MONTRÉAL.AI ACADEMY: ARTIFICIAL INTELLIGENCE 101 FIRST WORLD-CLASS OVERVIEW OF AI FOR ALL **VIP AI 101 CHEATSHEET**

A PREPRINT

Vincent Boucher[*]
MONTRÉAL.AI
Montreal, Quebec, Canada
info@montreal.ai

June 8, 2021

ABSTRACT

For the purpose of entrusting all sentient beings with powerful AI tools to learn, deploy and scale AI in order to enhance their prosperity, to settle planetary-scale problems and to inspire those who, with AI, will shape the 21st Century, **MONTRÉAL.AI** introduces this ***VIP AI 101 CheatSheet*** for All.

**MONTRÉAL.AI is preparing a global network of education centers.*
*****ALL OF EDUCATION, FOR ALL**. MONTRÉAL.AI is developing a teacher (**Saraswati AI**) and an agent learning to orchestrate synergies amongst academic disciplines (**Polymatheia AI**).*

Curated Open-Source Codes and Science: http://www.academy.montreal.ai/.

***K*eywords** AI-First · Artificial Intelligence · Deep Learning · Reinforcement Learning · Symbolic AI

1 AI-First

TODAY'S ARTIFICIAL INTELLIGENCE IS POWERFUL AND ACCESSIBLE TO ALL. AI is capable of transforming industries and opens up a world of new possibilities. **What's important is what you do with AI and how you embrace it**. To pioneer AI-First innovations advantages: start by exploring how to apply AI in ways never thought of.

The Emerging Rules of the AI-First Era: Search and Learning.

"***Search** and **learning** are general purpose methods that continue to scale with increased computation, even as the available computation becomes very great.*" — Richard Sutton in *The Bitter Lesson*

The Best Way Forward For AI[2].

"*... so far as I'm concerned, system 1 certainly knows language, understands language... system 2... it does involve certain manipulation of symbols... Gary Marcus ... Gary proposes something that seems very natural... **a hybrid architecture**... I'm influenced by him... if you look introspectively at the way the mind works... you'd get to that distinction between implicit and explicit... explicit looks like symbols.*" — Nobel Laureate Danny Kahneman at *AAAI-20 Fireside Chat with Daniel Kahneman* https://vimeo.com/390814190

In *The Next Decade in AI*[3], Gary Marcus proposes a hybrid, knowledge-driven, reasoning-based approach, centered around cognitive models, that could provide the substrate for a richer, more robust AI than is currently possible.

[*]Founding Chairman at MONTRÉAL.AI http://www.montreal.ai and QUÉBEC.AI http://www.quebec.ai.
[2]https://montrealartificialintelligence.com/aidebate/
[3]https://arxiv.org/abs/2002.06177v3

2 Getting Started

"*It takes a village to raise an AI that's ethical, robust, and trustworthy.*" — Gary Marcus

Figure 1: **AI DEBATE 2**: *Moving AI Forward*. Official Video: https://youtu.be/VOI3Bb3p4GM

The Measure of Intelligence (*Abstraction and Reasoning Corpus*[4]) https://arxiv.org/abs/1911.01547.

- **AI Paygrades** https://aipaygrad.es/.
- **CS231n Python Tutorial With Google Colab**[5].
- **HP TECH TAKES: Machine Learning - AI for Kids**[6].
- **Learn with Google AI** https://ai.google/education/.
- **Made With ML Topics** https://madewithml.com/topics/.
- **Papers with Datasets** https://paperswithcode.com/datasets.
- **One Place for Everything AI** https://aihub.cloud.google.com/.
- **Deep Learning Drizzle** https://deep-learning-drizzle.github.io.
- **Google Dataset Search** (Blog[7]) https://datasetsearch.research.google.com.
- **AI Literacy for K-12 School Children** https://aieducation.mit.edu/resources.
- **Learning resources from DeepMind** https://deepmind.com/learning-resources.
- **Papers With Code** (*Learn Python 3 in Y minutes*[8]) https://paperswithcode.com/state-of-the-art.

Tinker with neural networks in the browser with *TensorFlow Playground* http://playground.tensorflow.org/.

"*(Google) Dataset Search*[9] *has indexed almost 25 million of these datasets, giving you a single place to search for datasets and find links to where the data is.*" — Natasha Noy

[4] https://github.com/fchollet/ARC
[5] https://colab.research.google.com/github/cs231n/cs231n.github.io/blob/master/python-colab.ipynb
[6] https://www.hp.com/us-en/shop/tech-takes/computer-education-machine-learning-ai-for-kids
[7] https://blog.google/products/search/discovering-millions-datasets-web/
[8] https://learnxinyminutes.com/docs/python3/
[9] https://datasetsearch.research.google.com

2.1 In the Cloud

Colab [10]. Practice Immediately [11]. Labs[12]: Introduction to Deep Learning (MIT 6.S191)

- Free GPU compute via Colab `https://colab.research.google.com/notebooks/welcome.ipynb`.
- Colab can open notebooks directly from GitHub by simply replacing "*http://github.com*" with "*http://colab.research.google.com/github/*" in the notebook URL.
- Colab Pro `https://colab.research.google.com/signup`.

2.2 On a Local Machine

JupyterLab is an interactive development environment for working with notebooks, code and data [13].

- Install Anaconda `https://www.anaconda.com/download/` and launch '*Anaconda Navigator*'
- Update Jupyterlab and launch the application. Under Notebook, click on '*Python 3*'

IDE: Visual Studio Code `https://code.visualstudio.com/`.

"*If we truly reach AI, it will let us know.*" — Garry Kasparov

3 Deep Learning

Learning according to Mitchell (1997):

"*A computer program is said to learn from experience* **E** *with respect to some class of tasks* **T** *and performance measure* **P***, if its performance at tasks in* **T***, as measured by* **P***, improves with experience* **E***.*" — Tom Mitchell

After the **Historical AI Debate**[14]: "*Yoshua Bengio and Gary Marcus on the Best Way Forward for AI*" `https://montrealartificialintelligence.com/aidebate/`, there have been clarifications on the term "***deep learning***"[15].

"*Deep learning is inspired by neural networks of the brain to build learning machines which discover rich and useful internal representations, computed as a composition of learned features and functions.*" — Yoshua Bengio

"*DL is constructing networks of parameterized functional modules and training them from examples using gradient-based optimization.*" — Yann LeCun

"*... replace symbols by vectors and logic by continuous (or differentiable) functions.*" — Yann LeCun

Deep learning allows computational models that are composed of multiple processing layers to learn REPRESENTATIONS of (raw) data with multiple levels of abstraction[2]. At a high-level, neural networks are either encoders, decoders, or a combination of both[16]. Introductory course `http://introtodeeplearning.com`. See also Table 1.

Table 1: Types of Learning, by Alex Graves at NeurIPS 2018

Name	With Teacher	Without Teacher
Active	*Reinforcement Learning / Active Learning*	*Intrinsic Motivation / Exploration*
Passive	*Supervised Learning*	*Unsupervised Learning*

Deep learning assumes that the data was generated by the composition of factors potentially at multiple levels in a hierarchy[17]. Deep learning (*distributed representations + composition*) is a general-purpose learning procedure.

[10] `https://medium.com/tensorflow/colab-an-easy-way-to-learn-and-use-tensorflow-d74d1686e309`
[11] `https://colab.research.google.com/github/madewithml/practicalAI/`
[12] `https://colab.research.google.com/github/aamini/introtodeeplearning_labs`
[13] `https://blog.jupyter.org/jupyterlab-is-ready-for-users-5a6f039b8906`
[14] `https://www.zdnet.com/article/devils-in-the-details-in-bengio-marcus-ai-debate/`
[15] `https://www.zdnet.com/article/whats-in-a-name-the-deep-learning-debate/`
[16] `https://github.com/lexfridman/mit-deep-learning`
[17] `https://www.deeplearningbook.org`

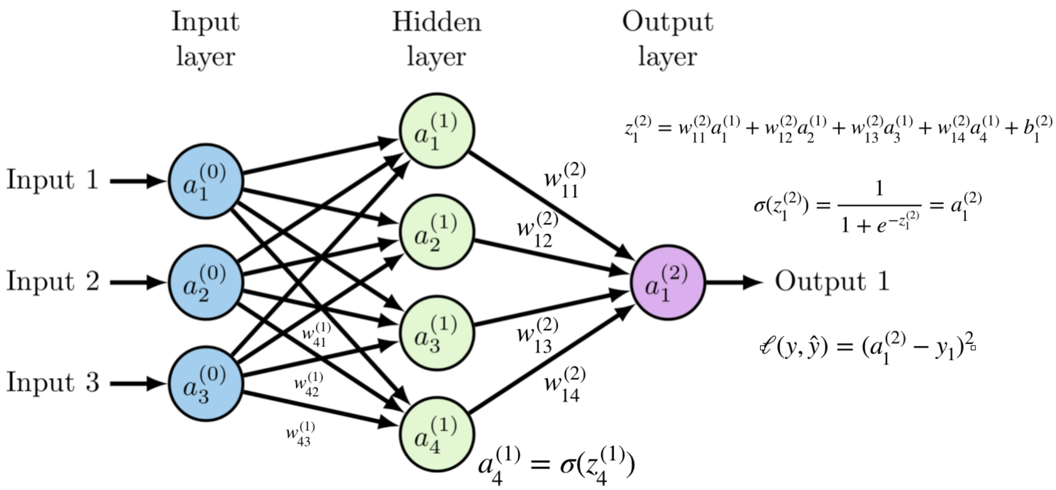

Figure 2: **Multilayer perceptron (MLP)**.

"*When you first study a field, it seems like you have to memorize a zillion things. You don't. What you need is to identify the 3-5 core principles that govern the field. The million things you thought you had to memorize are various combinations of the core principles.*" — J. Reed

"*1. Multiply things together*
2. Add them up
3. Replaces negatives with zeros
4. Return to step 1, a hundred times."
— Jeremy Howard

❖ Linear Algebra. Prof. Gilbert Strang[18].
❖ Dive into Deep Learning `http://d2l.ai`.
❖ Minicourse in Deep Learning with PyTorch[19].
❖ How to do Research At the MIT AI Lab (1988)[20].
❖ Introduction to Artificial Intelligence, Gilles Louppe[21].
❖ Energy-Based Models for Continual Learning, Anonymous[22].
❖ Fast and Easy Infinitely Wide Networks with Neural Tangents[23].
❖ Deep Learning. The full deck of (600+) slides, Gilles Louppe[24].
❖ Design Space for Graph Neural Networks, You et al.[25]. Code[26].
❖ These Lyrics Do Not Exist `https://theselyricsdonotexist.com`.
❖ AI and Wargaming, Goodman et al. `https://arxiv.org/abs/2009.08922`.
❖ Backward Feature Correction: How Deep Learning Performs Deep Learning[27].
❖ A Selective Overview of Deep Learning `https://arxiv.org/abs/1904.05526`.
❖ The Missing Semester of Your CS Education `https://missing.csail.mit.edu`.
❖ fastai: A Layered API for Deep Learning `https://arxiv.org/abs/2002.04688`.
❖ Thinking Fast and Slow in AI, Booch et al. `https://arxiv.org/abs/2010.06002`.
❖ Anatomy of Matplotlib `https://github.com/matplotlib/AnatomyOfMatplotlib`.
❖ Data project checklist `https://www.fast.ai/2020/01/07/data-questionnaire/`.

[18] `https://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/`
[19] `https://github.com/Atcold/pytorch-Deep-Learning-Minicourse`
[20] `http://dspace.mit.edu/bitstream/handle/1721.1/41487/AI_WP_316.pdf`
[21] `https://glouppe.github.io/info8006-introduction-to-ai/pdf/lec-all.pdf`
[22] `https://openreview.net/forum?id=j5d9qacxdZa`
[23] `https://ai.googleblog.com/2020/03/fast-and-easy-infinitely-wide-networks.html`
[24] `https://github.com/glouppe/info8010-deep-learning/raw/v2-info8010-2019/pdf/lec-all.pdf`
[25] `https://arxiv.org/abs/2011.08843`
[26] `https://github.com/snap-stanford/GraphGym`
[27] `https://arxiv.org/abs/2001.04413`

❖ Using Nucleus and TensorFlow for DNA Sequencing Error Correction, Colab Notebook[28].
❖ Machine Learning for Physicists https://machine-learning-for-physicists.org.
❖ Flow-edge Guided Video Completion, Gao et al. https://arxiv.org/abs/2009.01835.
❖ The world as a neural network, Vitaly Vanchurin https://arxiv.org/abs/2008.01540.
❖ Generalized Energy Based Models, Michael Arbel, Liang Zhou and Arthur Gretton, 2020[29].
❖ Representing Scenes as Neural Radiance Fields for View Synthesis. Mildenhall et al., 2020[30].
❖ PoseNet Sketchbook https://googlecreativelab.github.io/posenet-sketchbook/.
❖ The Neural Network, A Visual Introduction. Vivek Verma : https://youtu.be/UOvPeC8WOt8.
❖ Synthetic Data for Deep Learning, Sergey I. Nikolenko https://arxiv.org/abs/1909.11512.
❖ Removing people from complex backgrounds in real time using TensorFlow.js in the web browser[31].
❖ A Recipe for Training Neural Networks https://karpathy.github.io/2019/04/25/recipe/.
❖ TensorFlow Datasets: load a variety of public datasets into TensorFlow programs (Blog[32] | Colab[33]).
❖ Denoising Diffusion Probabilistic Models, Ho et al., 2020 https://arxiv.org/abs/2006.11239.
❖ The Markov-Chain Monte Carlo Interactive Gallery https://chi-feng.github.io/mcmc-demo/.
❖ NeurIPS 2019 Implementations https://paperswithcode.com/conference/neurips-2019-12.
❖ Involutive MCMC: a Unifying Framework, Neklyudov et al. https://arxiv.org/abs/2006.16653.
❖ Interpretable Machine Learning – A Brief History, State-of-the-Art and Challenges. Molnar et al., 2020[34].
❖ Politeness Transfer: A Tag and Generate Approach, Madaan et al. https://arxiv.org/abs/2004.14257.
❖ Algebra, Topology, Differential Calculus, and Optimization Theory For Computer Science and Machine Learning[35].
❖ How to Choose Your First AI Project https://hbr.org/2019/02/how-to-choose-your-first-ai-project.
❖ Technology Readiness Levels for Machine Learning Systems, Lavin et al. https://arxiv.org/abs/2101.03989.
❖ Blog | MIT 6.S191 https://medium.com/tensorflow/mit-introduction-to-deep-learning-4a6f8dde1f0c.
❖ A Fortran-Keras Deep Learning Bridge for Scientific Computing, Ott et al. https://arxiv.org/abs/2004.10652. GitHub[36].
❖ A Wholistic View of Continual Learning with Deep Neural Networks: Forgotten Lessons and the Bridge to Active and Open World Learning. Mundt et al., 2020[37].

3.1 Universal Approximation Theorem

The universal approximation theorem states that a feed-forward network with a single hidden layer containing a finite number of neurons can solve any given problem to arbitrarily close accuracy as long as you add enough parameters.

Neural Networks + Gradient Descent + GPU[38]:

- Infinitely flexible function: *Neural Network* (multiple hidden layers: Deep Learning)[39].
- All-purpose parameter fitting: *Backpropagation*[40][41]. Backpropagation is the key algorithm that makes training deep models computationally tractable and highly efficient[42]. The backpropagation procedure is nothing more than a practical application of the chain rule for derivatives.
- Fast and scalable: *GPU*.

