What Can We Learn About Innovation From the Theories That Drive Artificial Intelligence?

Christopher J. Hazard, PhD

Diveplane

HAZARDOUS SOFTWARE

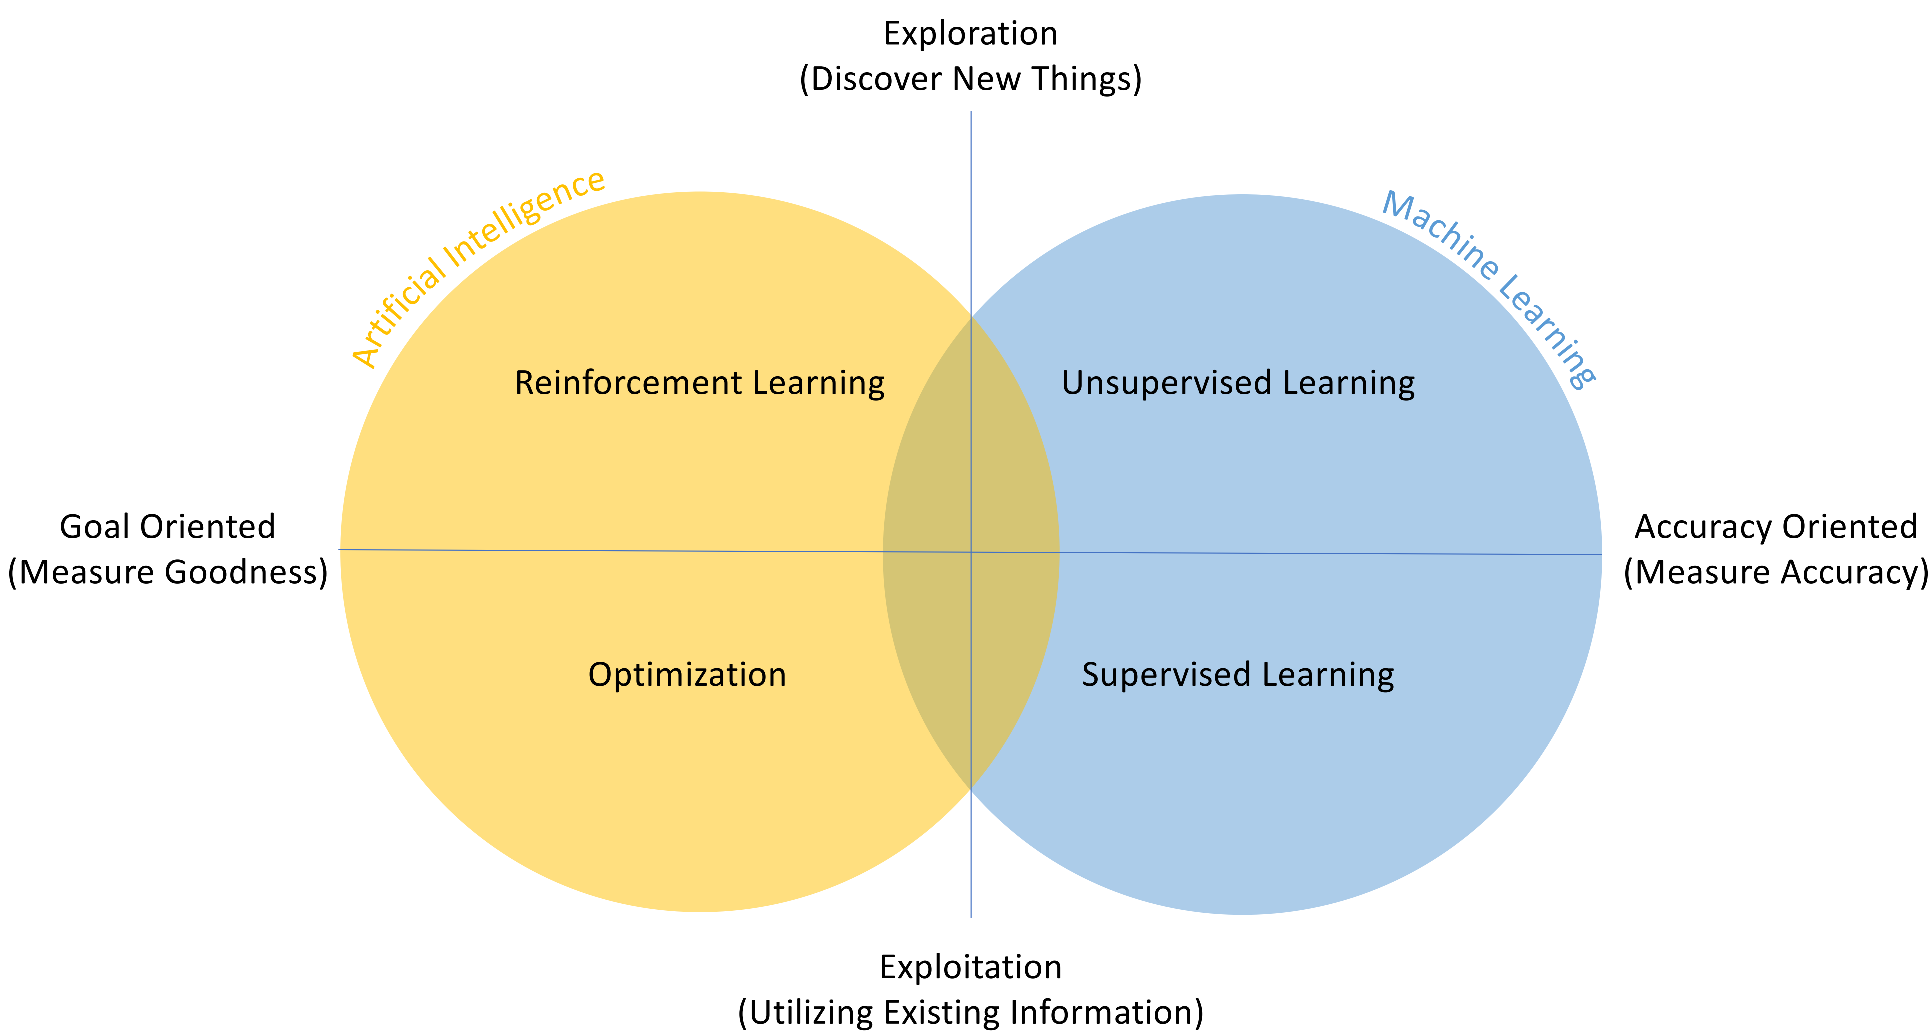
Exploration
(Discover New Things)
Artificial Intelligence
Machine Learning
Reinforcement Learning
Unsupervised Learning
Goal Oriented
(Measure Goodness)
Accuracy Oriented
(Measure Accuracy)
Optimization
Supervised Learning
Exploitation
(Utilizing Existing Information)

