Prediction in Natural and Artificial Intelligence

Keith L. Downing

Department of Computer Science
The Norwegian University of Science and Technology (NTNU)
Trondheim, Norway
keithd@ntnu.no

June 17, 2025

Charting a path from early artificial neural networks to the contemporary vision of the predictive brain, with rich forays into biology and evolution, this book explains the buzz about brains as engines of prediction.
(Andy Clark, University of Sussex)

MIT Press 2023

Downing's reach is omnidirectional. He connects the roots and new growth of deep learning with math, neuroscience, and evolutionary biology, ethology and computer science to show how intelligence emerged in animals and is emerging in machines.
(Josh Bongard, University of Vermont)

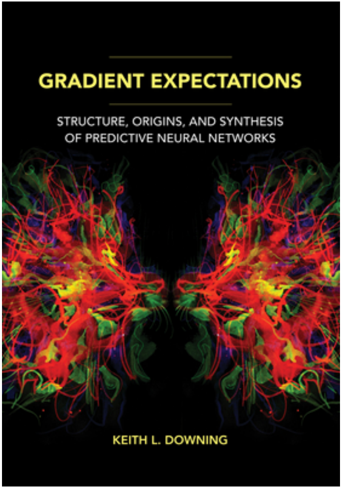

The Brain = A Prediction Machine ??

- The short punch line of this book is that brains are **foretelling devices**, and their **predictive** powers emerge from the various rhythms they perpetually generate...*Rhythms of the Brain* (Buzsaki, 2006)
- The capacity to **predict** the outcome of future events – critical to successful movement – is likely, the ultimate and most common of all global brain functions...*i of the Vortex* (Llinas, 2001)
- The mystery is, and remains, how mere matter manages to give rise to thinking, imagining, dreaming, and the whole smorgasboard of mentality, emotion and intelligent action....But there is an emerging clue...The clue can be summed up in a single word: **prediction**. To deal rapidly and fluently with an uncertain and noisy world, brains like ours have become **masters of prediction**...*Surfing Uncertainty* (Clark, 2016)
- ...the core task of all brains... is to regulate the organism's internal milieu – by responding to needs and, **better still, by anticipating needs and preparing to satisfy them before they arise**... ”Anticipatory regulation” replaces the more familiar ”homeostatic regulation”...*Principles of Neural Design* (Sterling & Laughlin, 2015)

Sensing and Acting in a Slow World

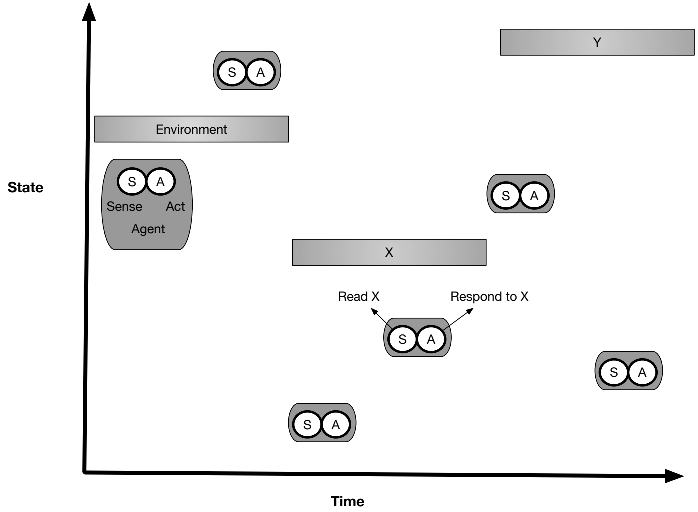

Life (and DEATH) in the Fast Lane

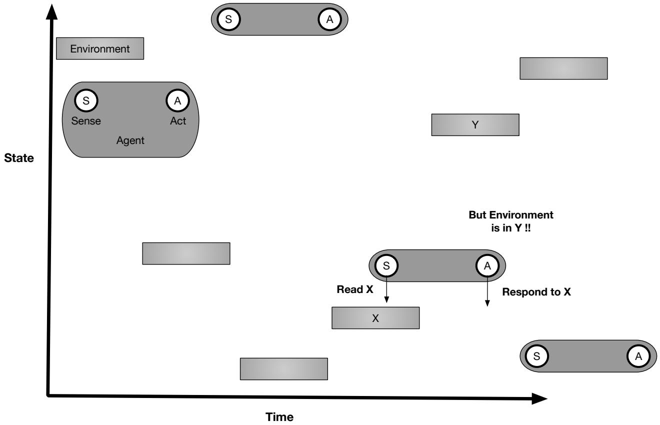

* For a **mobile** agent, relative frequency of environmental change increases.

Predict to Survive

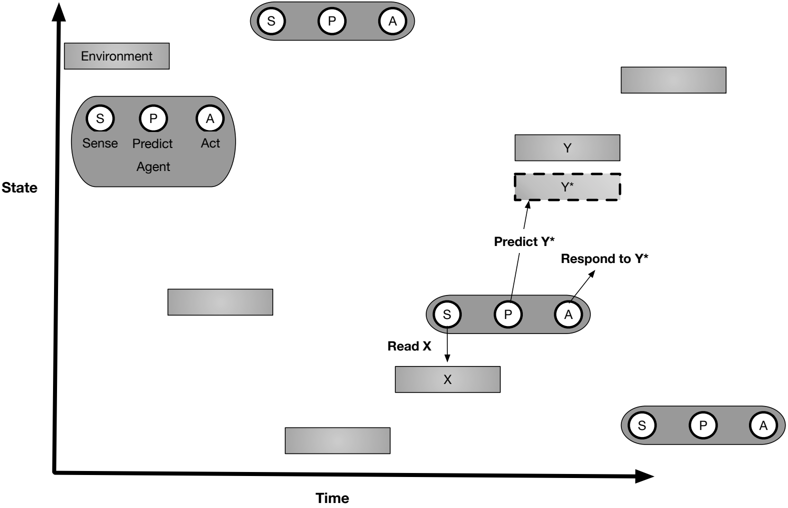

* Mobile agents need prediction.

The Brain-Eating Sea Squirt

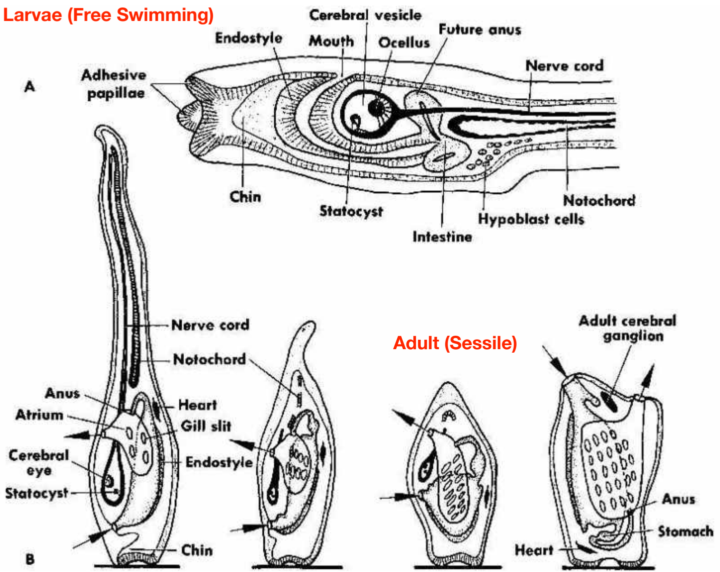

Implicit Prediction using Gradients

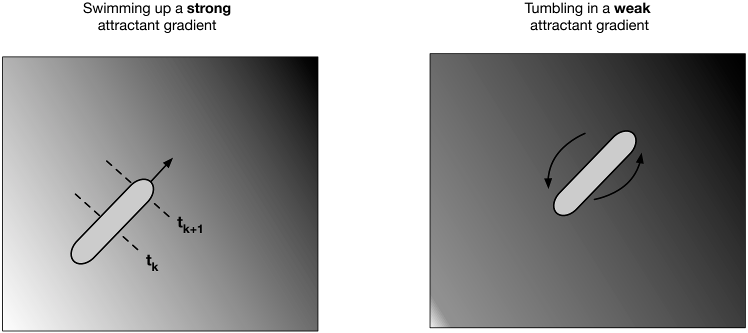

Gradient = $A(t_{k+1}) - A(t_k)$ A(t) = attractant read by head sensor at time t.

