

Executive Function: A Contrastive Value Policy for Resampling and Relabeling Perceptions via Hindsight Summarization?

Chris Lengerich, Ben Lengerich

November 29, 2021

Abstract

We develop the few-shot continual learning task from first principles and hypothesize an evolutionary motivation and mechanism of action for executive function as a contrastive value policy which resamples and relabels perception data via hindsight summarization to minimize attended prediction error, similar to an online prompt engineering problem. This is made feasible by the use of a memory policy and a pretrained network with inductive biases for a grammar of learning and is trained to maximize evolutionary survival. We show how this model of executive function can be used to implement hypothesis testing as a stream of consciousness and may explain observations of efficient human few-shot learning and neuroanatomy.

Contents

1	Derivation of the Predictive Learning Task	2
1.1	The Outer Loop: The Evolutionary Task	2
1.2	The Inner Loop: The Predictive Learning Task	2
1.3	The Continual Few-Shot Learning Constraint	2
1.4	What does it mean to learn?	2
1.4.1	Sparsity	2
1.4.2	Compositionality	3
1.4.3	Compositionality in Time	3
1.4.4	Compositionality in Abstraction	3
1.4.5	Communication is Composition in Time and Abstraction	3
1.4.6	Sparsity and Compositionality Are Necessary for Causality	3
1.4.7	Contrastive Attention to Prediction Loss Drives Data Resampling and Relabeling Needed For Causal Models	4
2	Biological Solutions	4
2.1	Prediction	5
2.2	Prediction Loss	5
2.2.1	Decreasing Prediction Loss	5
2.2.2	Increasing Prediction Loss	5
2.2.3	Acting to Minimize Expected Prediction Loss	5
2.3	Learning to Minimize Prediction Loss	6
2.3.1	What is Learning?	6
2.3.2	Taking actions in the world to resample data	6
2.3.3	Taking internal actions to resample data	6
2.4	Few-Shot Continual Learning	6
2.4.1	A Grammar of Learning?	6
2.4.2	Resampling and Relabeling Data Via Memory and Hindsight Summarization	6
2.4.3	Hypothesis testing via summarization	8
3	Conclusion	9

1 Derivation of the Predictive Learning Task

1.1 The Outer Loop: The Evolutionary Task

Consider the online learning task of an agent in an environment such that the agent has a sensor system that produces a stream of time-ordered perceptions, y_0, \dots, y_t . The agent lives completely within Plato’s Cave (Plato, 375 BC), such that there are no labels for perceptions apart from implicit labels from future values of the stream, however, there is an evolutionary selection process based on an external environment which deletes agents based on their lack of fitness. The task we assign the agent is that of evolution - for the genetic code of the agent’s species to continue to exist subject to this selection pressure.

1.2 The Inner Loop: The Predictive Learning Task

We consider in particular a *predictive learning agent* - the agent which has a learned probability density function, $p(y_t|y_{t-1}, \dots, y_0)$, a prediction sampling process which samples \hat{y}_k from $p(y_k|y_{t-1}, \dots, y_0)$ for a sampled $k \geq t$, and an action space, A , from which our agent draws actions conditioned on its predictions. In environmentally fit agents, we expect the following desiderata for the agent’s sampling policy and pdf:

1. Predictive: The agent learns to calibrate its expectations accurately to future perceptions such that it minimizes the prediction loss, $\hat{y}_k - y_k$
2. Controlled: The agent learn to predict calibrated ego-perceptions (perceptions which represent its own actions and thoughts) more often than other perceptions
3. Aligned: The agent learns to predict calibrated ego-perceptions grounded to its survival more often
4. Generalizable: The agent’s predictions fulfill 1-3 across many environments and data distributions

In contrast to the traditional TD Learning formulation of RL as a POMDP with a distinction between observations and rewards (Bellman, 1954; Sutton and Barto, 2018), we treat reward simply as attention to a specific type of prediction loss between predictions and observations which is learned via evolutionary selection, in line with the view of dopamine’s role as a generalized prediction loss signal within the brain rather than merely reward prediction (Gardner et al., 2018; Hohwy, 2014).

1.3 The Continual Few-Shot Learning Constraint

We also introduce a few-shot learning constraint of only being able to collect a few examples of data from the environment. Continual few-shot learning allows for the agent’s survival in environments with distribution shift and expensive data collection, such as the competitive natural world. Although we will use tools of generational learning (as in most current machine learning setups now) to train agents, we are primarily interested in agents that display intrinsically motivated few-shot hypothesis testing behavior which can efficiently build a tree of knowledge.

Note that if we solve Desiderata 1 in generality, we have made substantial progress towards Desiderata 2 + 3, as these are just a conditional distribution over the joint distribution learned in 1 and can be fine-tuned with an exploration/exploitation policy and evolutionary selection mechanism.

