Marino et al. [2020a] also applied normalizing flows across time in hierarchical models to assist in parameterizing dynamics, discussing connections to Friston's notion of generalized coordinates [Friston, 2008a]. Thus, multiple forms of spatial and temporal normalization may occur within cortex, allowing cortical columns to add and remove dependencies across space and time. Finally, multiple works within predictive coding have explored the use of prediction precision as a form of attention [Spratling, 2008, Feldman and Friston, 2010]. Increasing the precision of predictions modulates the gain of prediction errors in driving inference, leading to more precise inferred estimates. This may prove to be a useful technique in machine learning. The overall computational scheme, ignoring spatiotemporal normalization in cortex for simplicity, is shown in Figure 14.

Motor Dependencies Much like sensory input areas, there is a striking degree of low-level interneuron circuitry in motor output areas. A canonical example is central pattern generator (CPG) circuits, which, through local excitation and inhibition, give rise to coordinated muscle activation [Marder and Bucher, 2001]. These circuits provide a basis of "motor primitives," allowing muscle activations to be carried out in a lower-dimensional manifold rather than the entire combinatorial space. Given the anatomical similarity of these circuits with normalization circuitry in sensory areas, as well as the close relationship between adding (generation) and removing (inference) dependencies

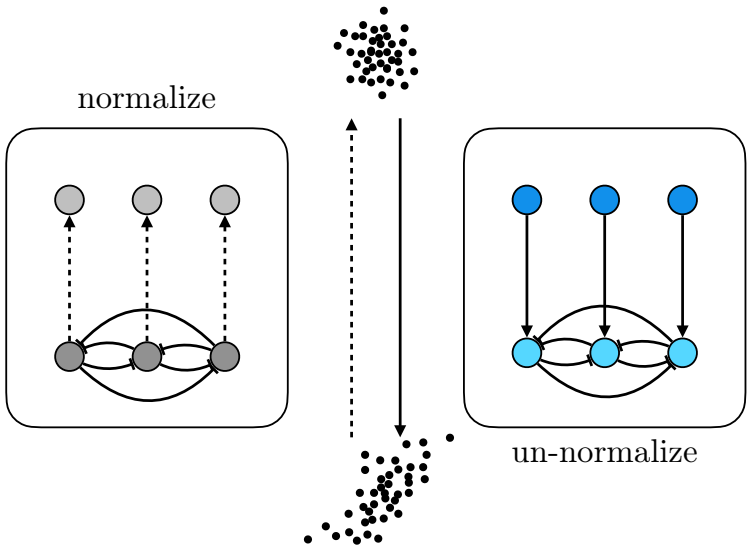

Figure 15: **Adding & Removing Dependencies with Normalizing Flows**. Normalizing flows provides a general mathematical framework for removing (left) or adding (right) probabilistic dependencies. Using lateral interactions, one can move between a normalized (top) or un-normalized (bottom) space. Normalized spaces have benefits for compression, whereas un-normalized spaces are more expressive. Neural systems may employ these transforms for sensory and motor processing.

in normalizing flows, it is possible that such low-level motor circuits are implementing flow-based distributions. Under this scheme, spinal circuits convert compressed, uncorrelated signals from higher-level motor areas into correlated low-level muscle activation. Similar efficiency arguments from sensory input areas [Atick and Redlich, 1990] are applicable to motor outputs, treating the spinal cord as a communication channel. Further, spatiotemporal dependencies in muscle activation may also improve the expressive capacity and generalization of motor routines. Indeed, multiple works have investigated improving control policies with normalizing flows across motor dimensions [Haarnoja et al., 2018, Tang and Agrawal, 2018, Ward et al., 2019], and Guerra and Marino [2020] investigated incorporating flow-based motor dependencies across time. Consistent with this setup, it has been empirically observed in neural systems that higher-level motor areas output transient update signals, which are converted to lower-level motor trajectories [Shalit et al., 2012]. While these findings are encouraging, more apparent benefits are likely to arise in high-dimensional action spaces, where the *curse of dimensionality* [Chen and Gopinath, 2001] makes it essential to model dependencies for effective exploration. For instance, current MuJoCo environments from OpenAI gym [Todorov et al., 2012, Brockman et al., 2016] contain $<$20 action dimensions, whereas it has been estimated that there are roughly 800 independent dimensions to human motor control [Powers, 1973]. Operating and learning efficiently in such large action spaces may require modeling spatiotemporal dependencies through some form of hierarchical or flow-based decomposition.