Example Domain: Food

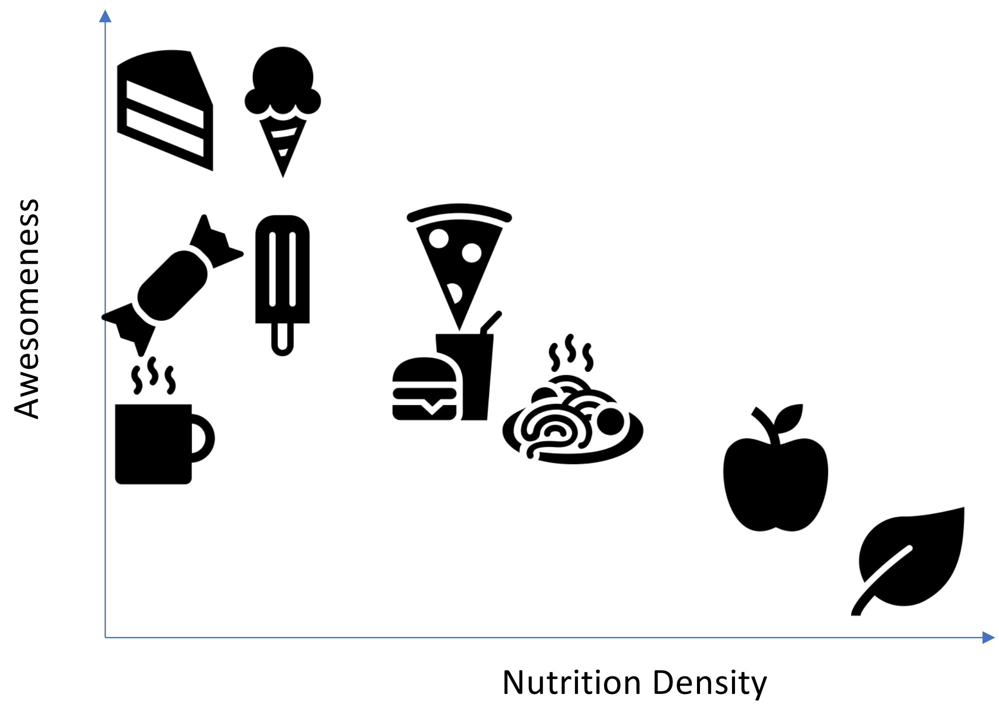

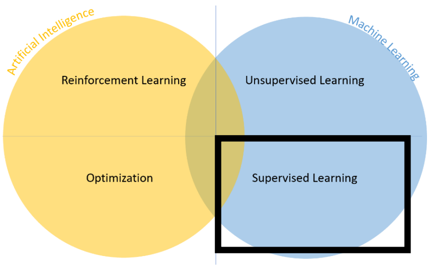

Supervised Learning

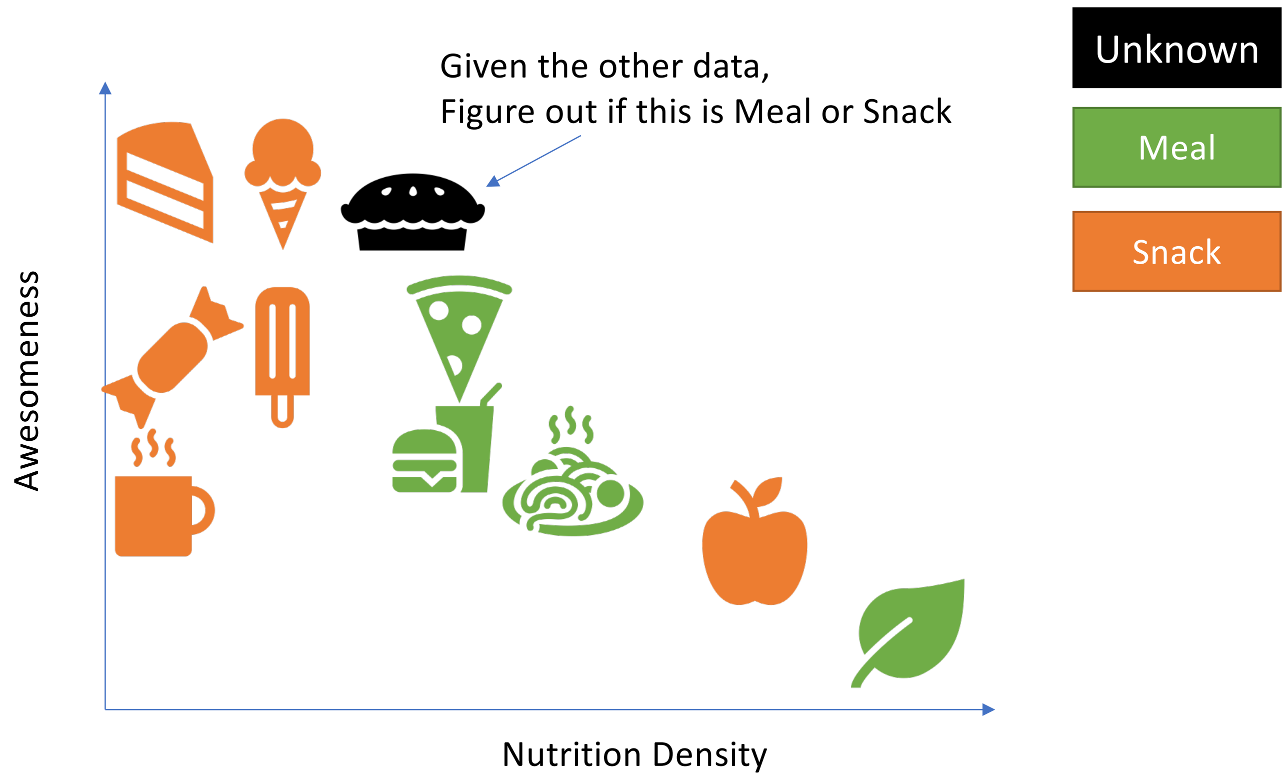

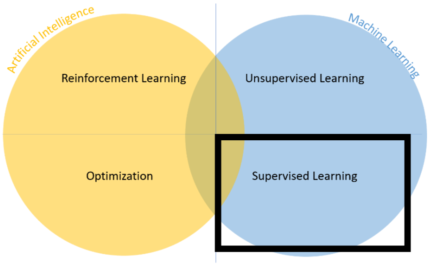

Supervised Learning:
Universal Function Approximators

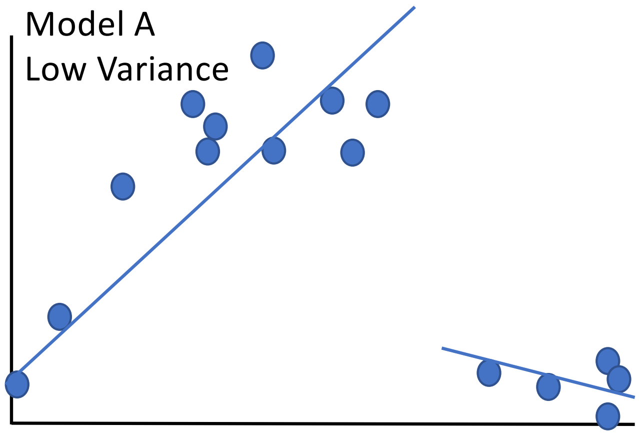

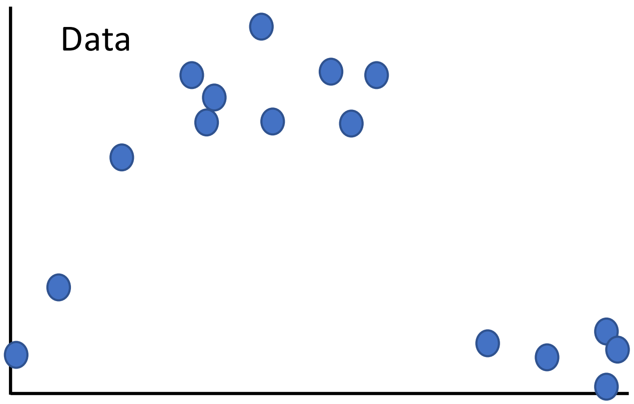

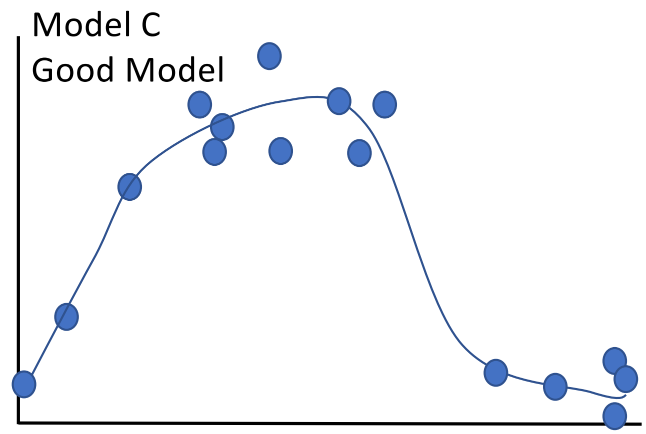

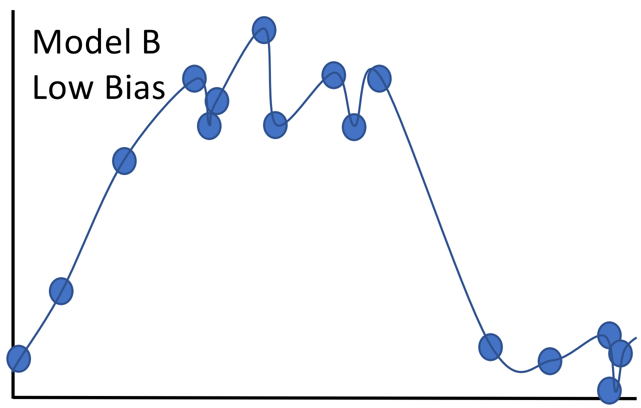

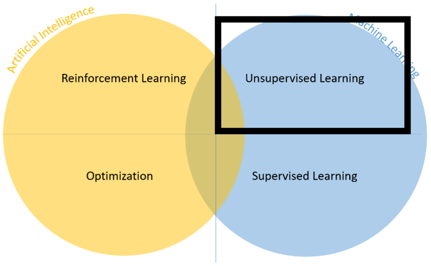

Unsupervised Learning

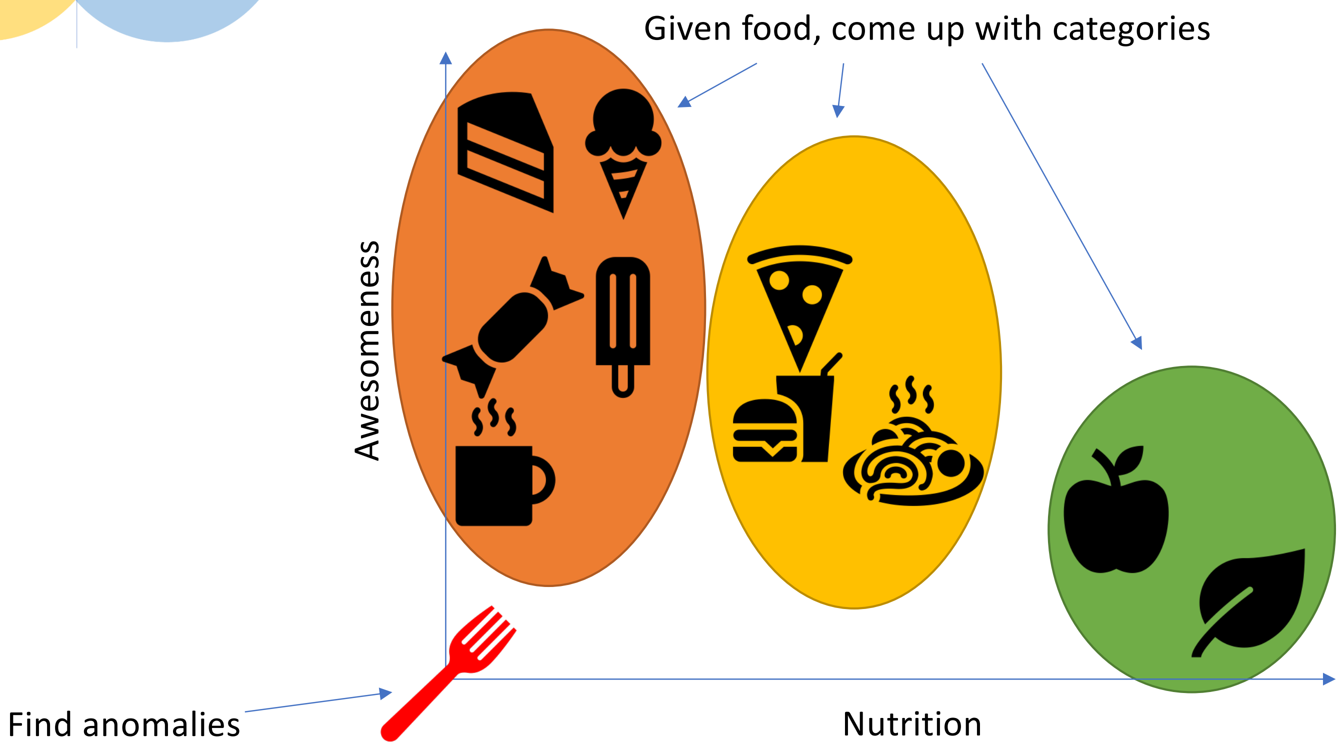

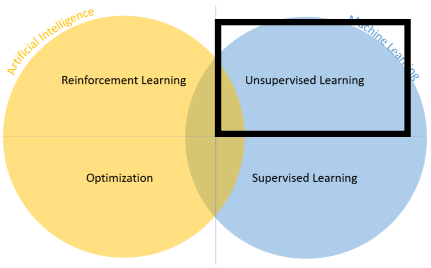

Unsupervised Learning:
Clustering and Anomaly Detection

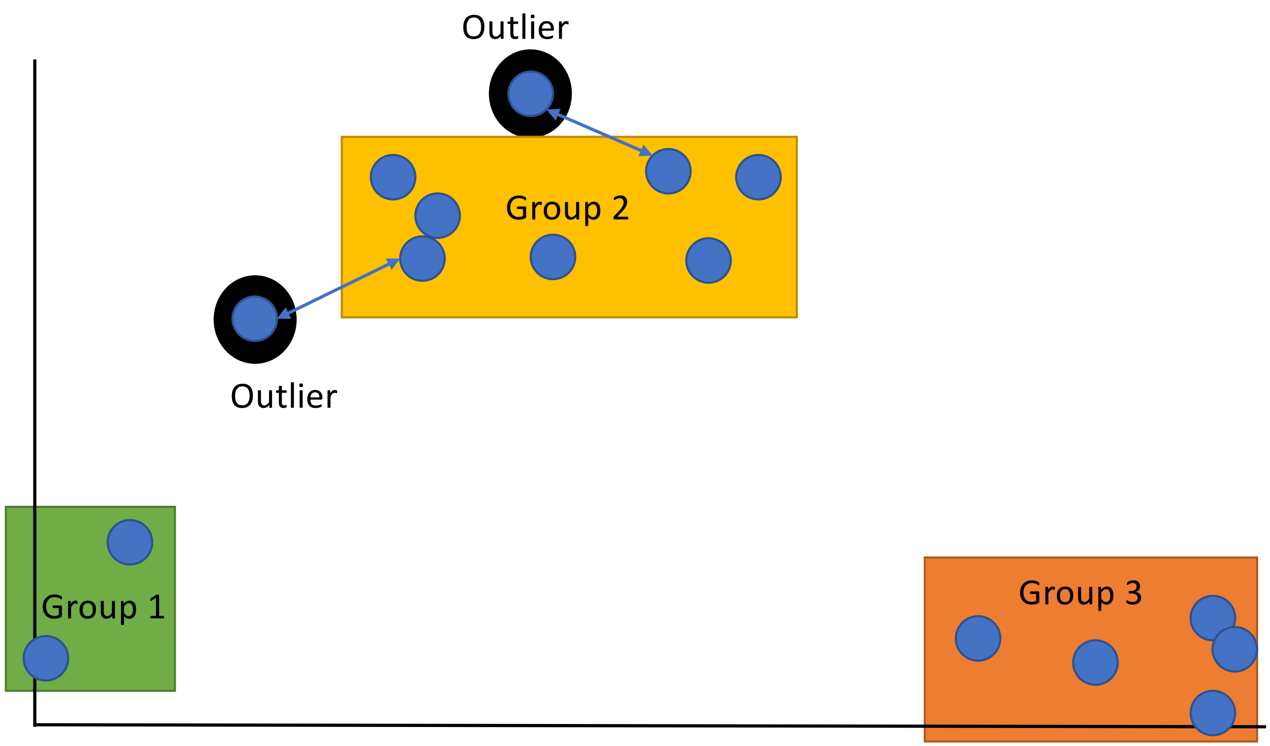

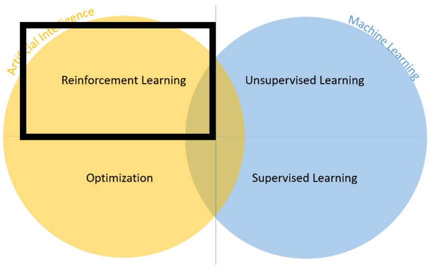

Reinforcement Learning

Objective: eat a highly nutritious meal

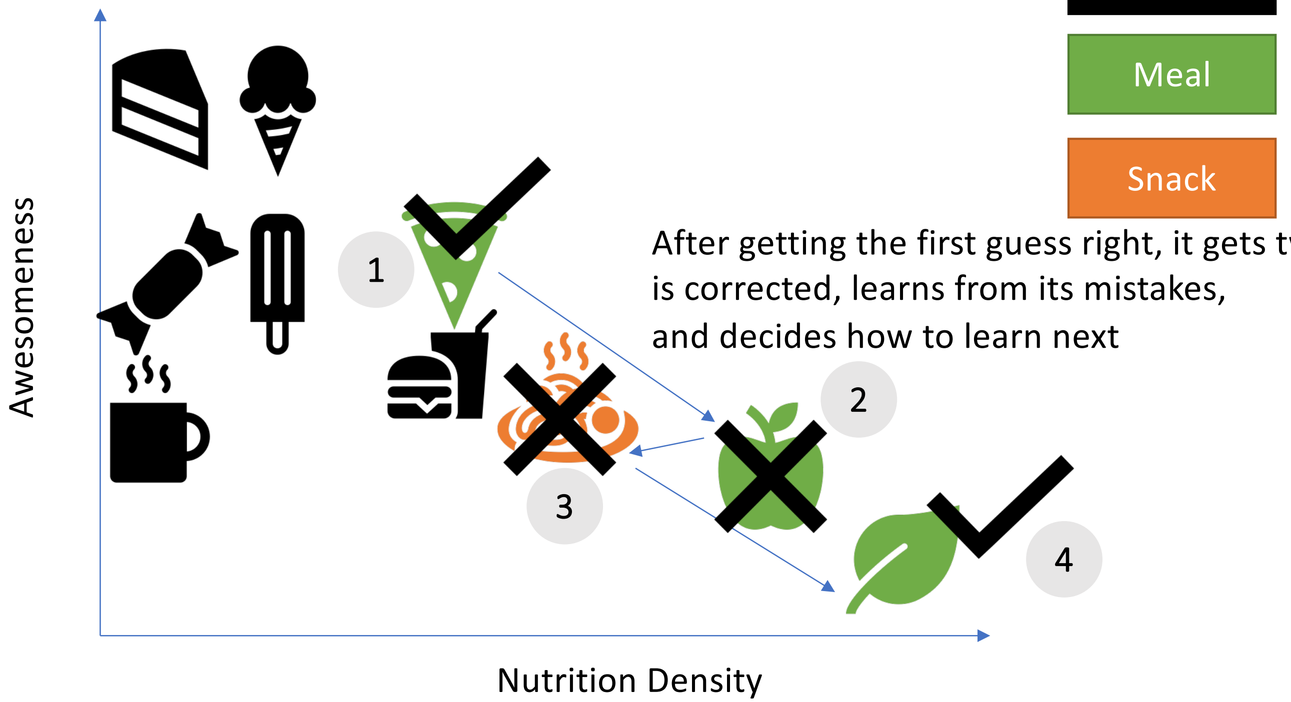

After getting the first guess right, it gets two wrong, is corrected, learns from its mistakes, and decides how to learn next