1.4 What does it mean to learn?

Arguably, the goal of learning is to produce a *world model* which is useful for prediction in environments with rapid change such as the competitive natural world (Ha and Schmidhuber, 2018). This requires the representations learned be *sparse, compositional and causal*.¹

1.4.1 Sparsity

Raw perception data is dense. However, in order to use it to predict, we would like to transform it into sparse² representations. Sparse representations are useful in that they can be converted into tokenized representations (vocabularies), which allow for compositionality and causal reasoning.

⁰The authors also thank MeiMei the golden retriever for insightful collaboration, in particular, deriving insights on human vs. dog data relabeling capabilities.

¹Note that although we would like world models to have accurate causal representations, even inaccurate causal models are useful for learning. Humans have been able to use approximate causal models of gravity effectively well before the discovery of Newton’s law.

²Sparsity is defined as the number of non-zeros in a matrix - a matrix with only a few non-zeros is sparse, while a matrix with many non-zeros is dense

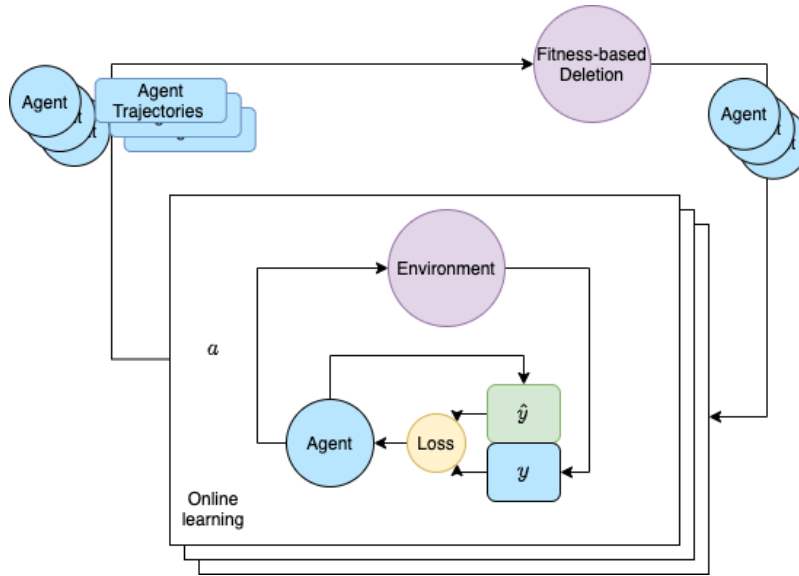


Figure 1: Learning Loops

1.4.2 Compositionality

We propose that it is particularly important to learn world models which have distributions of activations which are *compositional in time* and *compositional in abstractive depth*.

1.4.3 Compositionality in Time

We define composition in time as mapping an event A into an equivalent sequence of events $f(A) = a_0, a_1, \dots, a_n$ where a_{i+1} occurs after a_i , as well as the inverse mapping $g(a_0, a_1, \dots, a_n) = A$. This allows for the sequencing of events and long-horizon planning.

1.4.4 Compositionality in Abstraction

Composition in abstraction - that is, the ability to represent a concept either as an abstract idea (eg. “ $x = y$ ”) or a series of concrete examples (eg. “when $x = 1, y = 1$, when $x = 2, y = 2$, when $x = 3, y = 3$ ”) - allows for the communication of concrete representations to other neural networks. Mathematically, we would like to learn a mapping from an abstract concept A to a series of concrete examples: $f(A) = a_0, a_1, \dots, a_n$, and an inverse mapping $g(a_0, a_1, \dots, a_n) = A$. Note that compositionality in abstraction is a softer variant of hierarchical representation as multiple mappings may exist.

1.4.5 Communication is Composition in Time and Abstraction

Composition may also be recognizable as communication. For example:

$$\text{abstraction} \rightarrow 1\text{-}3 \text{ concrete examples or stories} \rightarrow \text{abstraction} \quad (1)$$

This is the process of two trained neural networks (human brains) communicating abstractions between each other’s networks using words. In the case of more shared parameters (common culture, familiarity), fewer words need to be said to transmit the same abstract idea between the two humans, and small changes in words can induce substantially different generalizations for humans (Luntz, 2007). As Chomsky has noted, the only indispensable operation of human grammar is that of Merge, which is not unlike compositionality in time and abstraction mentioned above (Chomsky, 1995).

Communication, however, does not only occur externally to other humans using language. It also occurs to ourselves (via our inner voice) and between layers in neural networks using non-word representations. After all, outputs from a layer L are latent variables which serve as inputs to layer $L + 1$. If you come from a machine learning background, you may notice that the diagram is simply the autoencoder setup shifted right by 1.

1.4.6 Sparsity and Compositionality Are Necessary for Causality

Sparsity, compositionality in time and compositionality in abstract representation are necessary properties of causal world models.

Although specifically how to learn causal world models is still an open research question, imparting inductive biases for sparse representations has been useful for improving neural network performance - key-value attention allows sparsity of communication in time (Vaswani et al., 2017), while skip connections allow sparsity of communication in abstractive depth (He et al., 2015), since it allows lower layer world models to communicate directly with higher layers of a network.

