DSC190: Machine Learning with Few Labels

"Standard Model" of ML

Zhiting Hu

Lecture 25, November 27, 2024

UC San Diego
HALICIOĞLU DATA SCIENCE INSTITUTE

Outline

"Standard Model" of ML

Presentations

- **Junyue Lin:** NeMo Guardrails: A Toolkit for Controllable and Safe LLM Applications with Programmable Rails
- **Charlie Gillet:** NeuralKart: A Real-Time Mario Kart 64 AI
- **Chojung (Angela) Tsai:** A Study on the Implementation Method of an Agent-Based Advanced RAG System Using Graph
- **Jason Dai:** Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena
- **Kevin Chan**: Discretization Drift in Two-Player Games
- **Ishaan Chadha**: Pearl: A Production-Ready Reinforcement Learning Agent
- **Asif Mahdin**: Solving Integrated Process Planning and Scheduling Problem via Graph Neural Network Based Deep Reinforcement Learning
- **Xiu Yuan**: Diffusion Policy

Experience of all kinds

Data examples

Type-2
diabetes is 90%
more common
than type-1

Rules/Constraints

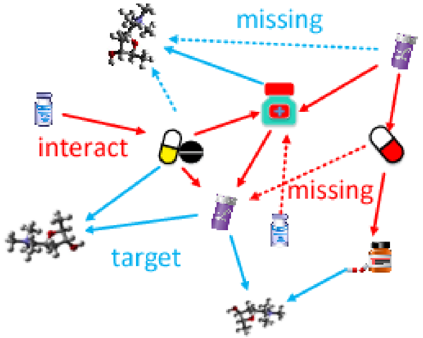

Knowledge graphs

Rewards

Auxiliary agents

Adversaries

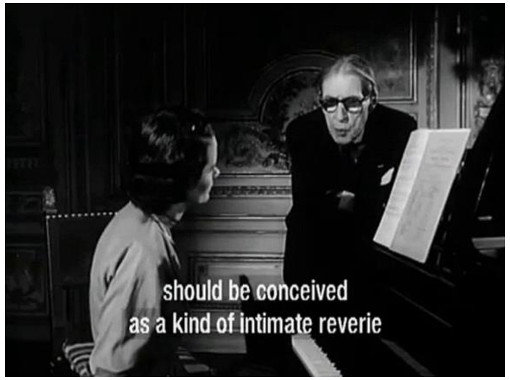

Master classes

…

- *And all combinations of such*
- *Interpolations between such*
- *…*

Human learning vs machine learning

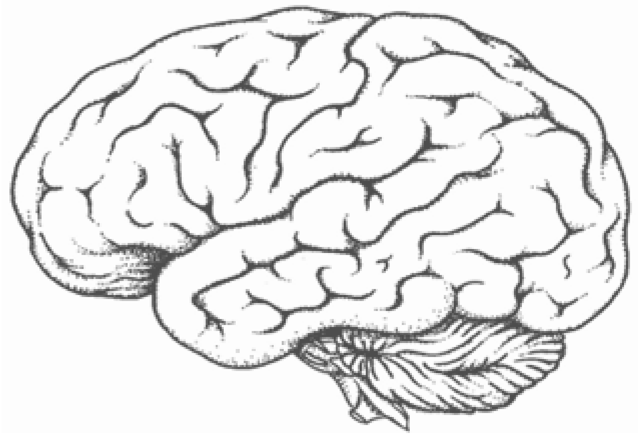

Type-2 diabetes is 90% more common than type-1

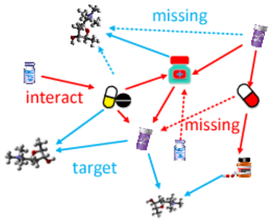

Data examples *Rules/Constraints* *Knowledge graphs* *Rewards* *Auxiliary agents*

