# Complementary Learning Theory, One Shot Learning and Contrastive Predictive Coding

Tom Dean, Stanford, April 24, 2021

Acknowledgements: Peter Samuel Chatain, William Ellsworth, Francis Eugene Lewis, Manan Ajay Shah, Yash Savani

# What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated

Dharshan Kumaran,[1,2,*] Demis Hassabis,[1,3,*] and James L. McClelland[4,*]

We update complementary learning systems (CLS) theory, which holds that intelligent agents must possess two learning systems, instantiated in mammalians in neocortex and hippocampus. The first gradually acquires structured knowledge representations while the second quickly learns the specifics of individual experiences. We broaden the role of replay of hippocampal memories in the theory, noting that replay allows goal-dependent weighting of experience statistics. We also address recent challenges to the theory and extend it by showing that recurrent activation of hippocampal traces can support some forms of generalization and that neocortical learning can be rapid for information that is consistent with known structure. Finally, we note the relevance of the theory to the design of artificial intelligent agents, highlighting connections between neuroscience and machine learning.

## Trends

Discovery of structure in ensembles of experiences depends on an interleaved learning process both in biological neural networks in neocortex and in contemporary artificial neural networks.

Recent work shows that once structured knowledge has been acquired in such networks, new consistent information can be integrated rapidly.

Both natural and artificial learning systems benefit from a second system that stores specific experiences, centred on the hippocampus in mammalians.

# Reinforcement Learning and Episodic Memory in Humans An Integrative Framework

Samuel J. Gershman[1] and Nathaniel D. Daw[2]

However, one challenge in the study of RL is computational: The simplicity of these tasks ignores important aspects of reinforcement learning in the real world: (*a*) State spaces are high-dimensional, continuous, and partially observable; this implies that (*b*) data are relatively sparse and, indeed, precisely the same situation may never be encountered twice; furthermore, (*c*) rewards depend on the long-term consequences of actions in ways that violate the classical assumptions that make RL tractable. A seemingly distinct challenge is that, cognitively, theories of RL have largely involved procedural and semantic memory, the way in which knowledge about action values or world models extracted gradually from many experiences can drive choice. This focus on semantic memory leaves out many aspects of memory, such as episodic memory, related to the traces of individual events. We suggest that these two challenges are related. The computational challenge can be dealt with, in part, by endowing RL systems with episodic memory, allowing them to (*a*) efficiently approximate value functions over complex state spaces, (*b*) learn with very little data, and (*c*) bridge long-term dependencies between actions and rewards.

## Global Neuronal Workspace model – Stanislas Dehaene & Bernard Baars

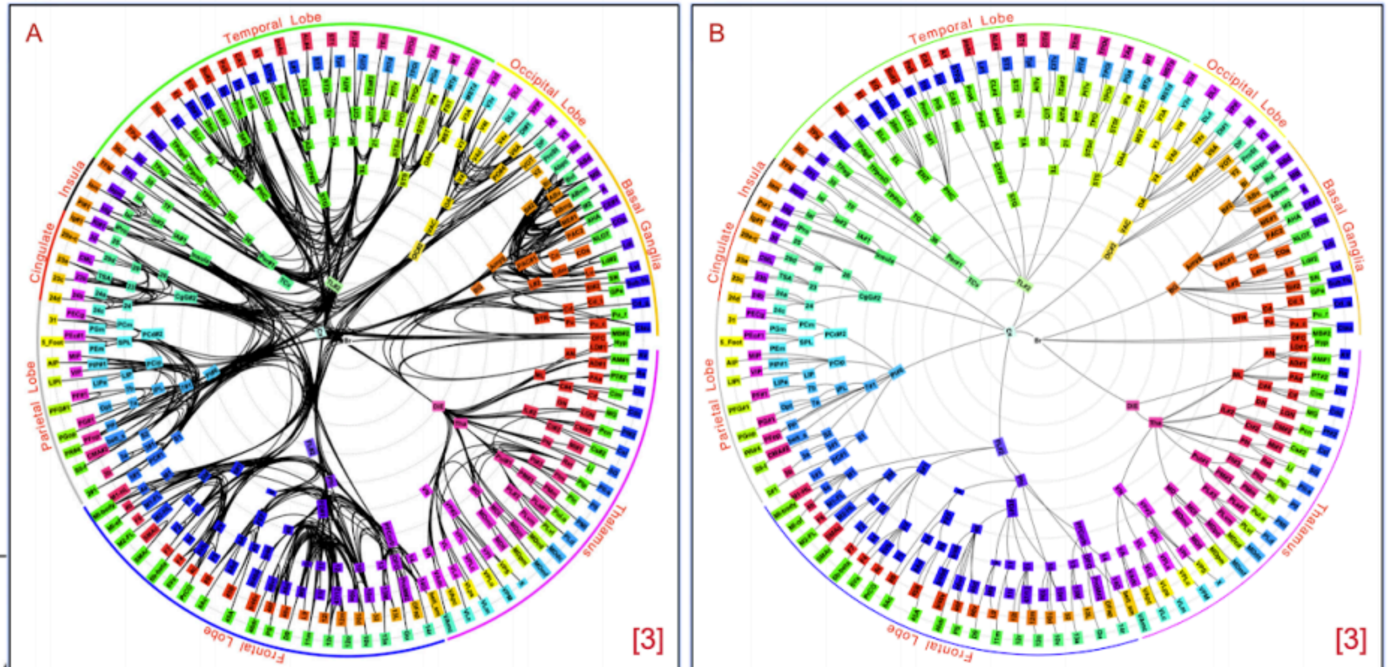


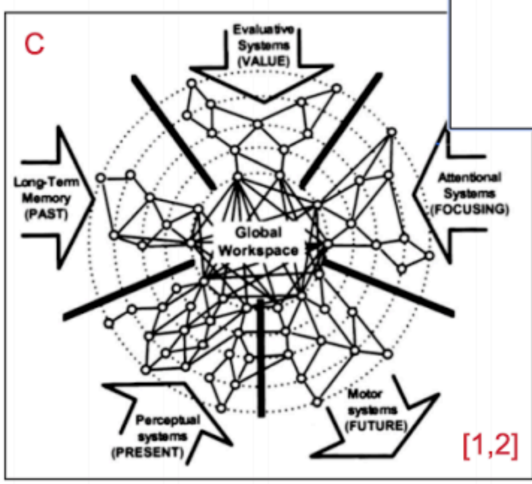



[1] Bernard Baars, Stan Franklin, and Thomas Ramsy. Global workspace dynamics: Cortical "binding and propagation" conscious contents. Frontiers in Psychology, 4:200, 2013.