7 Discussion

In this chapter, we connected the ideas presented in this thesis back to their origins in predictive coding. The two core techniques developed in this thesis, iterative amortized inference and sequential autoregressive flows, map onto aspects of hierarchical and spatiotemporal predictive coding, respectively. By connecting these techniques back to predictive coding, we arrived at a variety of possible

implied correspondences between machine learning and neuroscience. In particular,

- we identified the dendrites of pyramidal neurons as functionally analogous to (nonlinear) deep networks, and
- we identified lateral inhibition as implementing normalizing flows.

Placing pyramidal neuron dendrites in correspondence with deep networks departs from the traditional one-to-one analogy of biological and artificial neurons, raising a host of questions regarding dendritic computation and learning via backpropagation. Likewise, normalizing flows offers a more general framework for considering the normalization computations carried out by lateral inhibitory interactions found within multiple brain regions. We are hopeful that connecting these areas will provide new insights for both machine learning and neuroscience.

For practical reasons, we primarily focused on perception in discussing predictive coding and the ideas in this paper. This is a result of the fact that predictive coding was initially developed and studied in the context of generative models of sensory inputs [Srinivasan et al., 1982, Rao and Ballard, 1999, Friston, 2005]. However, if the cortical microcircuit implements a general-purpose modeling and inference algorithm, we should expect similar computations to be applicable to motor and prefrontal cortices. Over the past decade, Friston and colleagues have developed a range of exciting ideas, interpreting motor control as a process of proprioceptive prediction [Adams et al., 2013] and prefrontal cortex as performing hierarchical goal inference [Pezzulo et al., 2018]. Such ideas can be seen as modern extensions of early ideas in cybernetics [Wiener, 1948, MacKay, 1956, Powers, 1973], using motor control to correct for discrepancies (errors) between desired and actual outcomes. Indeed, as shown in Marino et al. [2020b], Guerra and Marino [2020], the *same* modeling and inference techniques used for perception can be applied to control. Nevertheless, further work is needed within machine learning to unify perception and control under a single formulation.

Acknowledgments

We thank Sam Gershman for comments on this manuscript, and we thank Karl Friston for useful early discussions related to these ideas. A condensed version of this paper appeared at the Real Neurons & Hidden Units workshop[2] at NeurIPS 2019. We thank the anonymous reviewers, as well as workshop organizers and participants, for their feedback.

References

David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.

Rick A Adams, Stewart Shipp, and Karl J Friston. Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218(3):611–643, 2013.

Siddharth Agrawal and Ambedkar Dukkipati. Deep variational inference without pixel-wise reconstruction. *arXiv preprint arXiv:1611.05209*, 2016.

Arjen Alink, Caspar M Schwiedrzik, Axel Kohler, Wolf Singer, and Lars Muckli. Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, 30(8):2960–2966, 2010.

Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3981–3989, 2016.

[2] https://openreview.net/group?id=NeurIPS.cc/2019/Workshop/Neuro_AI

W Ross Ashby. *An Introduction to Cybernetics*. Chapman and Hall, 1956.

B Atal and M Schroeder. Predictive coding of speech signals and subjective error criteria. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(3):247–254, 1979.

Joseph J Atick and A Norman Redlich. Towards a theory of early visual processing. *Neural computation*, 2(3):308–320, 1990.

Joseph J Atick and A Norman Redlich. What does the retina know about natural scenes? *Neural computation*, 4(2):196–210, 1992.

Stephen A Baccus, Bence P Ölveczky, Mihai Manu, and Markus Meister. A retinal circuit that computes object motion. *Journal of Neuroscience*, 28(27):6807–6817, 2008.