Adversaries

Master classes

…

- *And all combinations of such*
- *Interpolations between such*
- *…*

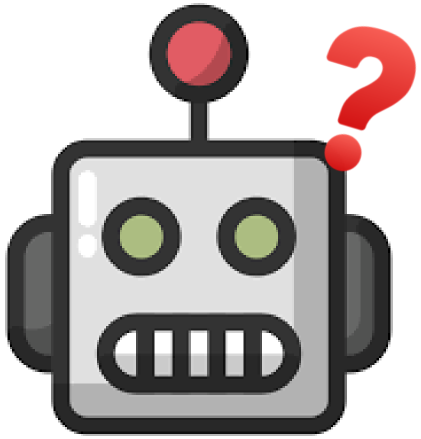

The zoo of ML/AI algorithms

maximum likelihood estimation

reinforcement learning as inference

data re-weighting

inverse RL

policy optimization

active learning

data augmentation

actor-critic

reward-augmented maximum likelihood

label smoothing

imitation learning

softmax policy gradient

adversarial domain adaptation

posterior regularization

GANs

constraint-driven learning

knowledge distillation

intrinsic reward

generalized expectation

prediction minimization

regularized Bayes

learning from measurements

energy-based GANs

weak/distant supervision

The zoo of ML/AI algorithms

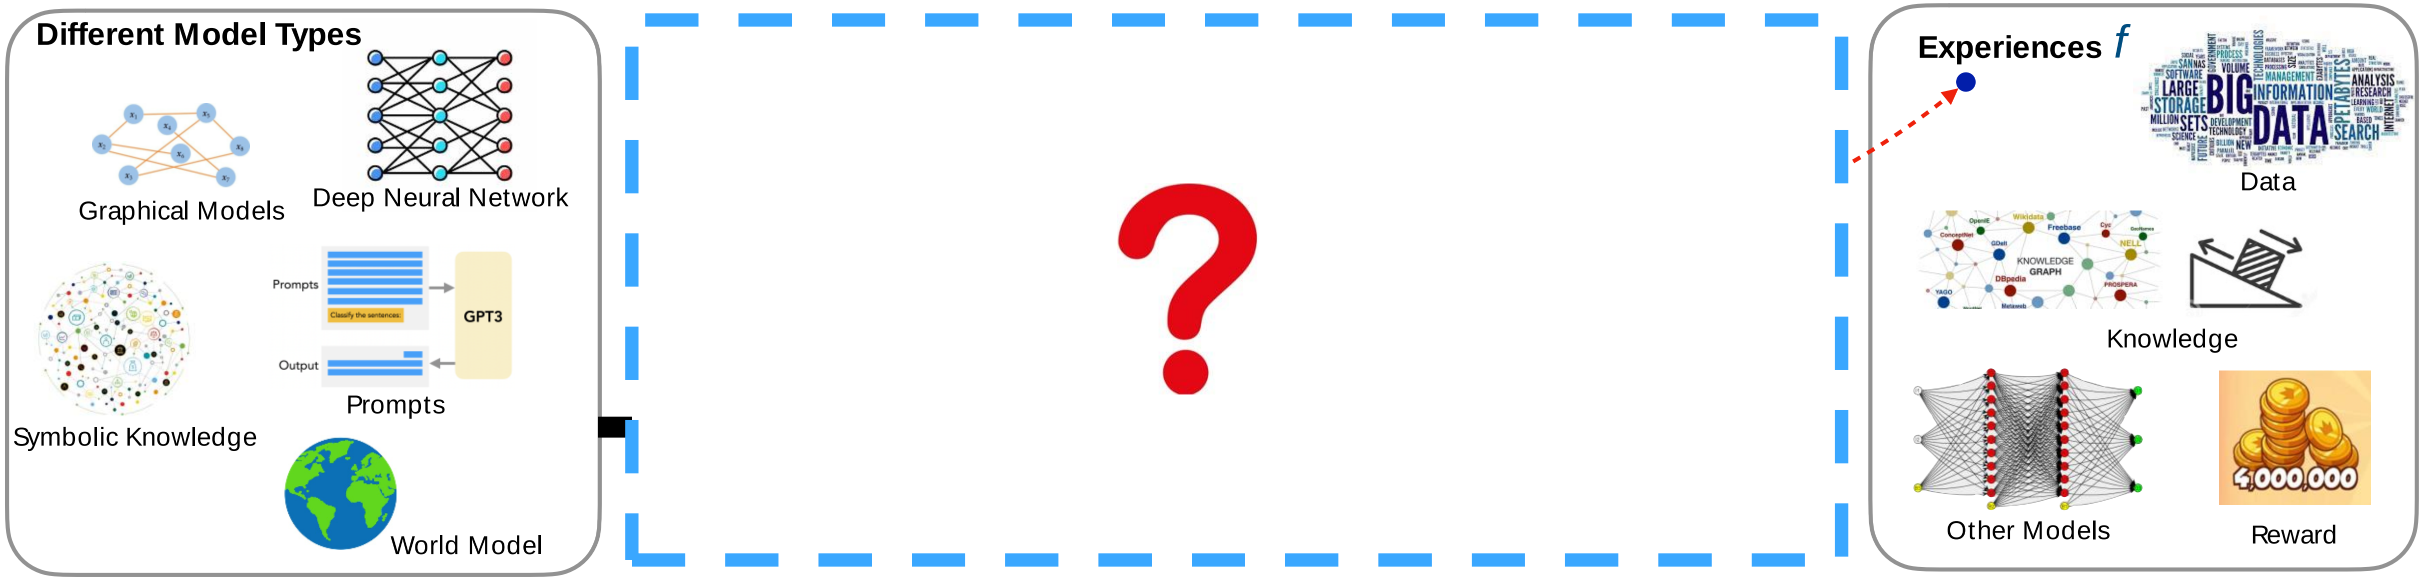

Standard Model in Physics

Diverse electro-magnetic theories

Maxwell's Eqns: original form

$e+\frac{df}{dx}+\frac{dg}{dy}+\frac{dh}{dz}=0$	(1)	Gauss' Law
$\mu\alpha=\frac{dH}{dy}-\frac{dG}{dz}$ $\mu\beta=\frac{dF}{dz}-\frac{dH}{dx}$ $\mu\gamma=\frac{dG}{dx}-\frac{dF}{dy}$	(2)	Equivalent to Gauss' Law for magnetism
$P=\mu\left(\gamma\frac{dy}{dt}-\beta\frac{dz}{dt}\right)-\frac{dF}{dt}-\frac{d\Psi}{dz}$ $Q=\mu\left(\alpha\frac{dz}{dt}-\gamma\frac{dx}{dt}\right)-\frac{dG}{dt}-\frac{d\Psi}{dy}$ $R=\mu\left(\beta\frac{dx}{dt}-\alpha\frac{dy}{dt}\right)-\frac{dH}{dt}-\frac{d\Psi}{dz}$	(3)	Faraday's Law (with the Lorentz Force and Poisson's Law)