[2] George Mashour, Pieter Roelfsema, Jean-Pierre Changeux, Stanislas Dehaene. Conscious processing global neuronal workspace hypothesis. Neuron, 105(5):776-798, 2020.

[3] Dharmendra Modha, Raghavendra Singh. Network architecture of the long-distance pathways in the macaque brain. The National Academy of Sciences, 107(30):13485-13490, 2010.

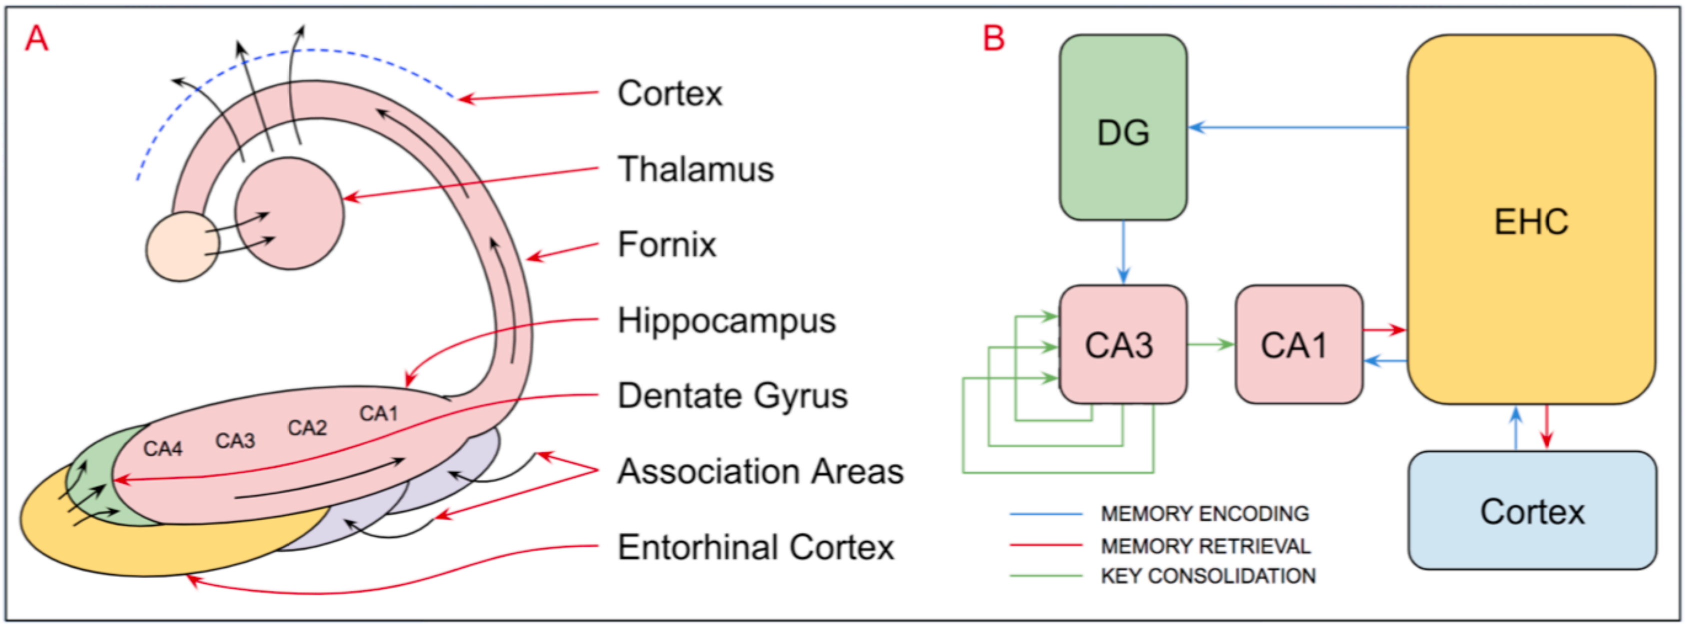


**See:** https://arxiv.org/pdf/1912.00421.pdf#page=13

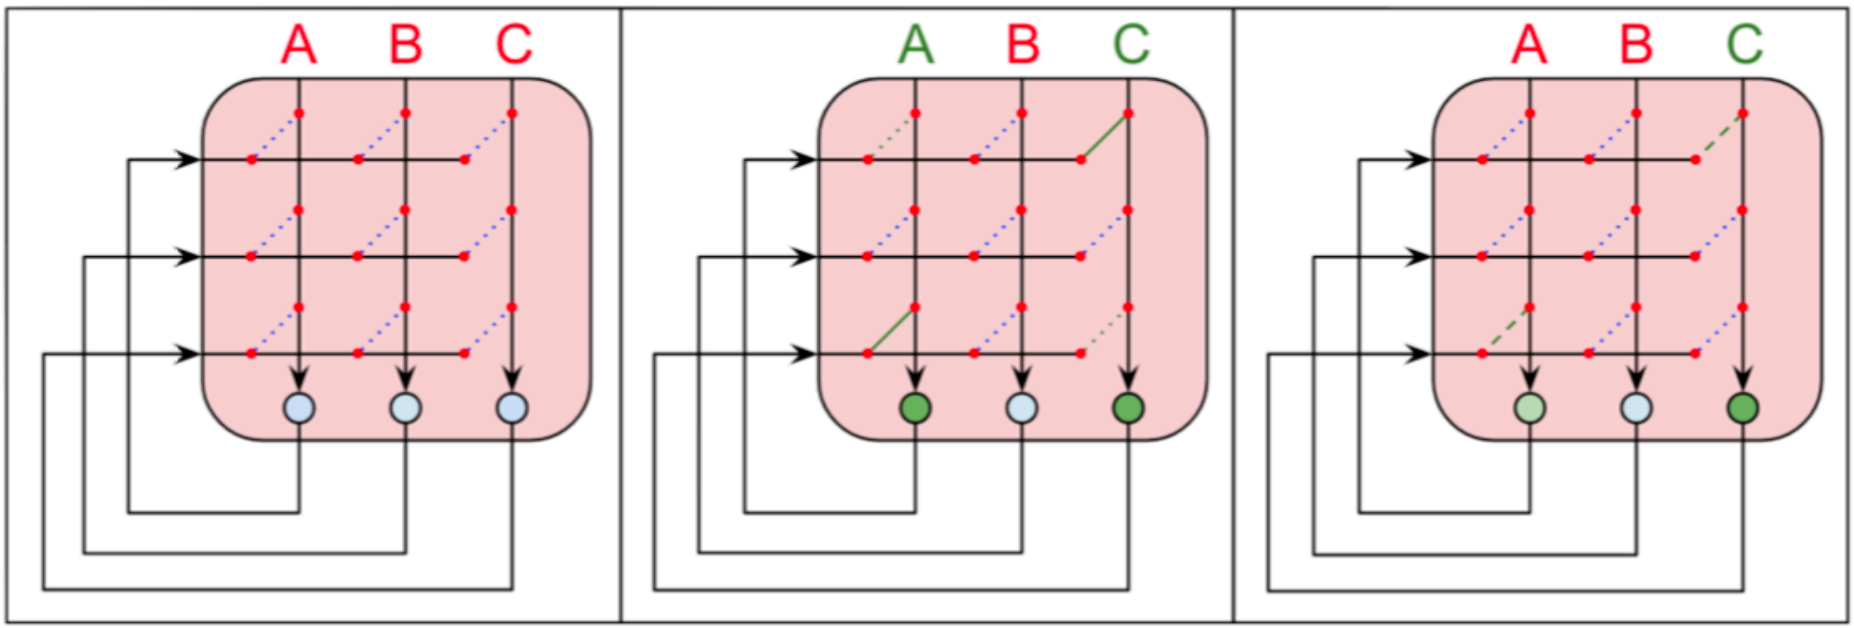


**See:** https://arxiv.org/pdf/1912.00421.pdf#page=14

# Hopfield Networks is All You Need

**Hubert Ramsauer**[*] **Bernhard Schäfl**[*] **Johannes Lehner**[*]
**Philipp Seidl**[*] **Michael Widrich**[*] **Lukas Gruber**[*] **Markus Holzleitner**[*]
**Milena Pavlović** **Geir Kjetil Sandve** **Victor Greiff**
**Günter Klambauer** **Johannes Brandstetter**[*] **Sepp Hochreiter**

The transformer and