1.4.7 Contrastive Attention to Prediction Loss Drives Data Resampling and Relabeling Needed For Causal Models

However, only having a sparse and compositional world model is not sufficient for it to be causal. For a world model to be causal it must be able to accurately model interventions and counterfactuals (Pearl, 2009). For example, “does sunlight catalyze the production of ATP in plants?” requires understanding how the intervention (sunlight) affects your world model of plant energy generation. Counterfactuals are a special type of intervention which is non-factual, for example, “without ATP synthase, would sunlight affect the production of ATP in plants?” If a model is able to answer a large number of these questions over a diversity of interesting interventions, we might consider that it has learned a causal model of the topic of sunlight and ATP.

Note that, with an existing world model of part of the input (say “sunlight”), it becomes easier to understand the question “does sunlight catalyze the production of ATP in plants?” If you didn’t know what sunlight means, on the other hand, your first instinct might be ask “what is sunlight?” to be able to systematically collect new training data about the part of the world model that was unknown. Similar contrastive prediction-loss motivated resampling has been shown to impart generalization with systematicity (Akyurek et al., 2021), especially when inputs are collected for contrastive interventions and outputs have contrastive predictions induced by the same contextual data (see Figure 2).

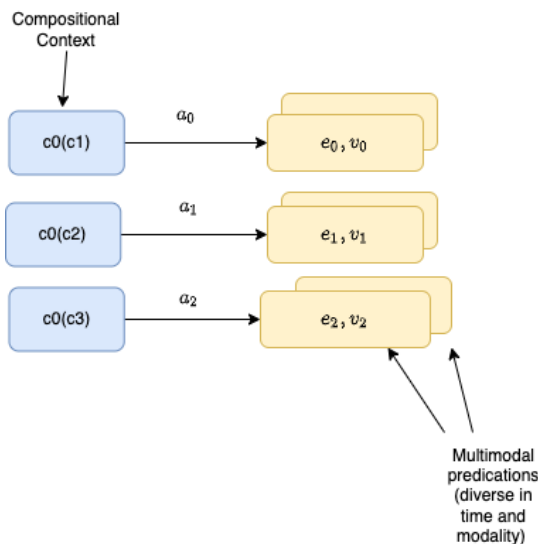


Figure 2: Compositional Contexts and Multi-Modal Predictions

Although there are many works describing useful inductive biases for neural network architectures trained on static training data, learning a policy for data labeling and data resampling over compositional representations needed for continual few-shot learning has not been well-addressed. To motivate this, we draw on observations of human learning, as humans are arguably experts in embodied continual few-shot learning that have behavior policies which have been pretrained by millenia of evolution.

2 Biological Solutions

Humans are efficient online learners who are often able to overcome the challenges of the continual learning environment and produce causal world models which are compositional in time and abstractive depth. Unlike many current machine learning models which run inference in serial, humans run learning, inference and control algorithms in parallel and may run training over the output perceptions from past training loops. Moreover, learning is tightly integrated with memory policies, unlike many contemporary RL agents.

2.1 Prediction

Humans are constantly making predictions about what perceptions they will experience in the future, \hat{y} . Predictions can be short-horizon (ie. where will my arm be in the next second) or long-horizon (what will I do in the next 5 years?). Predictions can be conscious or unconscious - at the conscious level, predictions are the goals, plans, tasks or imagination which drive behavior (see Figure 3). At the unconscious level, they may appear as dreams (Freud, 1900).

2.2 Prediction Loss

When predictions are compared against later perceptions they create prediction loss, $loss(y, \hat{y})$. One concrete form of loss is just the L1 distance, $|\hat{y} - y|$.

In practice, we have more than one possible prediction loss in focus at a specific time, attention to which we mediate via executive function. Although executive function is often defined as the ability to set goals, plan and have impulse control, here we hypothesize a simpler definition of executive function as stateful attention to a stack of predictions and perceptions, $loss(Y, \hat{Y}) = \sum a_i * loss(y_i, \hat{y}_i)$ where y_i is the i th perception on the stack, \hat{y}_i is its prediction and a_i is the attention weight of the pair. In classical RL terminology, this can be viewed as a long-horizon value policy for attention to predictive loss, which is then experienced as goals, plans and impulse control.

2.2.1 Decreasing Prediction Loss

Perceptions are subtly biased towards confirming predictions. Humans are prone to confirmation bias, whereby they create summaries which confirm our predictions without actually learning to control or predict accurately (Garrison and Hoskisson, 1989). At the subconscious level, we experience gestalt, subconsciously filling in missing parts of a whole (Wertheimer, 1923).