Dana H Ballard. Modular learning in neural networks. In *AAAI*, pages 279–284, 1987.

Horace B Barlow. The coding of sensory messages. *Current problems in animal behavior*, 1961.

Horace B Barlow. Unsupervised learning. *Neural computation*, 1(3):295–311, 1989.

Horace B Barlow, Tej P Kaushal, and Graeme J Mitchison. Finding minimum entropy codes. *Neural Computation*, 1(3):412–423, 1989.

Horace B Barlow et al. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234, 1961.

Andre Moraes Bastos, W Martin Usrey, Rick A Adams, George R Mangun, Pascal Fries, and Karl Friston. Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711, 2012.

Andre Moraes Bastos, Julien Vezoli, Conrado Arturo Bosman, Jan-Mathijs Schoffelen, Robert Oostenveld, Jarrod Robert Dowdall, Peter De Weerd, Henry Kennedy, and Pascal Fries. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85 (2):390–401, 2015.

John M Bekkers. Pyramidal neurons. *Current Biology*, 21(24):R975, 2011.

Anthony J Bell and Terrence J Sejnowski. The "independent components" of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, 1997.

Curtis C Bell. Memory-based expectations in electrosensory systems. *Current opinion in neurobiology*, 11(4):481–487, 2001.

Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

Yoshua Bengio. How auto-encoders could provide credit assignment in deep networks via target propagation. *arXiv preprint arXiv:1407.7906*, 2014.

Yoshua Bengio and Samy Bengio. Modeling high-dimensional discrete data with multi-layer neural networks. In *Advances in Neural Information Processing Systems*, pages 400–406, 2000.

David Beniaguev, Idan Segev, and Michael London. Single cortical neurons as deep artificial neural networks. *bioRxiv*, page 613141, 2020.

Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24): 10464–10472, 1998.

William Bialek, Fred Rieke, RR De Ruyter Van Steveninck, and David Warland. Reading a neural code. *Science*, 252(5014):1854–1857, 1991.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

János Brunner and János Szabadics. Analogue modulation of back-propagating action potentials enables dendritic hybrid signalling. *Nature communications*, 7:13033, 2016.

Matteo Carandini and David J Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, 2012.

Scott Saobing Chen and Ramesh A Gopinath. Gaussianization. In *Advances in neural information processing systems*, pages 423–429, 2001.

Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pages 4754–4765, 2018.

Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.

Rob Cornish, Anthony L Caterini, George Deligiannidis, and Arnaud Doucet. Relaxing bijectivity constraints with continuously indexed normalising flows. In *International Conference on Machine Learning*, 2020.

Elise N Covic and S Murray Sherman. Synaptic properties of connections between the primary and secondary auditory cortices in mice. *Cerebral Cortex*, 21(11):2425–2441, 2011.

Chris Cremer, Xuechen Li, and David Duvenaud. Inference suboptimality in variational autoencoders. In *International Conference on Machine Learning*, pages 1078–1086, 2018.

Francis Crick. The recent excitement about neural networks. *Nature*, 337(6203):129–132, 1989.

Milan Cvitkovic and Günther Koliander. Minimal achievable sufficient statistic learning. In *International Conference on Machine Learning*, pages 1465–1474, 2019.

Yang Dan, Joseph J Atick, and R Clay Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience*, 16(10):3351–3362, 1996.

Ishita Dasgupta, Eric Schulz, Noah D Goodman, and Samuel J Gershman. Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, 178:67–81, 2018.

Peter Dayan and Geoffrey E Hinton. Varieties of helmholtz machine. *Neural Networks*, 9(8):1385–1403, 1996.

Peter Dayan, Geoffrey E Hinton, Radford M Neal, and Richard S Zemel. The helmholtz machine. *Neural computation*, 7(5):889–904, 1995.

Roberto De Pasquale and S Murray Sherman. Synaptic properties of corticocortical connections between the primary and secondary visual cortical areas in the mouse. *Journal of Neuroscience*, 31 (46):16494–16506, 2011.

Gustavo Deco and Wilfried Brauer. Higher order statistical decorrelation without information loss. In *Advances in Neural Information Processing Systems*, pages 247–254, 1995.

Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. In *International Conference on Learning Representations*, 2015.

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations*, 2017.

Dawei W Dong and Joseph J Atick. Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in Neural Systems*, 6(2):159–178, 1995.

Rodney J Douglas, Kevan AC Martin, and David Whitteridge. A canonical microcircuit for neocortex. *Neural computation*, 1(4):480–488, 1989.

Kenji Doya, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao. *Bayesian brain: Probabilistic approaches to neural coding*. MIT press, 2007.

Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *arXiv preprint arXiv:1906.04032*, 2019.

Frederik Ebert, Chelsea Finn, Alex X Lee, and Sergey Levine. Self-supervised visual planning with temporal skip connections. In *Conference on Robot Learning*, 2017.

Bradley Efron and Carl Morris. Stein's estimation rule and its competitors—an empirical bayes approach. *Journal of the American Statistical Association*, 68(341):117–130, 1973.

Tobias Egner, Jim M Monti, and Christopher Summerfield. Expectation and surprise determine neural population responses in the ventral visual stream. *Journal of Neuroscience*, 30(49):16601–16608, 2010.

Steven J Eliades and Xiaoqin Wang. Neural substrates of vocalization feedback monitoring in primate auditory cortex. *Nature*, 453(7198):1102, 2008.

Harriet Feldman and Karl Friston. Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4, 2010.

Brendan J Frey, Geoffrey E Hinton, and Peter Dayan. Does the wake-sleep algorithm produce good density estimators? In *Advances in neural information processing systems*, pages 661–667, 1996.

Jerome H Friedman. Exploratory projection pursuit. *Journal of the American statistical association*, 82(397):249–266, 1987.

Karl Friston. A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456):815–836, 2005.

Karl Friston. Hierarchical models in the brain. *PLoS computational biology*, 4(11):e1000211, 2008a.

Karl Friston. Variational filtering. *NeuroImage*, 41(3):747–766, 2008b.

Karl Friston. The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13 (7):293–301, 2009.

Karl Friston. Does predictive coding have a future? *Nature neuroscience*, 21(8):1019, 2018.

Karl Friston, Jérémie Mattout, Nelson Trujillo-Barreto, John Ashburner, and Will Penny. Variational free energy and the laplace approximation. *Neuroimage*, 34(1):220–234, 2007.

Takako Fujioka, Laurelj Trainor, Edwardw Large, and Bernhard Ross. Beta and gamma rhythms in human auditory cortex during musical beat processing. *Annals of the New York Academy of Sciences*, 1169(1):89–92, 2009.

Samuel Gershman and Noah Goodman. Amortized inference in probabilistic reasoning. In *Proceedings of the Cognitive Science Society*, volume 36, 2014.

Samuel J Gershman. What does the free energy principle tell us about the brain? *arXiv preprint arXiv:1901.07945*, 2019.

Albert Gidon, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsi, Panayiota Poirazi, Martin Holtkamp, Imre Vida, and Matthew Evan Larkum. Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, 2020.

Charles D Gilbert and Mariano Sigman. Brain states: top-down influences in sensory processing. *Neuron*, 54(5):677–696, 2007.

Pascal Girard and Jean Bullier. Visual activity in area v2 during reversible inactivation of area 17 in the macaque monkey. *Journal of neurophysiology*, 62(6):1287–1302, 1989.

Pascal Girard, PA Salin, and Jean Bullier. Visual activity in areas v3a and v3 during reversible inactivation of area v1 in the macaque monkey. *Journal of Neurophysiology*, 66(5):1493–1503, 1991.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Daniel J Graham, Damon M Chandler, and David J Field. Can the theory of "whitening" explain the center-surround properties of retinal ganglion cell receptive fields? *Vision research*, 46(18): 2901–2913, 2006.

Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.

Karol Gregor, Ivo Danihelka, Andriy Mnih, Charles Blundell, and Daan Wierstra. Deep autoregressive networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1242–1250, 2014.

Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science*, 11(1):23–63, 1987.

