What is missing from current AI and how we may bridge the gap

Ce qui manque à l'IA actuelle et comment nous pouvons combler l’écart

Yoshua Bengio, April 19th, 2023, WSAI, Montreal

CIFAR

Université de Montréal

IVADO

Mila

What is missing from ChatGPT?

- Human-like reasoning
- Including causal discovery and causal reasoning
- Providing high-level explanation and sources of facts
- Generalizing far from its (huge) training set (OOD generalization)
- Generalizing well with less data (1000 human lives spent reading?)

Mila

What is human-like reasoning?

- Inference (believed to be approximately Bayesian)

- Composing pieces of knowledge (to explain, plan, imagine, learn)
  - Which requires factorizing knowledge into composable pieces

- Global Workspace Theory (GTW) bottleneck
  - Thought: like a short sentence, very few concepts
  - Hard stochastic choice: inference over a few latent elements
  - Train of thoughts: incomplete but very informative

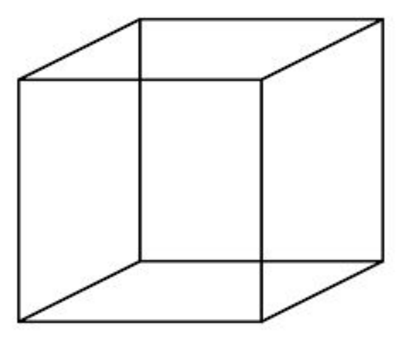

Necker cube is bistable

Unique cause for the gap: conscious processing?

Hypothesis:

This gap originates from a type of computation, knowledge representation and inference associated with conscious processing in humans, not yet mastered in AI

Mila

Conscious processing helps humans deal with odd settings

Faced with novel or rare situations, humans call upon conscious attention to combine on-the-fly the appropriate pieces of knowledge, to reason with them and imagine solutions.

⇨ we do not follow our habitual routines, we think hard to solve new problems.

Mila

System 1 vs. system 2 Cognition

2 systems (and categories of cognitive tasks):

Manipulates high-level / semantic concepts, which can be recombined combinatorially

System 1

- Intuitive, fast, **UNCONSCIOUS**, 1-step parallel, non-linguistic, habitual
- Implicit knowledge
- Current DL

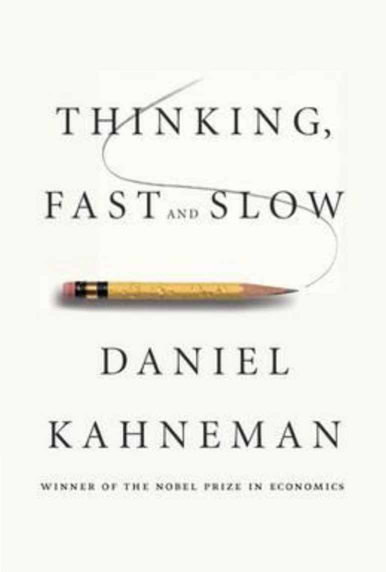

System 2

Slow, logical, sequential, **CONSCIOUS**, linguistic, algorithmic, planning, reasoning
Explicit knowledge
DL 2.0

Transfer to modified distribution: Beyond the iid assumption

- iid assumption too strong ⇨ poor out-of-distribution generalization
- relaxed assumptions: same causal dynamics, different state/intervention

Stochastic dynamical system

Initial conditions ⇨

⇨ Observed data

Mila

Causal model & OOD generalization

- Causal model vs regular distribution
  - joint over variables V and interventions I
  - family of distributions over V indexed by I
  - all sharing the same parameters!