"*You have relatively simple processing elements that are very loosely models of neurons. They have connections coming in, each connection has a weight on it, and that weight can be changed through learning.*" — Geoffrey Hinton

[28] https://colab.research.google.com/github/google/nucleus/blob/master/nucleus/examples/dna_sequencing_error_correction.ipynb
[29] https://arxiv.org/abs/2003.05033
[30] http://www.matthewtancik.com/nerf
[31] https://github.com/jasonmayes/Real-Time-Person-Removal
[32] https://medium.com/tensorflow/introducing-tensorflow-datasets-c7f01f7e19f3
[33] https://colab.research.google.com/github/tensorflow/datasets/blob/master/docs/overview.ipynb
[34] https://arxiv.org/abs/2010.09337
[35] https://drive.google.com/file/d/1sJvLQwxMyu89t2z4Zf9tD7O7efnbIUyB/view
[36] https://github.com/scientific-computing/FKB
[37] https://arxiv.org/abs/2009.01797
[38] http://wiki.fast.ai/index.php/Lesson_1_Notes
[39] http://neuralnetworksanddeeplearning.com/chap4.html
[40] https://github.com/DebPanigrahi/Machine-Learning/blob/master/back_prop.ipynb
[41] https://www.jeremyjordan.me/neural-networks-training/
[42] https://colah.github.io/posts/2015-08-Backprop/

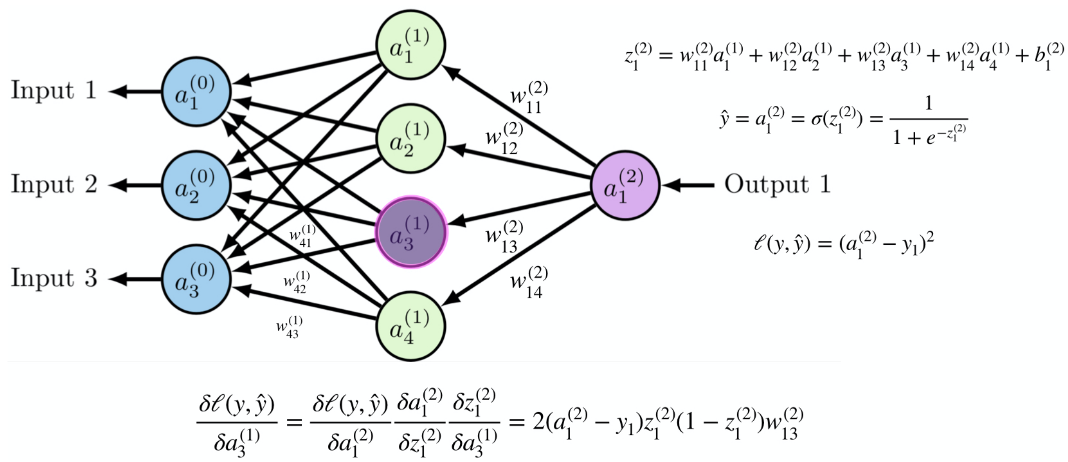

Figure 3: **All-purpose parameter fitting: Backpropagation**.

Deep learning : connect a dataset, a model, a cost function and an optimization procedure.

"*Deep learning has fully solved the curse of dimensionality. It vanished like an RNN gradient!*" — Ilya Sutskever

When a choice must be made, just feed the (raw) data to a deep neural network (Universal function approximators).

"*Here is my own deep: "**DEEP UNDERSTANDING**" with very clear definition: A mathematical object that supports reasoning across all 3 levels of the causal hierarchy.*" — Judea Pearl

3.2 Convolution Neural Networks (Useful for Images | Space)

Richer innate priors : *innateness that enables learning.*

A significant percentage of Deep Learning breakthroughs comes from reusable constructs and parameters sharing. The deep convolutional network is a construct that reuses weights in multiple locations (parameters sharing in space)[43].

"*Virtually all modern observers would concede that genes and experience work together; it is "nature and nurture", not "nature versus nurture". No nativist, for instance, would doubt that we are also born with specific biological machinery that allows us to learn. Chomsky's Language Acquisition Device should be viewed precisely as an innate learning mechanism, and nativists such as Pinker, Peter Marler (Marler, 2004) and myself (Marcus, 2004) have frequently argued for a view in which a significant part of a creature's innate armamentarium consists not of specific knowledge but of learning mechanisms, a form of **innateness that enables learning.***" — Gary Marcus, Innateness, AlphaZero, and Artificial Intelligence[44]

The deep convolutional network, inspired by Hubel and Wiesel's seminal work on early visual cortex, uses hierarchical layers of tiled convolutional filters to mimic the effects of receptive fields, thereby exploiting the local spatial correlations present in images[1]. See Figure 4. Demo `https://ml4a.github.io/demos/convolution/`.

A ConvNet is made up of Layers. Every Layer has a simple API: It transforms an input 3D volume to an output 3D volume with some differentiable function that may or may not have parameters[45]. Reading[46].

In images, local combinations of edges form motifs, motifs assemble into parts, and parts form objects[47][48].

Representation learning : *the language of neural networks.* The visual vocabulary of a convolutional neural network seems to emerge from low level features such as edges and orientations, and builds up textures, patterns and composites,

[43] `https://twitter.com/iamtrask/status/949439556499230720`
[44] `https://arxiv.org/abs/1801.05667`
[45] `http://cs231n.github.io/convolutional-networks/`
[46] `https://ml4a.github.io/ml4a/convnets/`
[47] `http://yosinski.com/deepvis`
[48] `https://distill.pub/2017/feature-visualization/`

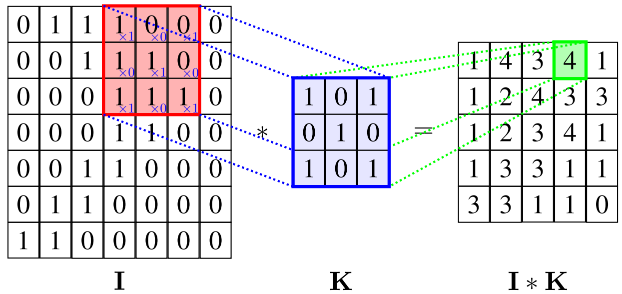

Figure 4: **2D Convolution**. Source: Cambridge Coding Academy

. . . and builds up even further into complete objects. This relates to Wittgenstein's "*language-game*" in ***Philosophical Investigations***[49], where a functional language emerge from simple tasks before defining a vocabulary[50].

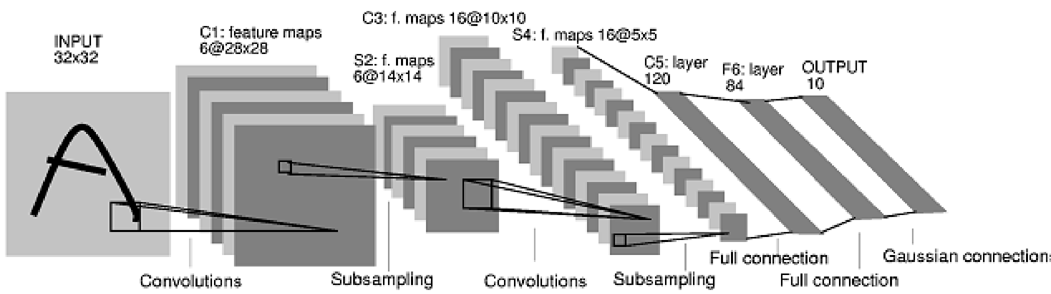

Figure 5: Architecture of LeNet-5, a Convolutional Neural Network. LeCun et al., 1998

"*DL is essentially a new style of programming – "differentiable programming" – and the field is trying to work out the reusable constructs in this style. We have some: convolution, pooling, LSTM, GAN, VAE, memory units, routing units, etc.*" — Thomas G. Dietterich

❖ Image Classification from Scratch[51].
❖ CS231N : Convolutional Neural Networks for Visual Recognition[52].
❖ Introduction to Graph Convolutional Network (GCN). Alfredo Canziani[53].
❖ Deep Plastic Surgery: Robust and Controllable Image Editing with Human-Drawn Sketches. Yang et al.[54].
❖ CNN Explainer: Learning Convolutional Neural Networks with Interactive Visualization. Wang et al. [55] [56].
❖ An Overview of Early Vision in InceptionV1 `https://distill.pub/2020/circuits/early-vision/`.
❖ Neural Voice Puppetry: Audio-driven Facial Reenactment. Thies et al. `https://arxiv.org/abs/1912.05566`.
❖ TensorSpace (`https://tensorspace.org`) offers interactive 3D visualizations of *LeNet*, *AlexNet* and *Inceptionv3*.

3.3 Recurrent Neural Networks (Useful for Sequences | Time)

Recurrent neural networks are networks with loops in them, allowing information to persist[57]. RNNs process an input sequence one element at a time, maintaining in their hidden units a 'state vector' that implicitly contains information about the history of all the past elements of the sequence[2]. For sequential inputs. See Figure 6.

[49] https://en.wikipedia.org/wiki/Philosophical_Investigations
[50] https://media.neurips.cc/Conferences/NIPS2018/Slides/Deep_Unsupervised_Learning.pdf
[51] https://colab.research.google.com/drive/1umJnCp8tZ7UDTYSQsuWdKRhqbHts38AC
[52] https://www.youtube.com/playlist?list=PLzUTmXVwsnXod6WNdg57Yc3zFx_f-RYsq
[53] https://atcold.github.io/pytorch-Deep-Learning/en/week13/13-3/
[54] https://arxiv.org/abs/2001.02890
[55] https://arxiv.org/abs/2004.15004
[56] http://poloclub.github.io/cnn-explainer/
[57] http://colah.github.io/posts/2015-08-Understanding-LSTMs/

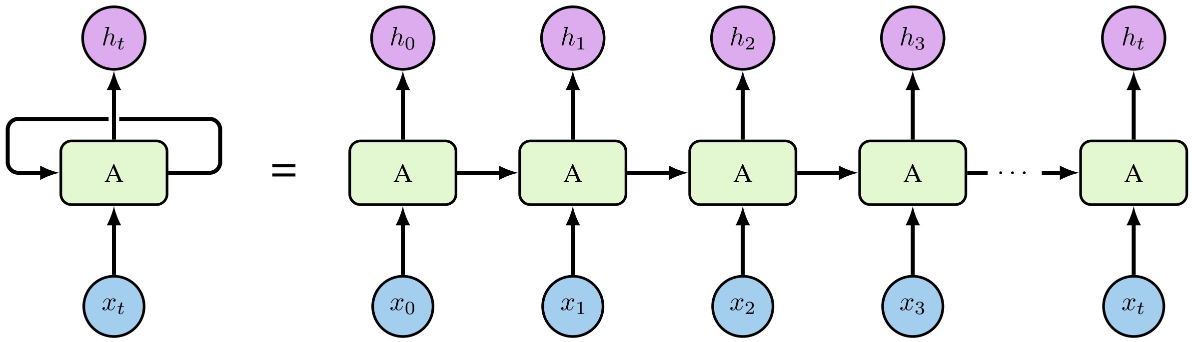

Figure 6: **RNN Layers Reuse Weights for Multiple Timesteps**.

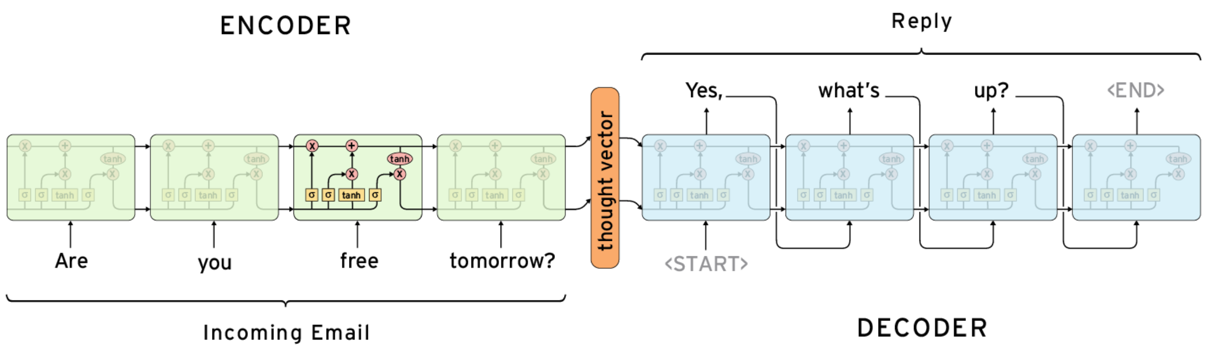

Figure 7: **Google *Smart Reply System* is built on a pair of recurrent neural networks**. Diagram by Chris Olah

"I feel like a significant percentage of Deep Learning breakthroughs ask the question "how can I reuse weights in multiple places?" – Recurrent (LSTM) layers reuse for multiple timesteps – Convolutional layers reuse in multiple locations. – Capsules reuse across orientation." — Andrew Trask

❖ CS224N : Natural Language Processing with Deep Learning[58].
❖ Long Short-Term-Memory (LSTM), Sepp Hochreiter and Jürgen Schmidhuber[59].
❖ The Unreasonable Effectiveness of Recurrent Neural Networks, blog (2015) by Andrej Karpathy[60].
❖ Understanding LSTM Networks `http://colah.github.io/posts/2015-08-Understanding-LSTMs/`.
❖ Can Neural Networks Remember? Slides by Vishal Gupta: `http://vishalgupta.me/deck/char_lstms/`.

3.4 Transformers

Transformers are generic, simples and exciting machine learning architectures designed to process a connected set of units (tokens in a sequence, pixels in an image, etc.) where the only interaction between units is through self-attention. Transformers' performance limit seems purely in the hardware (how big a model can be fitted in GPU memory)[61].

The fundamental operation of transformers is **self-attention** (a sequence-to-sequence operation, Figure 9): *an attention mechanism relating different positions of a single sequence in order to compute a representation of the same sequence*[62].

Let's call the input vectors (of dimension k) :

$$\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_t \tag{1}$$

Let's call the corresponding output vectors (of dimension k) :

$$\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_t \tag{2}$$

[58] `https://www.youtube.com/playlist?list=PLU40WL8Ol94IJzQtileLTqGZuXtGlLMP_`
[59] `https://www.bioinf.jku.at/publications/older/2604.pdf`
[60] `http://karpathy.github.io/2015/05/21/rnn-effectiveness/`
[61] `http://www.peterbloem.nl/blog/transformers`
[62] `https://lilianweng.github.io/lil-log/2018/06/24/attention-attention.html`

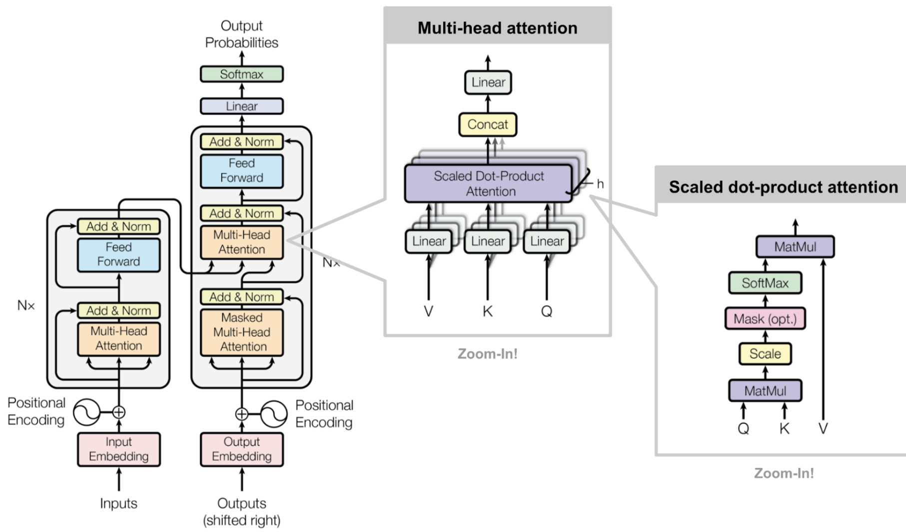

Figure 8: Attention Is All You Need. *Vaswani et al., 2017* : `https://arxiv.org/abs/1706.03762`.