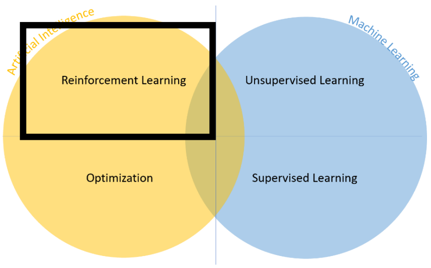

Reinforcement Learning:
Seeking Rewards, filling in Unknowns

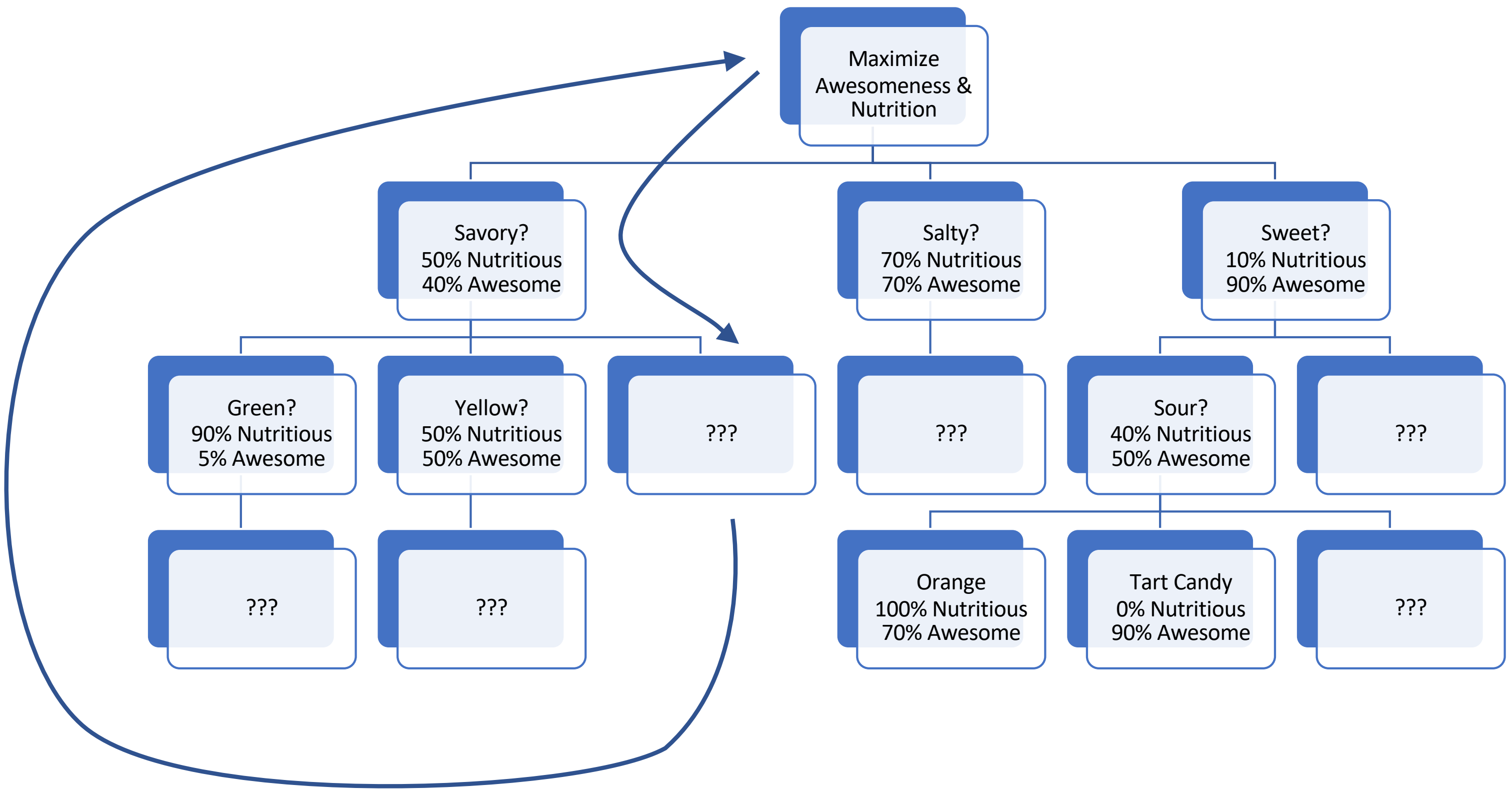

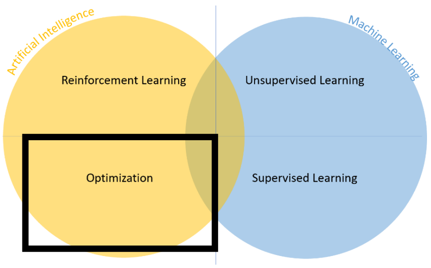

Optimization

Find the “best” meal

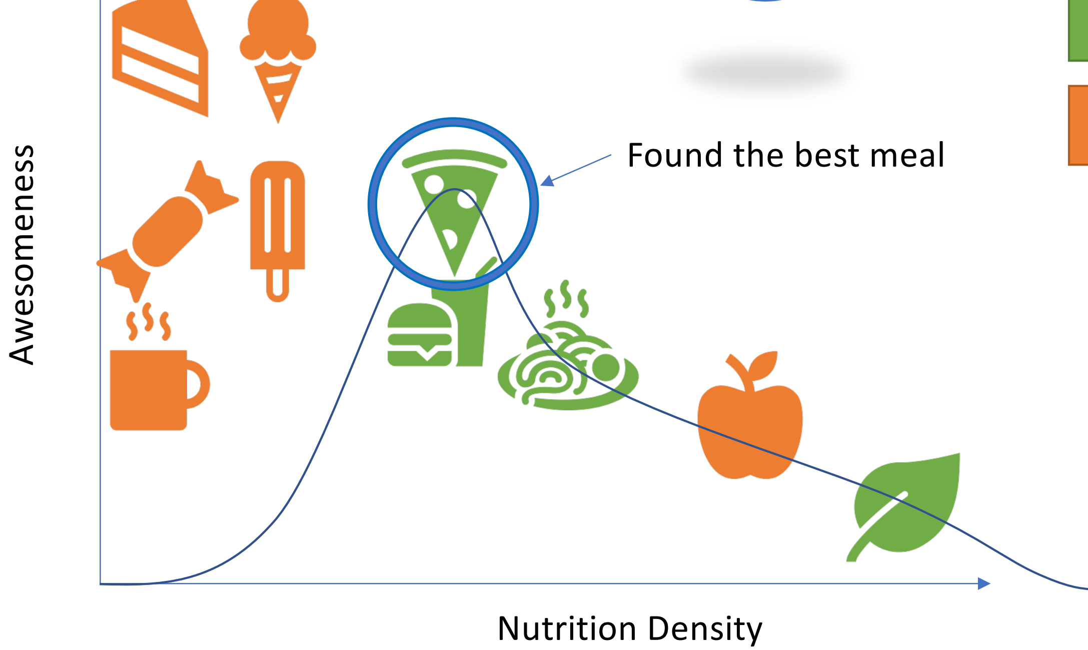

Unknown

Meal

Snack

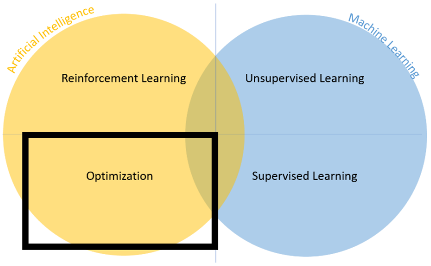

Optimization: Finding the Best

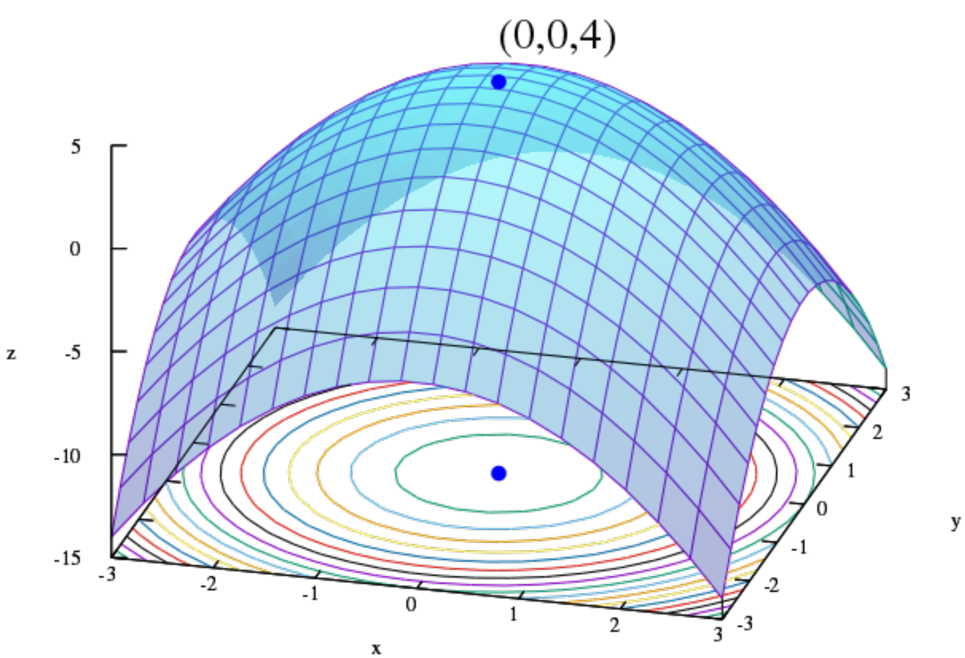

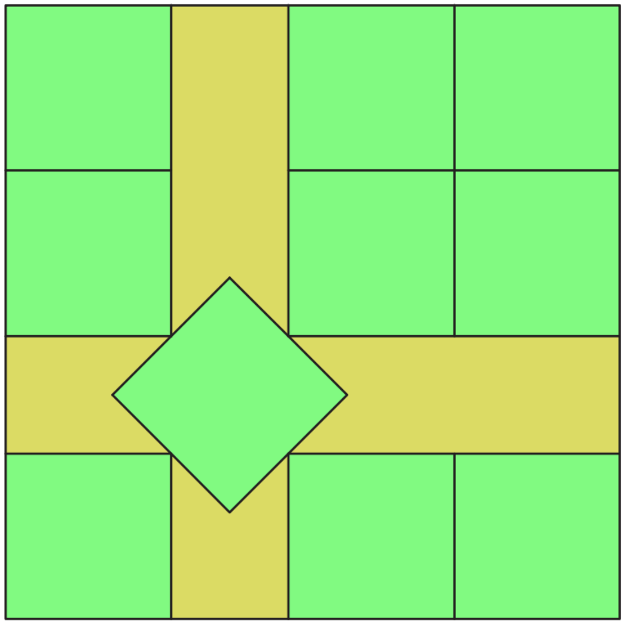

Innovation & Creativity

To make **new** and **valuable** things and ideas

Innovation & Creativity

Maximize Effectiveness

Minimize Expense

Minimize Complexity

To make **new** and **valuable** things and ideas

Maximize Surprisal

...using feedback

Filament Material	Voltage (Volts)	Power (Watts)	Thickness (Inches)	Length (Inches)	Gas	Pressure (Atm)	Lumens	Cost	Lifespan
Platinum	220	60	.0025	30	Air	.0005	400	$$$$	200 hours
Carbonized Bamboo	120	55	.0027	23.5	Air	.0002	250	$	1200 hours
Tungsten	120	100	.0018	22.8	Nitrogen	.7	1700	$	1000 hours
…	…	…	…	…	…	…	…	…	…