However, when prediction loss is high enough, we experience this as an uncomfortable feeling of cognitive dissonance (Kaaronen, 2018), and will take actions to reduce prediction loss by changing y or \hat{y} . Anticipated prediction error can be reduced by taking actions within the external environment to change y to match \hat{y} , while realized prediction error can be changed by modifying network weights so that \hat{y} becomes closer to y (“learning”), a behavior which is aided by using executive function to resample or relabel perceptions (changing y and \hat{y}) (“thinking”).

2.2.2 Increasing Prediction Loss

Although reducing prediction loss is a key drive, if we experience too little prediction loss, we also strive to increase it via curiosity-driven exploration, thus staying balanced within a “golden mean” of error (Plato, 375 BC) formed by a Markov blanket over our predictions.³

2.2.3 Acting to Minimize Expected Prediction Loss

When a prediction involves a ego-perception (ie. something that we will do in the future), in order to minimize expected predictive loss, we attempt to control our environment through action, a process called active inference (Friston et al., 2006) (Figure 3). This can be consciously shaped by summarization and visualization - in sports, positive self-talk and visualization is a technique which is correlated with improved performance by creating prediction loss with respect to target performance (Raalte and Vincent, 2017).

Both conscious and unconscious predictions motivate action, regardless of their origin - ideas that “pop into your head” tend to have connections to past events. For example, when you have just read the sentence “does sunlight catalyze the production of ATP in plants?”, you may be more likely to eat a salad tonight or go outside (Luntz, 2007).




Figure 3: Active Inference

³In regimes with high prediction loss, there is chaos, fear and uncomfortable uncertainty. In regimes with low prediction loss, boredom. In the middle is happiness. As in folk wisdom, “the journey is the reward.”

2.3 Learning to Minimize Prediction Loss

2.3.1 What is Learning?

The most straightforward way to minimize realized prediction loss is by updating weights or growing neural connections such that they directly change the world model so that the predictions are closer to perceptions. In artificial neural networks this occurs via backpropagation (Rumelhart et al., 1986) or many other optimization methods, and in the human brain, via Hebbian learning (Hebb, 1949).

2.3.2 Taking actions in the world to resample data

To assist with learning, as with most animals, we can resample data via physical movement which results in new contrastive perceptions (Friston and Kiebel, 2009). For example, at a cocktail party, you may move your head to look at a speaker’s lips to be able to better distinguish the utterances of the speaker from the background noise. We can also employ longer-term learning loops to resample data, for example, calling a friend for a second opinion on how to approach a situation. In this case, you acquire two summary labels for the situation (one from yourself, one from your friend). This contrastive data makes it easier to learn accurate causal world models.

2.3.3 Taking internal actions to resample data

Although taking actions is an effective way to gather contrastive data, it is also quite expensive and impractical in many settings. Shifting attention to novel internal perceptions allows us to sample new perceptions and predictions without motor control. Humans have learned to hack this process for pleasure and control. Attending to the sensations of your body which are easily predictable via meditation may be a method to consciously lower prediction loss by changing the contents of the prediction stack (Jamieson, 2015; Pagnoni, 2019). Mindfulness lowers reward prediction error and activates the putamen in practitioners (Kirk and Montague, 2015). Consciousness presence has been hypothesized to be due to top-down predictions successfully suppressing informative interoceptive signals (Seth et al., 2012).

2.4 Few-Shot Continual Learning

However, even with policies for data resampling, learning a causal world model from a handful of examples and random neural network weights tends to be infeasible. Humans, however, benefit from evolutionarily trained data encoded in the inductive biases of neuroanatomy. Learning policies which tap into these existing facilities are much faster than learning without.

2.4.1 A Grammar of Learning?

The Complementary Learning Systems theory (McClelland et al., 1995) proposes that the hippocampus encodes “fast” learning policies resulting in direct storage of incoming perceptions, while the neocortex replays these perceptions slowly (interspersed with other perceptions) to achieve a slow learning of abstractions and subconscious behaviors.

We propose that fast learning may be a behavior that uses the tools of summarization, active inference and multiple stateful passes through a pre-trained network (which activate latent abstractions) such that the resulting output data forms a batch which has sparse and compositional representations such that it can be more easily learned, a phenomenon we call learning a communicative mapping to a *grammar of learning*, similar to an online prompt engineering problem. Specifically, for any network θ which has been pre-trained on dataset D , there exists a data distribution D' such that an input $d' \in D'$ will produce a sparse, compositional distribution of outputs which can be resampled and relabeled contrastively, thus inducing efficient learning. We call this D' a grammar of learning for the network and the original grammar of learning without pre-training data a universal grammar of learning⁴. Given a novel input e , if the network can find a series of transformations to map e into a latent representation d' , then e can be more easily learned by association with d' .