The **self attention** operation takes a weighted average over all the input vectors :

$$\mathbf{y}_i = \sum_j w_{ij} \mathbf{x}_j \tag{3}$$

The weight w_{ij} is derived from a function over $\mathbf{x}_{\mathrm{i}}$ and $\mathbf{x}_{\mathrm{j}}$. The simplest option is the dot product (with softmax) :

$$w_{ij} = \frac{e^{\mathbf{x}_i^T \mathbf{x}_j}}{\sum_j e^{\mathbf{x}_i^T \mathbf{x}_j}} \tag{4}$$

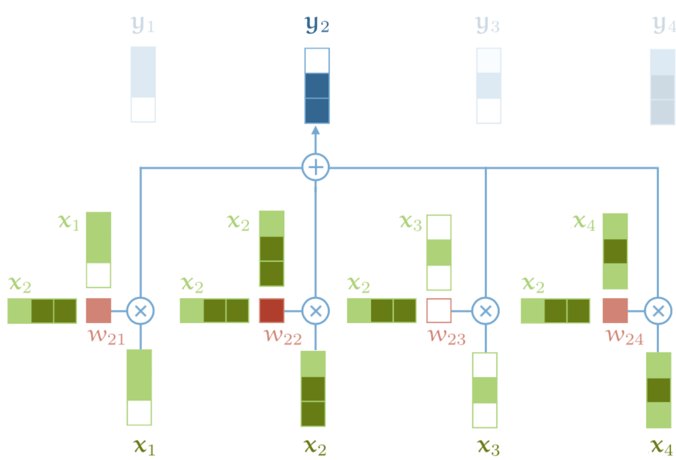

Figure 9: Self-attention. By *Peter Bloem* : `http://www.peterbloem.nl/blog/transformers`.

Transformers are Graph Neural Networks[63].

❖ The Transformer Family. By *Lilian Weng*[64].
❖ Transformers Notebooks. By *Hugging Face*[65].

[63] `https://graphdeeplearning.github.io/post/transformers-are-gnns/`
[64] `https://lilianweng.github.io/lil-log/2020/04/07/the-transformer-family.html`
[65] `https://github.com/huggingface/transformers/tree/master/notebooks`

❖ Text classification with Transformer. Colab[66].
❖ Making Transformer networks simpler and more efficient[67].
❖ Implementing a Transformer with PyTorch and PyTorch Lightnin. Colab[68].
❖ AttentioNN: All about attention in neural networks described as colab notebooks[69].
❖ Attention Is All You Need, Vaswani et al. `https://arxiv.org/abs/1706.03762`.
❖ Efficient Transformers: A Survey, Tay et al. `https://arxiv.org/abs/2009.06732`.
❖ How to train a new language model from scratch using Transformers and Tokenizers[70].
❖ Write With Transformer. By *Hugging Face*: `https://transformer.huggingface.co`.
❖ The Illustrated Transformer `http://jalammar.github.io/illustrated-transformer/`.
❖ How to generate text: using different decoding methods for language generation with Transformers[71].
❖ The annotated transformer (code) `http://nlp.seas.harvard.edu/2018/04/03/attention.html`.
❖ Attention and Augmented Recurrent Neural Networks `https://distill.pub/2016/augmented-rnns/`.
❖ Transformer model for language understanding. Tutorial showing how to write Transformer in TensorFlow 2.0[72].
❖ End-to-End Object Detection with Transformers, Carion et al. `https://arxiv.org/abs/2005.12872`. Colab[73].
❖ Transformer in TensorFlow 2.0 (code) `https://www.tensorflow.org/beta/tutorials/text/transformer`.

3.4.1 Natural Language Processing (NLP) | BERT: A New Era in NLP

BERT (Bidirectional Encoder Representations from Transformers)[6] is a *deeply bidirectional, unsupervised language representation*, pre-trained using only a plain text corpus (in this case, Wikipedia)[74].

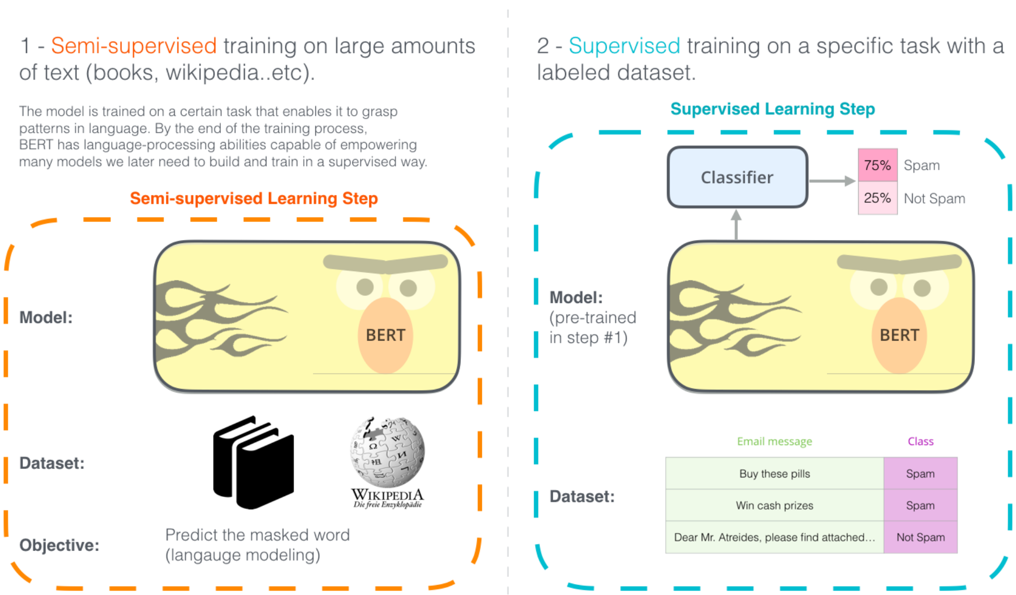

Figure 10: The two steps of how BERT is developed. Source `https://jalammar.github.io/illustrated-bert/`.

- Reading: Unsupervised pre-training of an LSTM followed by supervised fine-tuning[7].
- TensorFlow code and pre-trained models for BERT `https://github.com/google-research/bert`.
- Better Language Models and Their Implications[75].

[66] `https://colab.research.google.com/github/keras-team/keras-io/blob/master/examples/nlp/ipynb/text_classification_with_transformer.ipynb`
[67] `https://ai.facebook.com/blog/making-transformer-networks-simpler-and-more-efficient/`
[68] `https://colab.research.google.com/drive/1swXWW5sOLW8zSZBaQBYcGQkQ_Bje_bmI`
[69] `https://github.com/zaidalyafeai/AttentioNN`
[70] `https://huggingface.co/blog/how-to-train`
[71] `https://huggingface.co/blog/how-to-generate`
[72] `https://www.tensorflow.org/tutorials/text/transformer`
[73] `https://colab.research.google.com/drive/1rPm0-UrWHpJJRX9PsNb5SpzZiUlMh7wm`
[74] `https://ai.googleblog.com/2018/11/open-sourcing-bert-state-of-art-pre.html`
[75] `https://blog.openai.com/better-language-models/`

"*I think transfer learning is the key to general intelligence. And I think the key to doing transfer learning will be the acquisition of conceptual knowledge that is abstracted away from perceptual details of where you learned it from.*" — Demis Hassabis

❖ Towards a Conversational Agent that Can Chat About. . . Anything[76].
❖ How to Build OpenAI's GPT-2: "*The AI That's Too Dangerous to Release*"[77].
❖ A Primer in BERTology: What we know about how BERT works, Rogers et al., 2020[78].
❖ Play with BERT with your own data using TensorFlow Hub `https://colab.research.google.com/github/google-research/bert/blob/master/predicting_movie_reviews_with_bert_on_tf_hub.ipynb`.

3.5 Unsupervised Learning

True intelligence will require independent learning strategies.

"*Give a robot a label and you feed it for a second; teach a robot to label and you feed it for a lifetime.*" — Pierre Sermanet

Unsupervised learning is a paradigm for creating AI that learns without a particular task in mind: learning for the sake of learning[79]. It captures some characteristics of the joint distribution of the observed random variables (learn the underlying structure). The variety of tasks include density estimation, dimensionality reduction, and clustering.[4][80].

"*The unsupervised revolution is taking off!*" — Alfredo Canziani

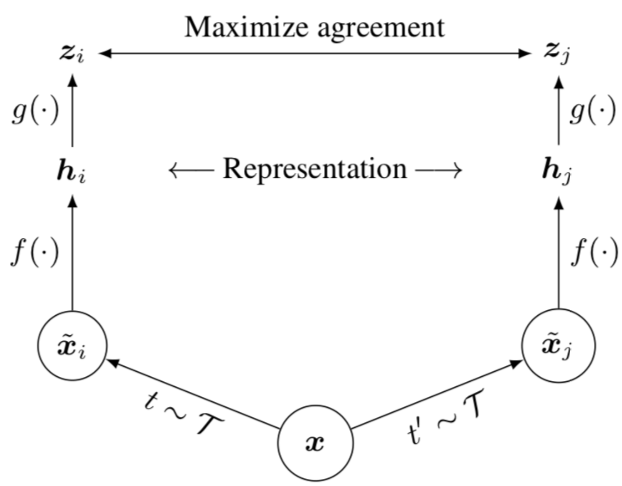

Figure 11: A Simple Framework for Contrastive Learning of Visual Representations, Chen et al., 2020

Self-supervised learning is derived form unsupervised learning where the data provides the supervision. E.g. Word2vec[81], a technique for learning vector representations of words, or word **embeddings**. An embedding is a mapping from discrete objects, such as words, to vectors of real numbers[82].

"*The next revolution of AI won't be supervised.*" — Yann LeCun

"*Self-supervised learning is a method for attacking unsupervised learning problems by using the mechanisms of supervised learning.*" — Thomas G. Dietterich

[76] `https://ai.googleblog.com/2020/01/towards-conversational-agent-that-can.html`
[77] `https://blog.floydhub.com/gpt2/`
[78] `https://arxiv.org/abs/2002.12327`
[79] `https://deepmind.com/blog/unsupervised-learning/`
[80] `https://media.neurips.cc/Conferences/NIPS2018/Slides/Deep_Unsupervised_Learning.pdf`
[81] `https://jalammar.github.io/illustrated-word2vec/`
[82] `http://projector.tensorflow.org`

❖ Self-Supervised Image Classification, Papers With Code[83].
❖ Self-supervised learning and computer vision, Jeremy Howard[84].
❖ Understanding Self-supervised Learning with Dual Deep Networks, Tian et al.[85]
❖ Momentum Contrast for Unsupervised Visual Representation Learning, He et al.[86]
❖ Data-Efficient Image Recognition with Contrastive Predictive Coding, Hénaff et al.[87]
❖ A Simple Framework for Contrastive Learning of Visual Representations, Chen et al.[88]
❖ CURL: Contrastive Unsupervised Representations for Reinforcement Learning, Srinivas et al.[89]
❖ FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence, Sohn et al.[90]
❖ Viewmaker Networks: Learning Views for Unsupervised Representation Learning, Tamkin et al.[91]
❖ Self-classifying MNIST Digits, Randazzo et al.: `https://distill.pub/2020/selforg/mnist/`.
❖ Self-Supervised Learning of Pretext-Invariant Representations, Ishan Misra, Laurens van der Maaten[92].

3.5.1 Generative Adversarial Networks

Simultaneously train two models: a generative model G that captures the data distribution, and a discriminative model D that estimates the probability that a sample came from the training data rather than G. The training procedure for G is to maximize the probability of D making a mistake. This framework corresponds to a minimax two-player game[3].

$$\min_{\theta_g} \max_{\theta_d} [\mathbb{E}_{\mathrm{x} \sim \mathrm{p_{data}(x)}} [log D_{\theta_d}(x)] + \mathbb{E}_{\mathrm{z} \sim \mathrm{p_z(z)}} [log(1 - D_{\theta_d}(G_{\theta_g}(z)))]] \tag{5}$$

"*What I cannot create, I do not understand.*" — Richard Feynman

Goodfellow et al. used an interesting analogy where the generative model can be thought of as analogous to a team of counterfeiters, trying to produce fake currency and use it without detection, while the discriminative model is analogous to the police, trying to detect the counterfeit currency. Competition in this game drives both teams to improve their methods until the counterfeits are indistiguishable from the genuine articles. See Figure 12.

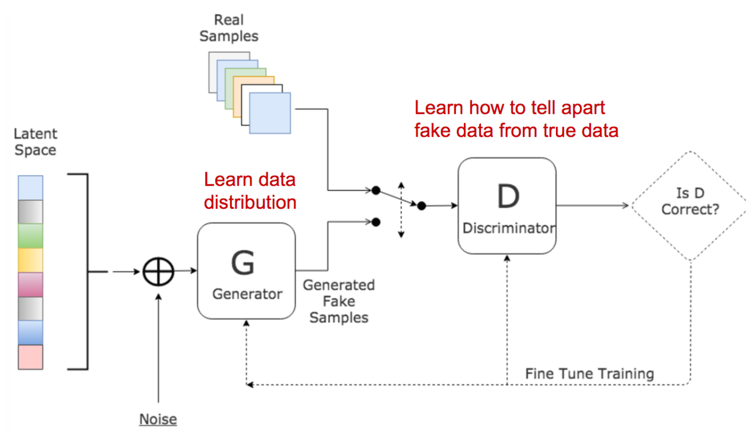

Figure 12: GAN: Neural Networks Architecture Pioneered by Ian Goodfellow at University of Montreal (2014).

StyleGAN: A Style-Based Generator Architecture for Generative Adversarial Networks

- Paper `http://stylegan.xyz/paper` | Code `https://github.com/NVlabs/stylegan`.

[83] `https://paperswithcode.com/task/self-supervised-image-classification`
[84] `https://www.fast.ai/2020/01/13/self_supervised/`
[85] `https://arxiv.org/abs/2010.00578`
[86] `https://arxiv.org/abs/1911.05722`
[87] `https://arxiv.org/abs/1905.09272`
[88] `https://arxiv.org/abs/2002.05709`
[89] `https://arxiv.org/abs/2004.04136`
[90] `https://arxiv.org/abs/2001.07685`
[91] `https://arxiv.org/abs/2010.07432`
[92] `https://arxiv.org/abs/1912.01991`

- **StyleGAN for art.** Colab `https://colab.research.google.com/github/ak9250/stylegan-art`.
- This Person Does Not Exist `https://thispersondoesnotexist.com`.
- Which Person Is Real? `http://www.whichfaceisreal.com`.
- This Resume Does Not Exist `https://thisresumedoesnotexist.com`.
- This Waifu Does Not Exist `https://www.thiswaifudoesnotexist.net`.
- Encoder for Official TensorFlow Implementation `https://github.com/Puzer/stylegan-encoder`.
- How to recognize fake AI-generated images. By Kyle McDonald[93].