⇨ Generalization over I = OOD generalization to distributions corresponding to unseen interventions

Mila

Learning & inference over causal model

- Effect of intervention I on P(V | I) = intervention surgery
- I may be observed or not or partially
- V may be partially observed
- Humans (and machines) can do and need probabilistic inference
- Surgery structure + causal constraints (temporality, marginal independence of causes) = inductive biases that can greatly help learn P(V, I) vs agnostic learning
- Other inductive biases (system 2), e.g., sparse causal graph

Bayesian causal modeling

- Markov Equivalence Class (MEC) = ambiguity even with infinite data
- Finite dataset D = more ambiguity about causal graph G
- Bayesian P(G | D) includes all sources of ambiguity
  - Converges to MEC as |D| ⇨ infinity
  - MEC = special case of P(G | D)
- P(G | D) may have exponential # modes ⇨ need powerful inference
- Full posterior: $$P(G, \theta, Z_1^T | X_1^T) \propto P(G)P(\theta|G) \prod_t P(Z_t|\theta, G)P(X_t|Z_t, \theta, G)$$

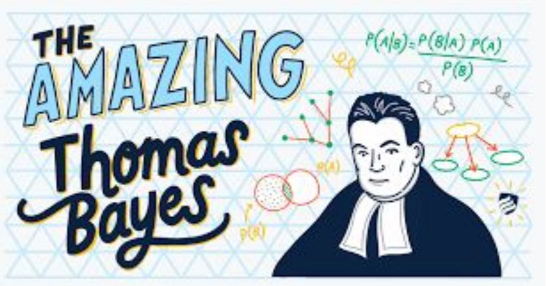

NeurIPS’2021
arXiv:2106.04399

Flow Network based Generative Models for Non-Iterative Diverse Candidate Generation

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, Yoshua Bengio




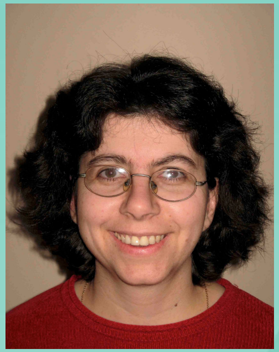
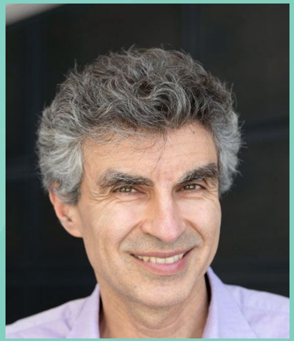

Extends amortized variational inference, but not using ELBO

Trains staged generator to sample X with probability proportional to R(x)

Multimodality & generalization

- Posterior inference P(Y|X)
- Y complicated (discrete/continuous), P(Y|X) highly multimodal
- GFN: decompose generation of Y|X into a sequence of stochastic steps

  # modes can grow exponentially with # steps

  If Y discrete, autoregressive generation = universal approximator
- Generalization power of ML: no need to visit all (X,Y) during training
  - Generalize from seen modes to unseen ones, unlike MCMC
- GFN = variational inference w/o ELBO

Neural Net Sequential Generation Policy

Deterministic environment:

$$s_{t+1} = T(s_t, a_t)$$

GFlowNet policy:

$$\pi(a_t|s_t) = P_F(s_{t+1}|s_t)$$

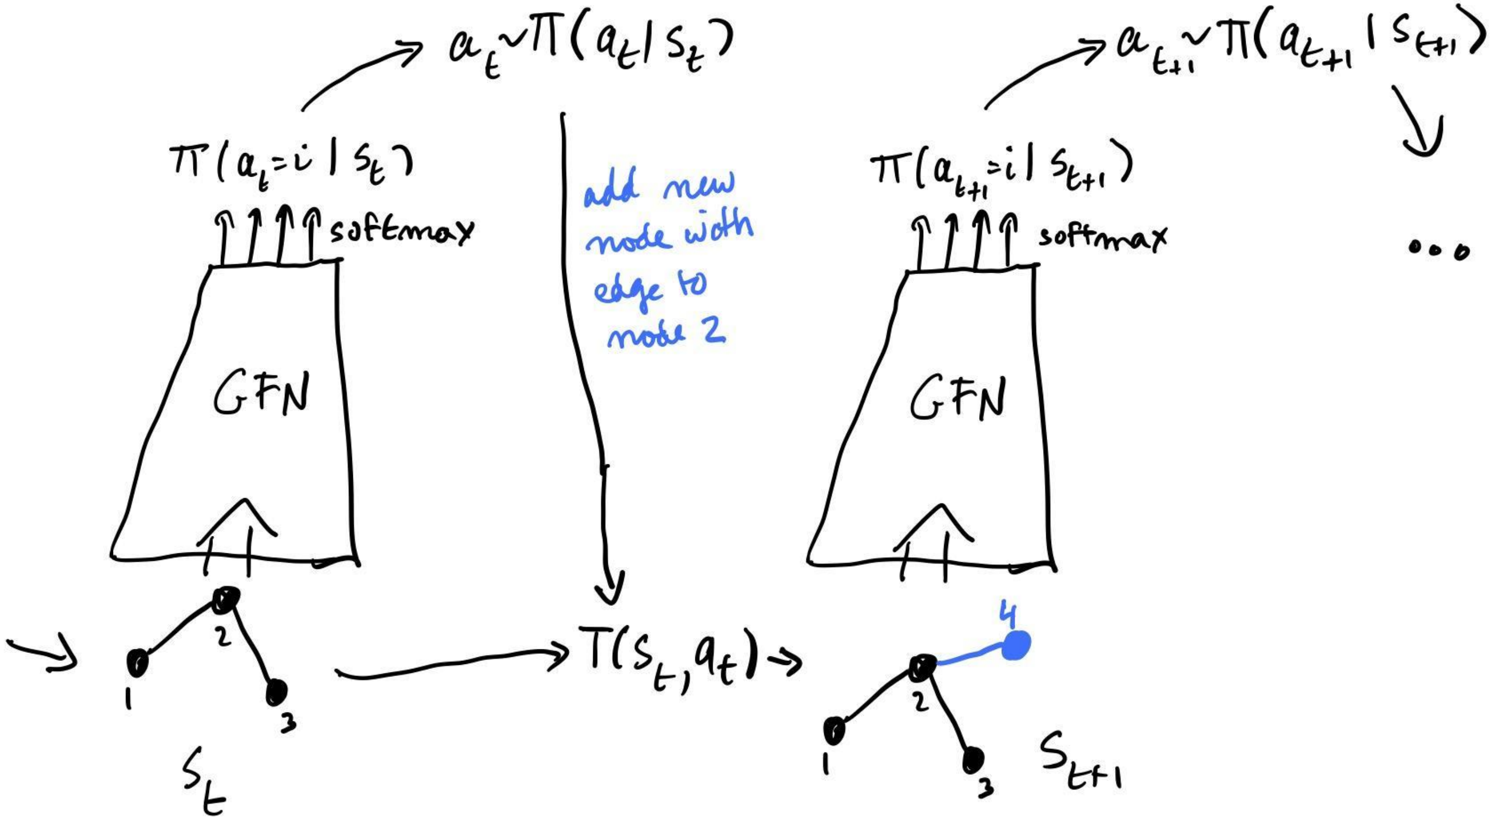

GFlowNets and Variational Inference

ICLR 2023, arXiv:2201.13259

Nikolay Malkin, Salem Lahlou, Tristan Deleu, Xu Ji, Edward Hu, Katie Everett, Dinghuai Zhang, Yoshua Bengio

       

A theory of continuous generative flow networks

arXiv:2301.12594

Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, Nikolay Malkin

DAG-GFN

- Generate a causal graph G sequentially while satisfying DAGness constraints exactly
- Generate latents Z | G (implicitly conditioning on data X)