BERT models pushed the performance on NLP tasks to new levels via their attention mechanism. We show that this attention mechanism is the update rule of a modern Hopfield network with continuous states. This new Hopfield network can store exponentially (with the dimension) many patterns, converges with one update, and has exponentially small retrieval errors. The number of stored patterns must be traded off against convergence speed and retrieval error.

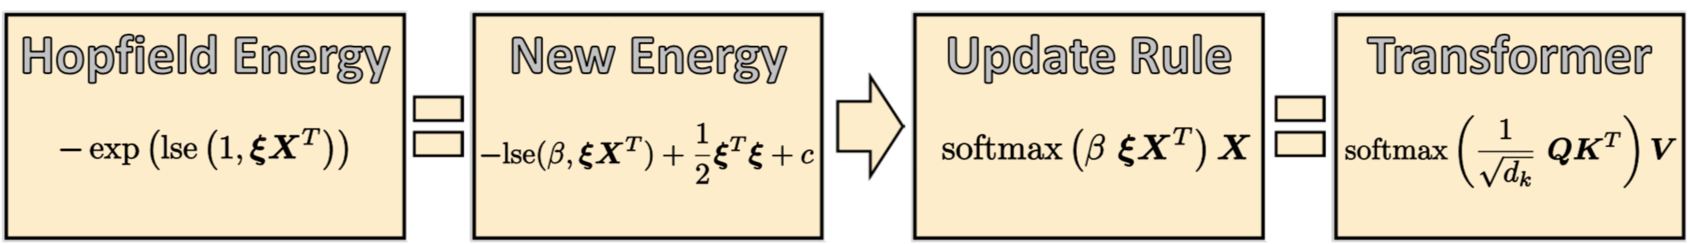


Figure 1: We generalized the energy of binary modern Hopfield networks for allowing continuous states while keeping convergence and storage capacity properties. We defined for the new energy also a new update rule that minimizes the energy. The new update rule is the attention mechanism of the transformer. Formulae are modified to express softmax as row vector as for transformers.