Jordan Guergiuev, Timothy P Lillicrap, and Blake A Richards. Biologically feasible deep learning with segregated dendrites. *arXiv preprint arXiv:1610.00161*, 2016.

Alex Guerra and Joseph Marino. Sequential autoregressive flow-based policies. In *ICML workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2020.

Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville. Pixelvae: A latent variable model for natural images. In *International Conference on Learning Representations*, 2017.

David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, pages 2450–2462, 2018.

Tuomas Haarnoja, Kristian Hartikainen, Pieter Abbeel, and Sergey Levine. Latent space policies for hierarchical reinforcement learning. In *International Conference on Machine Learning*, pages 1846–1855, 2018.

CW Harrison. Experiments with linear prediction in television. *Bell System Technical Journal*, 31 (4):764–783, 1952.

Jeff Hawkins and Sandra Blakeslee. *On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines*. Macmillan, 2004.

Donald O Hebb. The organization of behavior; a neuropsycholocigal theory. *A Wiley Book in Clinical Psychology.*, pages 62–78, 1949.

Geoffrey E. Hinton. How to do backpropagation in a brain. *NeurIPS Deep Learning Workshop*, 2007.

Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13. ACM, 1993.

Devon Hjelm, Ruslan R Salakhutdinov, Kyunghyun Cho, Nebojsa Jojic, Vince Calhoun, and Junyoung Chung. Iterative refinement of the approximate posterior for directed belief networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4691–4699, 2016.

John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

Toshihiko Hosoya, Stephen A Baccus, and Markus Meister. Dynamic predictive coding by the retina. *Nature*, 436(7047):71, 2005.

Chin-Wei Huang, Ahmed Touati, Laurent Dinh, Michal Drozdzal, Mohammad Havaei, Laurent Charlin, and Aaron Courville. Learnable explicit density for continuous latent space and variational inference. *arXiv preprint arXiv:1710.02248*, 2017.

Yanping Huang and Rajesh PN Rao. Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5):580–593, 2011.

Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.

Jeffry S Isaacson and Massimo Scanziani. How inhibition shapes cortical activity. *Neuron*, 72(2): 231–243, 2011.

Masao Ito. Cerebellar learning in the vestibulo–ocular reflex. *Trends in cognitive sciences*, 2(9): 313–321, 1998.

Janneke FM Jehee and Dana H Ballard. Predictive feedback can account for biphasic responses in the lateral geniculate nucleus. *PLoS Comput Biol*, 5(5):e1000373, 2009.

Friedrich W Johenning, Anne-Kathrin Theis, Ulrike Pannasch, Martin Rückl, Sten Rüdiger, and Dietmar Schmitz. Ryanodine receptor activation induces long-term plasticity of spine calcium dynamics. *PLoS biology*, 13(6):e1002181, 2015.

Ilenna Simone Jones and Konrad Paul Kording. Can single neurons solve mnist? the computational power of biological dendritic trees. *arXiv preprint arXiv:2009.01269*, 2020.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *NATO ASI SERIES D BEHAVIOURAL AND SOCIAL SCIENCES*, 89:105–162, 1998.

Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

Ryota Kanai, Yutaka Komura, Stewart Shipp, and Karl Friston. Cerebral hierarchies: predictive processing, precision and the pulvinar. *Phil. Trans. R. Soc. B*, 370(1668):20140169, 2015.

Georg B Keller and Thomas D Mrsic-Flogel. Predictive processing: a canonical cortical computation. *Neuron*, 100(2):424–435, 2018.

Georg B Keller, Tobias Bonhoeffer, and Mark Hübener. Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron*, 74(5):809–815, 2012.

Ann Kennedy, Greg Wayne, Patrick Kaifosh, Karina Alviña, LF Abbott, and Nathaniel B Sawtell. A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nature neuroscience*, 17(3):416–422, 2014.

Agnan Kessy, Alex Lewin, and Korbinian Strimmer. Optimal whitening and decorrelation. *The American Statistician*, 72(4):309–314, 2018.

Yoon Kim, Sam Wiseman, Andrew C Miller, David Sontag, and Alexander M Rush. Semi-amortized variational autoencoders. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.