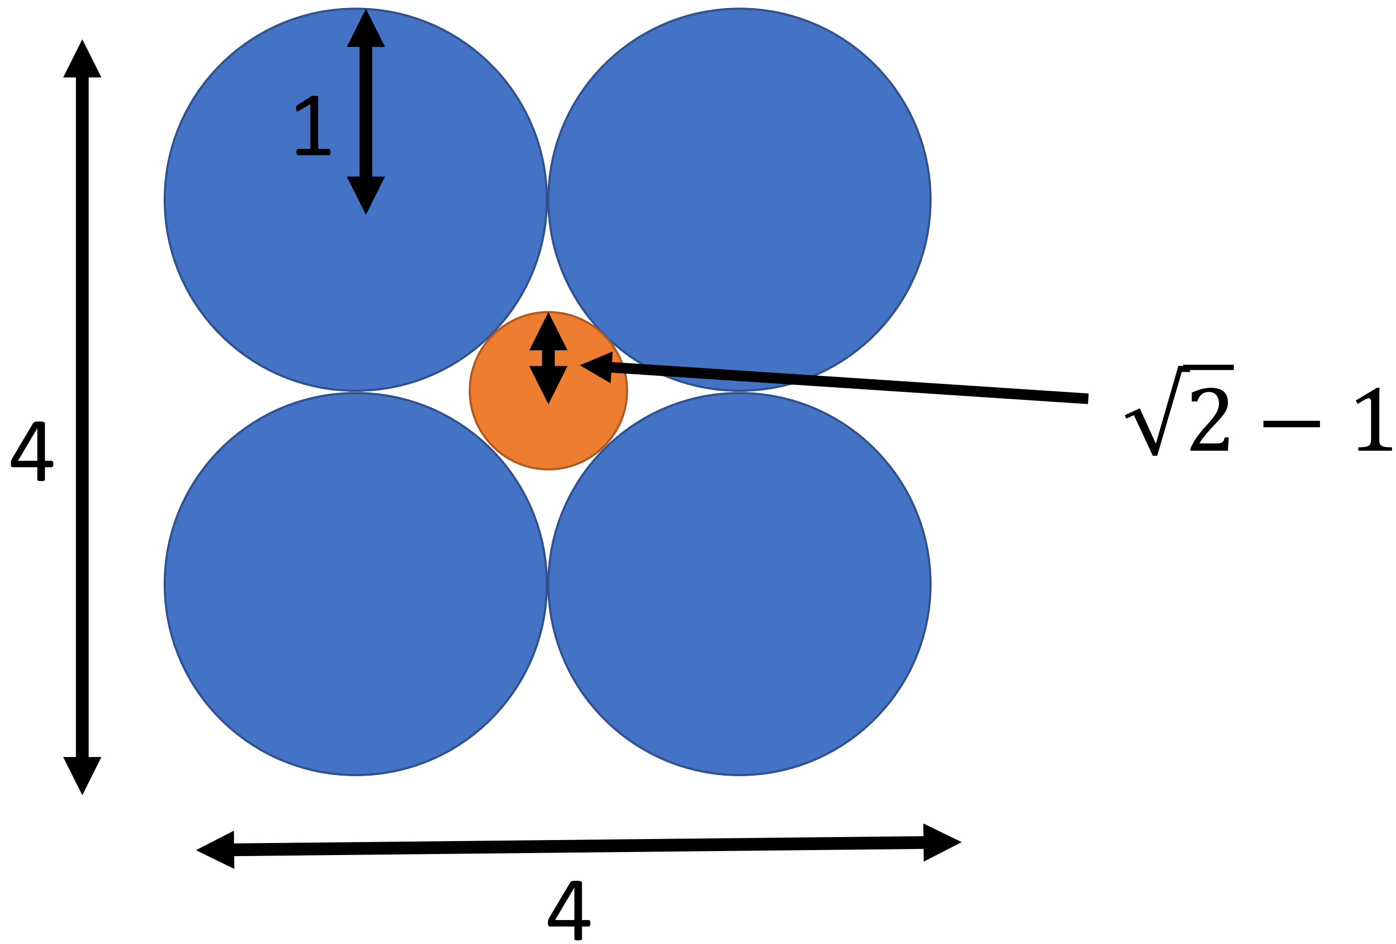
1
4
$\sqrt{2} - 1$
4

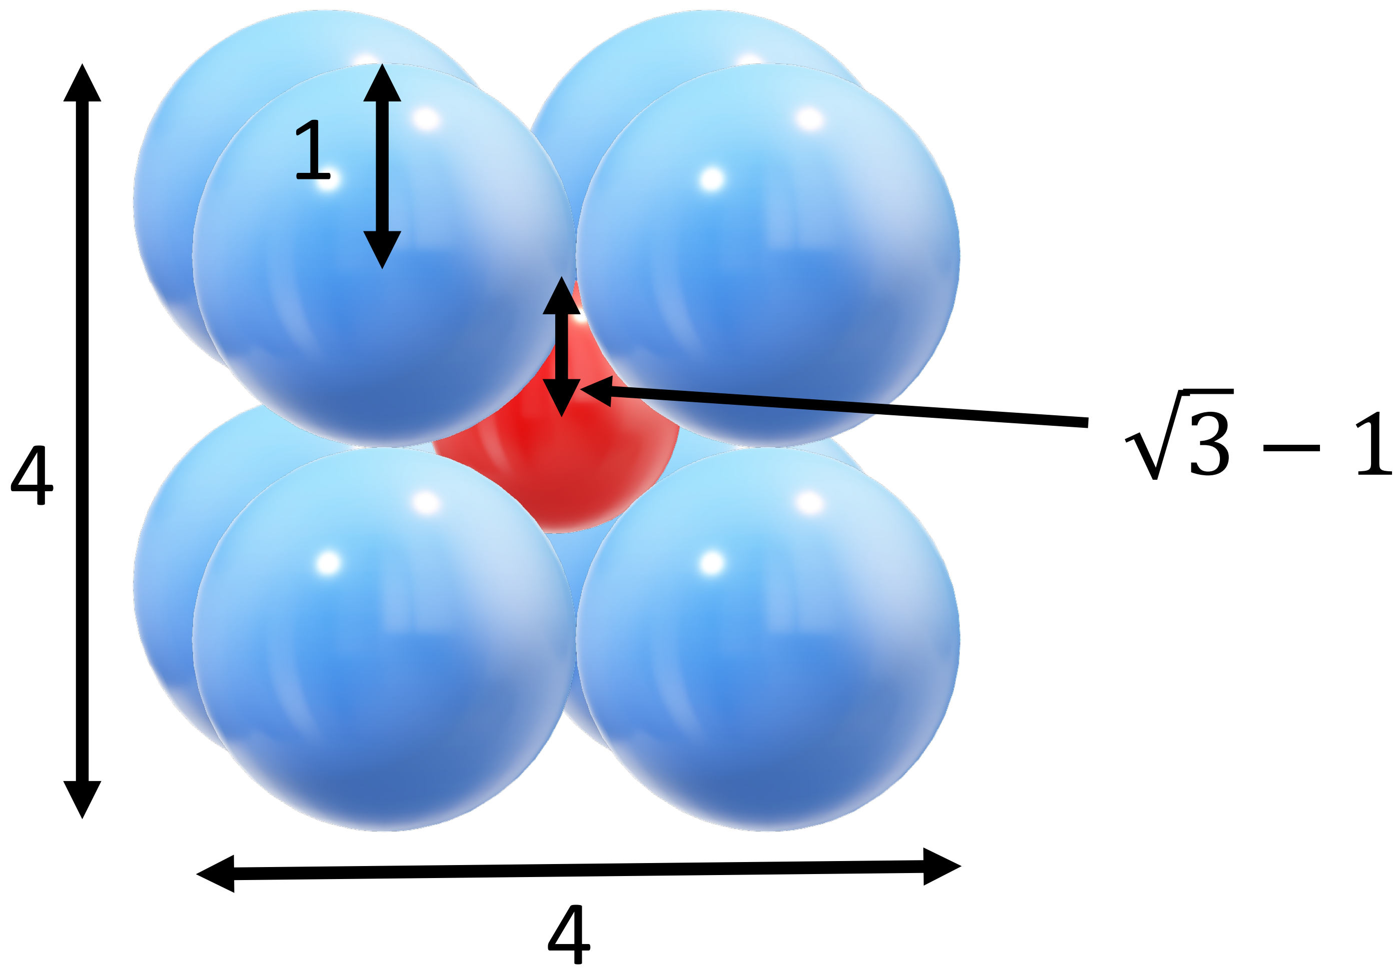
1
4
4
$\sqrt{3} - 1$

Dimensions	Diameter of Inner Sphere
1	$2\left(\sqrt{1}-1\right)=0$
4	$2\left(\sqrt{4}-1\right)=2$
9	$2\left(\sqrt{9}-1\right)=\mathbf{4}$
16	$2\left(\sqrt{16}-1\right)=6$
64	$2\left(\sqrt{64}-1\right)=14$

L_p Space / Minkowski Distance: $||x||_p = \left(\sum_{i \in \Xi} w_i x_i^p \right)^{1/p}$

A new L_0 "Norm": $\lim_{p \to 0} \left(\sum_{i=1}^{n} \frac{1}{n} x_i^p \right)^{1/p} = \left(\prod_{i=1}^{n} x_i \right)^{\frac{1}{n}}$