Gradients are simple, cheap (and amazingly accurate) predictors of future states: Future = Present + Gradient

Gradient = Derivative

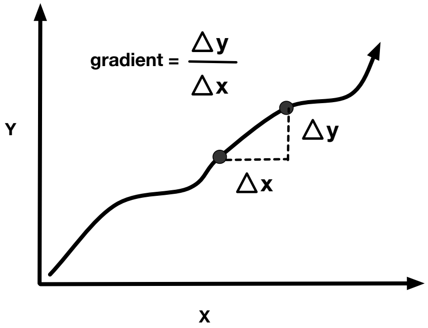

Situation	X	Y
Bacterial Foraging	Location	Nutrients
Finance	Time	Stock Price
Thermostat	Heat	Temperature
Deep Learning	Connection Weights	Output Error
Evolutionary Computation	Genotype	Fitness

Prediction via Gradients

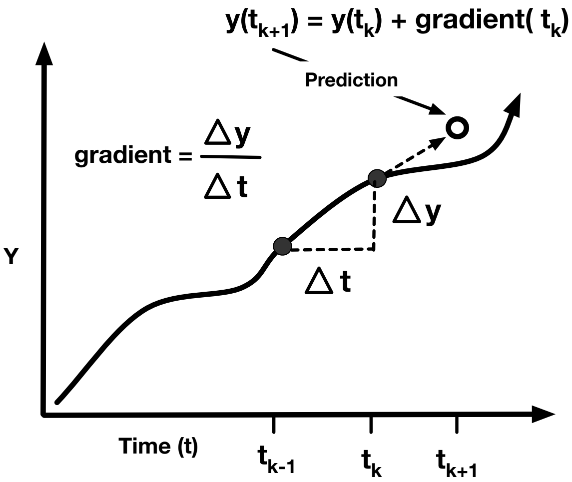

Gradient-Driven Behavior via Chemistry

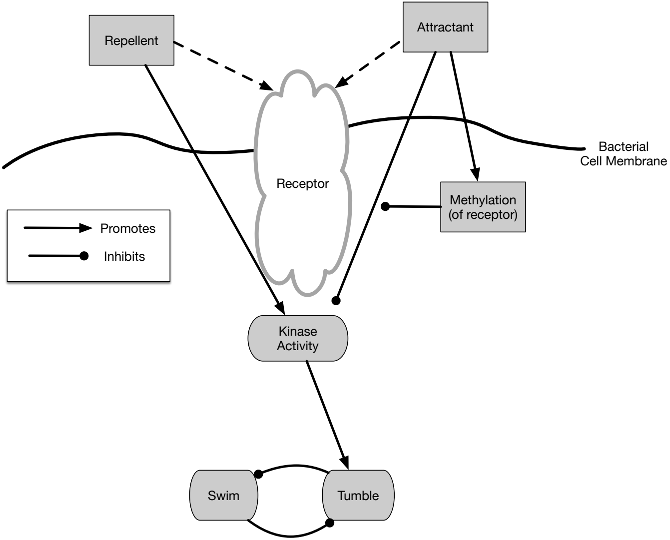

Bacterial Response to the Gradient

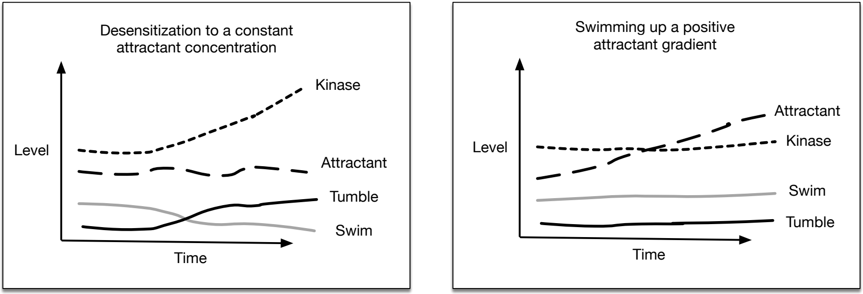

- Bacterium responds to the gradient of the attractant, not simply its current value.
- By moving, the bacterium encodes a spatial gradient as a temporal gradient, which it can then detect using receptors in a single location.

Temporal Differentiation via Habituation

Neurons can detect gradients too!

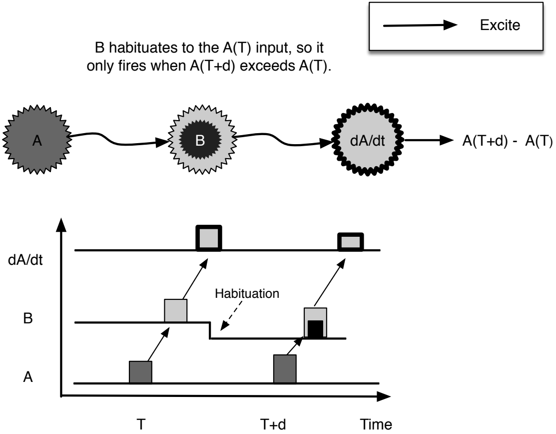

This is how (nematode worm) C. Elegans navigates. (Larsch et. al., 2015)

Temporal Differentiation via Delayed Inhibition

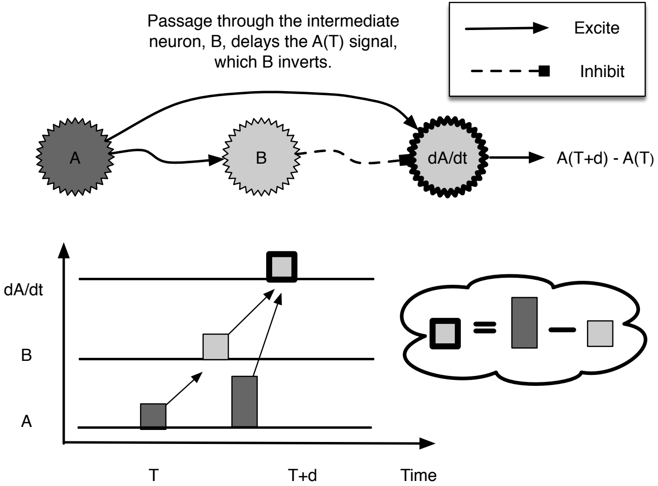

Temporal Differentiation via Synaptic Depression

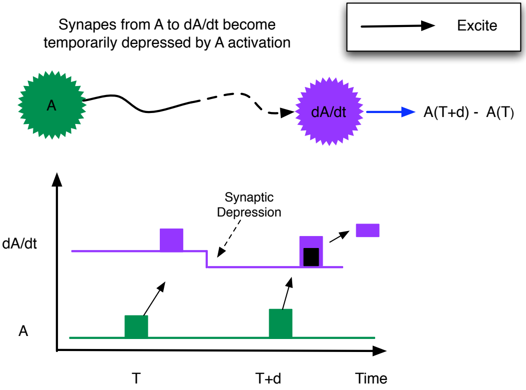

Tripp, B. and Eliasmith, C. (2010), Population Models of Temporal Differentiation, *Neural Computation*, 22, pp. 621-659.

Prediction via Averages

- A simple average (over time or space) can give a good prediction of a variable's next state ($S_{t+\triangle}$).
  - Time $\rightarrow$ History: $S_{t-k}, S_{t-k+1}, ... S_{t-1}, S_t$
  - Space $\rightarrow$ current value of the same variable in nearby regions, e.g. concentrations of a particular chemical in neighboring cells.
- Averaging $\rightarrow$ summing, integrating.
- Averages in space or time can also determine the **current** value S_t if it's unknown. These are also called *predictions* despite having a present rather than future tense.
- Neurons often average: they aggregate and scale signals over space and time (remember), and also leak (forget).
- Averages can both contribute to gradients and combine with gradients to support prediction.