2.4.2 Resampling and Relabeling Data Via Memory and Hindsight Summarization

The average human reaction time is 250ms and the human eye has a resolution of 576 megapixels, resulting in a data stream of at least 7GB/second. To have the compositional and sparse representations required for learning of causality, we must significantly reduce the dimensionality of this firehose of data. This done by hindsight summarization (Figure 4). Hindsight summarization has access to a memory and a grammar of learning to enable rapid non-linear data transformations via existing world models. Creating a hindsight summary changes the composition of the perception stack attended by executive function, creating new prediction losses which can be used directly as training data (Figure 5) or indirectly to produce further summaries, forming a stream of consciousness which quickly changes the composition of the entire perception stack.

⁴This is largely just an inductive bias for sparse and compositional input data

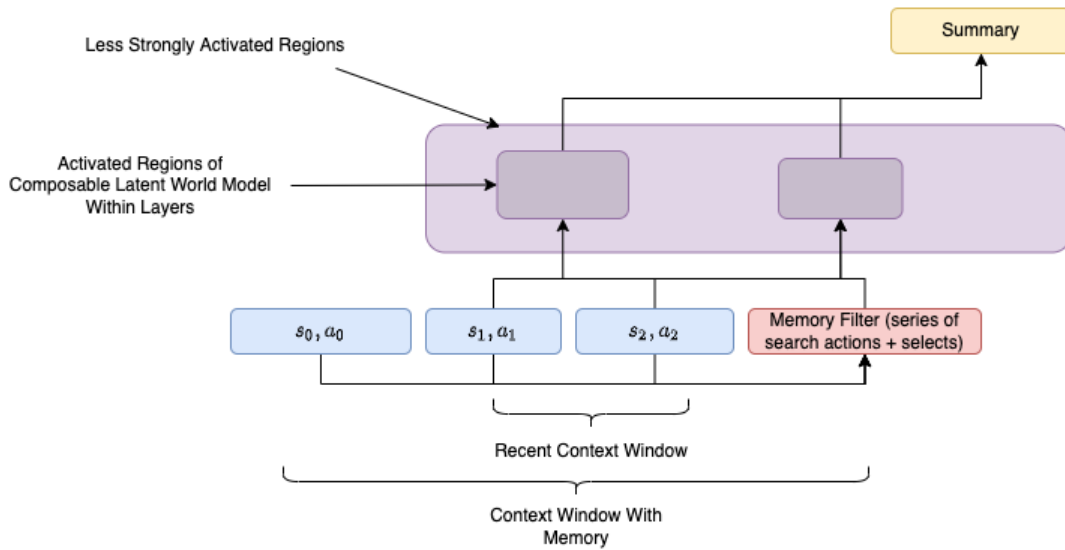


Figure 4: Relabeling Data for a Causal World Model Data Via Hindsight Summarization

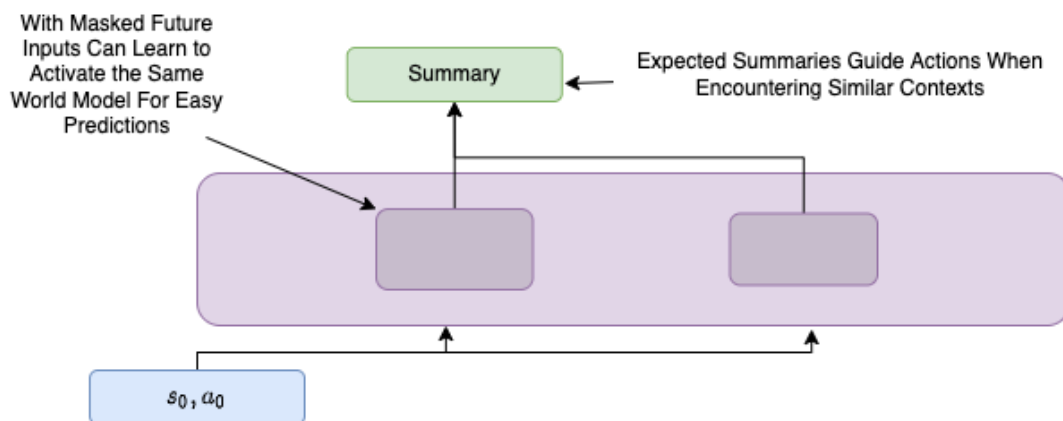


Figure 5: Continuous Training on Relabeled Data to Motivate Behavior To Connect Summaries to Future Actions

Streams of consciousness like beneficial self-talk (Raalte and Vincent, 2017) and cognitive behavioral therapy (Manjaly and Iglesias, 2020) may be learned behaviors to modify experienced prediction loss using hindsight summarization.

These can be complex and diverse. Adopting a growth mindset by framing unpredictable perceptions as a hypothesis test (so that any outcome is now expected with a certain probability) is correlated with reduced cognitive dissonance during experimentation for the authors. Making explicit predictions before reading (a form of data relabeling via increasing attention to prediction loss) has been shown to improve human learning (Thomas-Fair, 2005; Brod, 2021), and may be an example of consciously increasing prediction loss to focus learning mechanisms. They may also occur unconsciously - humans make long-range predictions at multiple levels of latent depth when reading (Caucheteux et al., 2021).⁵

Regions of the brain associated with executive function are larger in humans than other primates (Donahue et al., 2018), and humans tend to be better at few-shot learning than other animals, which may be related to the ability to set conscious plans (Fuster, 2017).

