

GFlowNets to accelerate scientific discovery with machine learning

Alex Hernández-García (he/il/él)

Mathematical and Scientific Machine Learning 2023 · ICERM, Providence, RI · June 6th 2023



Mila

Université 
de Montréal

Slides: alexhernandezgarcia.github.io/slides/msml23

Collaborators

- Nikita Saxena
 - Moksh Jain
 - Chenghao Liu
 - Yoshua Bengio
- Kolya Malkin
 - Salem Lahlou
 - Alexandra Volokhova
 - Emmanuel Bengio

Outline

- Part 1: Motivation: Why scientific discovery?
 - Challenges, limitations and opportunities for machine learning
- Part 2: A brief introduction to GFlowNets
- Part 3: Multi-fidelity active learning with GFlowNets


Motivation: Why scientific discovery?

Part 1



Why scientific discovery?

Climate change is a major challenge for humanity.




Consequences:

- Melting glaciers and polar ice
- Sea level rise
- Heatwaves
- Floods
- Droughts
- Wildfires
- ...

Modelled and observed global average temperatures in the last 2 millenia
(source graphic: [The Guardian](#).)

Why scientific discovery?


Climate change is a major challenge for humanity.



Median global warming across modelled scenarios. Adapted from IPCC Sixth Assessment Report, 2022

Why scientific discovery?

Climate change is a major challenge for humanity.



Median global warming across modelled scenarios. Adapted from IPCC Sixth Assessment Report, 2022

"The evidence is clear: the time for action is now." IPCC Sixth Assessment Report, 2022

Why scientific discovery?

"The time for action is now"

"Limiting global warming will require major transitions in the energy sector. This will involve a substantial reduction in fossil fuel use, widespread electrification, **improved energy efficiency, and use of alternative fuels (such as hydrogen).**" IPCC Sixth Assessment Report, 2022

"Net-zero CO₂ emissions from the industrial sector are challenging but possible. Reducing industry emissions will entail coordinated action throughout value chains to promote all mitigation options, including demand management, **energy and materials efficiency, circular material flows,** as well as abatement technologies and transformational changes in production processes." IPCC Sixth Assessment Report, 2022

Why scientific discovery?

"The time for action is now"

"Limiting global warming will require major transitions in the energy sector. This will involve a substantial reduction in fossil fuel use, widespread electrification, **improved energy efficiency, and use of alternative fuels (such as hydrogen).**" IPCC Sixth Assessment Report, 2022

"Net-zero CO2 emissions from the industrial sector are challenging but possible. Reducing industry emissions will entail coordinated action throughout value chains to promote all mitigation options, including demand management, **energy and materials efficiency, circular material flows,** as well as abatement technologies and transformational changes in production processes." IPCC Sixth Assessment Report, 2022


Mitigation of the climate crisis requires transformational changes in the energy and materials efficiency.

Traditional scientific discovery loop

The climate crisis demands accelerating scientific discoveries.

The **traditional pipeline** for scientific discovery (paradigms 1-3):

- relies on **highly specialised human expertise**,
- it is **time-consuming** and
- **financially and computationally expensive**.




Machine learning in the loop

The traditional scientific discovery loop is too slow.

A machine learning model can be:

- trained with data from *real-world* experiments and
- used to quickly and cheaply evaluate queries



Machine learning in the loop


The traditional scientific discovery loop is too slow.

A machine learning model can be:

- trained with data from *real-world* experiments and
- used to quickly and cheaply evaluate queries

A machine learning model
replacing real-world experiments
can *only* provide *linear* gains.

Not enough if the search space is very large (10^{180} stable materials)




Generative machine learning in the loop

Can we do better than *linear*?
An agent in the loop.

A machine learning **agent** in the loop can:

- learn structure from the available data,
- generalise to unexplored regions of the search space and
- build better queries




Generative machine learning in the loop

Can we do better than *linear*?
An agent in the loop.

A machine learning **agent** in the loop could (ideally):

- learn structure from the available data,
- generalise to unexplored regions of the search space and
- build better queries

A successful AL pipeline with an ML agent in the loop can provide *exponential* gains.




Generative machine learning in the loop

GFlowNet as the agent.