$\frac{d\gamma}{dy}-\frac{d\beta}{dz}=4\pi p'$ $\quad p'=p+\frac{df}{dt}$ $\frac{d\alpha}{dz}-\frac{d\gamma}{dx}=4\pi q'$ $\quad q'=q+\frac{dg}{dt}$ $\frac{d\beta}{dx}-\frac{d\alpha}{dy}=4\pi r'$ $\quad r'=r+\frac{dh}{dt}$	(4)	Ampère-Maxwell Law
$P=-\xi p \quad Q=-\xi q \quad R=-\xi r$		Ohm's Law
$P=kf \quad Q=kg \quad R=kh$		The electric elasticity equation (**E** = **D**/ε)
$\frac{de}{dt}+\frac{dp}{dx}+\frac{dq}{dy}+\frac{dr}{dz}=0$		Continuity of charge

Simplified w/ rotational symmetry

$$\nabla\cdot\mathbf{D}=\rho_V$$

$$\nabla\cdot\mathbf{B}=0$$

$$\nabla\times\mathbf{E}=-\frac{\partial\mathbf{B}}{\partial t}$$

$$\nabla\times\mathbf{H}=\frac{\partial\mathbf{D}}{\partial t}+\mathbf{J}$$

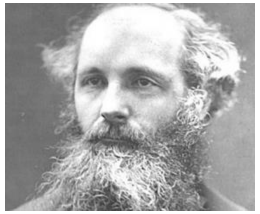

Further simplified w/ symmetry of special relativity

$$\varepsilon^{uvk\lambda}\partial_v F_{k\lambda}=0$$

$$\partial_v F^{uV}=\frac{4\pi}{c}j^u$$

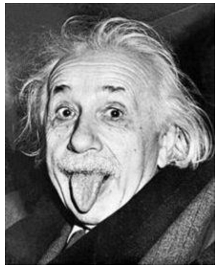

Standard Model w/ Yang-Mills theory and US(3) symmetry

$$\mathcal{L}_{gf}=-\frac{1}{2}\mathrm{Tr}(F^2)=-\frac{1}{4}F^{a\mu\nu}F^a_{\mu\nu}$$

Unification of fundamental forces?