Paul D King, Joel Zylberberg, and Michael R DeWeese. Inhibitory interneurons decorrelate excitatory cells to drive sparse code formation in a spiking model of v1. *Journal of Neuroscience*, 2013.

Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.

Durk P Kingma and Max Welling. Stochastic gradient vb and the variational auto-encoder. In *Proceedings of the International Conference on Learning Representations*, 2014.

Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016.

Konrad P Körding and Peter König. Supervised and unsupervised learning with two sites of synaptic integration. *Journal of computational neuroscience*, 11(3):207–215, 2001.

Rahul G Krishnan, Dawen Liang, and Matthew Hoffman. On the challenges of learning with inference networks on sparse, high-dimensional data. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 143–151, 2018.

Alex Krizhevsky and Geoffrey E Hinton. Learning multiple layers of features from tiny images. 2009.

Manoj Kumar, Mohammad Babaeizadeh, Dumitru Erhan, Chelsea Finn, Sergey Levine, Laurent Dinh, and Durk P Kingma. Videoflow: A flow-based generative model for video. In *International Conference on Learning Representations*, 2020.

Valero Laparra, Gustavo Camps-Valls, and Jesús Malo. Iterative gaussianization: from ica to random rotations. *IEEE transactions on neural networks*, 22(4):537–549, 2011.

Simon Laughlin. A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung c*, 36(9-10):910–912, 1981.

Yann LeCun, Yoshua Bengio, and Geoffrey E Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

Dong-Hyun Lee, Saizheng Zhang, Asja Fischer, and Yoshua Bengio. Difference target propagation. In *Joint european conference on machine learning and knowledge discovery in databases*, pages 498–515. Springer, 2015.

Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7: 13276, 2016.

Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, pages 1–12, 2020.

Michael London and Michael Häusser. Dendritic computation. *Annu. Rev. Neurosci.*, 28:503–532, 2005.

William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. In *International Conference on Learning Representations*, 2017.

William Lotter, Gabriel Kreiman, and David Cox. A neural network trained to predict future video frames mimics critical properties of biological neuronal responses and perception. *arXiv preprint arXiv:1805.10734*, 2018.

D M MacKay. The epistemological problem for automata. *Automata studies*, pages 235–252, 1956.

Eve Marder and Dirk Bucher. Central pattern generators and the control of rhythmic movements. *Current biology*, 2001.

Joseph Marino, Milan Cvitkovic, and Yisong Yue. A general method for amortizing variational filtering. In *Advances in Neural Information Processing Systems*, pages 7857–7868, 2018a.

Joseph Marino, Yisong Yue, and Stephan Mandt. Iterative amortized inference. In *International Conference on Machine Learning*, pages 3403–3412, 2018b.

Joseph Marino, Lei Chen, Jiawei He, and Stephan Mandt. Improving sequential latent variable models with autoregressive flows. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–16, 2020a.

Joseph Marino, Alexandre Piché, Alessandro Davide Ialongo, and Yisong Yue. Iterative amortized policy optimization. *arXiv preprint arXiv:2010.10670*, 2020b.

Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215, 1997.

Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.

Bartlett W Mel. The clusteron: toward a simple abstraction for a complex neuron. In *Advances in neural information processing systems*, 1992.

Hanno S Meyer, Daniel Schwarz, Verena C Wimmer, Arno C Schmitt, Jason ND Kerr, Bert Sakmann, and Moritz Helmstaedter. Inhibitory interneurons in a cortical column form hot zones of inhibition in layers 2 and 5a. *Proceedings of the National Academy of Sciences*, 108(40):16807–16812, 2011.

Travis Meyer and Carl R Olson. Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences*, 108(48):19401–19406, 2011.

Marvin Minsky and Seymour Papert. Perceptrons, 1969.

Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, pages 1791–1799, 2014.

VB Mountcastle, AL Berman, and PW Davies. Topographic organization and modality representation in first somatic area of cat's cerebral cortex by method of single unit analysis. *Am J Physiol*, 183 (464):10, 1955.

David Mumford. On the computational architecture of the neocortex. *Biological cybernetics*, 65(2): 135–145, 1991.