Hazard et al., DP TR 2019

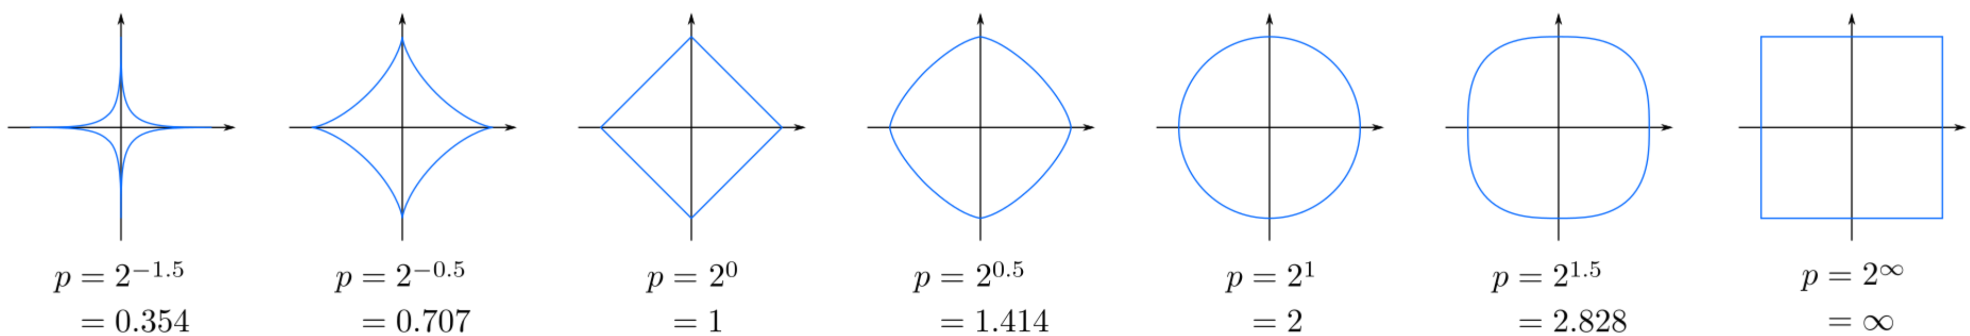

Original image by Waldyrious on Wikipedia

A Slower Speed of Light. Kortemeyer et al., FDG 2013

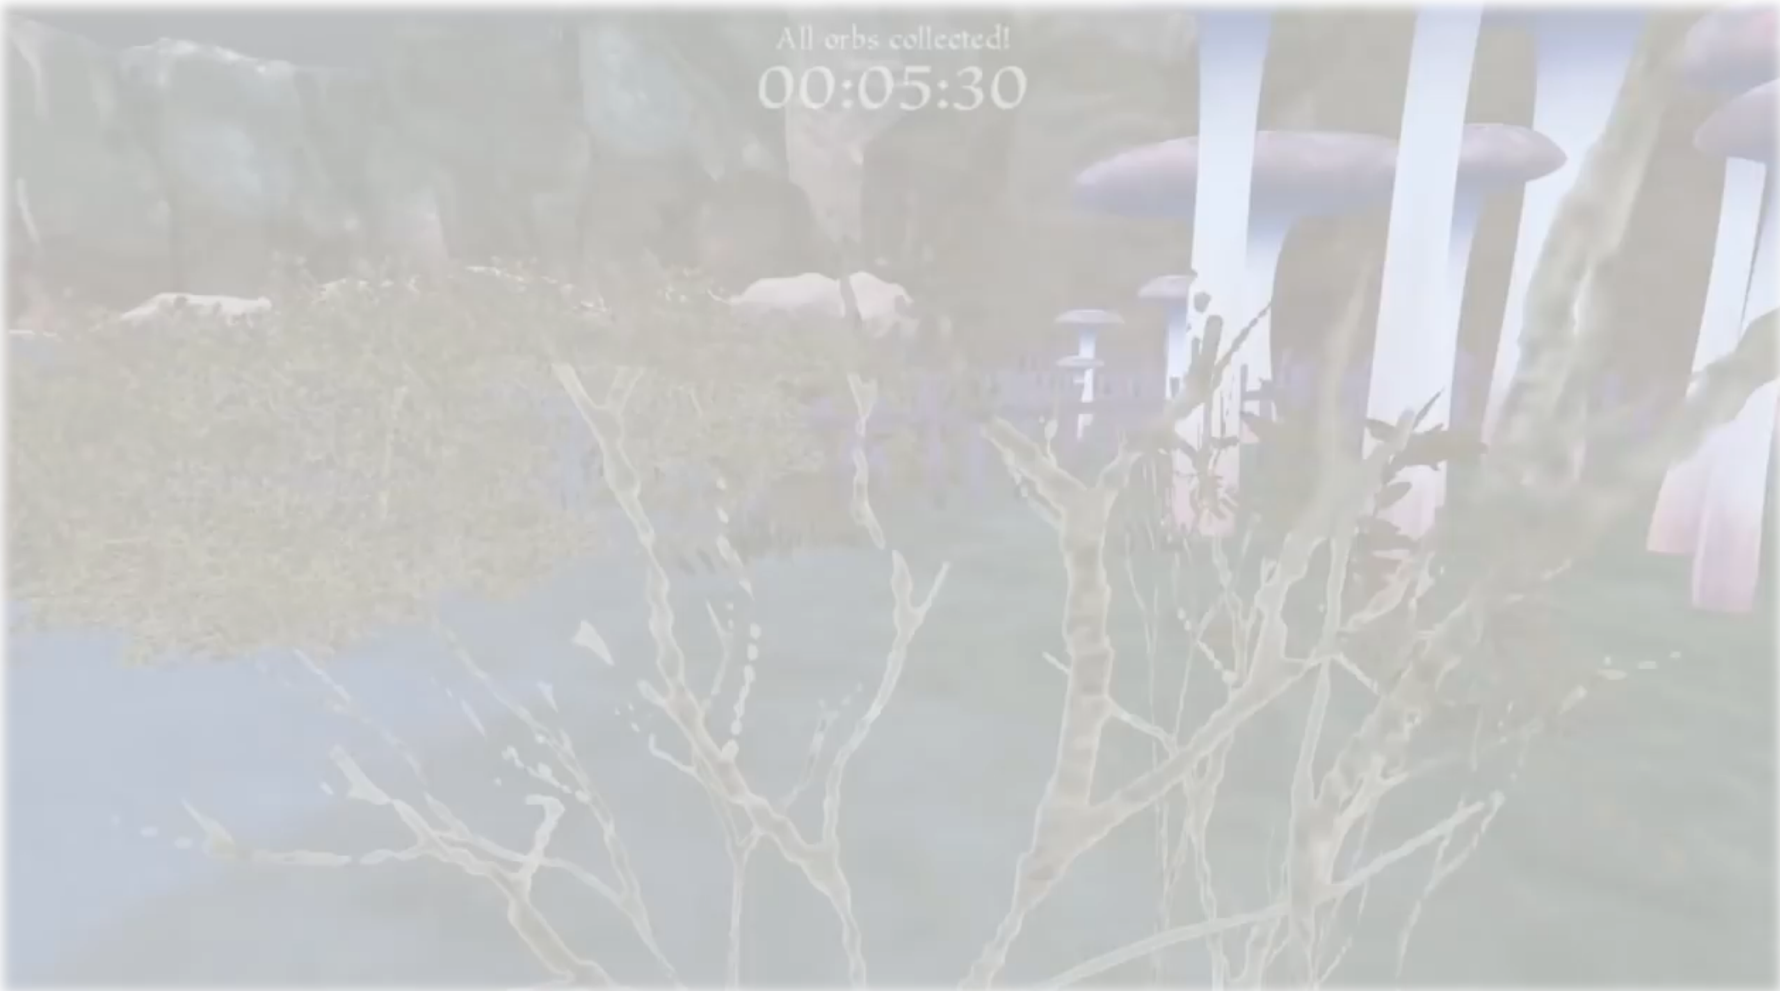

Nintendo: Mario Kart 8

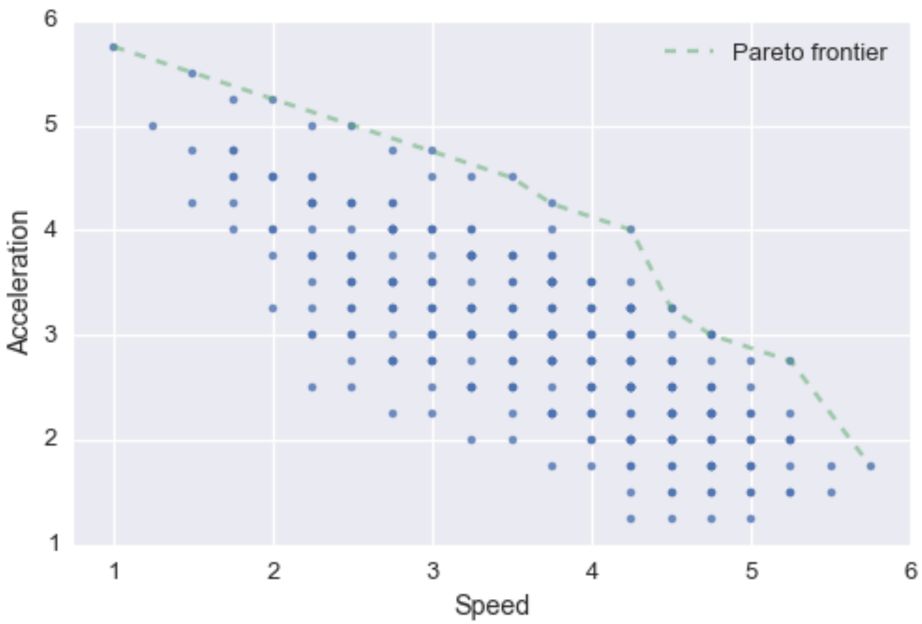

Character-Body-Tire	Weight	Speed	Acceleration	Handling	Traction
Baby Mario-Biddybuggy-Roller	1	1	5.75	5.5	4
Toad-Biddybuggy-Roller	1.5	1.5	5.5	5	3.75
Peach-Biddybuggy-Roller	2	2	5.25	4.5	3.5
Mario-Biddybuggy-Roller	2.5	2.5	5	4	3.25
Donkey Kong-Biddybuggy-Roller	3	3	4.75	3.5	3
Wario-Biddybuggy-Roller	3.5	3.5	4.5	3	2.75
Donkey Kong-Sports Bike-Roller	3.25	3.75	4.25	3.75	2
Wario-Sports Bike-Roller	3.75	4.25	4	3.25	1.75
Wario-Sports Bike-Wood	4	4.5	3.25	2.75	2.5
Wario-Biddybuggy-Slick	4.25	4.5	3.25	2.75	2
Donkey Kong-Sports Bike-Slick	4	4.75	3	3.5	1.25
Wario-Gold Standard-Roller	4.25	4.75	3	2.5	2
Wario-Sports Bike-Standard	4.25	4.75	3	3	2
Wario-Sports Bike-Slick	4.5	5.25	2.75	3	1
Wario-Gold Standard-Slick	5	5.75	1.75	2.25	1.25

Henry Hinnefeld: http://hinnefe2.github.io/python/tools/2015/09/21/mario-kart.html

Goodness Landscape
(projected to one dimension)

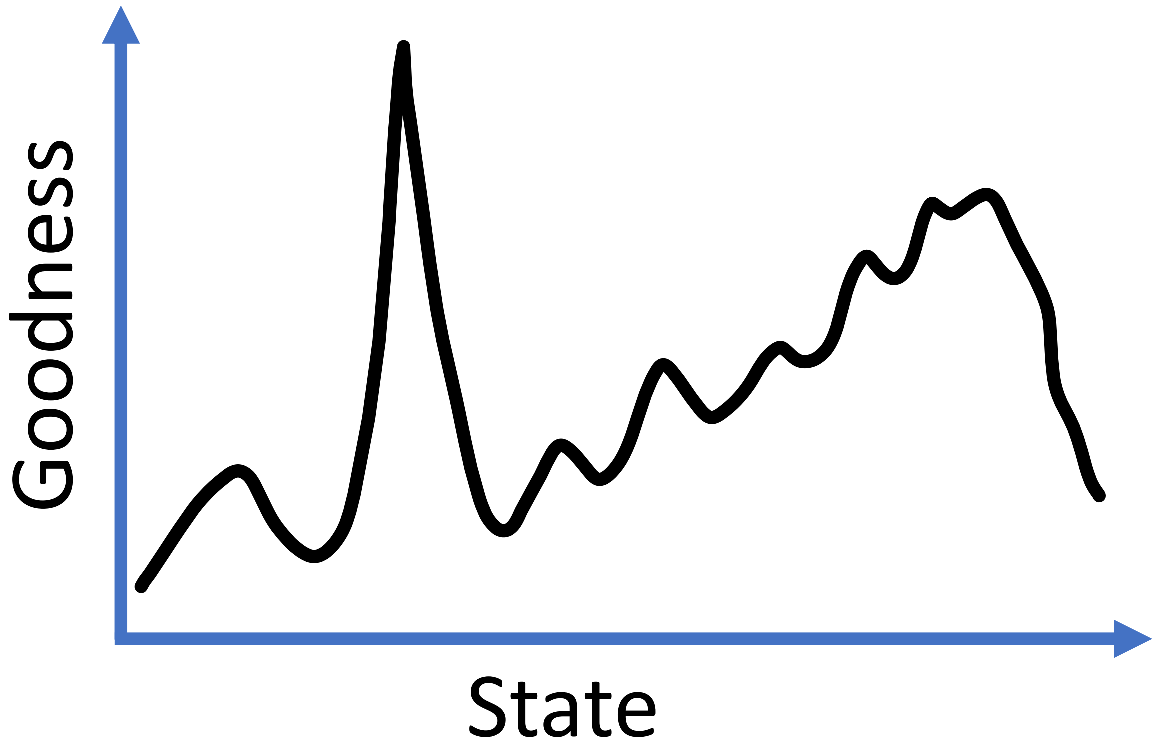

Sampling Goodness

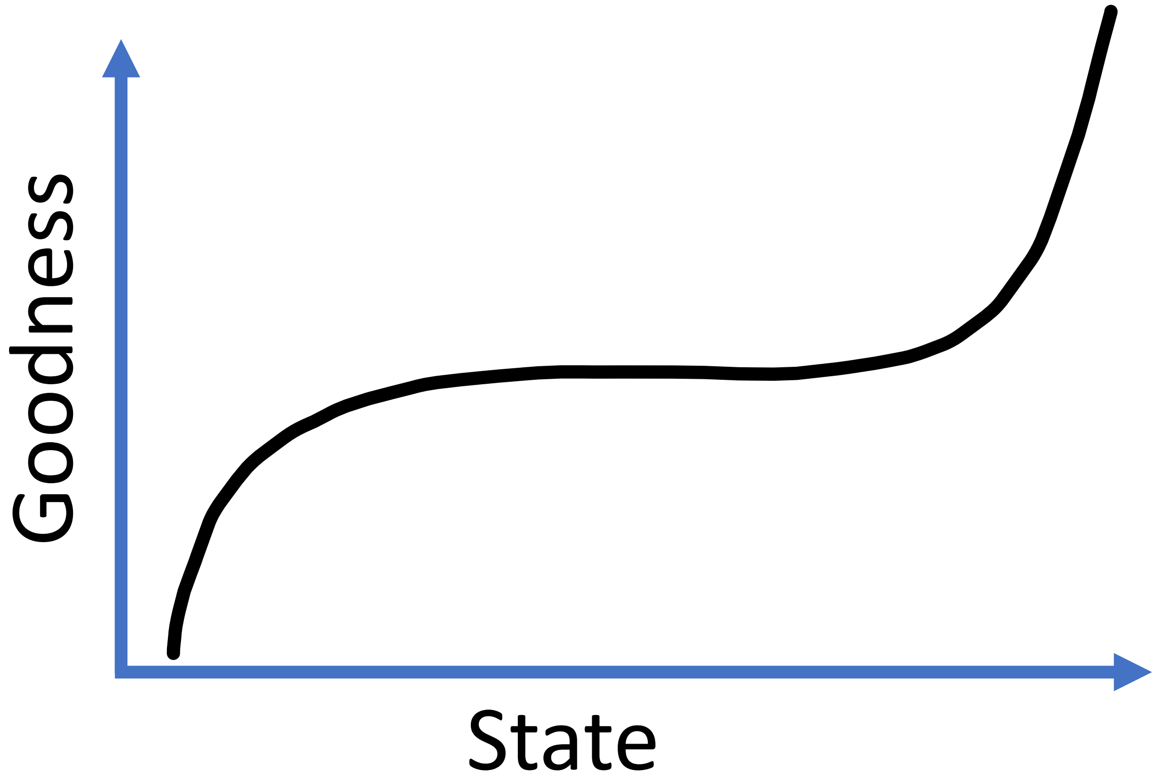

How Are Functions Fooled?