Predictions -vs- Goals

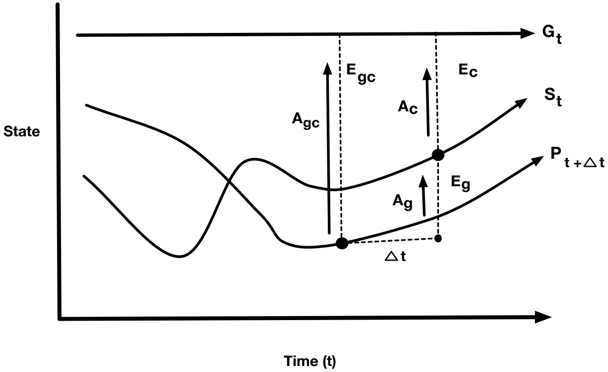

- G_t = Goal at time t.
- S_t = System state at time t.
- $P_{t+\triangle t}$ = Prediction (guess) at time t of state at time $t + \triangle t$.
- E = error of guess (g), control (c), control relative to guess (gc).
- A = actions to reduce errror of guess (g), control (c), both (gc).

Prediction ≈ Control

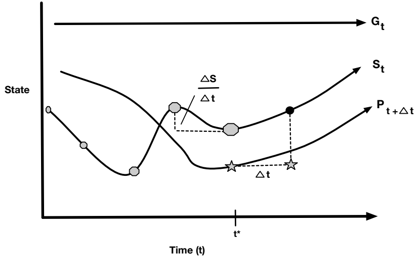

Prediction

$$E_{t+\triangle t} = \Gamma(G \vec{-} S_t) = \underbrace{k_g(G - S_t) + k_g \frac{\triangle(G - S_t)}{\triangle t}}_{gradient-based} + \underbrace{k_a \sum_{j=0}^{M} w_j(G - S_{t-j\triangle t})}_{average-based} \quad (1)$$

PID Control

$$u_t = k_p e_t + k_d \frac{\triangle e_t}{\triangle t} + k_i \sum_{j=0}^{t} e_j \quad (2)$$

Same neural circuits evolved for control could be reused for prediction.

Redundancy Reduction in Visual Perception

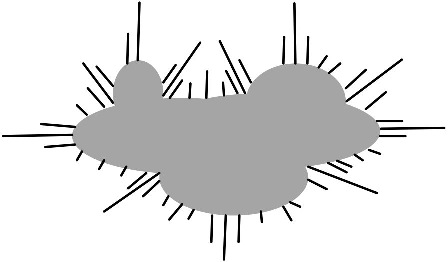

When we begin to consider perception as an information-handling process, it quickly becomes clear that much of the information received by any higher organism is redundant. Sensory events are highly interdependent in both space and time..... we can make better-than-chance inferences with respect to the prior and subsequent states of these receptors.. (Fred Attneave, 1954)

- *Efficient Coding* - Oliver(1952) - similar, for telecom.
- *Predictive coding* coined - Srinivasan et. al.(1982)

Predictive Coding in Neural Circuits

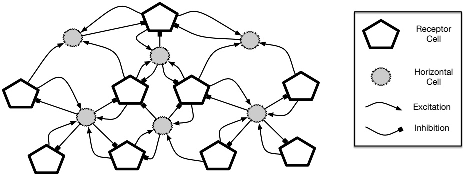

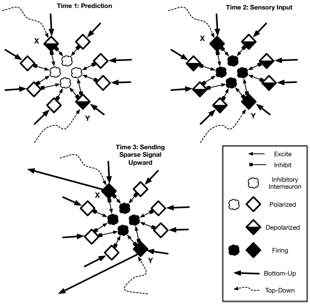

Srinivasan et. al. (1982)

Hawkins(2015)

Predictive Coding in Artificial Neural Networks

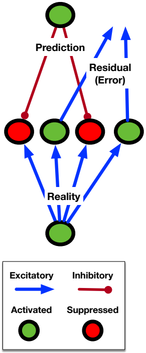

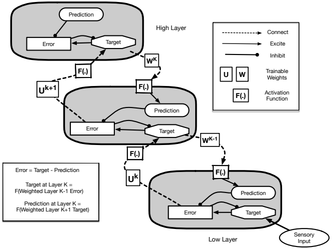

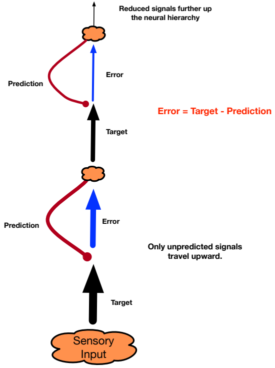

Neural Essence　　Rao + Ballard (1999)　　General Model

Predictive Coding Networks and Control

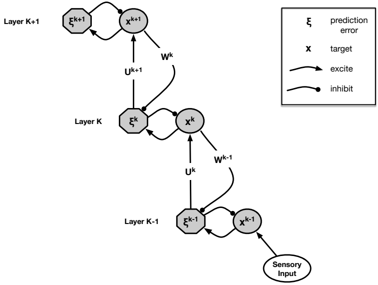

Layered Predictors

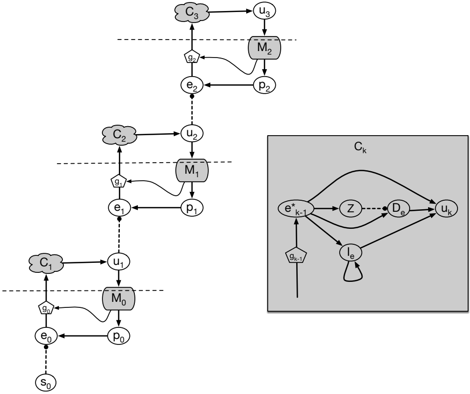

Hierarchy of PID Controllers

Prediction in Diverse Brain Regions

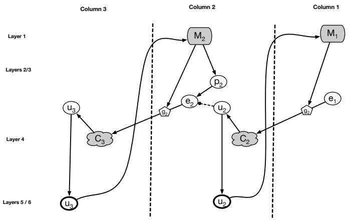

Neocortex

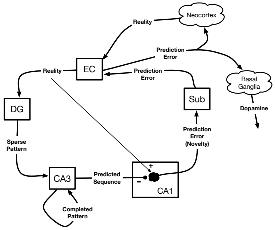

Hippocampus

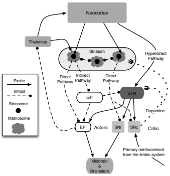

Basal Ganglia

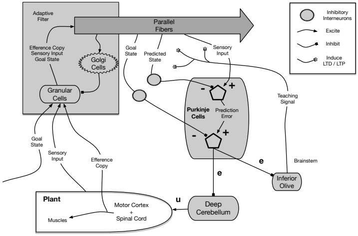

Cerebellum

Backpropagation: Cornerstone of Deep Learning

Data (Cases)

Features **Class**

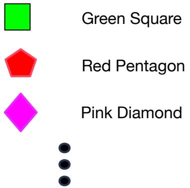

Complex Deep Learning problems require LOTS of data !!

Gradient Descent

$$\text{Gradient (G)} = \frac{\Delta \text{Error}}{\Delta \text{Weight}}$$

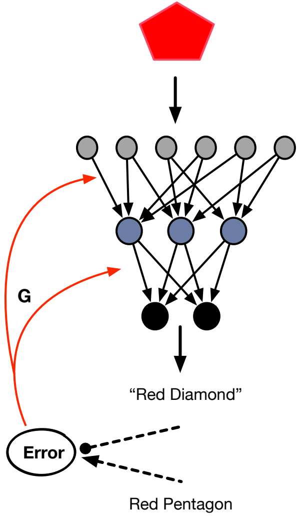

Gradient enables **prediction** of future error resulting from a weight change.