David Mumford. On the computational architecture of the neocortex. *Biological cybernetics*, 66(3): 241–251, 1992.

Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

Scott O Murray, Daniel Kersten, Bruno A Olshausen, Paul Schrater, and David L Woods. Shape perception reduces activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, 99(23):15164–15169, 2002.

Radford M Neal and Geoffrey E Hinton. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pages 355–368. Springer, 1998.

BM Oliver. Efficient coding. *The Bell System Technical Journal*, 31(4):724–750, 1952.

Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.

Bence P Ölveczky, Stephen A Baccus, and Markus Meister. Segregation of object and background motion in the retina. *Nature*, 423(6938):401–408, 2003.

Stephanie E Palmer, Olivier Marre, Michael J Berry, and William Bialek. Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22):6908–6913, 2015.

George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.

Lucas Parra, Gustavo Deco, and Stefan Miesbach. Redundancy reduction with information-preserving nonlinear maps. *Network: Computation in Neural Systems*, 6(1):61–72, 1995.

Gloria G Parras, Javier Nieto-Diego, Guillermo V Carbajal, Catalina Valdés-Baizabal, Carles Escera, and Manuel S Malmierca. Neurons along the auditory pathway exhibit a hierarchical organization of prediction error. *Nature communications*, 8(1):1–17, 2017.

Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3): 241–288, 1986.

Giovanni Pezzulo, Francesco Rigoli, and Karl J Friston. Hierarchical active inference: A theory of motivated control. *Trends in cognitive sciences*, 22(4):294–306, 2018.

Xaq Pitkow and Markus Meister. Decorrelation and efficient coding by retinal ganglion cells. *Nature neuroscience*, 15(4):628, 2012.

Panayiota Poirazi, Terrence Brannon, and Bartlett W Mel. Pyramidal neuron as two-layer neural network. *Neuron*, 37(6):989–999, 2003.

Alon Polsky, Bartlett W Mel, and Jackie Schiller. Computational subunits in thin dendrites of pyramidal cells. *Nature neuroscience*, 7(6):621, 2004.

Mohsen Pourahmadi. Covariance estimation: The glm and regularization perspectives. *Statistical Science*, pages 369–387, 2011.

William T Powers. *Behavior: The control of perception*. Aldine Chicago, 1973.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

Rajesh PN Rao. Correlates of attention in a model of dynamic visual recognition. In *Advances in neural information processing systems*, pages 80–86, 1998.

Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1), 1999.

Rajesh PN Rao and Terrence J Sejnowski. Predictive coding, cortical feedback, and spike-timing dependent plasticity. *Probabilistic models of the brain*, page 297, 2002.

Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, pages 14866–14876, 2019.

Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538, 2015.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the International Conference on Machine Learning*, pages 1278–1286, 2014.

Blake A Richards. personal communication.

Oren Rippel and Ryan Prescott Adams. High-dimensional probability estimation with deep density models. *arXiv preprint arXiv:1302.5125*, 2013.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 1986.

Mathieu Schiess, Robert Urbanczik, and Walter Senn. Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites. *PLoS computational biology*, 12(2):e1004638, 2016.

Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.

Uri Shalit, Nofya Zinger, Mati Joshua, and Yifat Prut. Descending systems translate transient cortical commands into a sustained muscle activation signal. *Cerebral cortex*, 22(8):1904–1914, 2012.

Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27 (3):379–423, 1948.

Jitendra Sharma, Alessandra Angelucci, and Mriganka Sur. Induction of visual orientation modules in auditory cortex. *Nature*, 404(6780):841, 2000.

S Murray Sherman and RW Guillery. The role of the thalamus in the flow of information to the cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1428):1695–1708, 2002.

Michael W Spratling. Reconciling predictive coding and biased competition models of cortical function. *Frontiers in computational neuroscience*, 2:4, 2008.

Mandyam Veerambudi Srinivasan, Simon Laughlin, and Andreas Dubs. Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1205):427–459, 1982.

David G Stork. Is backpropagation biologically plausible. In *International Joint Conference on Neural Networks*, volume 2, pages 241–246. IEEE Washington, DC, 1989.