# NEURAL PROGRAMMER-INTERPRETER

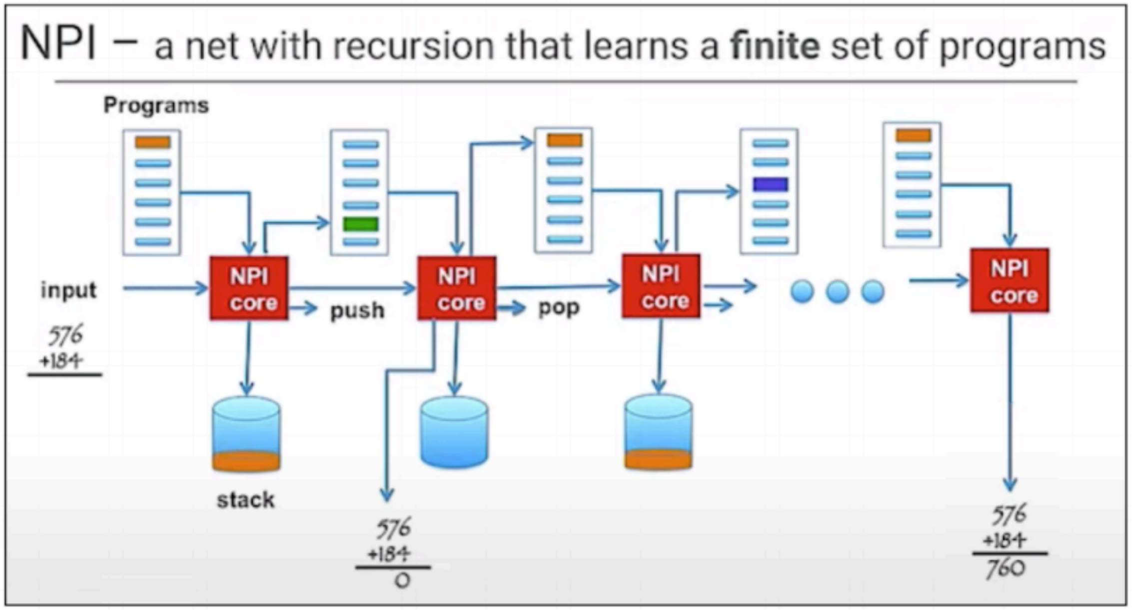


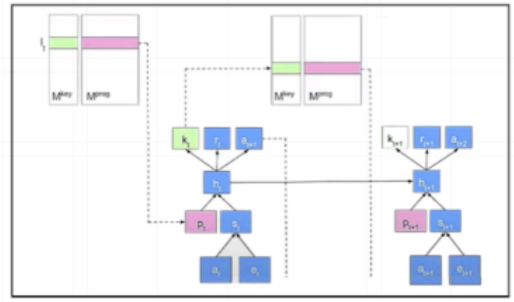

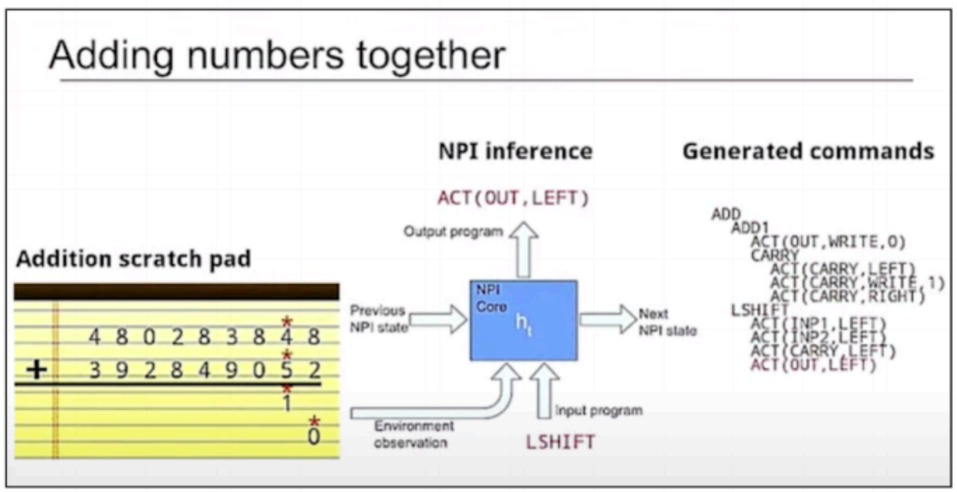


Scott E. Reed and Nando de Freitas. Neural programmer-interpreters. *CoRR*, arXiv:1511.06279, 2015.

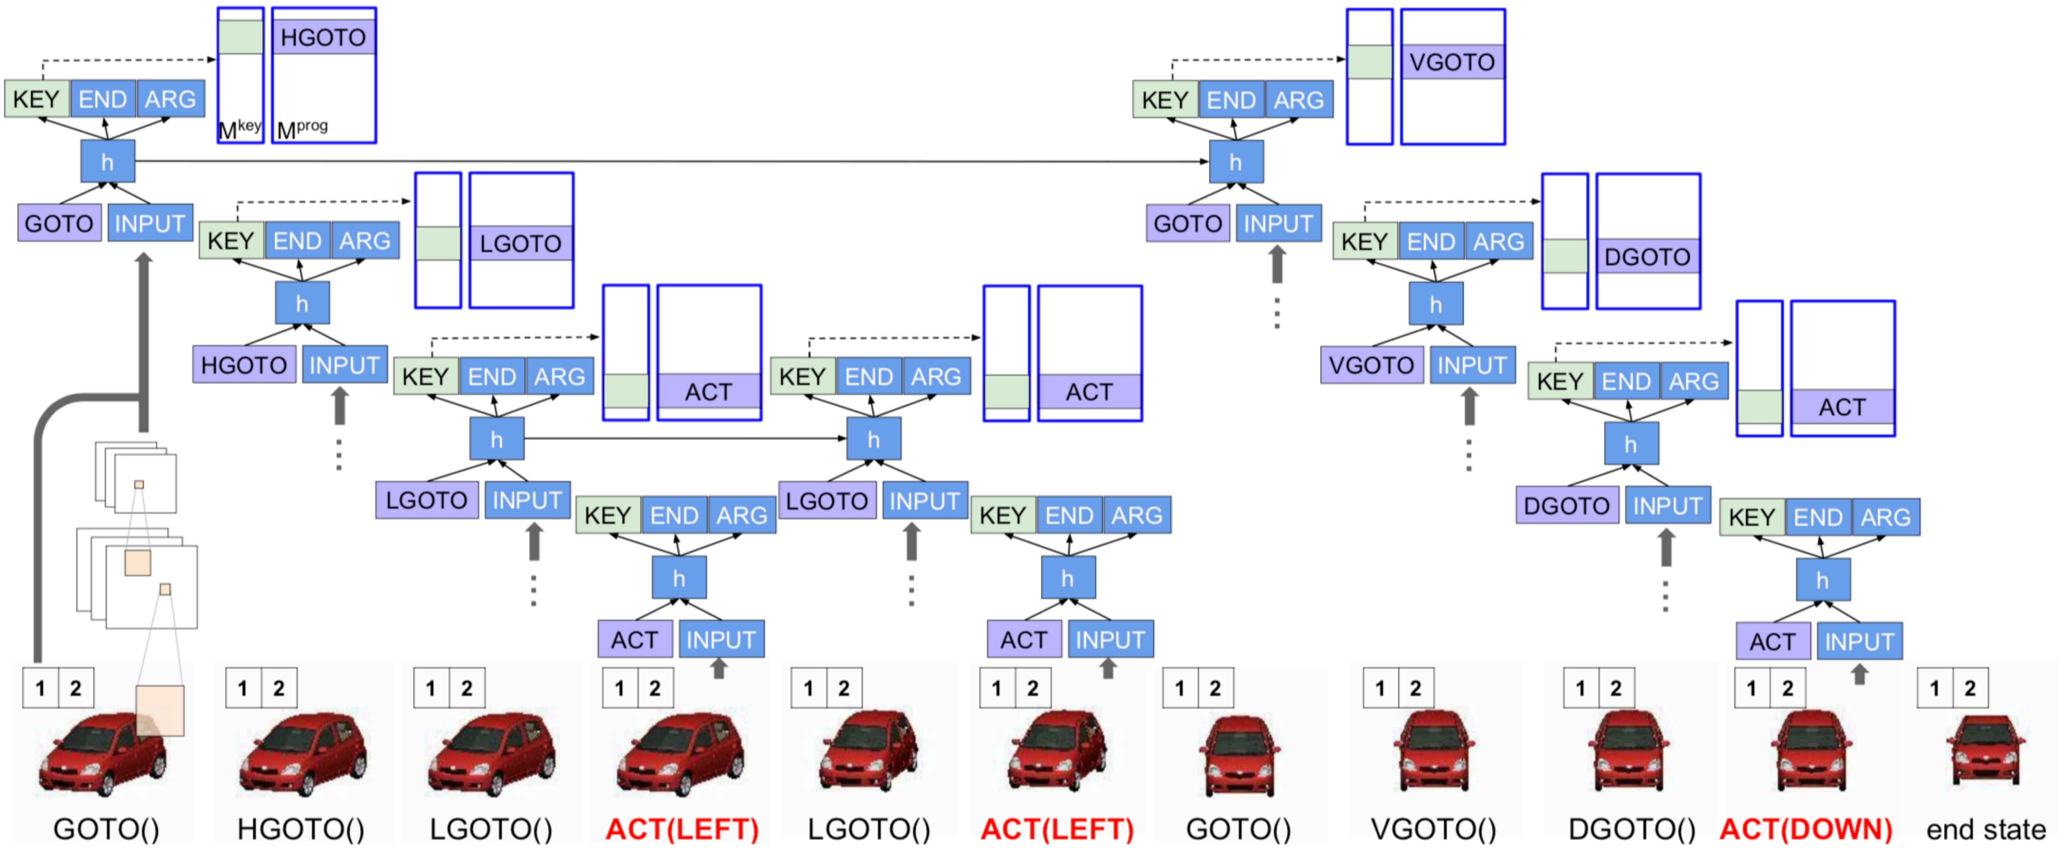
KEY
END
ARG
$M^{key}$
$M^{prog}$
HGOTO
h
GOTO
INPUT
KEY
END
ARG
LGOTO
h
HGOTO
INPUT
KEY
END
ARG
ACT
h
LGOTO
INPUT
KEY
END
ARG
h
ACT
INPUT
KEY
END
ARG
ACT
h
LGOTO
INPUT
KEY
END
ARG
h
ACT
INPUT
KEY
END
ARG
VGOTO
h
GOTO
INPUT
KEY
END
ARG
DGOTO
h
VGOTO
INPUT
KEY
END
ARG
ACT
h
DGOTO
INPUT
KEY
END
ARG
h
ACT
INPUT
1 2
GOTO()
HGOTO()
LGOTO()
ACT(LEFT)
LGOTO()
ACT(LEFT)
GOTO()
VGOTO()
DGOTO()
ACT(DOWN)
end state