Gradients: Global and Ubiquitous

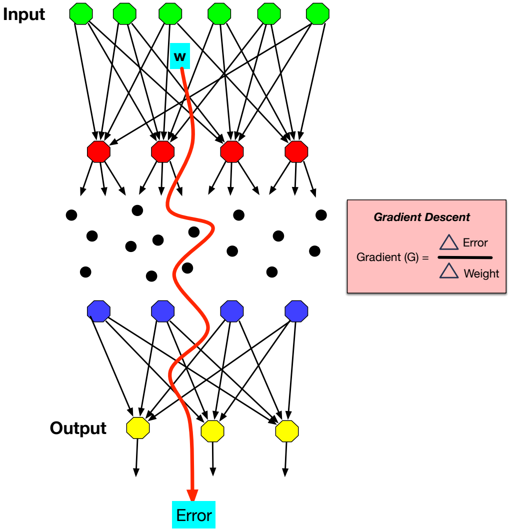

Learning based on these long-distance gradients → recent AI success. But not biologically plausible.

Descending the Error Gradient

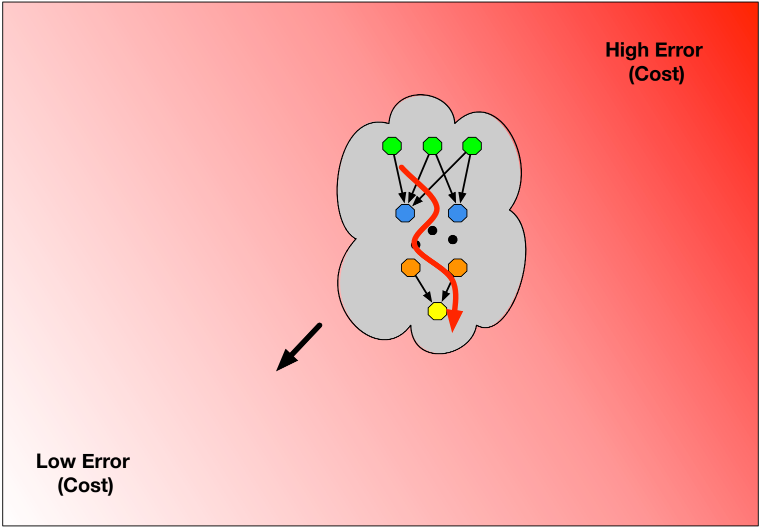

Weight modifications based on **global** (distant) network relationships $\longrightarrow$ System descends global error gradient. **Metric = Total Error**

Connectionism and Prediction: Energy Nets

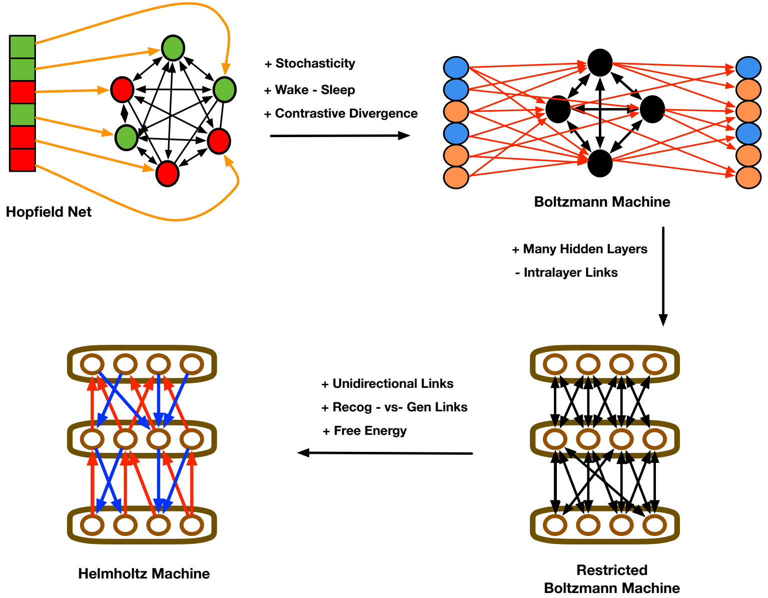

Alternating Phases: Interpretation + Prediction

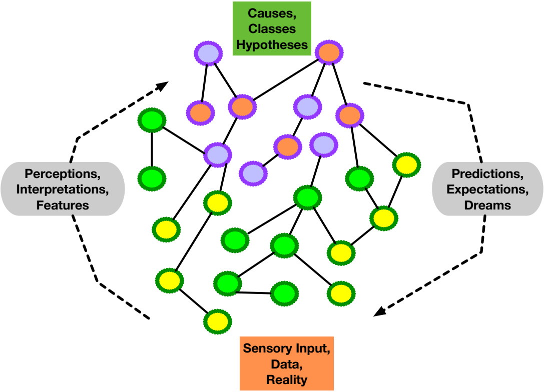

Wake-Sleep Training

- Wake phase (based on data)
- Sleep / Dream phase (based on model-generated patterns = predictions)
- Different variations in Boltzmann, Restricted Boltzmann and Helmholtz Machines.

Energy Metrics for Neural Networks

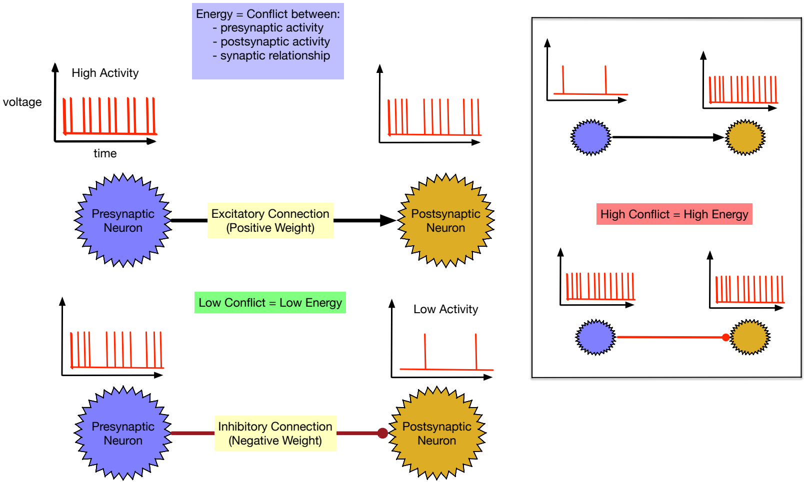

$$Energy = \sum_{pre,post} -X_{pre} X_{post} W_{pre \rightarrow post}$$

Energy Gradients

- Minimize energy instead of error.
- Each weight's contribution to energy is **only local**:

$$\frac{\partial Energy}{\partial W_{ab}} = \frac{\partial}{\partial W_{ab}} \sum_{j,k} -X_j X_k W_{jk} = -X_a X_b$$

where $X_j, X_k, W_{j,k} \in [-1, 1]$

- Learning = Adjusting weights to reduce energy.
- Learning = Descending the energy gradient.

$$\triangle W_{ab} = -\lambda \frac{\partial Energy}{\partial W_{ab}} = \lambda X_a X_b$$

where λ = learning rate

- This is very Hebbian, very biological..very ALIFE !!
- X_a, X_b and Energy take many forms (in different models), but learning remains Hebbian.