❖ GAN in Keras. Colab[94].
❖ 100,000 Faces Imagined by a GAN `https://generated.photos`.
❖ Introducing **TF-GAN**: A lightweight GAN library for TensorFlow 2.0[95].
❖ Generative Adversarial Networks (GANs) in 50 lines of code (PyTorch)[96].
❖ Few-Shot Adversarial Learning of Realistic Neural Talking Head Models[97].
❖ GANcraft: Unsupervised 3D Neural Rendering of Minecraft Worlds. Hao et al.[98].
❖ Wasserstein GAN `http://www.depthfirstlearning.com/2019/WassersteinGAN`.
❖ GANpaint Paint with GAN units `http://gandissect.res.ibm.com/ganpaint.html`.
❖ StyleGAN2 Distillation for Feed-forward Image Manipulation. Viazovetskyi et al.[99] Code[100].
❖ A Review on Generative Adversarial Networks: Algorithms, Theory, and Applications. Gui et al.[101].
❖ CariGANs: Unpaired Photo-to-Caricature Translation. Cao et al.: `https://cari-gan.github.io`.
❖ Infinite-resolution (CPPNs, GANs and TensorFlow.js) `https://thispicturedoesnotexist.com`.
❖ PyTorch pretrained BigGAN `https://github.com/huggingface/pytorch-pretrained-BigGAN`.
❖ GANSynth: Generate high-fidelity audio with GANs! Colab `http://goo.gl/magenta/gansynth-demo`.
❖ SC-FEGAN: Face Editing Generative Adversarial Network `https://github.com/JoYoungjoo/SC-FEGAN`.
❖ Demo of BigGAN in an official Colaboratory notebook (backed by a GPU) `https://colab.research.google.com/github/tensorflow/hub/blob/master/examples/colab/biggan_generation_with_tf_hub.ipynb`

3.5.2 Variational AutoEncoder

Variational Auto-Encoders[102] (VAEs) are powerful models for learning low-dimensional representations See Figure 13. Disentangled representations are defined as ones where a change in a single unit of the representation corresponds to a change in single factor of variation of the data while being invariant to others (Bengio et al. (2013).

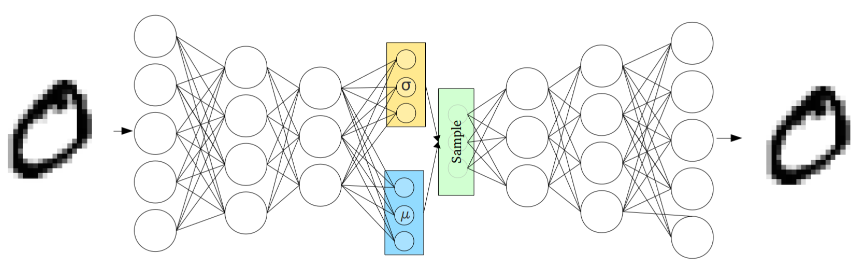

Figure 13: Variational Autoencoders (VAEs): Powerful Generative Models.

[93] `https://medium.com/@kcimc/how-to-recognize-fake-ai-generated-images-4d1f6f9a2842`
[94] `https://colab.research.google.com/drive/1CQ2XTMoUB7b9i9USUh4kp8BoCag1z-en`
[95] `https://medium.com/tensorflow/introducing-tf-gan-a-lightweight-gan-library-for-tensorflow-2-0-36d767e1abae`
[96] `https://medium.com/@devnag/generative-adversarial-networks-gans-in-50-lines-of-code-pytorch-e81b79659e3f`
[97] `https://arxiv.org/abs/1905.08233`
[98] `https://arxiv.org/abs/2104.07659`
[99] `https://arxiv.org/abs/2003.03581`
[100] `https://github.com/EvgenyKashin/stylegan2-distillation`
[101] `https://arxiv.org/abs/2001.06937`
[102] `https://arxiv.org/abs/1906.02691v2`

❖ Colab[103]: "*Debiasing Facial Detection Systems.*" ***AIEthics***
❖ NVAE: A Deep Hierarchical Variational Autoencoder, Arash Vahdat and Jan Kautz [104].
❖ Reading: Disentangled VAE's (DeepMind 2016) `https://arxiv.org/abs/1606.05579`.
❖ Slides: A Few Unusual Autoencoders `https://colinraffel.com/talks/vector2018few.pdf`.
❖ **MusicVAE**: Learning latent spaces for musical scores `https://magenta.tensorflow.org/music-vae`.
❖ Generative models in **Tensorflow 2** `https://github.com/timsainb/tensorflow2-generative-models/`.
❖ **SpaceSheet**: Interactive Latent Space Exploration with a Spreadsheet `https://vusd.github.io/spacesheet/`.

3.5.3 Capsule

Stacked Capsule Autoencoders. The inductive biases in this unsupervised version of capsule networks give rise to object-centric latent representations, which are learned in a self-supervised way—simply by reconstructing input images. Clustering learned representations is enough to achieve unsupervised state-of-the-art classification performance on MNIST (98.5%). Reference: blog by Adam Kosiorek.[105] Code[106].

Capsules learn *equivariant object representations* (applying any transformation to the input of the function has the same effect as applying that transformation to the output of the function).

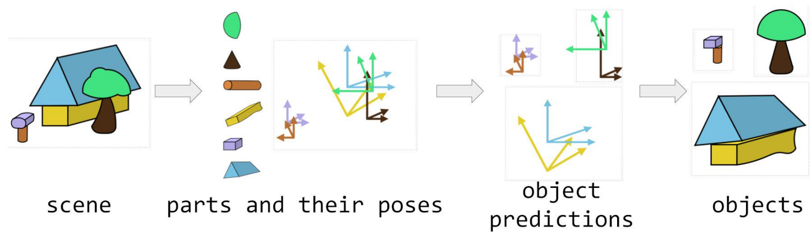

Figure 14: Stacked Capsule Autoencoders. Image source: Blog by Adam Kosiorek.

4 Autonomous Agents

We are on the dawn of *The Age of Artificial Intelligence.*

"*In a moment of technological disruption, leadership matters.*" — Andrew Ng

An **autonomous agent** is any device that perceives its environment and takes actions that maximize its chance of success at some goal. At the bleeding edge of AI, autonomous agents can learn from experience, simulate worlds and orchestrate meta-solutions. Here's an informal definition[107] of the *universal intelligence* of agent π[108]:

$$\Upsilon(\pi) := \sum_{\mu \in E} 2^{-K(\mu)} V_\mu^\pi \tag{6}$$

"*Intelligence measures an agent's ability to achieve goals in a wide range of environments.*" — Legg and Hutter, 2007

4.1 Deep Reinforcement Learning

Reinforcement learning (RL) studies how an agent can learn how to achieve goals in a complex, uncertain environment (Figure 15) [5]. Recent superhuman results in many difficult environments combine deep learning with RL (*Deep Reinforcement Learning*). See Figure 16 for a taxonomy of RL algorithms.

➙ Spinning Up in Deep RL - Proximal Policy Optimization (PPO), Colab Notebook[109].

[103] `https://colab.research.google.com/github/aamini/introtodeeplearning_labs/blob/master/lab2/Part2_debiasing_solution.ipynb`
[104] `https://arxiv.org/abs/2007.03898`
[105] `http://akosiorek.github.io/ml/2019/06/23/stacked_capsule_autoencoders.html`
[106] `https://github.com/google-research/google-research/tree/master/stacked_capsule_autoencoders`
[107] `https://arxiv.org/abs/0712.3329`
[108] Where μ is an environment, K is the Kolmogorov complexity function, E is the space of all computable reward summable environmental measures with respect to the reference machine U and the value function V_μ^π is the agent's "ability to achieve".
[109] `https://colab.research.google.com/drive/1piaU0x7nawRpSLKOTaCEdUG0KAR2OXku`

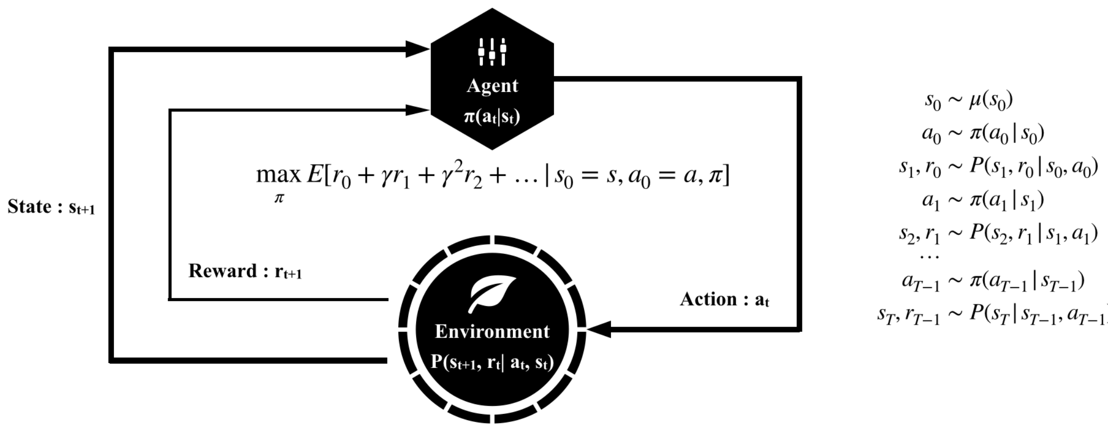

Figure 15: An Agent Interacts with an Environment.

❖ RL Tutorial, Behbahani et al.[110].
❖ An Opinionated Guide to ML Research[111].
❖ CS 188 : Introduction to Artificial Intelligence[112].
❖ Task-Agnostic Morphology Evolution, Hejna III et al.[113].
❖ Introduction to Reinforcement Learning by DeepMind[114].
❖ Isaac Gym `https://developer.nvidia.com/isaac-gym`.
❖ Discovering Reinforcement Learning Algorithms, Oh et al.[115].
❖ The NetHack Learning Environment, Küttler et al.[116] GitHub[117].
❖ "My Top 10 Deep RL Papers of 2019" by Robert Tjarko Lange[118].
❖ Lectures for UC Berkeley CS 285: Deep Reinforcement Learning[119].
❖ Behavior Priors for Efficient Reinforcement Learning, Tirumala et al.[120].
❖ Deep tic-tac-toe `https://zackakil.github.io/deep-tic-tac-toe/`.
❖ A Framework for Reinforcement Learning and Planning, Moerland et al. [121].
❖ Automatic Curriculum Learning For Deep RL: A Short Survey, Portelas et al.[122].
❖ ALLSTEPS: Curriculum-driven Learning of Stepping Stone Skills, Xie et al. [123].
❖ Decoupling Representation Learning from Reinforcement Learning, Stooke et al.[124] Code[125]
❖ Chip Placement with Deep Reinforcement Learning `https://arxiv.org/abs/2004.10746`.
❖ RL Unplugged: Benchmarks for Offline Reinforcement Learning, Gulcehre et al.[126] GitHub[127].
❖ CS 287: Advanced Robotics[128]. `https://people.eecs.berkeley.edu/~pabbeel/cs287-fa19/`.
❖ SCC: an efficient deep reinforcement learning agent mastering the game of StarCraft II, Wang et al.[129].

[110] https://github.com/eemlcommunity/PracticalSessions2020/blob/master/rl/EEML2020_RL_Tutorial.ipynb
[111] http://joschu.net/blog/opinionated-guide-ml-research.html
[112] https://inst.eecs.berkeley.edu/~cs188/fa18/
[113] https://arxiv.org/abs/2102.13100
[114] https://www.youtube.com/watch?v=2pWv7GOvuf0&list=PLqYmG7hTraZDM-OYHWgPebj2MfCFzFObQ
[115] https://arxiv.org/abs/2007.08794
[116] https://arxiv.org/abs/2006.13760
[117] https://github.com/facebookresearch/nle
[118] https://roberttlange.github.io/posts/2019/12/blog-post-9/
[119] https://www.youtube.com/playlist?list=PL_iWQOsE6TfURIIhCrlt-wj9ByIVpbfGc
[120] https://arxiv.org/pdf/2010.14274.pdf
[121] https://arxiv.org/abs/2006.15009
[122] https://arxiv.org/abs/2003.04664
[123] https://www.cs.ubc.ca/~van/papers/2020-allsteps/
[124] https://arxiv.org/abs/2009.08319
[125] https://github.com/astooke/rlpyt/tree/master/rlpyt/ul
[126] https://arxiv.org/abs/2006.13888
[127] https://github.com/deepmind/deepmind-research/tree/master/rl_unplugged
[128] https://people.eecs.berkeley.edu/~pabbeel/cs287-fa19/exam/cs287-fa19-exam-study-handout.pdf
[129] https://arxiv.org/abs/2012.13169

❖ Combining Deep Reinforcement Learning and Search for Imperfect-Information Games, Brown et al.[130].
❖ MDP Homomorphic Networks: Group Symmetries in Reinforcement Learning, Elise van der Pol et al.[131].
❖ Too many cooks: Bayesian inference for coordinating multi-agent collaboration, Wang et al.[132] GitHub[133].
❖ One Policy to Control Them All: Shared Modular Policies for Agent-Agnostic Control, Huang et al.[134]. Code[135].
❖ Decentralized Reinforcement Learning: Global Decision-Making via Local Economic Transactions, Chang et al. [136].

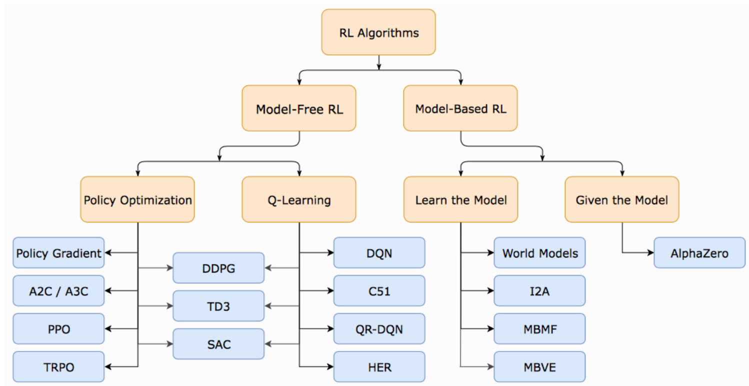

Figure 16: A Taxonomy of RL Algorithms. Source: Spinning Up in Deep RL by Achiam et al. | OpenAI

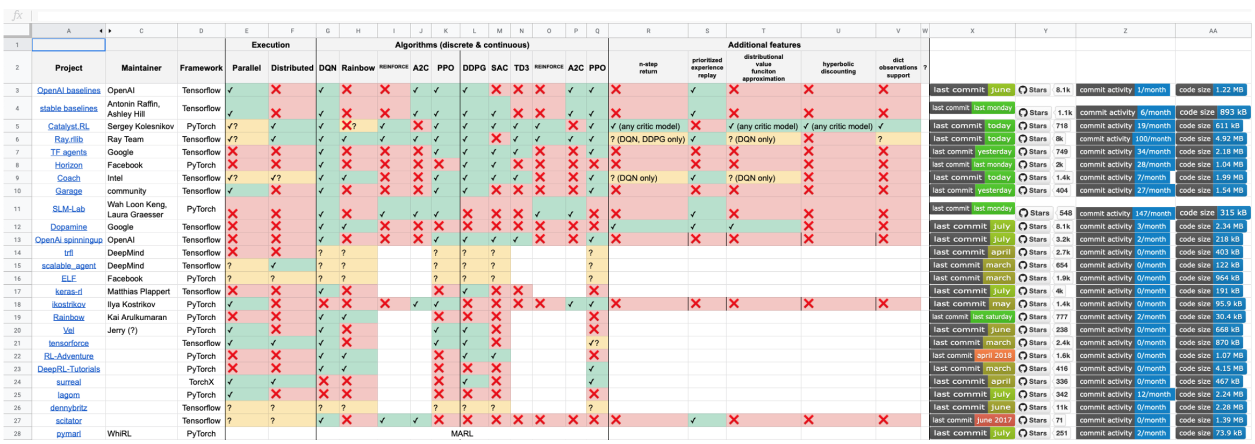

Figure 17: **Open-Source RL Algorithms** `https://docs.google.com/spreadsheets/d/1EeFPd-XIQ3mq_9snTlAZSsFY7Hbnmd7P5bbT8LPuMn0/`