=

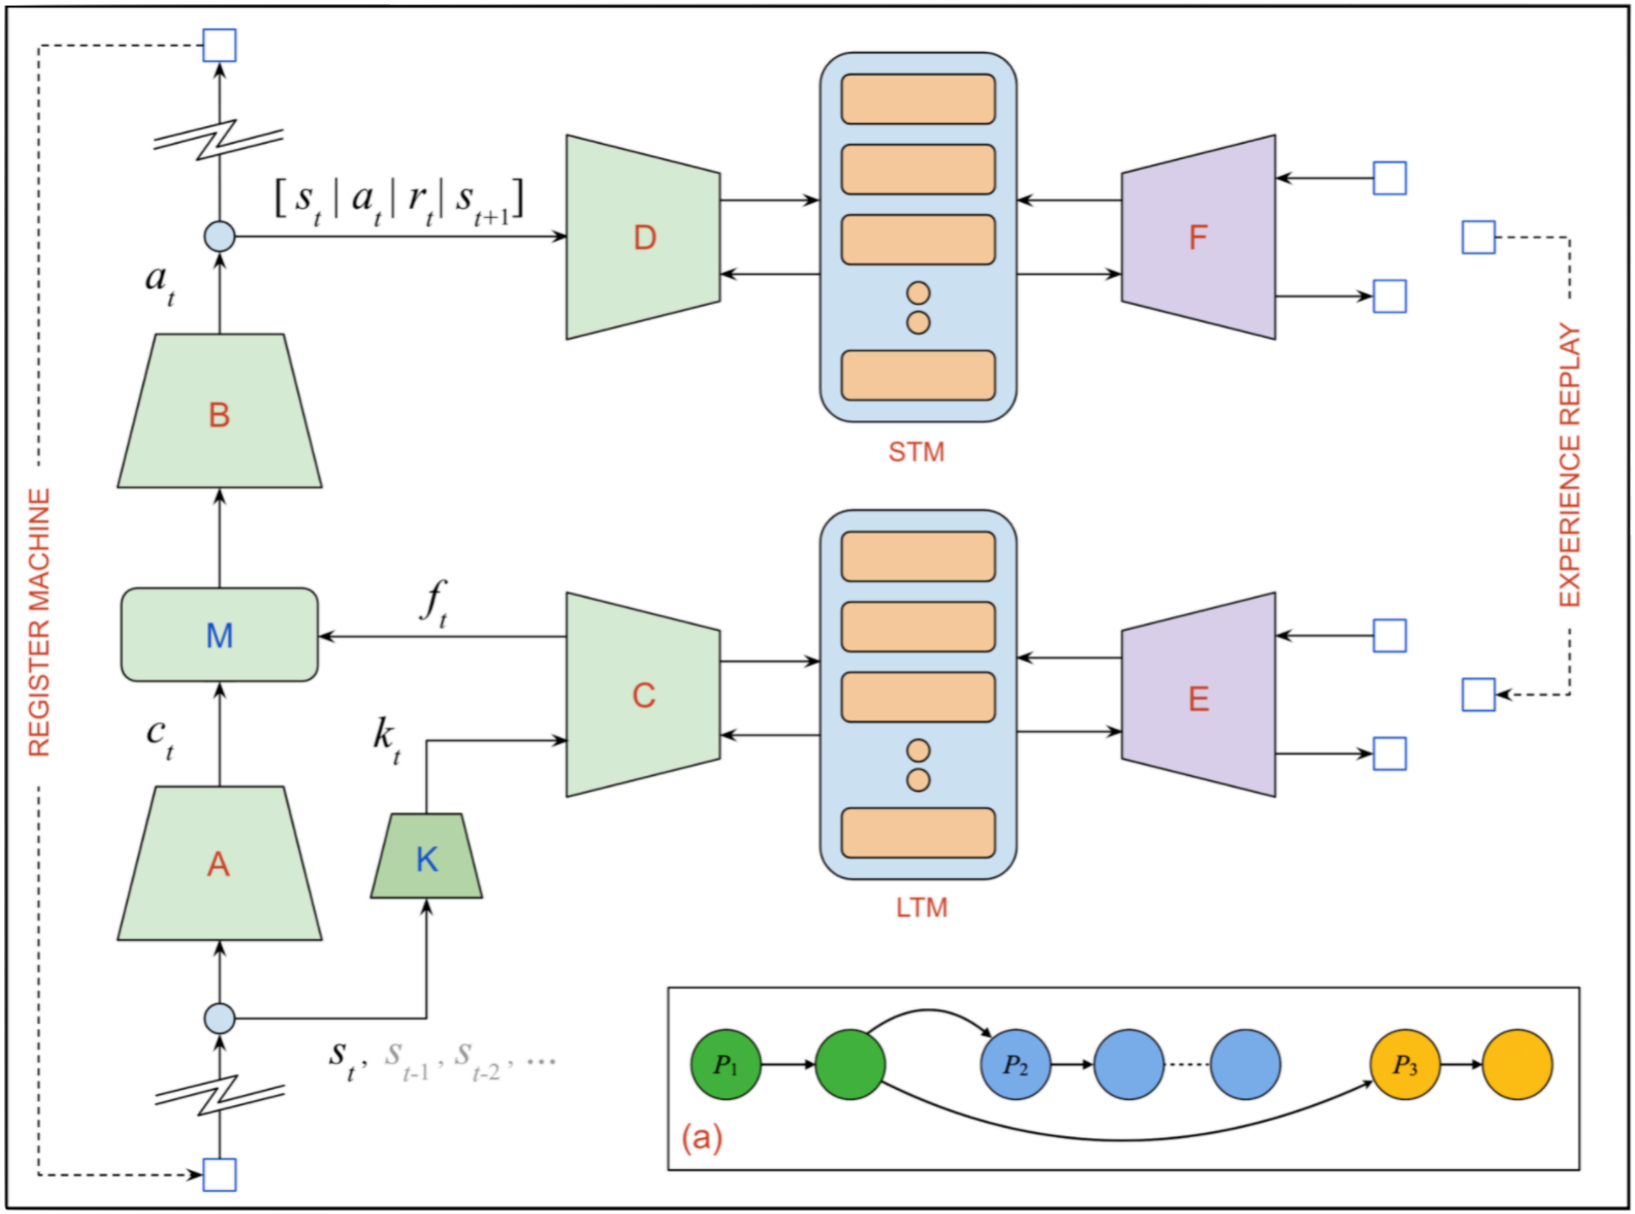


**See:** https://arxiv.org/pdf/1912.00421.pdf#page=23

# Modern Hopfield Networks for Few- and Zero-Shot Reaction Prediction

**Philipp Seidl** [1 2] **Philipp Renz** [1 2] **Natalia Dyubankova** [3] **Paulo Neves** [3] **Jonas Verhoeven** [3] **Jörg K. Wegner** [3]
**Sepp Hochreiter** [1 2 4] **Günter Klambauer** [1 2]

In drug discovery, a desired property can be the inhibition of a disease or a virus and in material science, thermal stability. From the design idea of the molecule, a virtual molecule is constructed, which enables to simulate or to predict the molecule's properties by the means of computational methods

However, to eventually test its hypothetical properties, the molecule has to be made physically available through chemical synthesis. The chemical synthesis problem, that is, how to assemble a given molecule with a series of chemical reactions, is a multi-step process with many possible choices at each step and, hence, highly complex. New molecules only come into physical existence if their synthesis route is known, otherwise, they are just an idea.

# Modern Hopfield Networks for Few- and Zero-Shot Reaction Prediction

**Philipp Seidl** [1,2] **Philipp Renz** [1,2] **Natalia Dyubankova** [3] **Paulo Neves** [3] **Jonas Verhoeven** [3] **Jörg K. Wegner** [3]
**Sepp Hochreiter** [1,2,4] **Günter Klambauer** [1,2]

An essential step in the discovery of new drugs and materials is the synthesis of a molecule that exists so far only as an idea to test its biological and physical properties. While computer-aided design of virtual molecules has made large progress, computer-assisted synthesis planning (CASP) to realize physical molecules is still in its infancy and lacks a performance level that would enable large-scale molecule discovery. CASP supports the search for multi-step synthesis routes, which is very challenging due to high branching factors in each synthesis step and the hidden rules that govern the reactions. The central and repeatedly applied step in CASP is reaction prediction, for which machine learning methods yield the best performance. We propose a novel reaction prediction approach that uses a deep learning architecture with modern Hopfield networks (MHNs) that is optimized by contrastive learning. An MHN is an associative memory that can store and retrieve chemical reactions in each layer of a deep learning architecture. We show that our MHN contrastive learning approach enables few- and zero-shot learning for reaction prediction which, in contrast to previous methods, can deal with rare, single, or even no training example(s) for a reaction.

# Modern Hopfield Networks for Few- and Zero-Shot Reaction Prediction

**Philipp Seidl** [1 2] **Philipp Renz** [1 2] **Natalia Dyubankova** [3] **Paulo Neves** [3] **Jonas Verhoeven** [3] **Jörg K. Wegner** [3]
**Sepp Hochreiter** [1 2 4] **Günter Klambauer** [1 2]