Descending the Free-Energy Gradient

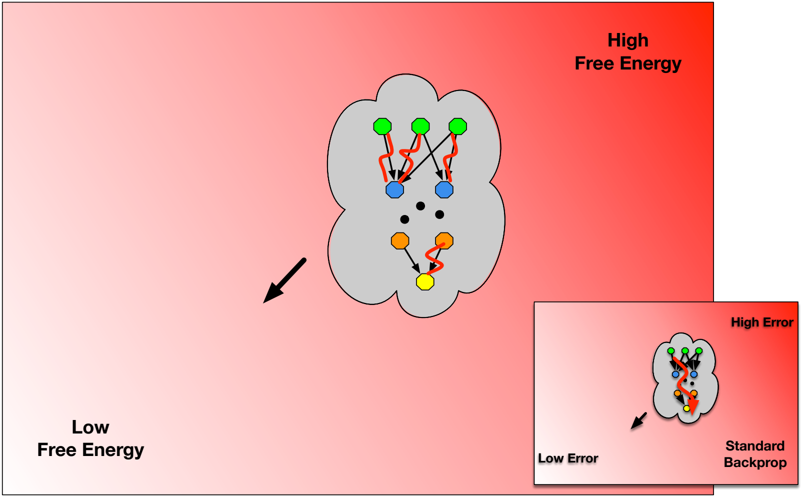

Weight modifications based on **local** network relationships $\longrightarrow$ System moves down the global free-energy gradient.

Predictive Coding instead of Backprop

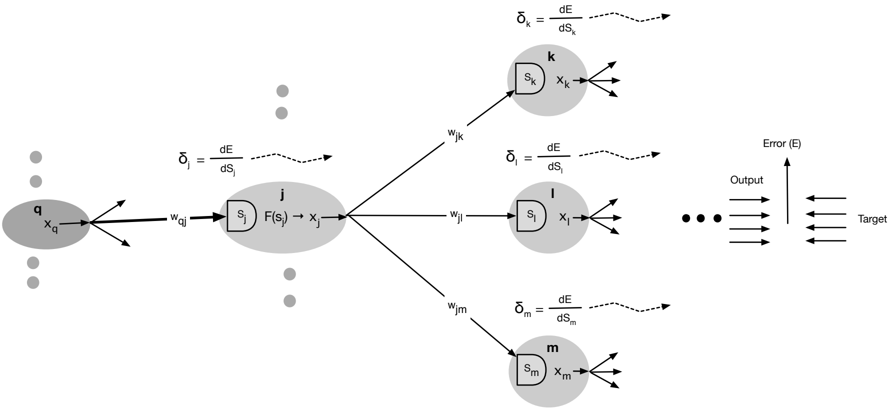

- Replacing δ (a long-distance gradient) with ξ (a prediction error) yields an energy network with respectable classification abilities.
- Rafal Bogacz et. al. (2017, 2019, 2021) - they do the math...and code.

Prediction Enables Supervised Learning

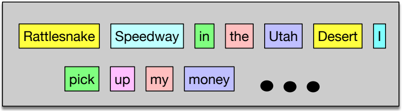

Features	Label
Rattlesnake Speedway in	the
Speedway in the	Utah
in the Utah	Desert
the Utah Desert	I

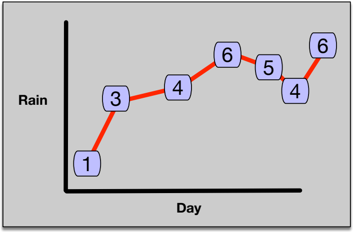

Features			Label
1	3	4	6
3	4	6	5
4	6	5	4
6	5	4	6

Prediction = easy way to generate datasets

Reinforcement Learning

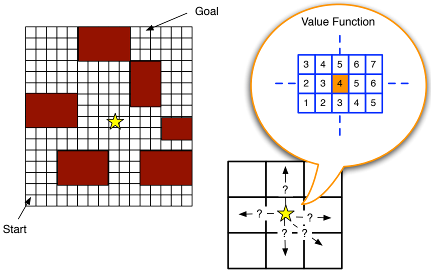

- Do series of actions (policy) to get from start to goal.
- Receive **intermittent** feedback (i.e. reward)
- Over **many trials**, learn a good policy.
- RL uses experience to generate data cases (features → target):

  Model: (situation, action) → (next situation, reward)
  Critic: situation → value
  Actor (Policy): situation → action
  * All targets are **dynamic estimates**.

Prediction of Future Reward

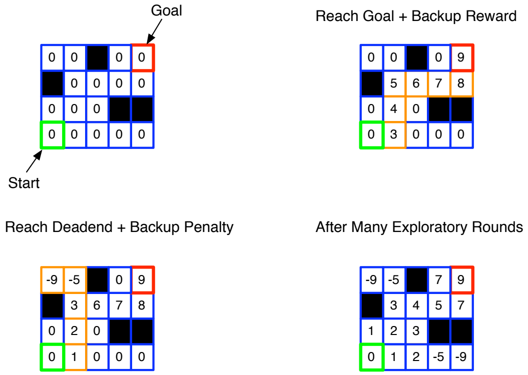

- V(s) = Value of state s = predicted cumulative reward from s to a goal state.

Temporal Differencing (RL Variant)

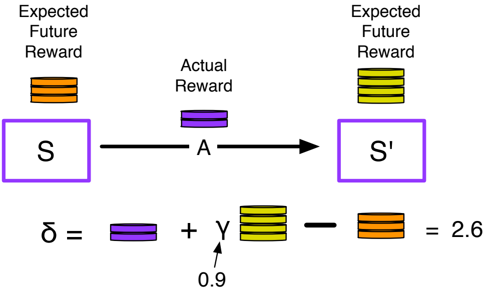

$$\triangle \mathbf{V}(\mathbf{S}) = \lambda \delta$$

where $\lambda =$ learning rate, $\gamma =$ discount factor, and $\delta =$ TD Error

Bootstrapped Prediction Improvement

Prediction of Sum Reward($t_k \rightarrow t_{Final}$) updated by

Prediction of Sum Reward($t_{k+1} \rightarrow t_{Final}$)

Adaptation driven by **gradients of predictions**.

TD-Gammon (Tesauro, 1995): RL + NN

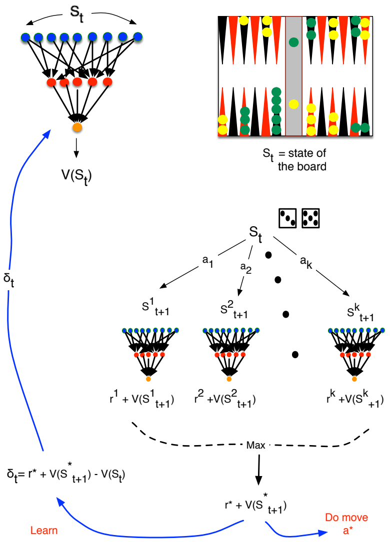

Net learns value function, V(s), using **Predicted-reward gradient:** $\frac{\triangle V(s)}{\triangle w}$

AlphaGo (DeepMind, 2016): RL + Deep Nets

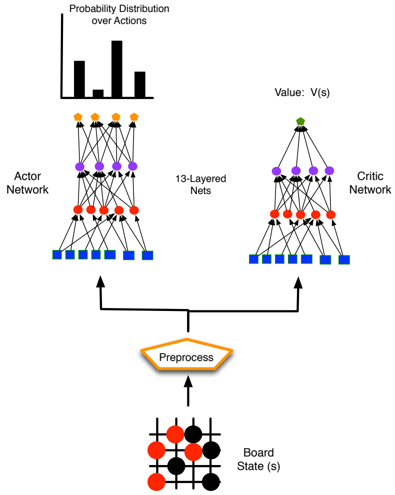

Taking Tesauro's (1995) work to a higher level with deep nets.