4.1.1 Model-Free RL | Value-Based

The goal in RL is to train the agent to maximize the discounted sum of all future rewards $\boldsymbol{R_t}$, called the return:

$$R_t = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots \tag{7}$$

The Q-function captures the expected total future reward an agent in state $\boldsymbol{s}$ can receive by executing a certain action $\boldsymbol{a}$:

$$Q(s, a) = E[R_t] \tag{8}$$

[130] `https://arxiv.org/abs/2007.13544`
[131] `https://arxiv.org/abs/2006.16908`
[132] `https://arxiv.org/abs/2003.11778`
[133] `https://github.com/rosewang2008/gym-cooking`
[134] `https://arxiv.org/abs/2007.04976`
[135] `https://huangwl18.github.io/modular-rl`
[136] `https://arxiv.org/abs/2007.02382`

The optimal policy should choose the action ***a*** that maximizes ***Q(s,a)***:

$$\pi^*(s) = argmax_a Q(s,a) \tag{9}$$

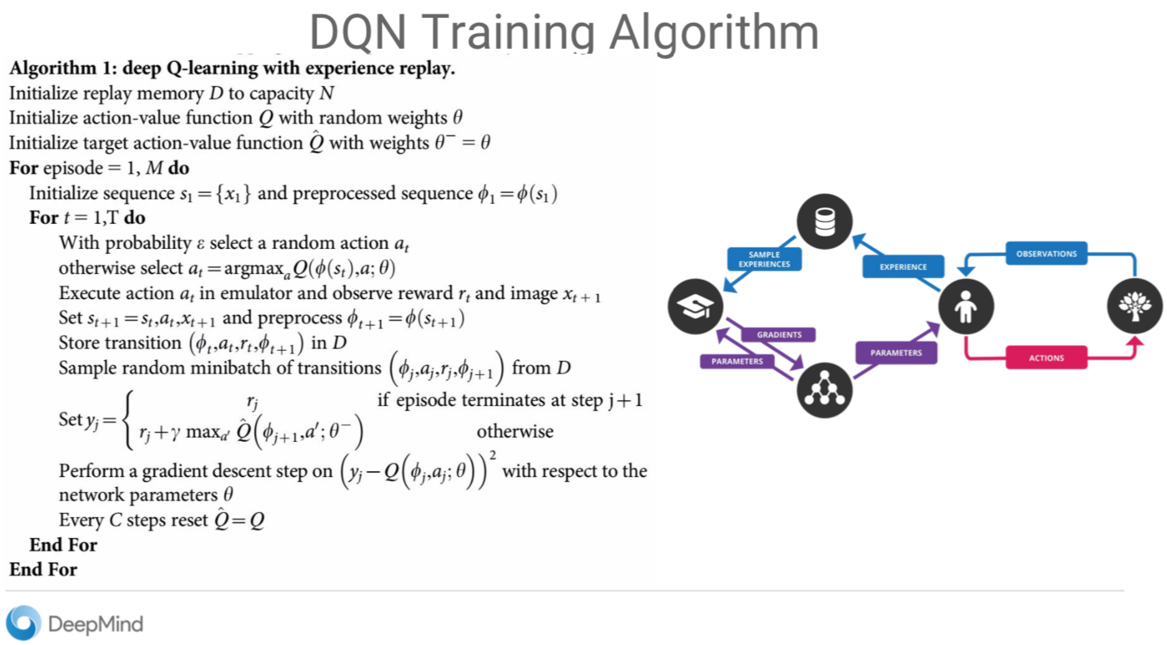

Figure 18: **DQN Training Algorithm.** Volodymyr Mnih, Deep RL Bootcamp

- **Q-Learning**: *Playing Atari with Deep Reinforcement Learning* (DQN). Mnih et al, 2013[10]. See Figure 18.

"*There's no limit to intelligence.*" — David Silver

❖ Q-Learning in enormous action spaces via amortized approximate maximization, de Wiele et al.[137].
❖ TF-Agents (DQN Tutorial) | Colab `https://colab.research.google.com/github/tensorflow/agents`.

4.1.2 Model-Free RL | Policy-Based

An RL agent learns the stochastic policy function that maps state to action and act by sampling policy.

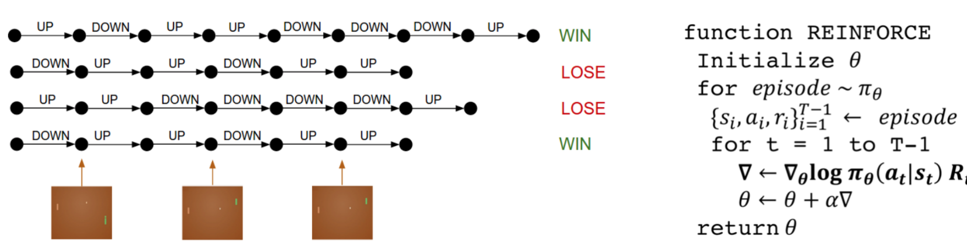

Figure 19: Policy Gradient Directly Optimizes the Policy.

Run a policy for a while (code: `https://gist.github.com/karpathy/a4166c7fe253700972fcbc77e4ea32c5`):

$$\tau = (s_0, a_0, r_0, s_1, a_1, r_1, \ldots, s_{T-1}, a_{T-1}, r_{T-1}, s_T) \tag{10}$$

[137] `https://arxiv.org/abs/2001.08116`

Increase probability of actions that lead to high rewards and decrease probability of actions that lead to low rewards:

$$\nabla_\theta E_\tau[R(\tau)] = E_\tau\left[\sum_{t=0}^{T-1} \nabla_\theta \log \pi(a_t|s_t,\theta) R(\tau)\right] \tag{11}$$

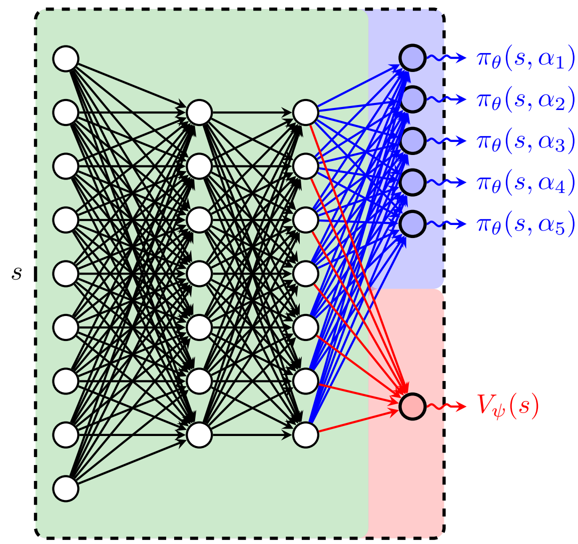

Figure 20: **Asynchronous Advantage Actor-Critic (A3C)**. Source: Petar Velickovic

- **Policy Optimization**: *Asynchronous Methods for Deep Reinforcement Learning* (A3C). Mnih et al, 2016[8].
- **Policy Optimization**: *Proximal Policy Optimization Algorithms* (PPO). Schulman et al, 2017[9].

❖ Phasic Policy Gradient. Cobbe et al.[138] Code[139].
❖ Deep Reinforcement Learning for Playing 2.5D Fighting Games. Li et al.[140].
❖ rlpyt: A Research Code Base for Deep Reinforcement Learning in PyTorch. Adam Stooke, Pieter Abbee[141].

4.1.3 Model-Based RL

In Model-Based RL, the agent generates predictions about the next state and reward before choosing each action.

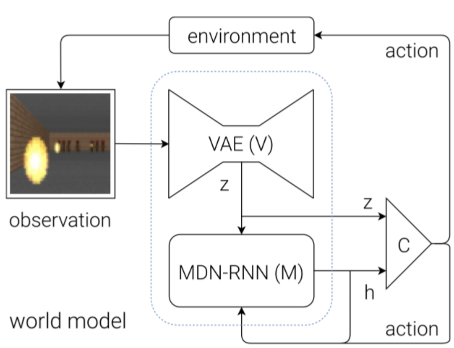


```
def rollout(controller):
  ''' env, rnn, vae are '''
  ''' global variables  '''
  obs = env.reset()
  h = rnn.initial_state()
  done = False
  cumulative_reward = 0
  while not done:
    z = vae.encode(obs)
    a = controller.action([z, h])
    obs, reward, done = env.step(a)
    cumulative_reward += reward
    h = rnn.forward([a, z, h])
  return cumulative_reward
```

Figure 21: World Model's Agent consists of: Vision (V), Memory (M), and Controller (C). | Ha et al, 2018[11]

[138] https://arxiv.org/abs/2009.04416
[139] https://github.com/openai/phasic-policy-gradient
[140] https://arxiv.org/abs/1805.02070
[141] https://arxiv.org/abs/1909.01500

- **Learn the Model**: *Recurrent World Models Facilitate Policy Evolution* (World Models[142]). The world model agent can be trained in an unsupervised manner to learn a compressed spatial and temporal representation of the environment. Then, a compact policy can be trained. See Figure 21. Ha et al, 2018[11].
- **Learn the Model**: *Learning Latent Dynamics for Planning from Pixels* `https://planetrl.github.io/`.
- **Given the Model**: *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm* (AlphaZero). Silver et al, 2017[14]. AlphaGo Zero Explained In One Diagram[143].

❖ Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model. Schrittwieser et al.[144]. Pseudocode[145].

4.1.4 Toward a General AI-Agent Architecture: SuperDyna (*General Dyna-style RL Agent*)

"*Intelligence is the computational part of the ability to predict and control a stream of experience.*" — Rich Sutton

SuperDyna.[146] The ambition: a general AI agent for Artificial Biological Reinforcement Learning.

1. Interact with the world: sense, update state and take an action
2. Learn from what just happened: see what happened and learn from it
3. Plan: (while there is time remaining in this time step) imagine hypothetical states and actions you might take
4. Discover : curate options and features and measure how well they're doing

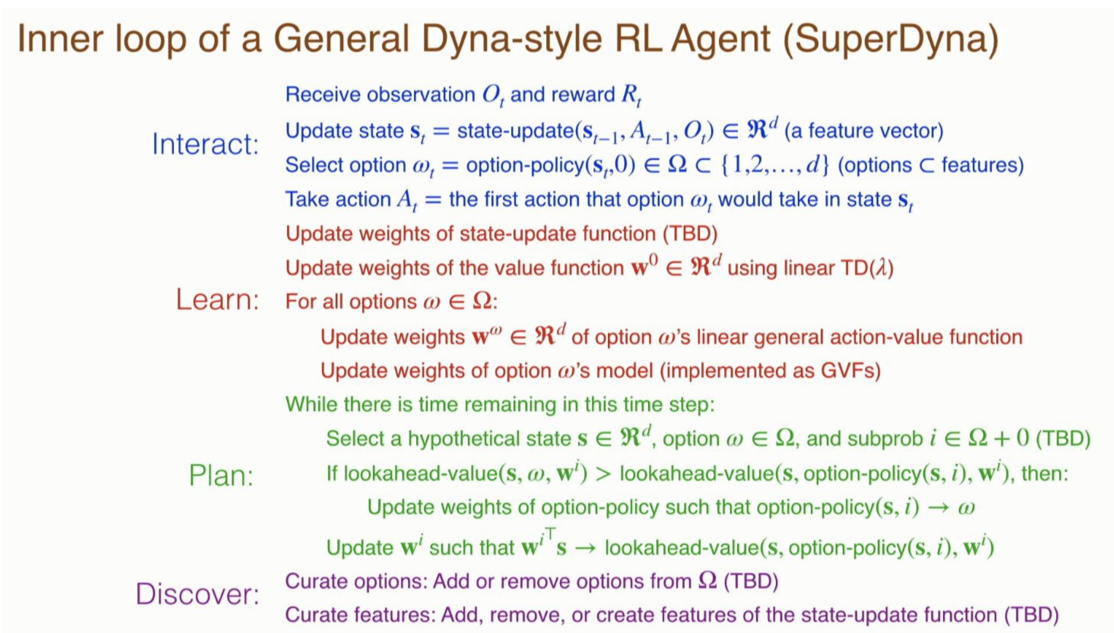

Figure 22: Inner Loop of a General Dyna-Style RL Agent (**SuperDyna**).

The first complete and scalable general AI-agent architecture that has all the most important capabilities and desiderata:

- Acting, learning, planning, model-learning, subproblems, and options.
- Function approximation, partial observability, non-stationarity and stochasticity.
- Discovery of state features, and thereby of subproblems, options and models.
- All feeding back to motivate new, more-abstract features in a virtuous cycle of discovery.

[142] `https://worldmodels.github.io`
[143] `https://applied-data.science/static/main/res/alpha_go_zero_cheat_sheet.png`
[144] `https://arxiv.org/abs/1911.08265`
[145] `https://arxiv.org/src/1911.08265v2/anc/pseudocode.py`
[146] `https://insidehpc.com/2020/02/video-toward-a-general-ai-agent-architecture/`

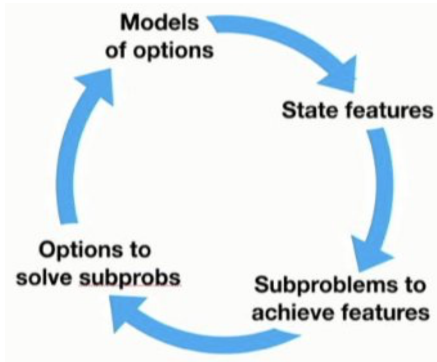

Figure 23: **SuperDyna:** Virtuous cycle of discovery.

Presentation by Richard Sutton (starts at 15 min.)[147].

"*In practice, I work primarily in reinforcement learning as an approach to artificial intelligence. I am exploring ways to represent a broad range of human knowledge in an empirical form–that is, in a form directly in terms of experience–and in ways of reducing the dependence on manual encoding of world state and knowledge.*" — Richard S. Sutton

4.1.5 Improving Agent Design

Via Reinforcement Learning: Blog[148]. arXiv[149]. ASTool `https://github.com/hardmaru/astool/`.
Via Evolution: Video[150]. Evolved Creatures `http://www.karlsims.com/evolved-virtual-creatures.html`.