Analogues to simple patterns of hindsight summarization for non-embodied agents can also be seen in SOTA approaches to self-supervised learning, such as data2vec (Baevski et al., 2022), where summaries can be considered to be the target latent representations of the teacher, and the future masking of Figure 5 is analogous to the input masking in the student model in Baevski et al. (2022). However, data2vec does not use multiple steps of conditional computation, as humans do.

Hindsight summarization can also be compared to other hindsight schemes such as HER (Andrychowicz et al., 2018), however summarization is a learned path function over the past trajectories rather than a deterministic function of the last state, as in HER. Unlike generalized hindsight (Li et al., 2020), hindsight summarization encodes information about a trajectory of observations, not just rewards, and runs on a much faster timescale, possibly summarizing a sequence of perceptions which span intervening training epochs.

2.4.3 Hypothesis testing via summarization

Sequential steps of hindsight summarization with intervening training and counterfactual sampling may be recognizable as the Scientific Method, a well-known meta-learning algorithm for scientific discovery. See Figure 6 for how we might model this stream of consciousness and use insights to inform future behavior.

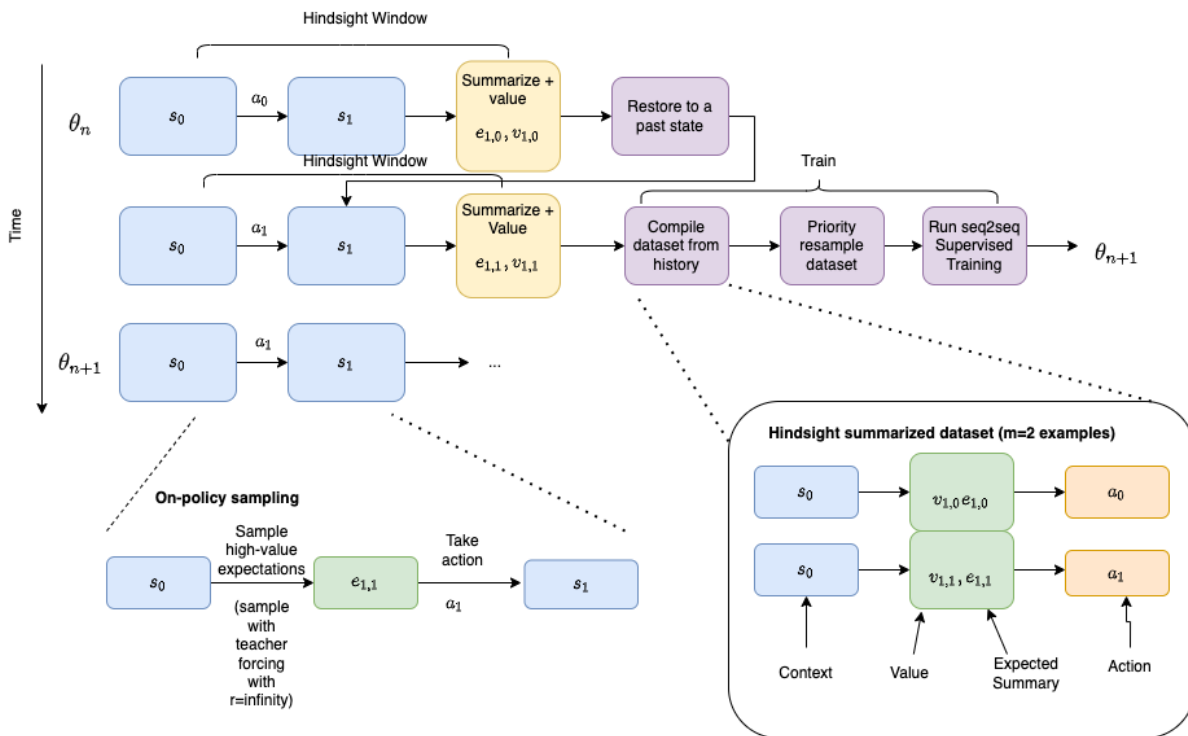


Figure 6: Hypothesis Testing

⁵Eric Hayot, a comparative literature professor, once told one of the authors to “kick the ladder” - that is, after finishing a draft, rewrite it with the benefit of hindsight so that you weave anticipation of the summary into the piece, and remove the confusing scratch writings you used to bootstrap the summary for yourself as an author. In this way, the reader’s executive function naturally anticipates the conclusion with less cognitive dissonance.