Greg Stuart, Nelson Spruston, Bert Sakmann, and Michael Häusser. Action potential initiation and backpropagation in neurons of the mammalian cns. *Trends in neurosciences*, 20(3):125–131, 1997.

Greg J Stuart and Bert Sakmann. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, 367(6458):69, 1994.

Christopher Summerfield, Tobias Egner, Matthew Greene, Etienne Koechlin, Jennifer Mangels, and Joy Hirsch. Predictive codes for forthcoming perception in the frontal cortex. *Science*, 314(5803): 1311–1314, 2006.

Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.

Esteban G Tabak and Cristina V Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.

Yunhao Tang and Shipra Agrawal. Boosting trust region policy optimization by normalizing flows policy. *arXiv preprint arXiv:1809.10326*, 2018.

Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE, 2012.

Jakub M Tomczak and Max Welling. Improving variational auto-encoders using householder flow. *arXiv preprint arXiv:1611.09630*, 2016.

Benigno Uria, Iain Murray, and Hugo Larochelle. A deep and tractable density estimator. In *International Conference on Machine Learning*, pages 467–475, 2014.

Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in Neural Information Processing Systems*, 33, 2020.

Rianne van den Berg, Leonard Hasenclever, Jakub M Tomczak, and Max Welling. Sylvester normalizing flows for variational inference. *arXiv preprint arXiv:1803.05649*, 2018.

Gerben van den Broeke. What auto-encoders could learn from brains. Master's thesis, Aalto University, 2016.

Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *9th ISCA Speech Synthesis Workshop*, pages 125–125, 2016a.

Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756, 2016b.

Aaron van den Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International Conference on Machine Learning*, pages 3915–3923, 2018.

Hermann Von Helmholtz. *Handbuch der physiologischen Optik*, volume 9. Voss, 1867.

Catherine Wacongne, Etienne Labyt, Virginie van Wassenhove, Tristan Bekinschtein, Lionel Naccache, and Stanislas Dehaene. Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences*, 108(51):20754–20759, 2011.

Kevin S Walsh, David P McGovern, Andy Clark, and Redmond G O'Connell. Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the New York Academy of Sciences*, 1464(1):242, 2020.

Patrick Nadeem Ward, Ariella Smofsky, and Avishek Joey Bose. Improving exploration in soft-actor-critic with normalizing flows policies. *ICML Workshop on Invertible Neural Nets and Normalizing Flows*, 2019.

Paul Werbos. Beyond regression:" new tools for prediction and analysis in the behavioral sciences. *Ph. D. dissertation, Harvard University*, 1974.

James CR Whittington and Rafal Bogacz. An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation*, 2017.

Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7): 560–576, 2003.

Norbert Wiener. The interpolation, extrapolation and smoothing of stationary time series. ndrc report, 1942.

Norbert Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*. MIT press, 1948.

Stephen R Williams and Greg J Stuart. Backpropagation of physiological spike trains in neocortical pyramidal neurons: implications for temporal coding in dendrites. *Journal of Neuroscience*, 2000.

Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019.

Xiaohui Xie and H Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation*, 15(2):441–454, 2003.

Ruihan Yang, Yibo Yang, Joseph Marino, and Stephan Mandt. Hierarchical autoregressive modeling for neural video compression. *arXiv preprint arXiv:2010.10258*, 2020.

Ilker Yildirim, Tejas D Kulkarni, Winrich A Freiwald, and Joshua B Tenenbaum. Efficient and robust analysis-by-synthesis in vision: A computational framework, behavioral tests, and modeling neuronal representations. In *Thirty-Seventh Annual Conference of the Cognitive Science Society*, volume 4, 2015.

Anthony M Zador, Brenda J Claiborne, and Thomas H Brown. Nonlinear pattern separation in single hippocampal neurons with active dendritic membrane. In *Advances in neural information processing systems*, pages 51–58, 1992.

Pawel Zmarz and Georg B Keller. Mismatch receptive fields in mouse visual cortex. *Neuron*, 92(4): 766–772, 2016.