- Exploit spurious correlations in random features
  - 200 coin flips: 6 in a row
- Exploit irregular boundaries
  - Incorrect margins
  - Incorrect slope
  - Irregular shape
- Simpson’s Paradox / Wrong Features

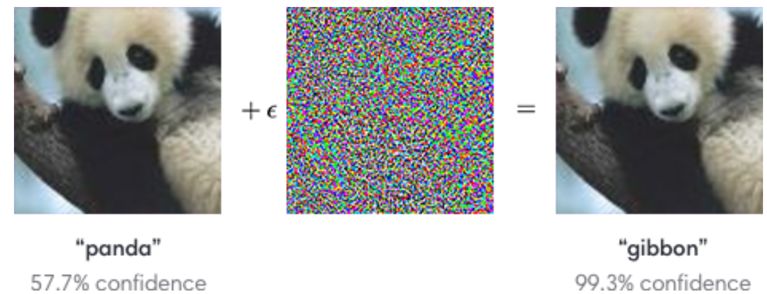

Goodfellow et al., ICMR 2015

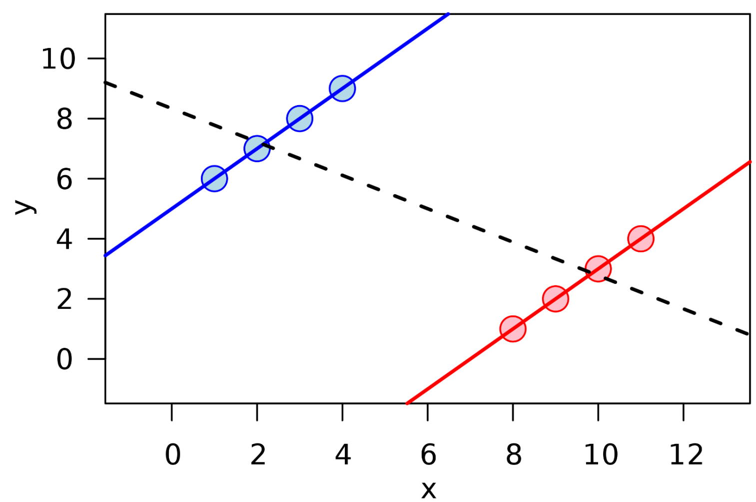

Data vs Games

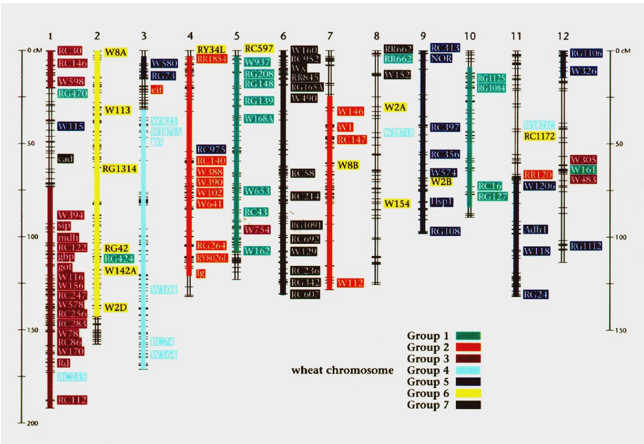

Wheat Genome

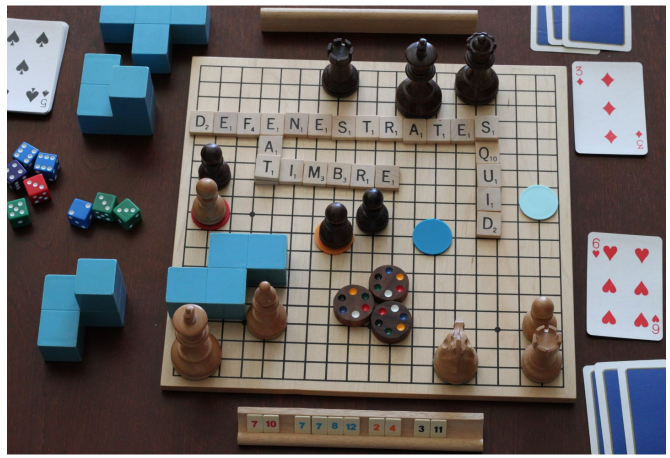

Calvinball/Nomic with Hazard

Starcraft 2 – Blizzard

Google Image Labeler

INMAST – Hazardous Software, 2017

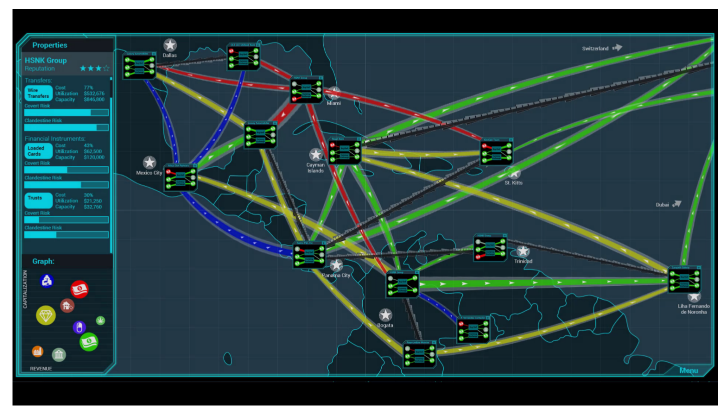

What Are you Optimizing For?

Goal	Example Technique	Requires	Benefits	Drawbacks
Maximize expected value	MCTS	Data	Great results without adversary	Not strong vs formidable / creative adversary
Minimize expected regret	MCCRM	Knowledge of causality and uncertainty	Unlikely to lose or lose by much, will do well vs adversary	Need to codify what are and are not rules / causal
Minimize maximum loss (minmax)	Nash Equilibrium (or other solution concept)	Knowledge of causality and uncertainty fully characterized	Won't lose except by chance	Often higher computational complexity, will not take advantage of weak adversaries

Data vs Game: Resources Spent on Defense

- ~20-30%

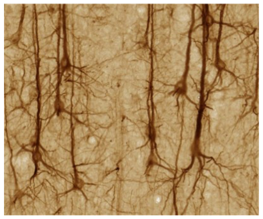

brainmaps.org

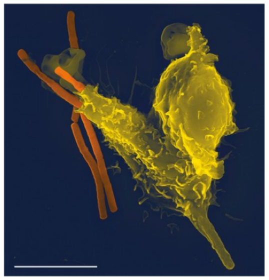

Volker Brinkmann

- ~3-8% (increasing?)

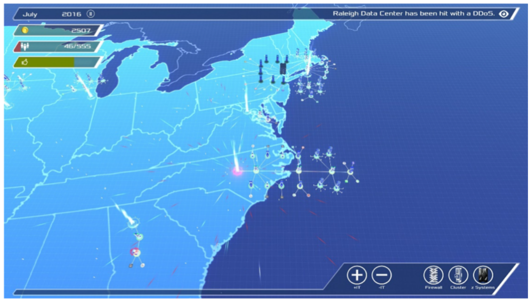

- ~1%

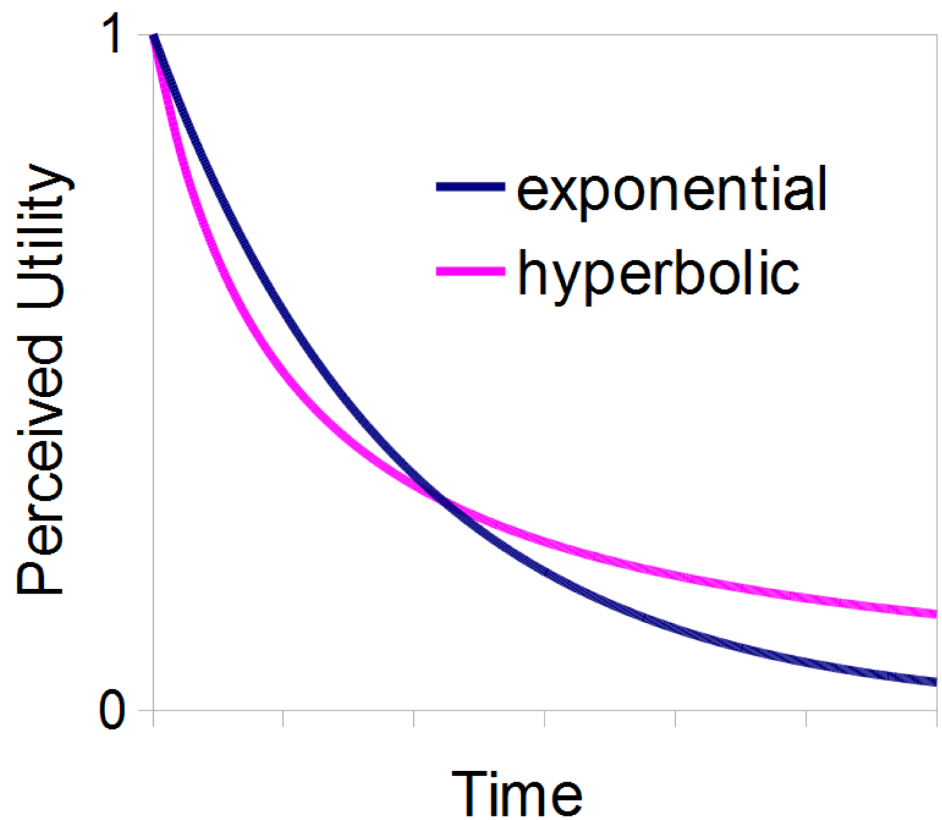
1
0
Perceived Utility
Time
exponential
hyperbolic

Measuring discount factor by choice

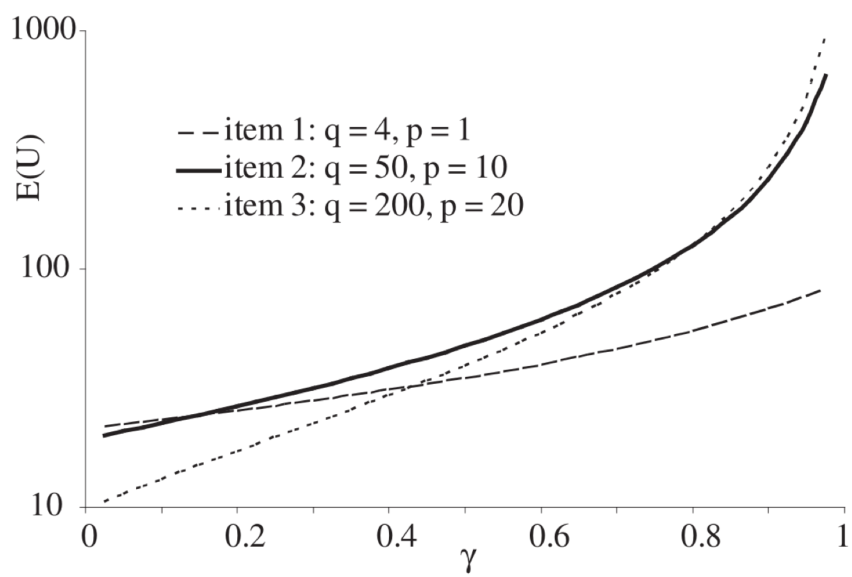

Hazard & Singh, TKDE, 2010

Time Preference and Switching Cost

- Why do some technologies get adopted? E.g., TCP and UDP dominate when more capable technologies exist such as SCTP

- Time preference, switching costs, and trend following scales the number of early adopters required

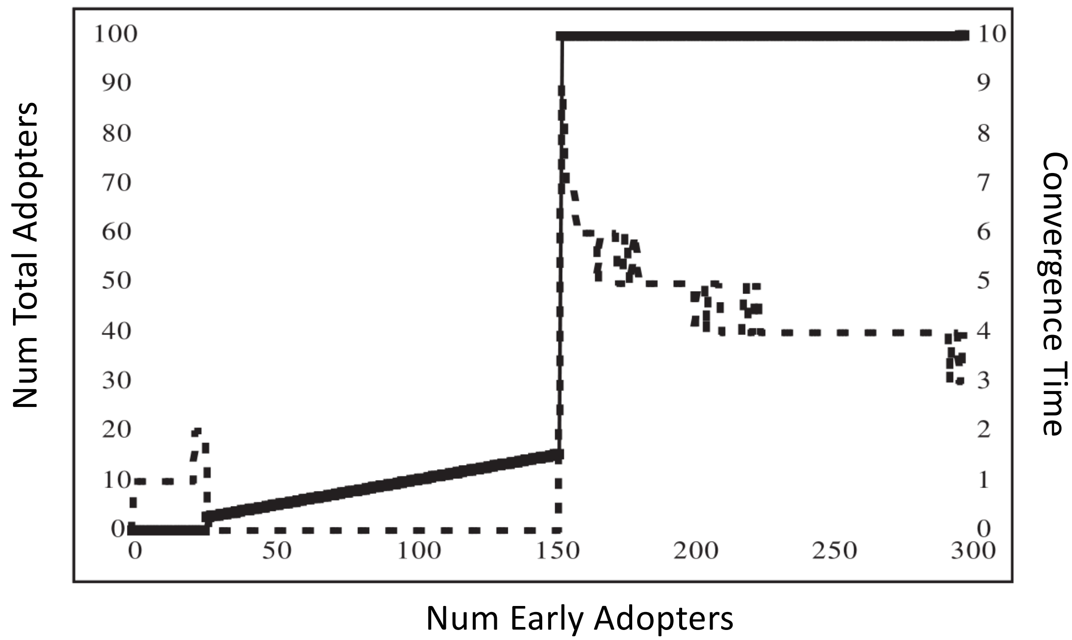

Hazard & Wurman, ICEC, 2007

Minority Game: The Path Less Taken

Challet et al., Oxford Press, 2005

- El Farol Bar problem
- Hard to find valuable unknowns in large population of smart agents
- Related to No Free Lunch Theorem: know the data