3 Conclusion

We have defined the few-shot continual learning task from first principles and hypothesized an evolutionary motivation and mechanism of action for executive function as a contrastive value policy which resamples and relabels perception data via hindsight summarization to minimize attended prediction error, similar to an online prompt engineering problem. This is enabled by the use of a memory policy and a pretrained network with inductive biases for a grammar of learning and is trained to maximize long-term survival. We have shown how this model of executive function can be used to implement hypothesis testing and may explain observations of few-shot learning and neuroanatomy.

There is ample room for future work. Hypothesis testing is only one example of a stream of consciousness that can be implemented using this architecture. In contrast, humans are experts at using complicated patterns of hindsight summarization and intervening training and action for various meta-learning strategies. As a simple example, how to implement imitation learning where a student learns not only to copy a teacher, but who to trust within a competitive environment?

The parallels between human learning and machine learning also remain to be explored more rigorously. Although we have shown correlations in empirical observations and a hypothesized mechanism of action, and have proved the efficacy of the architectures above on real-world text RL problems (redacted), to show causality and biologically relevant mechanisms requires additional study. We anticipate isomorphism in the objective functions, training data and behavior of in silico and in vitro algorithms but not necessarily the specific architectures. For example, the brain implements information processing algorithms in spiking neurons and has evolved distinctive regions that may not necessarily be seen in machine learning models. We do, however, see concrete opportunities to use isomorphism in purpose and algorithmic structure to better inform our understanding of learning in both humans and machines.

References

- Ekin Akyürek, Afra Feyza Akyürek, and Jacob Andreas. Learning to Recombine and Resample Data for Compositional Generalization. *arXiv:2010.03706 [cs]*, June 2021. URL <http://arxiv.org/abs/2010.03706>. arXiv: 2010.03706.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight Experience Replay, 2018. [eprint: 1707.01495](https://arxiv.org/abs/1707.01495).
- Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language. February 2022. URL <https://arxiv.org/abs/2202.03555v1>.
- Richard Bellman. The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6):503–515, November 1954. ISSN 0002-9904, 1936-881X. URL <https://projecteuclid.org/journals/bulletin-of-the-american-mathematical-society/volume-60/issue-6/The-theory-of-dynamic-programming/bams/1183519147.full>. Publisher: American Mathematical Society.
- Garvin Brod. Predicting as a learning strategy. *Psychonomic Bulletin & Review*, 28(6):1839–1847, December 2021. ISSN 1531-5320. doi: 10.3758/s13423-021-01904-1. URL <https://doi.org/10.3758/s13423-021-01904-1>.
- Charlotte Caucheteux, Alexandre Gramfort, and Jean-Remi King. Long-range and hierarchical language predictions in brains and algorithms. *arXiv:2111.14232 [cs, q-bio]*, November 2021. URL <http://arxiv.org/abs/2111.14232>. arXiv: 2111.14232.
- Noam Chomsky. *The Minimalist Program*. MIT Press, Cambridge, MA, USA, September 1995. ISBN 978-0-262-03229-2.
- Chad J. Donahue, Matthew F. Glasser, Todd M. Preuss, James K. Rilling, and David C. Van Essen. Quantitative assessment of prefrontal cortex in humans relative to nonhuman primates. *Proceedings of the National Academy of Sciences*, 115(22):E5183–E5192, May 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1721653115. URL <https://www.pnas.org/content/115/22/E5183>. Publisher: National Academy of Sciences Section: PNAS Plus.
- Sigmund Freud. *The Interpretation of Dreams*. Empire Books, Place of publication not identified, 1900. ISBN 978-1-61949-232-5.

⁵The authors thank Volodymyr Kuleshov, Pang Wei Koh, Lili Chen, Michael Janner, Andrej Karpathy, Laura Deming, Minn Kim, Maarten Bosma, Ruoxi Wang, Chip Huyen, Kanjun Qiu, Eugene Lengerich, Rebecca Lengerich and Charles Antle for insightful feedback on drafts of this work. Empirical experiments were run on OpenAI’s fine-tuning API.