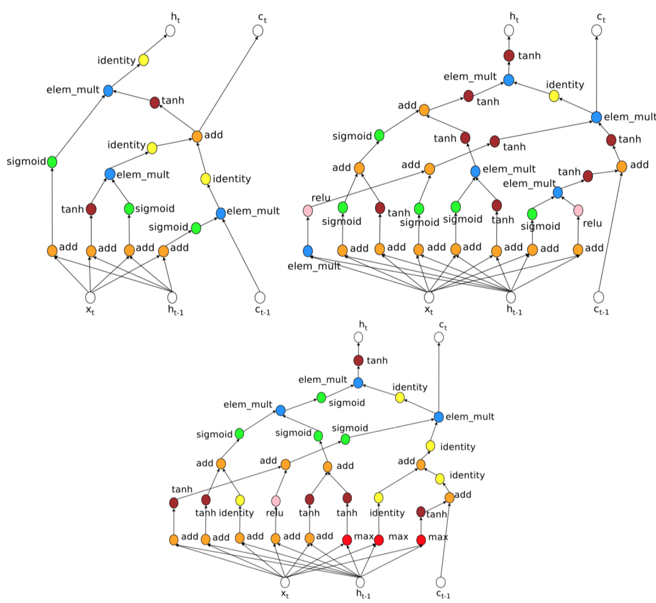

Figure 24: A comparison of the original LSTM cell vs. two new good generated. Top left: LSTM cell. [19]

"*The future of high-level APIs for AI is... a problem-specification API. Currently we only search over network weights, thus "problem specification" involves specifying a model architecture. In the future, it will just be: "tell me what data you have and what you are optimizing"*." — François Chollet

[147] `https://slideslive.com/38921889/biological-and-artificial-reinforcement-learning-4`
[148] `https://designrl.github.io`
[149] `https://arxiv.org/abs/1810.03779`
[150] `https://youtu.be/JBgG_VSP7f8`

❖ Teacher algorithms for curriculum learning of Deep RL in continuously parameterized environments[151].

4.1.6 OpenAI Baselines

High-quality implementations of reinforcement learning algorithms `https://github.com/openai/baselines`.

➺ Colab Notebook `https://colab.research.google.com/drive/1amdIQaHWyc8Av_DoM5yFYHyYvyqD5BZX`.

4.1.7 Google Dopamine and A Zoo of Agents

Dopamine is a research framework for fast prototyping of reinforcement learning algorithms.[152].

A Zoo of Atari-Playing Agents: Code[153], Blog[154] and Colaboratory notebook `https://colab.research.google.com/github/uber-research/atari-model-zoo/blob/master/colab/AtariZooColabDemo.ipynb`.

4.1.8 TRFL : TensorFlow Reinforcement Learning

TRFL ("truffle"): a library of reinforcement learning building blocks `https://github.com/deepmind/trfl`.

4.1.9 *bsuite* : Behaviour Suite for Reinforcement Learning

A collection of experiments that investigate core capabilities of RL agents `http://github.com/deepmind/bsuite`.

4.2 Evolution Strategies (ES)

In her Nobel Prize in Chemistry 2018 Lecture "*Innovation by Evolution: Bringing New Chemistry to Life*" (Nobel Lecture)†[155], Prof. Frances H. Arnold said :

> "*Nature ... invented life that has flourished for billions of years. (...) Equally awe-inspiring is the process by which Nature created these enzyme catalysts and in fact everything else in the biological world. The process is evolution, the grand diversity-generating machine that created all life on earth, starting more than three billion years ago. (...)* ***evolution executes a simple algorithm of diversification and natural selection, an algorithm that works at all levels of complexity from single protein molecules to whole ecosystems.***" — Prof. Frances H. Arnold

➺ Demo: ES on LunarLanderContinuous-v2. Colab Notebook[156]. Python Code[157]

Evolution and neural networks proved a potent combination in nature.

> "*Evolution is a slow learning algorithm that with the sufficient amount of compute produces a human brain.*" — Wojciech Zaremba

Natural evolutionary strategy **directly evolves the weights of a DNN** and performs competitively with the best deep reinforcement learning algorithms, including deep Q-networks (DQN) and policy gradient methods (A3C)[21].

Neuroevolution, which harnesses evolutionary algorithms to optimize neural networks, enables capabilities that are typically unavailable to gradient-based approaches, including learning neural network building blocks, architectures and even the algorithms for learning[12].

> "*. . . evolution — whether biological or computational — is inherently creative, and should routinely be expected to surprise, delight, and even outwit us.*" — The Surprising Creativity of Digital Evolution, Lehman et al.[22]

The ES algorithm is a “guess and check” process, where we start with some random parameters and then repeatedly:

1. Tweak the guess a bit randomly, and

[151] `https://arxiv.org/abs/1910.07224`
[152] `https://github.com/google/dopamine`
[153] `https://github.com/uber-research/atari-model-zoo`
[154] `https://eng.uber.com/atari-zoo-deep-reinforcement-learning/`
[155] `https://onlinelibrary.wiley.com/doi/epdf/10.1002/anie.201907729`
[156] `https://colab.research.google.com/drive/1PpYYaihoJWiszZh1vhKXvmN2X9KnLA7i`
[157] `https://drive.google.com/file/d/1YlKNorK21GMffz-29omEB7q_iLYRlmXD/view?usp=sharing`

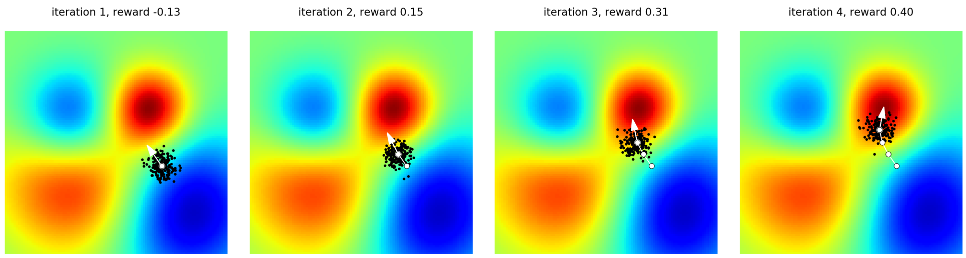

Figure 25: `https://colab.research.google.com/github/karpathy/randomfun/blob/master/es.ipynb`.

2. Move our guess slightly towards whatever tweaks worked better.

Neural architecture search has advanced to the point where it can outperform human-designed models[13].

"*Caterpillar brains LIQUIFY during metamorphosis, but the butterfly retains the caterpillar's memories!*" — M. Levin

"Open-ended" algorithms are algorithms that endlessly create. Brains and bodies evolve together in nature.

"We're machines," says Hinton. "*"We're just produced biologically (...)"* — Katrina Onstad, Toronto Life

❖ Evolution Strategies[158].
❖ VAE+CPPN+GAN[159].
❖ Demo: ES on CartPole-v1[160].
❖ Spiders Can Fly Hundreds of Miles Riding the Earth's Magnetic Fields[161].
❖ AutoML-Zero: Evolving Machine Learning Algorithms From Scratch, Real et al.[162] Code[163].
❖ A Visual Guide to ES `http://blog.otoro.net/2017/10/29/visual-evolution-strategies/`.
❖ Growing Neural Cellular Automata, Mordvintsev et al. `https://distill.pub/2020/growing-ca/`.
❖ **Xenobots** A scalable pipeline for designing reconfigurable organisms, Kriegman et al.[164]. Learn[165]. Evolve[166].

4.3 Self Play

Silver et al.[15] introduced an algorithm based solely on reinforcement learning, without human data, guidance or domain knowledge. Starting tabula rasa (and being its own teacher!), AlphaGo Zero achieved superhuman performance. AlphaGo Zero showed that algorithms matter much more than big data and massive amounts of computation.

"*Self-Play is Automated Knowledge Creation.*" — Carlos E. Perez

Self-play mirrors similar insights from coevolution. Transfer learning is the key to go from self-play to the real world[167].

"*Open-ended self play produces: Theory of mind, negotiation, social skills, empathy, real language understanding.*" — Ilya Sutskever, Meta Learning and Self Play

❖ How To Build Your Own MuZero AI Using Python[168].
❖ AlphaGo - The Movie | Full Documentary `https://youtu.be/WXuK6gekU1Y`.

[158] `https://lilianweng.github.io/lil-log/2019/09/05/evolution-strategies.html`
[159] `https://colab.research.google.com/drive/1_OoZ3z_C5Jl5gnxDOE9VEMCTs-Fl8pvM`
[160] `https://colab.research.google.com/drive/1bMZWHdhm-mT9NJENWoVewUks7cGV10go`
[161] `https://www.cell.com/current-biology/fulltext/S0960-9822(18)30693-6`
[162] `https://arxiv.org/abs/2003.03384`
[163] `https://github.com/google-research/google-research/tree/master/automl_zero`
[164] `https://www.pnas.org/content/early/2020/01/07/1910837117`
[165] `https://cdorgs.github.io`
[166] `https://github.com/skriegman/reconfigurable_organisms`
[167] `http://metalearning-symposium.ml`
[168] `https://medium.com/applied-data-science/how-to-build-your-own-muzero-in-python-f77d5718061a`

❖ TensorFlow.js Implementation of DeepMind's AlphaZero Algorithm for Chess. Live Demo[169]. | Code[170].
❖ An open-source implementation of the AlphaGoZero algorithm `https://github.com/tensorflow/minigo`.
❖ ELF OpenGo: An Open Reimplementation of AlphaZero, Tian et al.: `https://arxiv.org/abs/1902.04522`.

4.4 Multi-Agent Populations

"*We design a Theory of Mind neural network – a ToMnet – which uses meta-learning to build models of the agents it encounters, from observations of their behaviour alone.*" — Machine Theory of Mind, Rabinowitz et al.[25]

Cooperative Agents. Learning to Model Other Minds, by OpenAI[24], is an algorithm which accounts for the fact that other agents are learning too, and discovers self-interested yet collaborative strategies. Also: OpenAI Five[171].

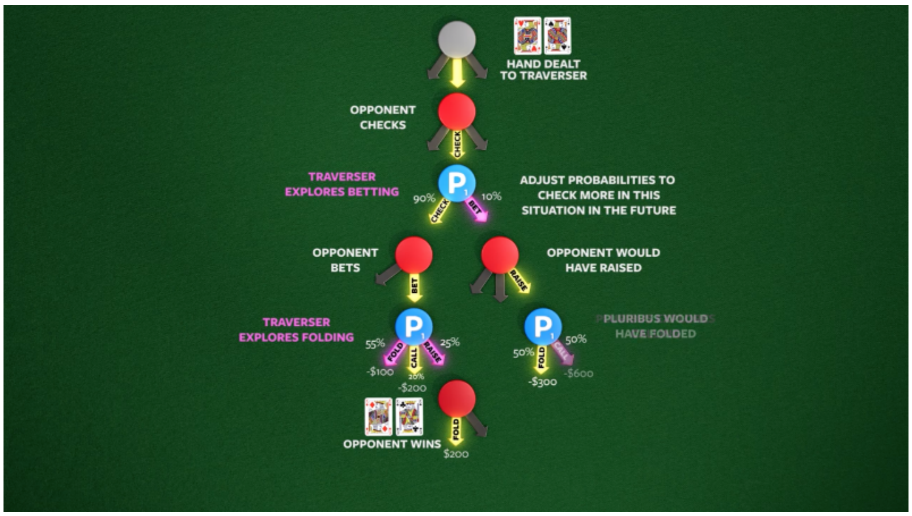

Figure 26: **Facebook, Carnegie Mellon build first AI that beats pros in 6-player poker** `https://ai.facebook.com/blog/pluribus-first-ai-to-beat-pros-in-6-player-poker`

"*Artificial Intelligence is about recognising patterns, Artificial Life is about creating patterns.*" — Mizuki Oka et al.

Active Learning Without Teacher. In *Intrinsic Social Motivation via Causal Influence in Multi-Agent RL*, Jaques et al. (2018) `https://arxiv.org/abs/1810.08647` propose an intrinsic reward function designed for multi-agent RL (MARL), which awards agents for having a causal influence on other agents' actions. Open-source implementation [172].

"*Open-ended Learning in Symmetric Zero-sum Games,*" Balduzzi et al.: `https://arxiv.org/abs/1901.08106`

❖ Lenia and Expanded Universe, Bert Wang-Chak Chan `https://arxiv.org/abs/2005.03742`.
❖ Neural MMO v1.3: A Massively Multiagent Game Environment for Training and Evaluating Neural Networks, Suarezet al.[173] Project Page `https://jsuarez5341.github.io`, Video[174] and Slides[175].
❖ Neural MMO: A massively multiagent env. for simulations with many long-lived agents. Code[176] and 3D Client[177].

4.5 Deep Meta-Learning

Learning to Learn[16].

[169] `https://frpays.github.io/lc0-js/engine.html`
[170] `https://github.com/frpays/lc0-js/`
[171] `https://blog.openai.com/openai-five/`
[172] `https://github.com/eugenevinitsky/sequential_social_dilemma_games`
[173] `https://arxiv.org/abs/2001.12004`
[174] `https://youtube.com/watch?v=DkHopV1RSxw`
[175] `https://docs.google.com/presentation/d/1tqm_Do9ph-duqqAlx3r9lI5Nbfb9yUfNEtXk1Qo4zSw/edit?usp=sharing`
[176] `https://github.com/openai/neural-mmo`
[177] `https://github.com/jsuarez5341/neural-mmo-client`

"*The notion of a neural "architecture" is going to disappear thanks to meta learning.*" — Andrew Trask

❖ Stanford CS330: Multi-Task and Meta-Learning, Finn et al., 2019[178].
❖ Meta Learning Shared Hierarchies[18] (*The Lead Author is in High School!*).
❖ Causal Reasoning from Meta-reinforcement Learning `https://arxiv.org/abs/1901.08162`.
❖ Meta-Learning through Hebbian Plasticity in Random Networks, Elias Najarro and Sebastian Risi, 2020[179].
❖ Meta-Learning Symmetries by Reparameterization, Zhou et al., 2020 `https://arxiv.org/abs/2007.02933`.

4.5.1 *MAML*: Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks

The goal of *model-agnostic meta-learning for fast adaptation of deep networks* is to train a model on a variety of learning tasks, such that it can solve new learning tasks using only a small number of training samples[20].

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}\left(f_{\theta_i'}\right) \tag{12}$$

A meta-learning algorithm takes in a distribution of tasks, where each task is a learning problem, and it produces a quick learner — a learner that can generalize from a small number of examples[17].

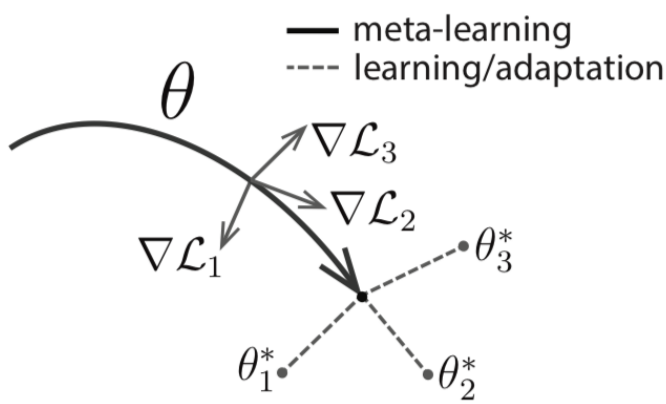

Figure 27: Diagram of Model-Agnostic Meta-Learning (MAML)

❖ How to Train MAML (Model-Agnostic Meta-Learning)[180].
❖ Meta-Learning with Implicit Gradients `https://arxiv.org/abs/1909.04630`.
❖ Colaboratory reimplementation of MAML (Model-Agnostic Meta-Learning) in TF 2.0[181].
❖ Torchmeta: A Meta-Learning library for PyTorch[182] `https://github.com/tristandeleu/pytorch-meta`.

4.5.2 The Grand Challenge for AI Research | *AI-GAs*: AI-Generating Algorithms, an Alternate Paradigm for Producing General Artificial Intelligence

In *AI-GAs: AI-generating algorithms, an alternate paradigm for producing general artificial intelligence*[183], Jeff Clune describes an exciting path that ultimately may be successful at producing general AI. The idea is to create an AI-generating algorithm (AI-GA), which automatically learns how to produce general AI.