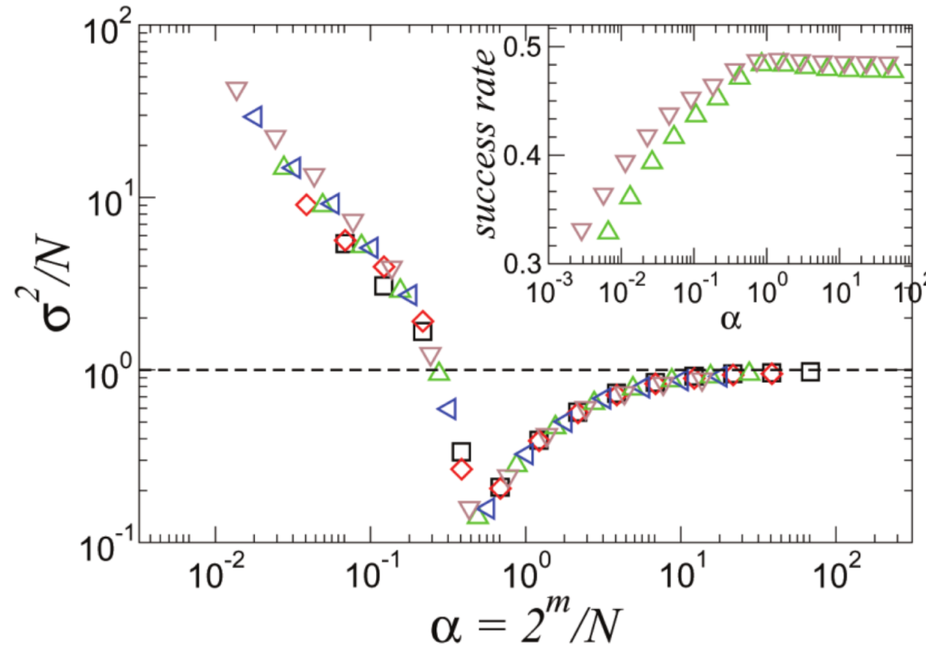

Esteban & Moro, '04

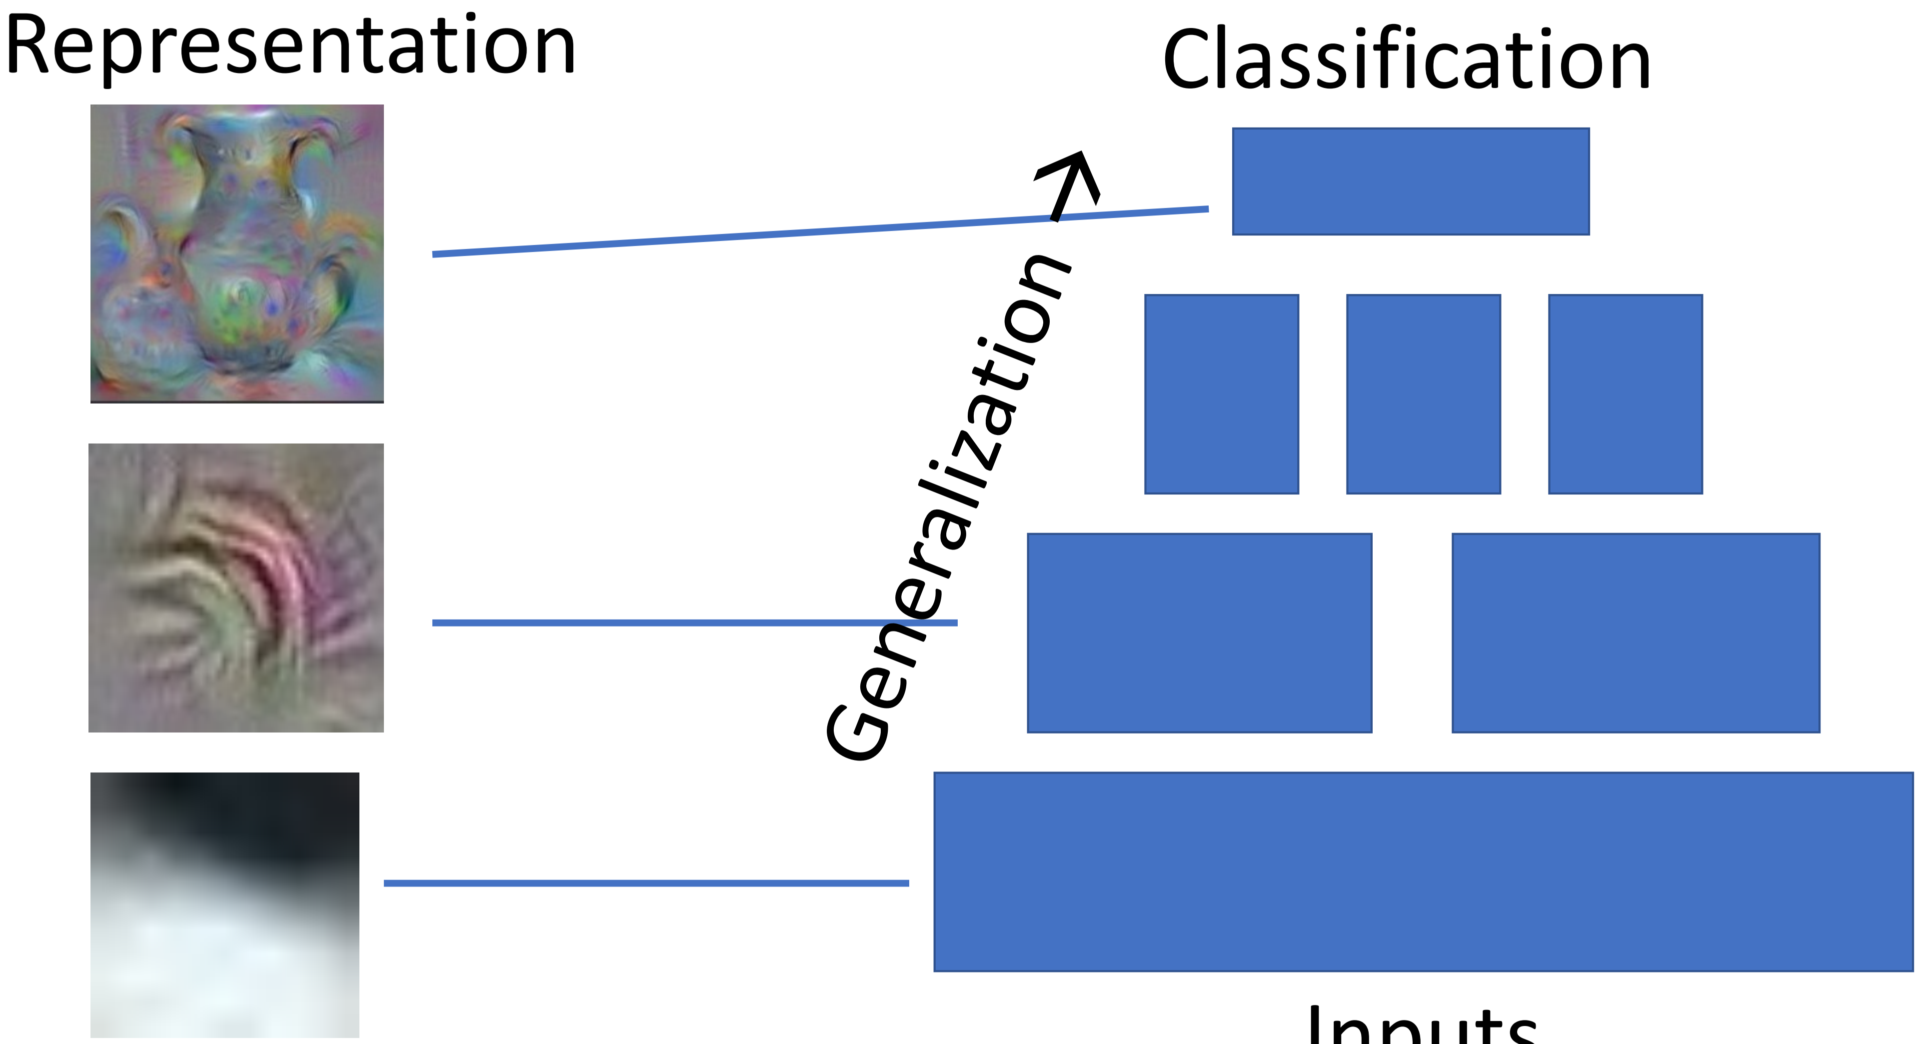

Yosinski et al., ICML DL 2015

What if we flatten a neural network?
Memorization without generalization

Neurons

Softmax

Inputs

Input

Logical conjunction: need a value for each combination of values (exponential!)

$$\sum_{j=1}^{m} w_j \sigma \left(\sum_{i=1}^{n} a_{ij} x_i \right) \approx \prod_{i=1}^{n} x_i$$

Output Weights

Input Scale

Lin, Tegmark, Rolnick, J Stat Physics, 2017

Desirability Index

Harrington, IQC, 1965

$$\sqrt[k]{\prod_{i=1}^{k} d_i(f_i(X_1, \ldots, X_n, 0))}$$

- Multicriteria optimization for innovating in chemistry, and chemical and mechanical engineering Trautmann, Drug Design Workshop, 2009

- Gaming and strategy

Point Recon, Hazardous Software, 2013

Generalized Diversity Index & Generalized Mean

$$^{q}D = \frac{1}{\sqrt[q-1]{\sum_{i=1}^{S} p_i p_i^{q-1}}}$$

$$M_p = \left(\frac{1}{n} \sum_{i=1}^{n} x_i^p \right)^{\frac{1}{p}}$$

$$^{q}D = \frac{1}{M_{q-1}}$$

$$^{1}D = \frac{1}{\prod_{i=1}^{R} p_i^{p_i}} = \exp\left(-\sum_{i=1}^{R} p_i \ln(p_i) \right)$$

$$M_0 = \sqrt[n]{\prod_{i=1}^{n} x_i}$$

Surprisal & Shannon Information

- Self-information: information of outcome of random event
- Surprisal: $-\log_2 P(x_i)$
- Information: Expected surprisal
- Information gain, KL-divergence, cross-entropy

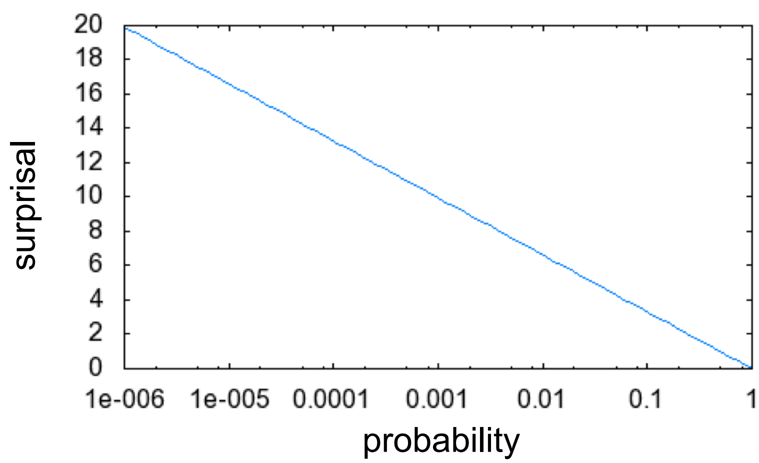

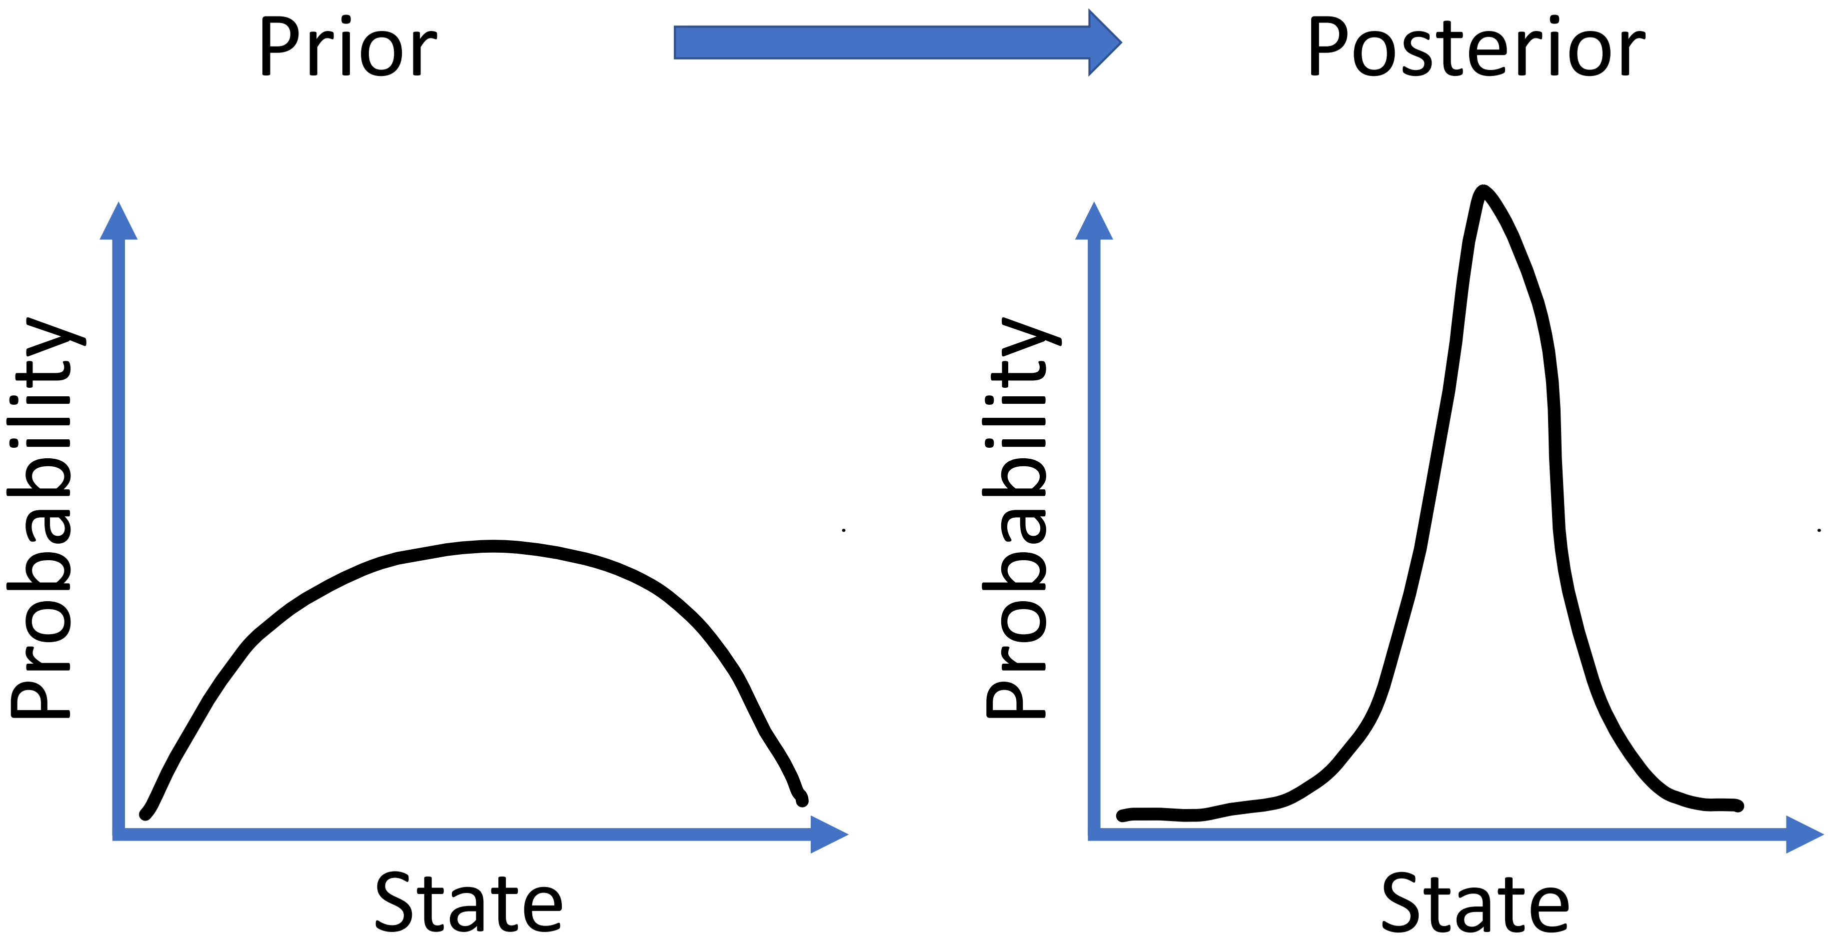
Prior
Posterior
Probability
State
Probability
State