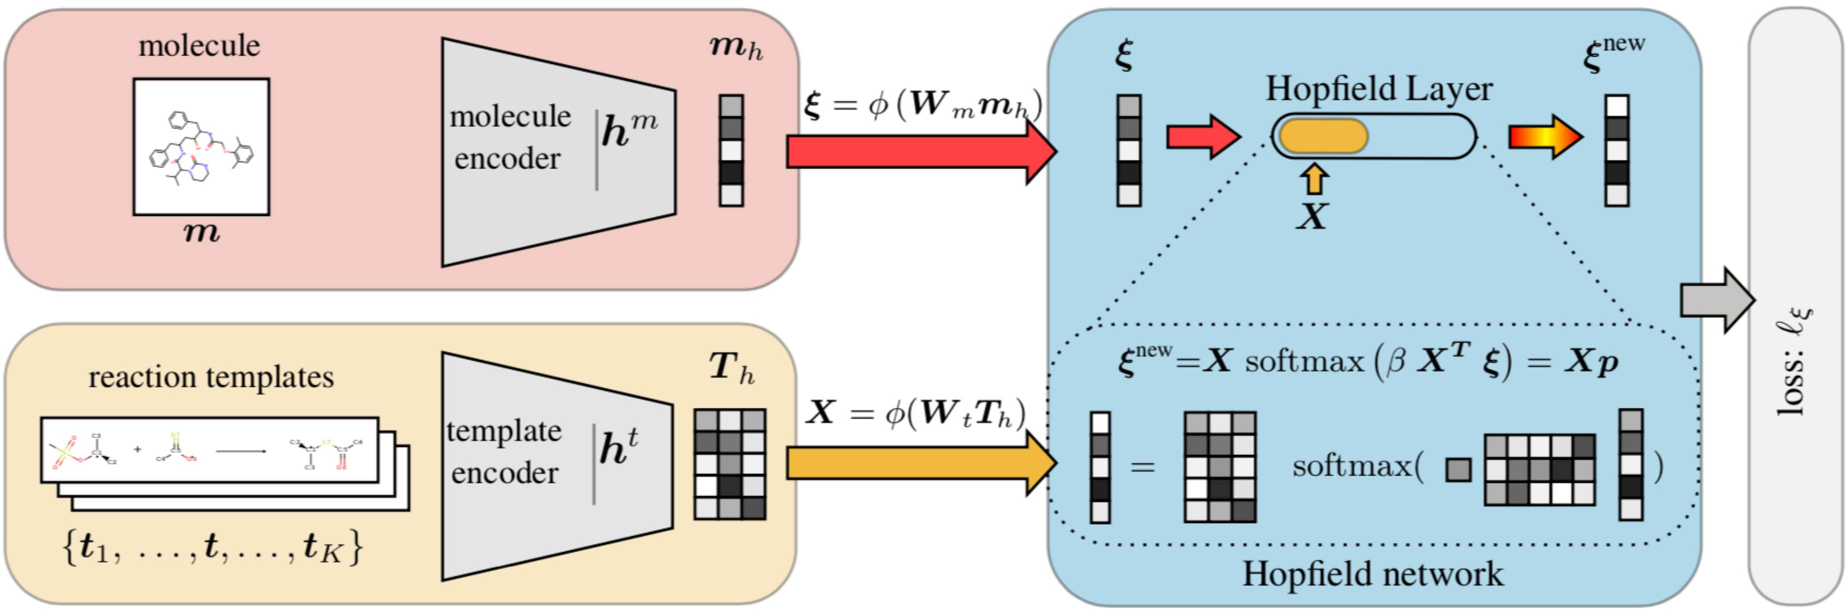


*Figure 2.* Simplified schematic representation of our approach. Standard approaches only encode the molecule and predict a fixed set of templates, in our MHN-based approach, the templates are also encoded and transformed to stored-pattern via the template encoder. The Hopfield-layer learns to associate the encoded input molecule, the state pattern $\boldsymbol{\xi}$, with the memory of encoded templates, the stored patterns $\boldsymbol{X}$.

In the CPC architecture, each input image is first divided into a grid of overlapping patches $\boldsymbol{x}_{i,j}$, where $i, j$ denote the location of the patch. Each patch is encoded with a neural network $f_\theta$ into a single vector $\boldsymbol{z}_{i,j} = f_\theta(\boldsymbol{x}_{i,j})$. To make predictions, a masked convolutional network $g_\phi$ is then applied to the grid of feature vectors. The masks are such that the receptive field of each resulting *context vector* $c_{i,j}$ only includes feature vectors that lie above it in the image (i.e. $c_{i,j} = g_\phi(\{\boldsymbol{z}_{u,v}\}_{u \le i,v})$). The prediction task then consists of predicting 'future' feature vectors $z_{i+k,j}$ from current context vectors $c_{i,j}$, where $k > 0$. The predictions are made linearly: given a context vector $c_{i,j}$, a prediction length $k > 0$, and a prediction matrix $W_k$, the predicted feature vector is $\hat{\boldsymbol{z}}_{i+k,j} = \boldsymbol{W}_k \boldsymbol{c}_{i,j}$.

The quality of this prediction is then evaluated using a contrastive loss. Specifically, the goal is to correctly recognize the target $z_{i+k,j}$ among a set of randomly sampled feature vectors $\{z_l\}$ from the dataset. We compute the probability assigned to the target using a softmax, and rate this probability using the usual cross-entropy loss. Summing this loss over locations and prediction offsets, we arrive at the CPC objective as defined in (van den Oord et al., 2018):

$$\mathcal{L}_{\text{CPC}} = -\sum_{i,j,k} \log p(\boldsymbol{z}_{i+k,j} | \hat{\boldsymbol{z}}_{i+k,j}, \{\boldsymbol{z}_l\})$$
$$= -\sum_{i,j,k} \log \frac{\exp(\hat{\boldsymbol{z}}_{i+k,j}^T \boldsymbol{z}_{i+k,j})}{\exp(\hat{\boldsymbol{z}}_{i+k,j}^T \boldsymbol{z}_{i+k,j}) + \sum_l \exp(\hat{\boldsymbol{z}}_{i+k,j}^T \boldsymbol{z}_l)}$$

The *negative samples* $\{\boldsymbol{z}_l\}$ are taken from other locations in the image and other images in the mini-batch. This loss is called InfoNCE as it is inspired by Noise-Contrastive Estimation (Gutmann & Hyvärinen, 2010; Mnih & Kavukcuoglu, 2013) and has been shown to maximize the mutual information between $\boldsymbol{c}_{i,j}$ and $\boldsymbol{z}_{i+k,j}$ (van den Oord et al., 2018).

- Surprisingly, this representation also surpasses supervised ResNets when given the entire ImageNet dataset (+3.2% Top-1 accuracy). Alternatively, our classifier is able to match fully-supervised ones while only using half of the labels.