Three Pillars are essential for the approach: (1) **Meta-learning architectures**, (2) **Meta-learning the learning algorithms themselves**, and (3) **Generating effective learning environments**.

- **The First Pillar**, meta-learning architectures, could potentially discover the building blocks : *convolution, recurrent layers, gradient-friendly architectures, spatial tranformers, etc*.

[178] `http://youtube.com/playlist?list=PLoROMvodv4rMC6zfYmnD7UG3LVvwaITY5`
[179] `https://arxiv.org/abs/2007.02686`
[180] `https://medium.com/towards-artificial-intelligence/how-to-train-maml-model-agnostic-meta-learning-90aa093f8e46`
[181] `https://colab.research.google.com/github/mari-linhares/tensorflow-maml/blob/master/maml.ipynb`
[182] `https://medium.com/pytorch/torchmeta-a-meta-learning-library-for-pytorch-f76c2b07ca6d`
[183] `https://arxiv.org/abs/1905.10985`

- **The Second Pillar**, meta-learning learning algorithms, could potentially learn the building blocks : *intelligent exploration, auxiliary tasks, efficient continual learning, causal reasoning, active learning, etc*.
- **The Third Pillar**, generating effective and fully expressive learning environments, could learn things like : *co-evolution / self-play, curriculum learning, communication / language, multi-agent interaction, etc*.

On Earth,

"*(. . .) a remarkably simple algorithm (Darwinian evolution) began producing solutions to relatively simple environments. The 'solutions' to those environments were organisms that could survive in them. Those organism often created new niches (i.e. environments, or opportunities) that could be exploited. Ultimately, that process produced all of the engineering marvels on the planet, such as jaguars, hawks, and the human mind.*" — Jeff Clune

Turing Complete (universal computer) : an encoding that enables the creation any possible learning algorithm.
Darwin Complete : an environmental encoding that enables the creation of any possible learning environment.

❖ **Learning to Continually Learn**. Beaulieu et al. https://arxiv.org/abs/2002.09571. Code[184].
❖ **Self-Organizing Intelligent Matter: A blueprint for an AI generating algorithm**. Anonymous et al.[185]

"*We propose an artificial life framework of interacting neural elements as a basis of an AI generating algorithm.*" — Anonymous[186]

❖ **Fully Differentiable Procedural Content Generation through Generative Playing Networks**. Bontrageret et al.[187]

5 Symbolic AI

❖ Generative Neurosymbolic Machines. Jindong Jiang, Sungjin Ahn[188]
❖ Neural Module Networks for Reasoning over Text. Gupta et al.[189] Code.[190]
❖ **Neurosymbolic AI: The 3rd Wave.** Artur d'Avila Garcez and Luis Lamb[191]
❖ Inductive Logic Programming: Theory and methods. Muggleton, S.; De Raedt, L.[192]
❖ (*Original FOIL paper*) Learning Logical Definitions from Relations. J.R. Quinlan.[193]
❖ Neural-Symbolic Learning and Reasoning: A Survey and Interpretation. Besold et al.[194]
❖ On neural-symbolic computing: suggested readings on foundations of the field. Luis Lamb[195].
❖ Neuro-symbolic A.I. is the future of artificial intelligence. Here's how it works. Luke Dormehl[196].
❖ Dimensions of Neural-symbolic Integration - A Structured Survey. Sebastian Bader, Pascal Hitzler[197].
❖ Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective. Lamb et al.[198].

"*The paper was inspired by the **AIDebate**, Gary Marcus writings, the AAAI2020 Firechat with Daniel Kahneman, and surveys not only our work, but the work of many in these AI fields.*" — Luis Lamb

[184] https://github.com/uvm-neurobotics-lab/ANML
[185] https://openreview.net/forum?id=160xFQdp7HR
[186] https://openreview.net/forum?id=160xFQdp7HR
[187] https://arxiv.org/abs/2002.05259
[188] https://arxiv.org/abs/2010.12152
[189] https://arxiv.org/abs/1912.04971
[190] https://nitishgupta.github.io/nmn-drop
[191] https://arxiv.org/abs/2012.05876
[192] https://doi.org/10.1016/0743-1066(94)90035-3
[193] https://link.springer.com/article/10.1023/A:1022699322624
[194] https://arxiv.org/abs/1711.03902
[195] https://twitter.com/luislamb/status/1218575842340634626
[196] https://www.digitaltrends.com/cool-tech/neuro-symbolic-ai-the-future/
[197] https://arxiv.org/abs/cs/0511042
[198] https://arxiv.org/abs/2003.00330

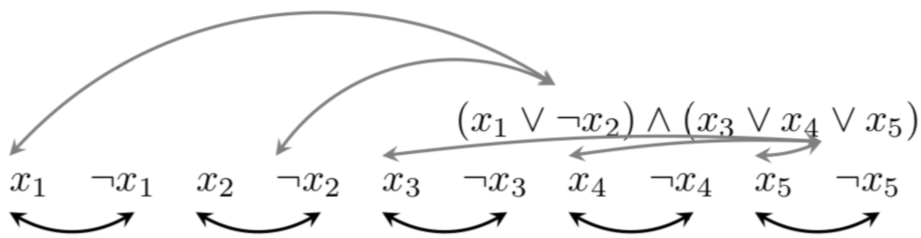

Figure 28: Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective. Lamb et al.

❖ DDSP: Differentiable Digital Signal Processing. Engel et al. Blog[199], Colab[200], Paper[201] and Code[202].
❖ The compositionality of neural networks: integrating symbolism and connectionism. Hupkes et al.[203]
❖ Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective. Lamb et al.[204]
❖ A computing procedure for quantification theory. M .Davis, H. Putnam. J. of ACM, Vol. 7, pp. 201-214, 1960
❖ Discovering Symbolic Models from Deep Learning with Inductive Biases, Cranmer et al.[205]. Blog and code[206].
❖ Symbolic Pregression: Discovering Physical Laws from Distorted Video. Silviu-Marian Udrescu, Max Tegmark[207]
❖ (*Workshop series on neurosymbolic AI*) Neural-Symbolic Integration. Hitzler et al. `http://neural-symbolic.org`
❖ Graph Colouring Meets Deep Learning: Effective Graph Neural Network Models for Combinatorial Problems. Lemos et al. `https://arxiv.org/abs/1903.04598`.
❖ Neural-Symbolic Relational Reasoning on Graph Models: Effective Link Inference and Computation from Knowledge Bases. Lemos et al. `https://arxiv.org/abs/2005.02525`.
❖ Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning. Garcez et al. `https://arxiv.org/abs/1905.06088`
❖ Differentiable Reasoning on Large Knowledge Bases and Natural Language. Minervini et al.[208] Open-source neuro-symbolic reasoning framework, in TensorFlow `https://github.com/uclnlp/gntp`.
❖ (*Original ILP foundational work*) Automatic Methods of Inductive Inference, Plotkin G.D. PhD thesis, University of Edinburgh, 1970 `https://era.ed.ac.uk/bitstream/handle/1842/6656/Plotkin1972.pdf;sequence=1`.

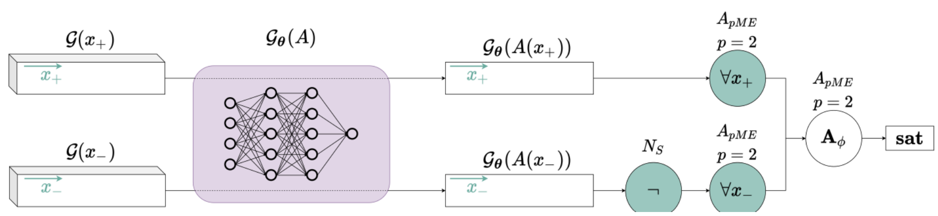

Figure 29: Logic Tensor Networks (*with explanatory examples using TensorFlow 2*). Badreddine et al.

6 Environments

Platforms for training autonomous agents.

"*Run a physics sim long enough and you'll get intelligence.*" — Elon Musk

[199] `http://magenta.tensorflow.org/ddsp`
[200] `http://g.co/magenta/ddsp-demo`
[201] `http://g.co/magenta/ddsp-paper`
[202] `http://github.com/magenta/ddsp`
[203] `https://arxiv.org/abs/1908.08351`
[204] `https://arxiv.org/abs/2003.00330`
[205] `https://arxiv.org/abs/2006.11287`
[206] `https://astroautomata.com/paper/symbolic-neural-nets/`
[207] `https://arxiv.org/abs/2005.11212`
[208] `https://arxiv.org/abs/1912.10824`

6.1 OpenAI Gym

"Situation awareness is the perception of the elements in the environment within a volume of time and space, and the comprehension of their meaning, and the projection of their status in the near future." — Endsley (1987)

The OpenAI Gym https://gym.openai.com/ (Blog[209] | GitHub[210]) is a toolkit for developing and comparing reinforcement learning algorithms. What makes the gym so great is a common API around environments.

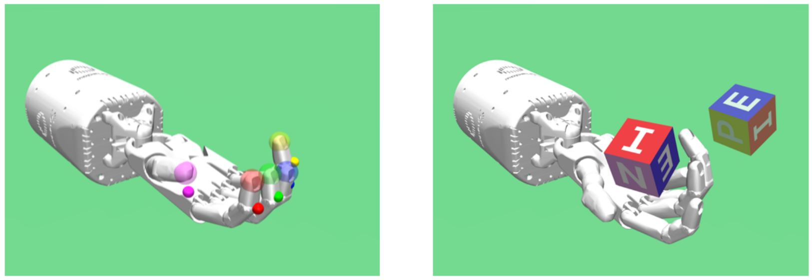

Figure 30: Robotics Environments https://blog.openai.com/ingredients-for-robotics-research/

"By framing the approach within the popular OpenAI Gym framework, design firms can create more realistic environments – for instance, incorporate strength of materials, safety factors, malfunctioning of components under stressed conditions, and plug existing algorithms into this framework to optimize also for design aspects such as energy usage, easy-of-manufacturing, or durability." — David Ha[211]

➾ Getting Started with the OpenAI Gym, Colab Notebook[212]

How to create new environments for Gym[213]. **Minimal example with code and agent** *(evolution strategies on foo-v0):*

1. Download *gym-foo* https://drive.google.com/file/d/1r2A8J9CJjIQNwss246gATeD0LLMzpUT-/view?usp=sharing
2. cd gym-foo
3. pip install -e .
4. python ES-foo.py

He're another more difficult (*for the agent!*) new environment for Gym *(evolution strategies on foo-v3):*

1. Download *gym-foo-v3*[214]
2. cd gym-foo-v3
3. pip install -e .
4. python ES-foo-v3.py

➾ Create a New Environment (foo) from Scratch, Colab Notebook[215]

❖ OpenAI Gym Environment for Trading[216].
❖ Fantasy Football AI Environment https://github.com/njustesen/ffai.
❖ Create custom gym environments from scratch — A stock market example[217].

[209] https://blog.openai.com/openai-gym-beta/
[210] https://github.com/openai/gym
[211] https://designrl.github.io
[212] https://colab.research.google.com/drive/1fBDH7xfpwH9SKj5J9TAH9XOTGJF61vJZ
[213] https://github.com/openai/gym/blob/master/docs/creating-environments.md
[214] https://drive.google.com/file/d/1cGncsXJ56UUKCO9MaRWJVTnxiQEnLuxS/view?usp=sharing
[215] https://colab.research.google.com/drive/1hXW5hQn1MO4kjgc2W2wjyTwDcId5QGCD
[216] https://github.com/hackthemarket/gym-trading
[217] https://towardsdatascience.com/creating-a-custom-openai-gym-environment-for-stock-trading-be532be3910e

❖ Spot Mini Mini OpenAI Gym Environment. Maurice Rahme, blog[218] et code[219].
❖ IKEA Furniture Assembly Environment `https://clvrai.github.io/furniture/`.
❖ Minimalistic Gridworld Environment `https://github.com/maximecb/gym-minigrid`.
❖ **DoorGym**: A Scalable Door Opening Environment and Baseline Agent, Urakami et al., 2019[220].
❖ **gym-gazebo2**, a toolkit for reinforcement learning using ROS 2 and Gazebo, Lopez et al., 2019[221].
❖ OFFWORLD GYM Open-access physical robotics environment for real-world reinforcement learning[222].
❖ **RecSim NG**: Toward Principled Uncertainty Modeling for Recommender Ecosystems, Mladenov et al., 2021[223].
❖ Safety Gym: environments to evaluate agents with safety constraints `https://github.com/openai/safety-gym`.
❖ **Gym-ANM**: Reinforcement Learning Environments for Active Network Management Tasks in Electricity Distribution Systems, Robin Henry, Damien Ernst, 2021[224].
❖ **TensorTrade:** An open source reinforcement learning framework for training, evaluating, and deploying robust trading agents `https://github.com/tensortrade-org/tensortrade`.

6.2 DeepMind Lab

DeepMind Lab: A customisable 3D platform for agent-based AI research `https://github.com/deepmind/lab`.

- DeepMind Control Suite `https://github.com/deepmind/dm_control`.
- Convert DeepMind Control Suite to OpenAI Gym Envs `https://github.com/zuoxingdong/dm2gym`.

6.3 Unity ML-Agents

Unity ML Agents allows to create environments where intelligent agents (*Single Agent*, *Cooperative and Competitive Multi-Agent* and *Ecosystem*) can be trained using RL, neuroevolution, or other ML methods `https://unity3d.ai`.

- Announcing ML-Agents Unity Package v1.0! Mattar et al.[225].
- Getting Started with Marathon Environments for Unity ML-Agents[226] `https://github.com/Unity-Technologies/marathon-envs`.
- Arena: A General Evaluation Platform and Building Toolkit for Multi-Agent Intelligence[227].

❖ Unity Robotics Hub `https://github.com/Unity-Technologies/Unity-Robotics-Hub`.

6.4 AI Habitat

AI Habitat enables training of embodied AI agents (virtual robots) in a highly photorealistic and efficient 3D simulator, before transferring the learned skills to reality. By Facebook AI Research `https://aihabitat.org/`.

Why the name Habitat? Because that's where AI agents live!

6.5 POET: Paired Open-Ended Trailblazer

Diversity is the premier product of evolution. Endlessly generate increasingly complex and diverse learning environments[228]. Open-endedness could generate learning algorithms reaching human-level intelligence[23].

- Implementation of the POET algorithm `https://github.com/uber-research/poet`.

[218] `https://moribots.github.io/project/spot-mini-mini`
[219] `https://github.com/moribots/spot_mini_mini`
[220] `https://arxiv.org/abs/1908.01887`
[221] `https://arxiv.org/abs/1903.06278`
[222] `https://gym.offworld.ai`
[223] `https://arxiv.org/abs/2103.08057`
[224] `https://arxiv.org/abs/2103.07932`
[225] `https://blogs.unity3d.com/2020/05/12/announcing-ml-agents-unity-package-v1-0/`
[226] `https://towardsdatascience.com/gettingstartedwithmarathonenvs-v0-5-0a-c1054a0b540c`
[227] `https://arxiv.org/abs/1905.08085`
[228] `https://eng.uber.com/poet-open-ended-deep-learning/`

- Enhanced POET: Open-Ended Reinforcement Learning through Unbounded Invention of Learning Challenges and their Solutions. Wang et al., 2020 https://arxiv.org/abs/2003.08536. Code[229].