Corpse Party
Chapter 1 Infirmary

Corpse Party
Chapter 1 Infirmary

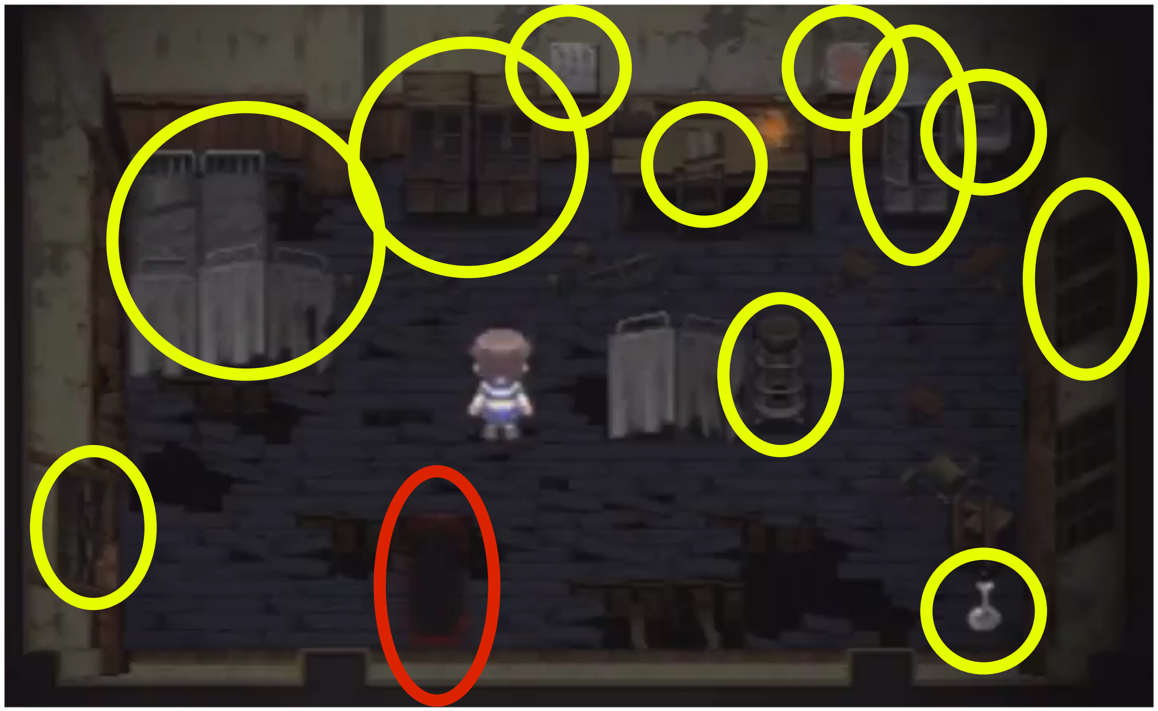

Infirmary Flow

take match from furnace

try door

try match

try door

try match

get rubbing alcohol

try door

exit

- Actual branching factor: 12
- Perceived branching factor: 11
- Exaggerated expectation

  [Hilbert, PSYCHOL BULL '12]

  - P(progress | revisit item) higher than anticipated

Infirmary Surprisal

- Player unsure of what to do, so assume uniform distribution over new possibilities:

  Q(X) ≈ 1/11, Q(Repeat) ≈ 0 => ~3.5 bits

- Correct distribution over possibilities, minimizing assumptions: P(X) = 1/12

$$D_{KL}(P||Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)}$$

Q(repeat) ≈ 0 means

1/12 * log((1/12) / 0) = 1/12 * ln(∞) = ∞

Massive surprisal if assume no repeat actions advance game

Measuring Complexity By Decision Information Rate

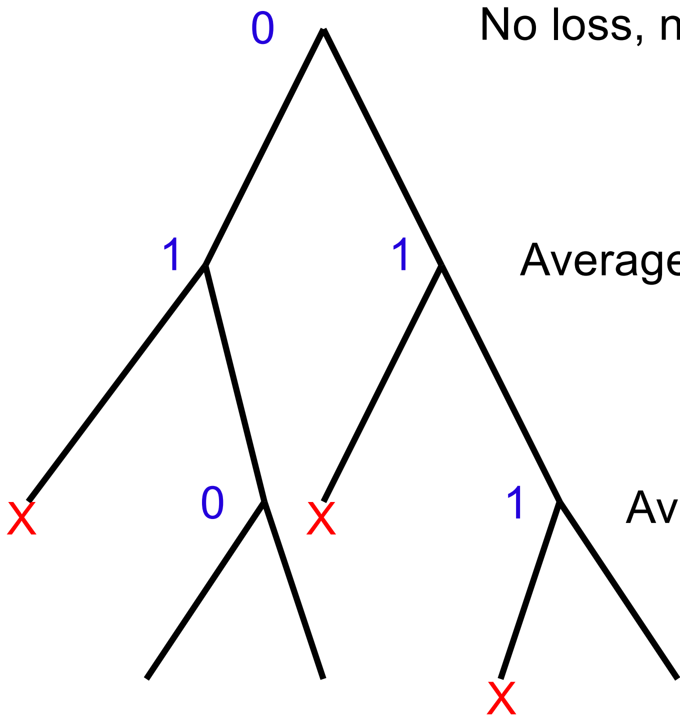

No loss, no information

Average 1 bit of information

Average 0.5 bits of information

3 out of 6 paths fail

1.5 bits of total information to succeed

1.5 bits / 2 steps = 0.75 bits per step to succeed

Combining Information Theory & Game Theory

- Maximum Entropy Correlated Equilibria
  (Ortiz et al., 2007)
- Measure information gain between player strategy and optimal
- Just add stochasticity!
  - Rock, Paper, Scissors:
    - 1/3 rock, 1/3 paper, 1/3 scissors
    - 1/4 rock, 1/4 paper, 1/2 scissors
- The value of soothsayers and randomness
- Robust sampling (e.g., Bayesian Optimization, MCCFR)

Peoples of the Steppe

Ambiguity of Strategy Via Information Theory: Maximum Difficulty

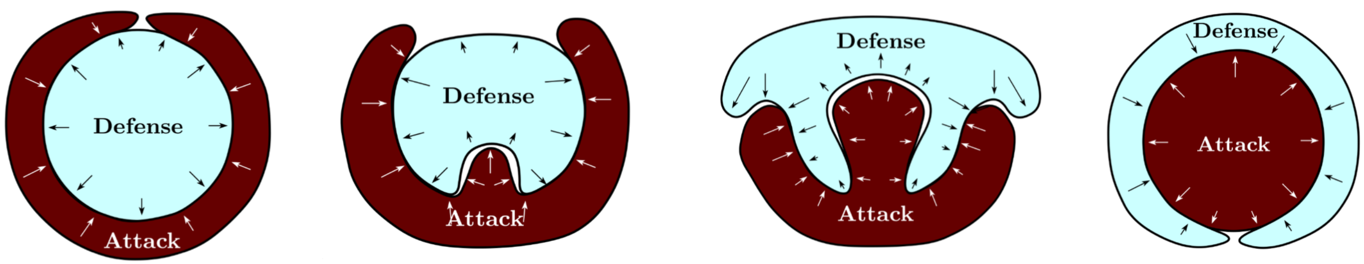

Fortification Honeypot Sampling Adaption

Pavlovic, Proc 2011 ACM New Sec Paradigms Workshop

Nomads → Pirates → Intellectual Property (Industrial Revolution) → Illicit Networks & Well-funded Startups

History Is Generalized & Compressed

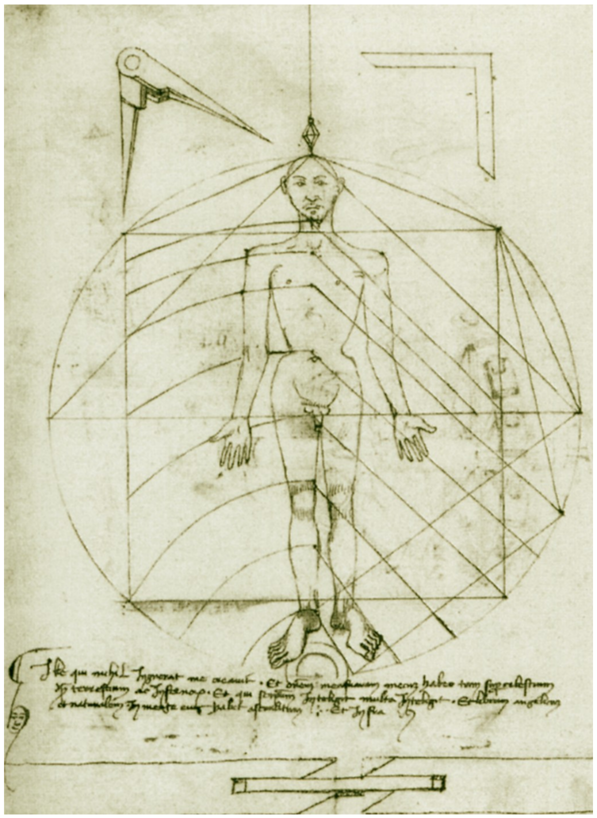

~1420, Taccola

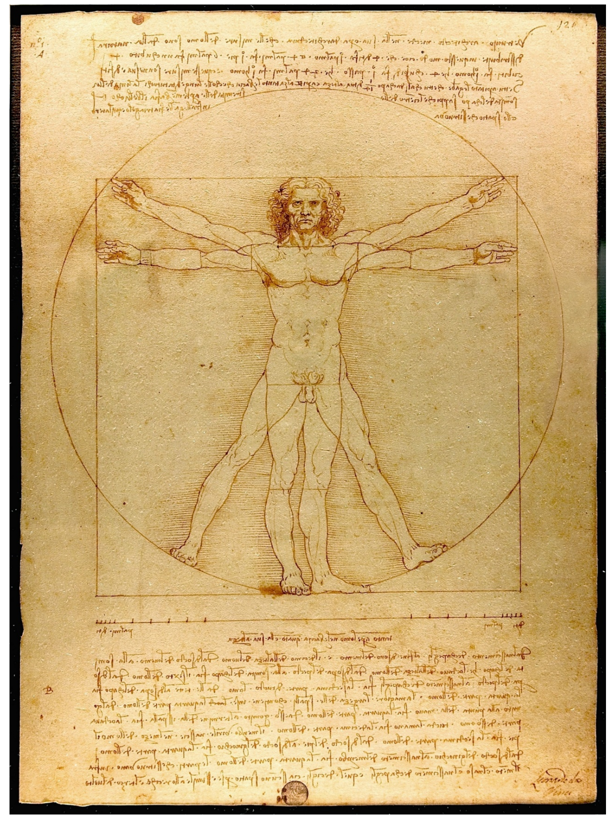

1490, da Vinci

A Formula for Measuring Creativity of a Solution

$$C(x, A, v_1, \dots, v_n) = \underset{a \in A}{min} \; D_{KL}(x|a) - \big(I(x) - I(a)\big) + \frac{1}{n}\sum_{i=1}^{n}(\ln v_i(x) - \ln v_i(a))$$

Compare to closest | Relative Novelty | Relative Complexity | Relative Desirability

x : configuration
A : set of known configuration
v_i : value function

Thanks!