Gradients of Deep Reinforcement Learning

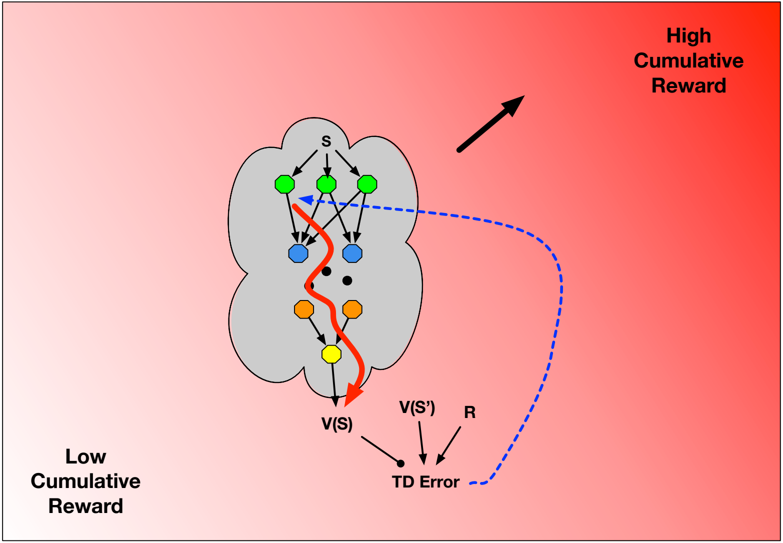

Learning via long gradients ($\frac{\triangle V(s)}{\triangle w}$) **modulated by** TD Error.

Biologically Realistic Deep RL

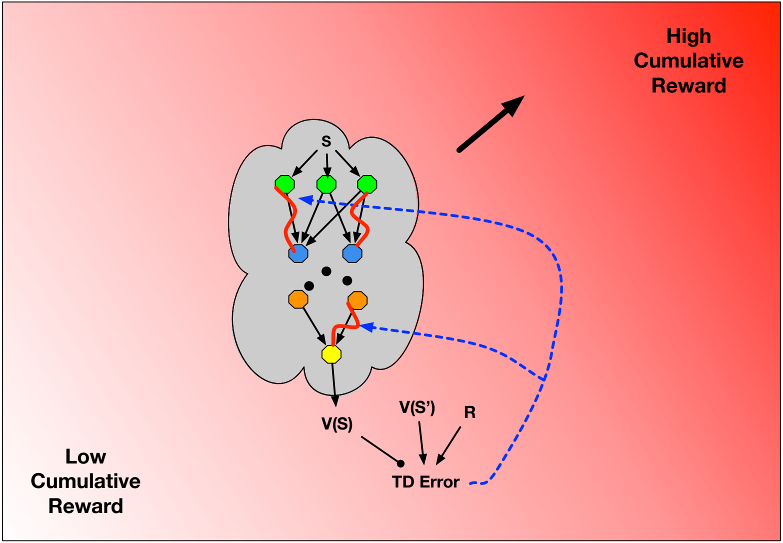

Hebbian learning (local) modulated by TD Error.

RL in the Basal Ganglia

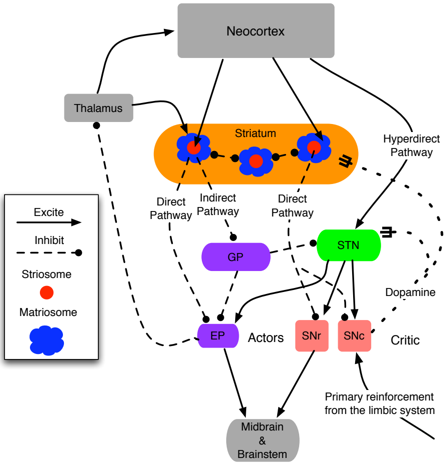

Neural computation of TD Error: $\delta = V(S_t) + R_t - V(S_{t-1})$

- Excitatory Inputs to SNc: V(S_t) (hyperdirect path) + R_t (limbic system)
- Inhibitory Inputs to SNc: V(S_{t-1}) via direct path.

As shown earlier....

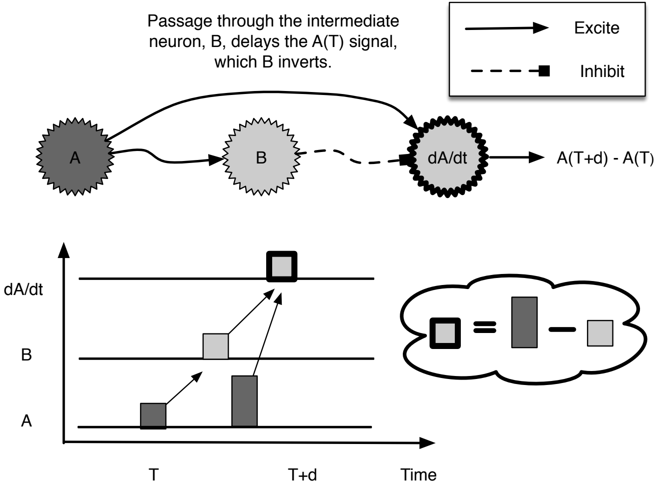

Monte Carlo Tree Search: Policy Bootstrapping

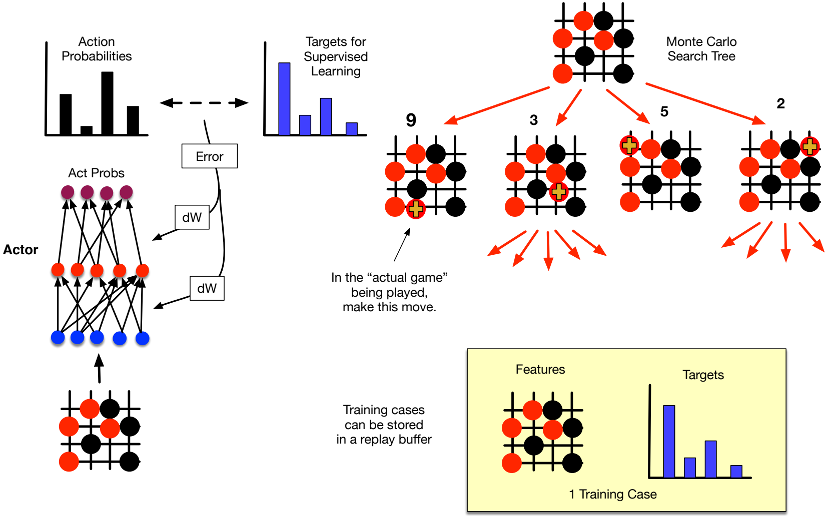

Actor / policy used to build tree, and tree is basis for updating the actor.

DeepMind's Self-Play: More Policy Bootstrapping

- Pre-trained on catalog of expert games
- **Self-play** from intermediate level upwards.
- Beat human world champion.

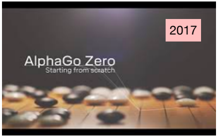

- No expert knowledge needed.
- **Self-play from scratch.**
- Became world champ.
- Beat AlphaGo 100-0

- Generalized AlphaGo Zero to other games.
- Became world champ at them.
- Increased Artificial General Intelligence (AGI)

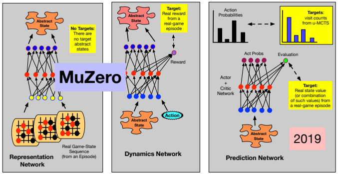

- Model free + ignorant of game rules.
- Learns abstract model of the game.
- Equals Alpha Zero in chess, go, shogi.
- Superhuman level on Atari arcade games.

- Many of the harder proofs require constructions.
- Constructions require considerable creative insight and/or search

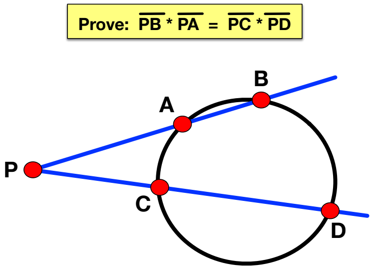

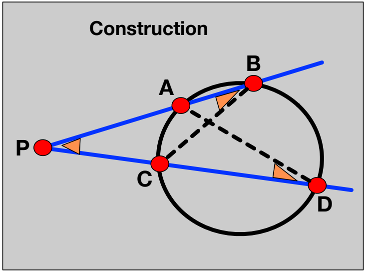

- $\measuredangle PDA \cong \measuredangle PBC$ (same subtended arc)
- $\triangle PDA \sim \triangle PBC$ (3 equal angles)
- $\frac{PB}{PC} = \frac{PD}{PA}$ (since similar triangles)
- $PB \star PA = PC \star PD$ (rearrangement) Q.E.D.