7 Deep-Learning Hardware

Figure 31: Edge TPU - Dev Board https://coral.ai/products/dev-board/

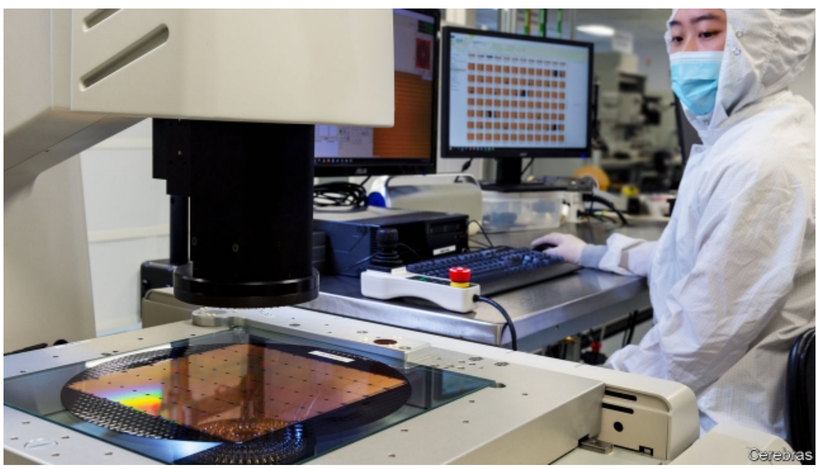

Figure 32: The world's largest chip : Cerebras Wafer Scale Engine https://www.cerebras.net

❖ Which GPU(s) to Get for Deep Learning, by Tim Dettmers[230].
❖ A Full Hardware Guide to Deep Learning, by Tim Dettmers[231].
❖ Jetson Nano. A small but mighty AI computer to create intelligent systems[232].
❖ Build AI that works offline with Coral Dev Board, Edge TPU, and TensorFlow Lite, by Daniel Situnayake[233].

[229] http://github.com/uber-research/poet
[230] http://timdettmers.com/2019/04/03/which-gpu-for-deep-learning/
[231] http://timdettmers.com/2018/12/16/deep-learning-hardware-guide/
[232] https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-nano/
[233] https://medium.com/tensorflow/build-ai-that-works-offline-with-coral-dev-board-edge-tpu-and-tensorflow-lite-70

8 Deep-Learning Software

8.1 TensorFlow

TensorFlow Hub is a library for reusable ML modules https://www.tensorflow.org/hub. Tutorials[234].
TensorFlow.js allows machine learning to happen within the web browser https://www.tensorflow.org/js/.

- TF-Coder https://goo.gle/3gwTbB6.
- TensorFlow Lite for Microcontrollers[235].
- Intro to Keras for Researchers. Colab[236].
- Introduction to Keras for Engineers. Colab[237].
- TensorBoard in Jupyter Notebooks[238]. Colab[239].
- TensorFlow 2.0 + Keras Crash Course. Colab[240].
- tf.keras (TensorFlow 2.0) for Researchers: Crash Course. Colab[241].
- TensorFlow Tutorials https://www.tensorflow.org/tutorials.
- Exploring helpful uses for BERT in your browser with TensorFlow.js[242].
- TensorFlow 2.0: basic ops, gradients, data preprocessing and augmentation, training and saving. Colab[243].

8.2 PyTorch

- PyTorch primer. Colab[244].
- Get started with PyTorch, Cloud TPUs, and Colab[245].
- MiniTorch https://minitorch.github.io/index.html
- Effective PyTorch https://github.com/vahidk/EffectivePyTorch
- PyTorch internals http://blog.ezyang.com/2019/05/pytorch-internals/
- PyTorch Lightning Bolts https://github.com/PyTorchLightning/pytorch-lightning-bolts

9 AI Art | A New Day Has Come in Art Industry

The code (*art-DCGAN*) for the first artificial intelligence artwork ever sold at Christie's auction house (Figure 33) is a modified implementation of DCGAN focused on generative art: https://github.com/robbiebarrat/art-dcgan.

- **The Creative AI Lab** https://creative-ai.org.
- **TensorFlow Magenta**. An open source research project exploring the role of ML in the creative process.[246].
- **Magenta Studio**. A suite of free music-making tools using machine learning models![247].
- **Style Transfer Tutorial** https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/r2/tutorials/generative/style_transfer.ipynb

[234] https://www.tensorflow.org/hub/tutorials
[235] https://www.tensorflow.org/lite/microcontrollers
[236] https://colab.research.google.com/drive/1qKPITTI879YHTxbTgYW_MAWMHFkbOBIk
[237] https://colab.research.google.com/drive/1lWUGZarlbORaHYUZlF9muCgpPl8pEvve
[238] https://www.tensorflow.org/tensorboard/tensorboard_in_notebooks
[239] https://colab.research.google.com/github/tensorflow/tensorboard/blob/master/docs/tensorboard_in_notebooks.ipynb
[240] https://colab.research.google.com/drive/1UCJt8EYjlzCs1H1d1X0iDGYJsHKwu-NO
[241] https://colab.research.google.com/drive/14CvUNTaX1OFHDfaKaaZzrBsvMfhCOHIR
[242] https://blog.tensorflow.org/2020/03/exploring-helpful-uses-for-bert-in-your-browser-tensorflow-js.html
[243] https://colab.research.google.com/github/zaidalyafeai/Notebooks/blob/master/TF_2_0.ipynb
[244] https://colab.research.google.com/drive/1DgkVmi6GksWOByhYVQpyUB4Rk3PUq0Cp
[245] https://medium.com/pytorch/get-started-with-pytorch-cloud-tpus-and-colab-a24757b8f7fc
[246] https://magenta.tensorflow.org
[247] https://magenta.tensorflow.org/studio

Figure 33: On October 25, 2018, the first AI artwork ever sold at Christie's auction house fetched USD 432,500.

- **AI x AR Paper Cubes** `https://experiments.withgoogle.com/paper-cubes`.
- **Photo Wake-Up** `https://grail.cs.washington.edu/projects/wakeup/`.
- **COLLECTION**. AI Experiments `https://experiments.withgoogle.com/ai`.

"*The Artists Creating with AI Won't Follow Trends; THEY WILL SET THEM.*" — The House of Montréal.AI Fine Arts

❖ Tuning Recurrent Neural Networks with Reinforcement Learning[248].
❖ **MuseNet**. Generate Music Using Many Different Instruments and Styles![249].
❖ Infinite stream of machine generated art. Valentin Vieriu `https://art42.net`.
❖ Deep Multispectral Painting Reproduction via Multi-Layer, Custom-Ink Printing. Shi et al.[250].
❖ Discovering Visual Patterns in Art Collections with Spatially-consistent Feature Learning. Shen et al.[251].
❖ Synthesizing Programs for Images using Reinforced Adversarial Learning, Ganin et al., 2018[252]. Agents[253].

10 AI Macrostrategy: Aligning AGI with Human Interests

Montréal.AI Governance: Policies at the intersection of AI, Ethics and Governance.

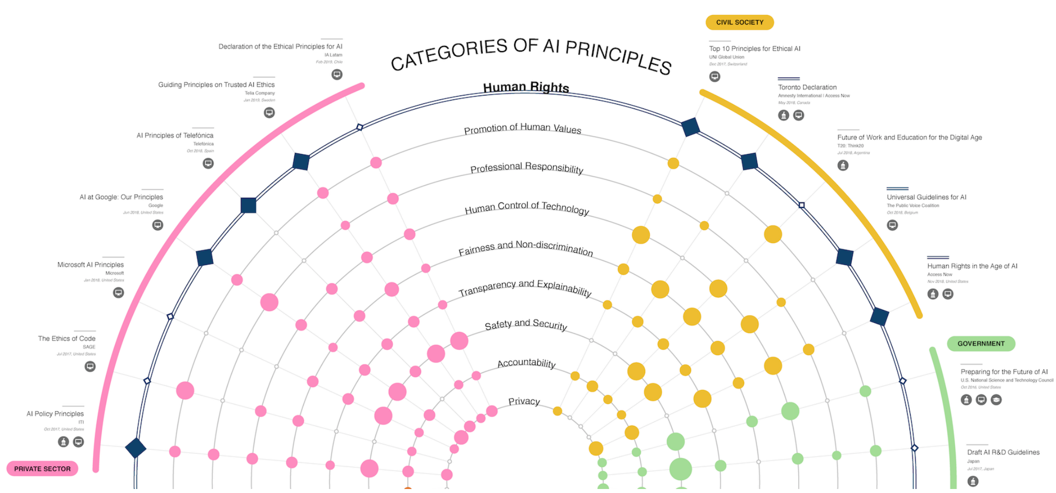

Figure 34: A Map of Ethical and Right-Based Approaches `https://ai-hr.cyber.harvard.edu/primp-viz.html`

[248] `https://magenta.tensorflow.org/2016/11/09/tuning-recurrent-networks-with-reinforcement-learning`
[249] `https://openai.com/blog/musenet/`
[250] `http://people.csail.mit.edu/liangs/papers/ToG18.pdf`
[251] `https://arxiv.org/pdf/1903.02678.pdf`
[252] `http://proceedings.mlr.press/v80/ganin18a.html`
[253] `https://github.com/deepmind/spiral`

"*(AI) will rank among our greatest technological achievements, and everyone deserves to play a role in shaping it.*" — Fei-Fei Li

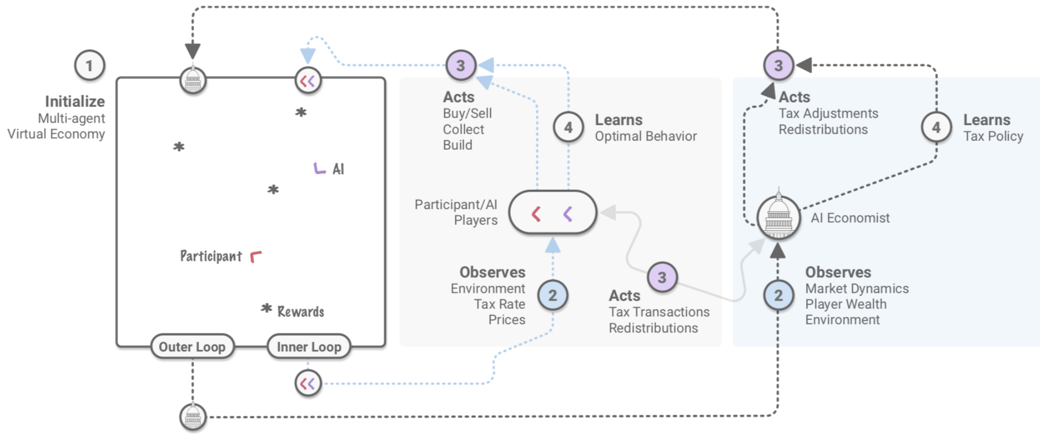

Figure 35: The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies. Zheng et al. https://arxiv.org/abs/2004.13332

❖ **AI Index**. http://aiindex.org.
❖ **The State of AI Report**. https://www.stateof.ai/.
❖ **Malicious AI Report**. https://arxiv.org/pdf/1802.07228.pdf.
❖ **Artificial Intelligence and Human Rights**. https://ai-hr.cyber.harvard.edu.
❖ **The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies**, Zheng et al.[254]. Blog[255].
❖ **Ethically Aligned Design, First Edition**[256]. From Principles to Practice https://ethicsinaction.ieee.org.
❖ **ADDRESS PREPARED BY POPE FRANCIS FOR THE PLENARY ASSEMBLY OF THE PONTIFICAL ACADEMY FOR LIFE**[257].

"*It's springtime for AI, and we're anticipating a long summer.*" — Bill Braun

References

[1] Mnih et al. Human-Level Control Through Deep Reinforcement Learning. In *Nature* 518, pages 529–533. 26 February 2015. https://storage.googleapis.com/deepmind-media/dqn/DQNNaturePaper.pdf

[2] Yann LeCun, Yoshua Bengio and Geoffrey Hinton. Deep Learning. In *Nature* 521, pages 436–444. 28 May 2015. https://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf

[3] Goodfellow et al. Generative Adversarial Networks. *arXiv preprint arXiv:1406.2661*, 2014. https://arxiv.org/abs/1406.2661

[4] Yoshua Bengio, Andrea Lodi, Antoine Prouvost. Machine Learning for Combinatorial Optimization: a Methodological Tour d'Horizon. *arXiv preprint arXiv:1811.06128*, 2018. https://arxiv.org/abs/1811.06128

[5] Brockman et al. OpenAI Gym. 2016. https://gym.openai.com

[6] Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*, 2018. https://arxiv.org/abs/1810.04805

[7] Dai et al. Semi-supervised Sequence Learning. *arXiv preprint arXiv:1511.01432*, 2015. https://arxiv.org/abs/1511.01432

[8] Mnih et al. Asynchronous Methods for Deep Reinforcement Learning. *arXiv preprint arXiv:1602.01783*, 2016. https://arxiv.org/abs/1602.01783

[254] https://arxiv.org/abs/2004.13332
[255] https://blog.einstein.ai/the-ai-economist/
[256] https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf
[257] http://w2.vatican.va/content/francesco/en/speeches/2020/february/documents/papa-francesco_20200228_accademia-perlavita.html

[9] Schulman et al. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017. https://arxiv.org/abs/1707.06347

[10] Mnih et al. Playing Atari with Deep Reinforcement Learning. *DeepMind Technologies*, 2013. https://www.cs.toronto.edu/~vmnih/docs/dqn.pdf

[11] Ha et al. Recurrent World Models Facilitate Policy Evolution. *arXiv preprint arXiv:1809.01999*, 2018. https://arxiv.org/abs/1809.01999

[12] Kenneth et al. Designing neural networks through neuroevolution. In *Nature Machine Intelligence* VOL 1, pages 24–35. January 2019. https://www.nature.com/articles/s42256-018-0006-z.pdf

[13] So et al. The Evolved Transformer. *arXiv preprint arXiv:1901.11117*, 2019. https://arxiv.org/abs/1901.11117

[14] Silver et al. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *arXiv preprint arXiv:1712.01815*, 2017. https://arxiv.org/abs/1712.01815

[15] Silver et al. AlphaGo Zero: Learning from scratch. In *DeepMind's Blog*, 2017. https://deepmind.com/blog/alphago-zero-learning-scratch/

[16] Andrychowicz et al. Learning to learn by gradient descent by gradient descent. *arXiv preprint arXiv:1606.04474*, 2016. https://arxiv.org/abs/1606.04474

[17] Nichol et al. Reptile: A Scalable Meta-Learning Algorithm. 2018. https://blog.openai.com/reptile/

[18] Frans et al. Meta Learning Shared Hierarchies. *arXiv preprint arXiv:1710.09767*, 2017. https://arxiv.org/abs/1710.09767

[19] Zoph and Le, 2017 Neural Architecture Search with Reinforcement Learning. *arXiv preprint arXiv:1611.01578*, 2017. https://arxiv.org/abs/1611.01578

[20] Finn et al., 2017 Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv preprint arXiv:1703.03400*, 2017. https://arxiv.org/abs/1703.03400

[21] Salimans et al. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. 2017. https://blog.openai.com/evolution-strategies/

[22] Lehman et al. The Surprising Creativity of Digital Evolution: A Collection of Anecdotes from the Evolutionary Computation and Artificial Life Research Communities. *arXiv preprint arXiv:1803.03453*, 2018. https://arxiv.org/abs/1803.03453

[23] Wang et al. Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions. *arXiv preprint arXiv:1901.01753*, 2019. https://arxiv.org/abs/1901.01753

[24] Foerster et al. Learning to Model Other Minds. 2018. https://blog.openai.com/learning-to-model-other-minds/

[25] Rabinowitz et al. Machine Theory of Mind. *arXiv preprint arXiv:1802.07740*, 2018. https://arxiv.org/abs/1802.07740