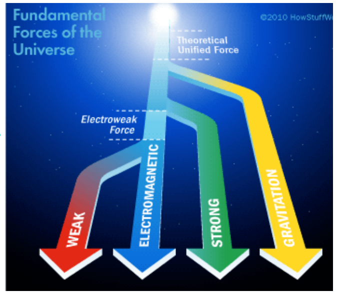

1861 1910s 1970s ?

Toward a 'Standard Model' of Machine Learning

Zhiting Hu[†,*], Eric P. Xing[‡,⋄,♮,**]
[†] Halıcıoğlu Data Science Institute, University of California San Diego, San Diego, USA
[‡] Machine Learning Department, Carnegie Mellon University, Pittsburgh, USA
[♮] Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE
[⋄] Petuum Inc., Pittsburgh, USA

[Hu & Xing, Harvard Data Science Review, 2022]: https://arxiv.org/abs/2108.07783

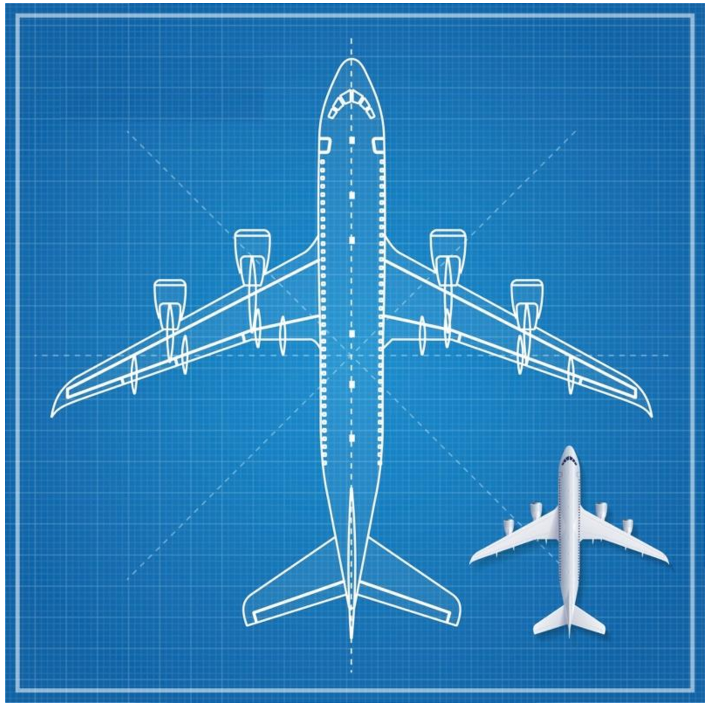

$$\min_{q,\theta} \; -\mathbb{E} + \mathbb{D} - \mathbb{H}$$

Experience Divergence Uncertainty

A “Standard Model” of ML

$$\min_{q,\,\theta} -\alpha\mathbb{H}(q) + \beta\mathbb{D}\left(q(\boldsymbol{t}), p_\theta(\boldsymbol{t})\right) - \mathbb{E}_{q(\boldsymbol{t})}\left[f(\boldsymbol{t})\right]$$

3 terms:

Uncertainty
(self-regularization)
e.g., Shannon entropy

Uncertainty

Divergence
(fitness)
e.g., Cross Entropy

Teacher
$q(\boldsymbol{t})$

Student
$p_\theta(\boldsymbol{t})$

Experiences
(exogenous regularizations)
e.g., data examples, rules

Textbook $f(\boldsymbol{t})$

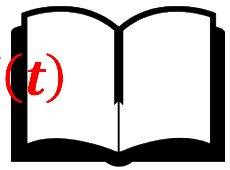

Check out more: Toward a ‘Standard Model’ of Machine Learning

Presentations

Questions?