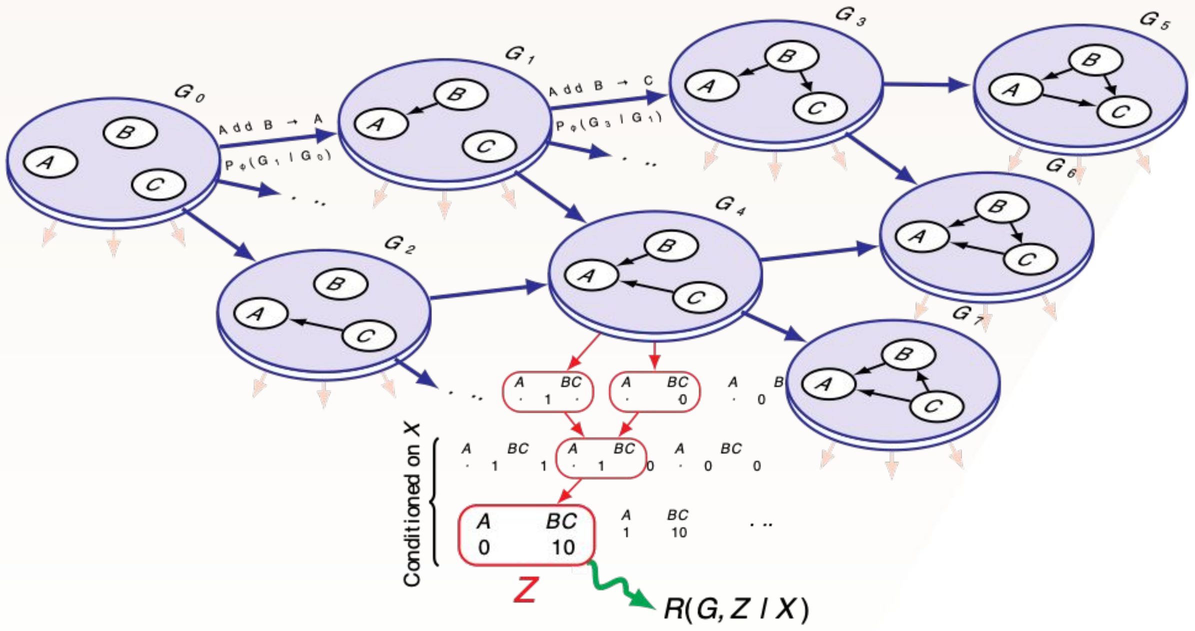

Mila

Bayesian Structure Learning with Generative Flow Networks

UAI'2022, arXiv:2202.13903

Tristan Deleu, António Góis, Chris Emezue, Mansi Rankawat, Simon Lacoste-Julien, Stefan Bauer, Yoshua Bengio







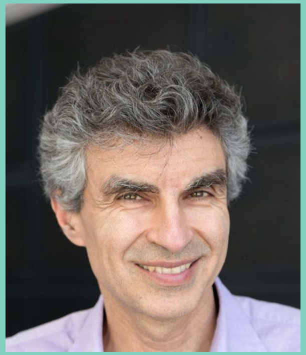

Model-based machine learning

- Separate inference machine (answer questions) from model
- Model's optimal capacity << inference machine's optimal capacity

⇨ Standard end-to-end deep learning to fit data confounds both
⇨ Overfits world model, underfits inference machine
Cannot incorporate causality inductive biases

- N.B.: inference machine can be trained with unlimited # queries to world model (like GFNs)

⇨ Bayesian model-based ML: learn P(model | data)

"Capacity" = $E_{P(\text{model}|\text{data})}[-\log P(\text{model})]$ = expected description length

⇨ Learning as well as question-answering become probabilistic inference problems

Mila

GFlowNet Tutorial

https://tinyurl.com/gflownet-tutorial

Papers at NeurIPS 2021, ICML 2022, UAI 2022, NeurIPS 2022, ICLR 2023 + ArXiv

Thank you!

Merci!