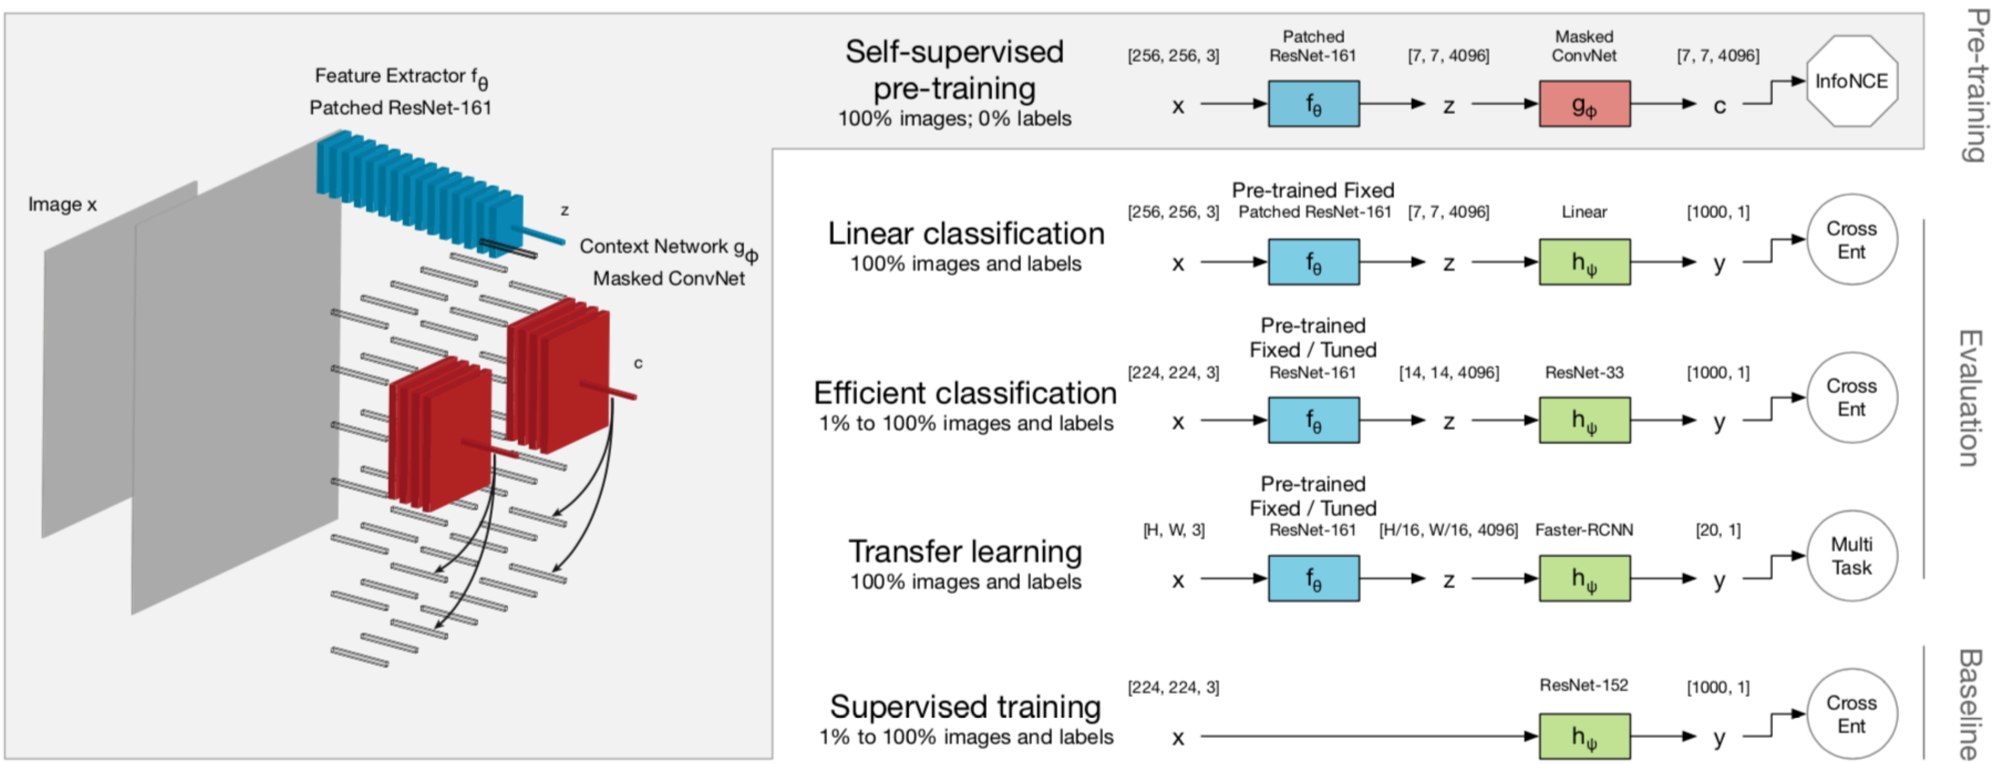


*Figure 2.* Overview of the framework for semi-supervised learning with Contrastive Predictive Coding. Left: unsupervised pre-training with the spatial prediction task (See Section 2.1). First, an image is divided into a grid of overlapping patches. Each patch is encoded independently from the rest with a feature extractor (blue) which terminates with a mean-pooling operation, yielding a single feature vector for that patch. Doing so for all patches yields a field of such feature vectors (wireframe vectors). Feature vectors above a certain level (in this case, the center of the image) are then aggregated with a context network (red), yielding a row of context vectors which are used to linearly predict features vectors below. Right: using the CPC representation for a classification task. Having trained the encoder network, the context network (red) is discarded and replaced by a classifier network (green) which can be trained in a supervised manner. In some experiments, we also fine-tune the encoder network (blue) for the classification task. When applying the encoder to cropped patches (as opposed to the full image) we refer to it as a *patched* ResNet in the figure.

| | |
|---|---|
| [1] | Bernard Baars, Stan Franklin, and Thomas Ramsy. Global workspace dynamics: Cortical "binding and propagation" enables conscious contents. *Frontiers in Psychology*, 4:200, 2013. |
| [2] | Samuel J. Gershman and Nathaniel D. Daw. Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Reviews of Psychology*, 68:101-128, 2017. |
| [3] | Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, and Aäron van den Oord. Data-efficient image recognition with contrastive predictive coding. *CoRR*, arXiv:1905.09272, 2019. |
| [4] | Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512--534, 2016. |
| [5] | George A. Mashour, Pieter Roelfsema, Jean-Pierre Changeux, and Stanislas Dehaene. Conscious processing and the global neuronal workspace hypothesis. *Neuron*, 105(5):776-798, 2020. |
| [6] | Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Sandve, Victor Greiff, David Kreil, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. Hopfield networks is all |
| [7] | Philipp Seidl, Philipp Renz, Natalia Dyubankova, Paulo Neves, Jonas Verhoeven, Jörg K. Wegner, Sepp Hochreiter, and Gunter Klambauer. Modern hopfield networks for few- and zero-shot reaction prediction. *CoRR*, arXiv:2104.03279, 2021. |
| [8] | Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *CoRR*, arXiv:1711.00937, 2017. |
| [9] | Representation Learning with Contrastive Predictive Coding. Aaron van den oord and yazhe li and oriol vinyals. *CoRR*, arXiv:1807.03748, 2018. |