- Karl Friston and Stefan Kiebel. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521):1211–1221, May 2009. doi: 10.1098/rstb.2008.0300. URL <https://royalsocietypublishing.org/doi/10.1098/rstb.2008.0300>. Publisher: Royal Society.
- Karl Friston, James Kilner, and Lee Harrison. A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1-3):70–87, July 2006. ISSN 09284257. doi: 10.1016/j.jphysparis.2006.10.001. URL <https://linkinghub.elsevier.com/retrieve/pii/S092842570600060X>.
- Joaquín M. Fuster. Chapter 1 - Prefrontal Executive Functions Predict and Preadapt. In Elkhonon Goldberg, editor, *Executive Functions in Health and Disease*, pages 3–19. Academic Press, San Diego, January 2017. ISBN 978-0-12-803676-1. doi: 10.1016/B978-0-12-803676-1.00001-5. URL <https://www.sciencedirect.com/science/article/pii/B9780128036761000015>.
- Matthew P. H. Gardner, Geoffrey Schoenbaum, and Samuel J. Gershman. Rethinking dopamine as generalized prediction error. *Proceedings of the Royal Society B: Biological Sciences*, 285(1891):20181645, November 2018. ISSN 0962-8452. doi: 10.1098/rspb.2018.1645. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6253385/>.
- James W. Garrison and Kenneth Hoskisson. Confirmation Bias in Predictive Reading. *The Reading Teacher*, 42(7):482–486, 1989. ISSN 0034-0561. URL <https://www.jstor.org/stable/20200195>. Publisher: [Wiley, International Reading Association].
- David Ha and Jürgen Schmidhuber. World Models. *arXiv:1803.10122 [cs, stat]*, March 2018. doi: 10.5281/zenodo.1207631. URL <http://arxiv.org/abs/1803.10122>. arXiv: 1803.10122.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, December 2015. URL <http://arxiv.org/abs/1512.03385>. arXiv: 1512.03385.
- D. O. Hebb. *The organization of behavior; a neuropsychological theory*. The organization of behavior; a neuropsychological theory. Wiley, Oxford, England, 1949. Pages: xix, 335.
- Jakob Hohwy. *The predictive mind*. 2014. ISBN 978-0-19-176635-0. OCLC: 940606476.
- Graham Jamieson. A unified theory of hypnosis and meditation states: the interoceptive predictive coding approach. January 2015. doi: 10.13140/2.1.1447.5844.
- Roope Oskari Kaaronen. A Theory of Predictive Dissonance: Predictive Processing Presents a New Take on Cognitive Dissonance. *Frontiers in Psychology*, 9:2218, 2018. ISSN 1664-1078. doi: 10.3389/fpsyg.2018.02218. URL <https://www.frontiersin.org/article/10.3389/fpsyg.2018.02218>.
- Ulrich Kirk and P. Read Montague. Mindfulness meditation modulates reward prediction errors in a passive conditioning task. *Frontiers in Psychology*, 6, 2015. ISSN 1664-1078. URL <https://www.frontiersin.org/article/10.3389/fpsyg.2015.00090>.
- Alexander C. Li, Lerrel Pinto, and Pieter Abbeel. Generalized Hindsight for Reinforcement Learning. *arXiv:2002.11708 [cs, stat]*, February 2020. URL <http://arxiv.org/abs/2002.11708>. arXiv: 2002.11708.
- Frank Luntz. *Words That Work: It's Not What You Say, It's What People Hear*. 2007. URL https://www.amazon.com/dp/B00JJ9QL0W/ref=dp-kindle-redirect?_encoding=UTF8&btkr=1.
- Zina-Mary Manjaly and Sandra Iglesias. A Computational Theory of Mindfulness Based Cognitive Therapy from the “Bayesian Brain” Perspective. *Frontiers in Psychiatry*, 11:404, 2020. ISSN 1664-0640. doi: 10.3389/fpsyg.2020.00404.
- James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, July 1995. ISSN 1939-1471, 0033-295X. doi: 10.1037/0033-295X.102.3.419. URL <http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.102.3.419>.
- Giuseppe Pagnoni. The contemplative exercise through the lenses of predictive processing: A promising approach. *Progress in Brain Research*, 244:299–322, 2019. ISSN 1875-7855. doi: 10.1016/bs.pbr.2018.10.022.
- Judea Pearl. *CAUSALITY, 2nd Edition, 2009*. 2009. URL <http://bayes.cs.ucla.edu/B00K-2K/>.
- Plato. *Republic, Book 7*. 375 BC. URL <http://www.perseus.tufts.edu/hopper/text?doc=Perseus%3Atext%3A1999.01.0168%3Abook%3D7&force=y>.
- Judy L. Van Raalte and Andrew Vincent. Self-Talk in Sport and Performance, March 2017. URL <https://oxfordre.com/psychology/view/10.1093/acrefore/9780190236557.001.0001/acrefore-9780190236557-e-157>. ISBN: 9780190236557.

- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, October 1986. ISSN 1476-4687. doi: 10.1038/323533a0. URL <https://www.nature.com/articles/323533a0>. Number: 6088 Publisher: Nature Publishing Group.
- Anil Seth, Keisuke Suzuki, and Hugo Critchley. An Interoceptive Predictive Coding Model of Conscious Presence. *Frontiers in Psychology*, 2:395, 2012. ISSN 1664-1078. doi: 10.3389/fpsyg.2011.00395. URL <https://www.frontiersin.org/article/10.3389/fpsyg.2011.00395>.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018. ISBN 0-262-03924-9.
- Ursula Thomas-Fair. The Power of Prediction: Using Prediction Journals to increase comprehension in kindergarten. page 17, 2005.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need, 2017. eprint: 1706.03762.
- Max Wertheimer. Untersuchungen zur Lehre von der Gestalt. II. *Psychologische Forschung*, 4(1):301–350, January 1923. ISSN 1430-2772. doi: 10.1007/BF00410640. URL <https://doi.org/10.1007/BF00410640>.