AlphaGeometry: Deduction + Construction

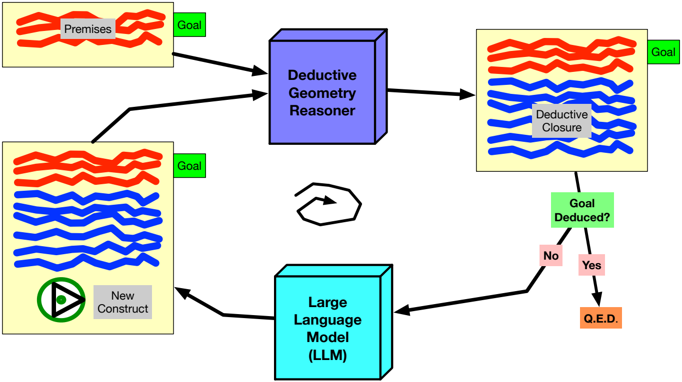

- Computing deductive closure $\approx$ a classification task: recognizing all consequences of the given facts.
- Construction = A Generative process indicating deeper understanding of domain.
- Google Deepmind (2024)

Generating Training Data: Constructor Bootstrapping

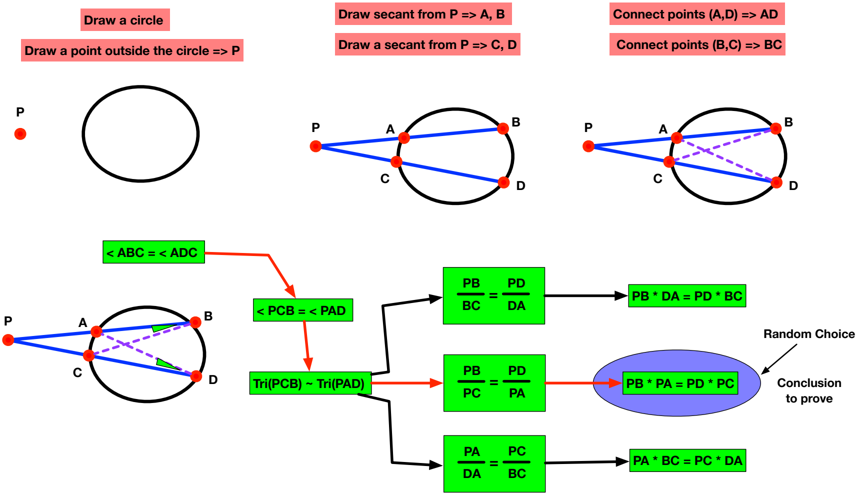

- P,A,B,C,D, $\overline{PB}$, $\overline{PA}$, $\overline{PC}$, and $\overline{PD}$ are all part of the conclusion.
- But $\overline{AD}$ and $\overline{BC}$ are not → They are viewed as **constructions**.

A Training Case

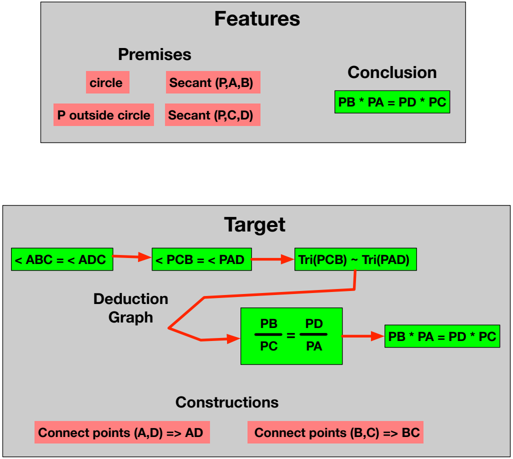

- AlphaGeometry generates millions of these !!
- Trains an LLM Transformer.

Diminishing Human Contributions

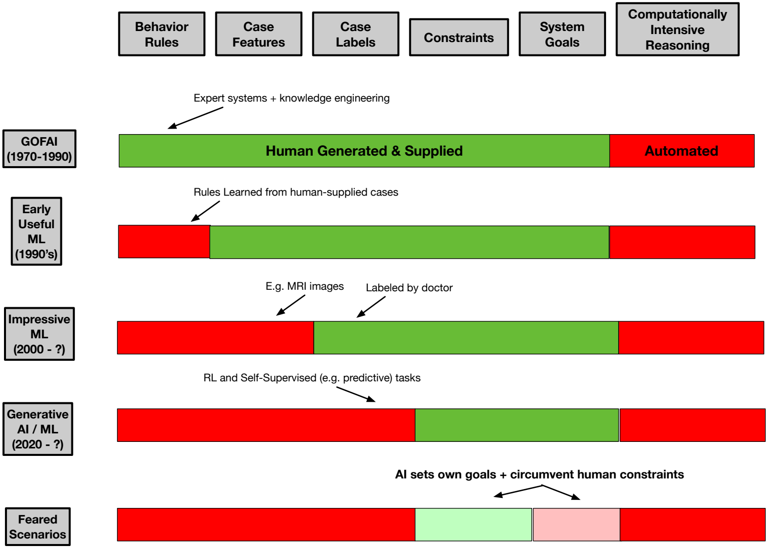

All this bootstrapping is booting humans out of the loop.

Emergent Intelligence in Large Language Models

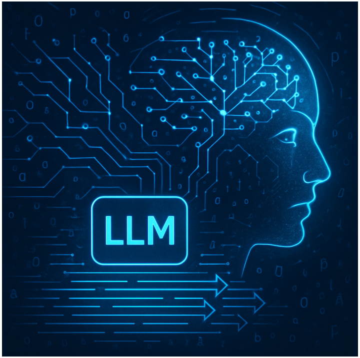

- LLMs are trained for a predictive task:
  What is the next word (or token)?
- Prediction is a cheap (no human labels) way to do supervised learning.
- From this, intelligent behavior has emerged.
- Is prediction the foundation of emergent intelligence?
- Can this prediction-based emergent intelligence achieve AGI?

My prompt: Please predict the main events that will happen in the next few minutes, hours and days when given the following scenario:

> *Marlon has just been released from prison after serving 4 years for armed robbery. As he is walking down the street, he notices that the owner of a jewelry shop has rushed out to his car and driven away, apparently for an emergency. He left the door open to the shop, and nobody is inside. Marlon has a wife and 3 young children at home, and they are desperate for the income that they hope he can earn after his prison release...*

ChatGPT: Given this scenario, we can project likely developments based on character motivations, past behavior, and situational factors. Here's a breakdown of what could happen in the next few minutes, hours, and days, taking into account different potential paths:

ChatGPT's Short-Term Predictions

Marlon is faced with a powerful temptation. He's alone, recently released, likely under pressure to provide for his family, and standing at the open door of a vulnerable jewelry store. Several possibilities could unfold:

Temptation Wins (High-Risk Path):

- Marlon enters the store, quickly grabs valuables, and flees.
- Surveillance footage likely exists.
- Silent alarm or nearby witnesses may be triggered.
- This would constitute a parole violation and new charges.

Resists Temptation (Redemptive Path):

- He hesitates but chooses to walk away.
- Demonstrates growth or desire to change.
- Might report the open store to authorities, earning trust.

Gray Area / Hesitation:

- He steps inside to look, perhaps grappling with the decision.
- A call or message from his wife or children might influence his choice.
- A bystander or returning owner could interrupt.

ChatGPT's Medium-Term Predictions

If He Robs the Store:

- Police may identify him via video.
- A warrant could be issued quickly.
- He may attempt to hide, reach out to his family, or plan an escape.
- Panic and regret may set in, especially if he wanted to go straight.

If He Walks Away or Reports the Situation:

- Police might take a statement.
- He could earn credibility with his parole officer.
- The experience might reinforce his resolve to change.
- He might reflect on how close he came to falling back.

ChatGPT's Long-Term Predictions

If He Committed a Crime:

- Arrest is highly likely within 1–2 days.
- Media might report on the "repeat offender."
- He faces a return to prison, deeper family hardship.
- His relationship with his wife and children could suffer.

If He Did Not:

- He may face continued hardship, job hunting, or parole check-ins.
- Social services, job programs, or support networks might play a role.
- This event could become a turning point in staying clean.
- Trust might grow between him and his family if they learn of his decision.

ChatGPT's Conclusion

- If Marlon is written as a realistic, multi-dimensional character — not just a stereotype — then I think he pauses. Maybe he walks in out of curiosity or reflex. But the weight of his kids, the memory of prison, and the chance to do better... they pull him back.
- He might close the door himself. Maybe he calls someone. Maybe he just keeps walking, sweating, but proud.
- Because the harder choice — especially when no one's watching — is usually the right one. And that's what redemption actually looks like.

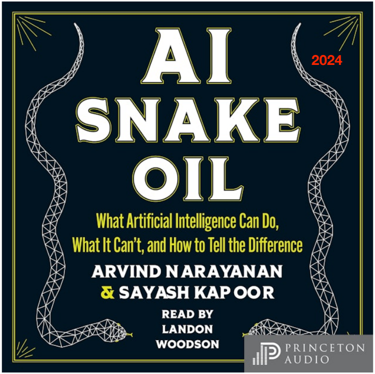

Predictive Tasks
-vs-
Generative Tasks

Classic AI Predictive Tasks (and Failures):

- Will this person pay back their bank loan?
- Should we hire this person?
- Will this person commit another crime?

Would an LLM do a better job?

Generative Tasks: Great AI Success

Expanding AI's Scope of Competence

Modern AI systems are very specific and brittle.

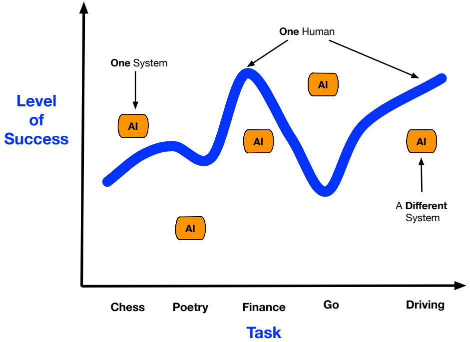

- **Artificial Narrow Intelligence (ANI)** = today's AI systems; very good at a few things.
- **Artificial General Intelligence (AGI)** = a single AI equal to humans on many (most) tasks.
- **Artificial Super Intelligence (ASI)** = a single AI better than humans at many (most) tasks.

Acquiring Commonsense

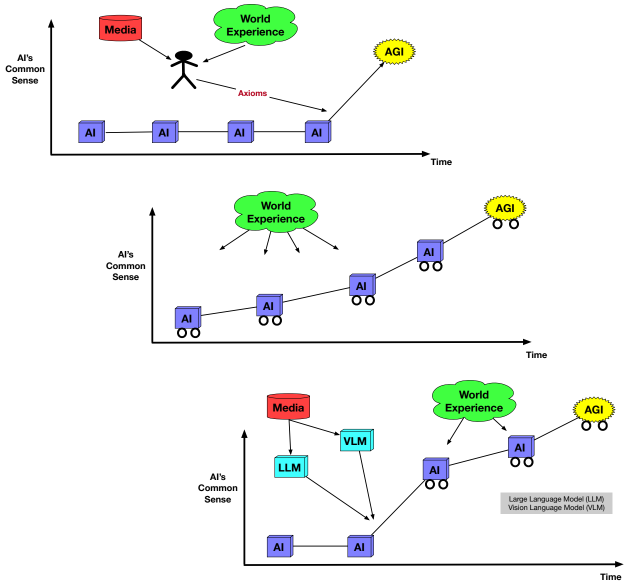

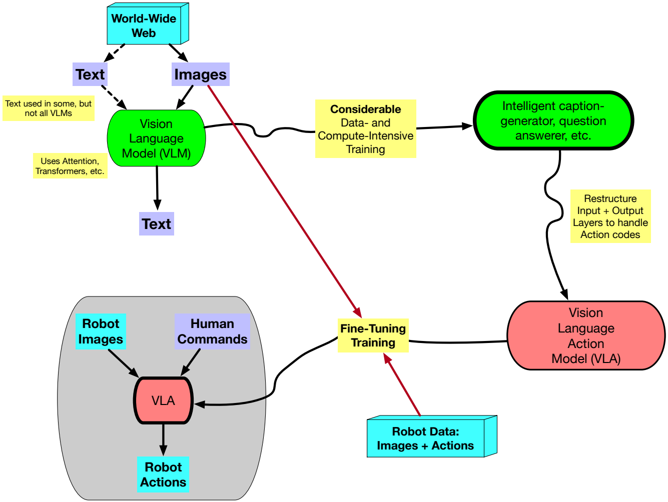

Google DeepMind (2023)

My Evolving Trauma

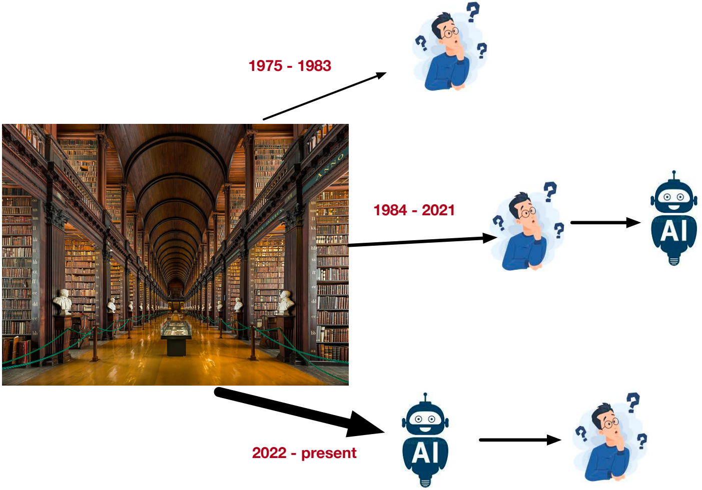

- Prediction: What will it do?
- Explanation: Why did it do it?

Ubiquity of Prediction

Prediction may not be the basis of all intelligence, but it pops up a lot in neuroscience, cognitive science and AI.

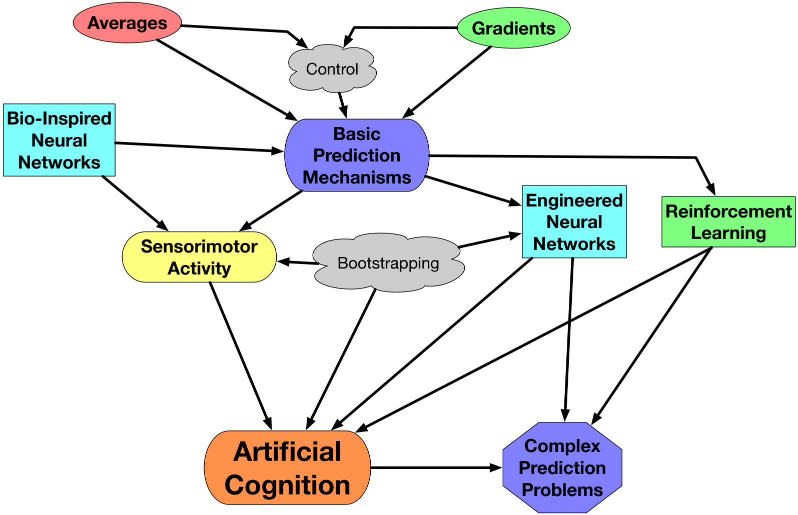

Yogi Berra - American Baseball Legend

It's tough to make predictions.... especially about the future.

The future ain't what it used to be.