A machine learning **agent** in the loop could (ideally):

- learn structure from the available data,
- generalise to unexplored regions of the search space and
- build better queries

A successful AL pipeline with an ML agent in the loop can provide *exponential* gains.



Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

→ Need for **efficient search and generalisation** of underlying structure.

Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

→ Need for **efficient search and generalisation** of underlying structure.

Challenge: underspecification of objective functions or metrics.

Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

→ Need for **efficient search and generalisation** of underlying structure.

Challenge: underspecification of objective functions or metrics.

→ Need for **diverse** candidates.

Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

→ Need for **efficient search and generalisation** of underlying structure.

Challenge: underspecification of objective functions or metrics.

→ Need for **diverse** candidates.

Limitation: Reinforcement learning and MCMC methods are good at optimisation but poor at mode mixing.

Machine learning for scientific discovery

Challenges and limitations of existing methods

Challenge: very large search spaces.

→ Need for **efficient search and generalisation** of underlying structure.

Challenge: underspecification of objective functions or metrics.


→ Need for **diverse** candidates.

Limitation: Reinforcement learning and MCMC methods are good at optimisation but poor at mode mixing.

→ Need for **multi-modal** optimisation.

A brief introduction to GFlowNets

Part 2



GFlowNet in a nutshell

Given a reward or objective function $R(x)$, GFlowNet can be seen a generative model trained to sample objects $x \in X$ according to a sampling policy $\pi(x)$ proportional to the reward $R(x)$: $\pi(x) \propto R(x)$

GFlowNet in a nutshell

Given a reward or objective function $R(x)$, GFlowNet can be seen a generative model trained to sample objects $x \in X$ according to a sampling policy $\pi(x)$ proportional to the reward $R(x)$: $\pi(x) \propto R(x)$

→ Sampling proportionally to the reward function induces multi-modal search and diversity.


GFlowNet in a nutshell

Given a reward or objective function $R(x)$, GFlowNet can be seen a generative model trained to sample objects $x \in X$ according to a sampling policy $\pi(x)$ proportional to the reward $R(x)$: $\pi(x) \propto R(x)$

→ Sampling proportionally to the reward function induces multi-modal search and diversity.


The policy $\pi_{\theta}(x)$ is modelled by a deep neural network, parameterised by θ , thus providing amortised inference.

→ Amortised inference can be thought of as *exploration with memory*, which induces systematic generalisation.




GFlowNet in a nutshell

- Objects $x \in X$ are constructed through a sequence of steps τ from an action space A .
- At each step of the trajectory $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_f)$, we get a partially constructed object s in state space S .
- This induces a directed acyclic graph (DAG) $G = (S, A)$, with all possible constructions in the domain.



GFlowNet in a nutshell


- Objects $x \in X$ are constructed through a sequence of steps τ from an action space A .
- At each step of the trajectory $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_f)$, we get a partially constructed object s in state space S .
- This induces a directed acyclic graph (DAG) $G = (S, A)$, with all possible constructions in the domain.



This terminology is reminiscent of reinforcement learning.



An intuitive toy example

Task: find arrangements of Tetris pieces on the board that minimise the empty space.



An intuitive toy example


Task: find arrangements of Tetris pieces on the board that minimise the empty space.




This task resembles designing DNA sequences or molecules or materials via fragments, with the objective of optimising certain properties.

An intuitive toy example


Task: find arrangements of Tetris pieces on the board that minimise the empty space.



State space S




Action space A




An intuitive toy example


Task: find arrangements of Tetris pieces on the board that minimise the empty space.



State space S



Action space A



An intuitive toy example

Task: find arrangements of Tetris pieces on the board that minimise the empty space.



Score: 0/12



Score: 4/12



Score: 8/12



Score: 12/12

An intuitive toy example

Task: find arrangements of Tetris pieces on the board that minimise the empty space.



Score: 12/12



Score: 12/12



Score: 12/12



Score: 12/12




Score: 12/12

The *reward function* of this task has multiple modes. With a larger board and more pieces, the number of combinations and modes grow exponentially and the task of efficiently finding them is non-trivial for machine learning models.

GFlowNet flows

The edges or transitions in the DAG can be quantified by their *flow*.

- Analogous to water-flow in pipes.
- Trajectory Flow $F(\tau)$ denotes probability mass assigned to trajectory τ .
- State Flow $F(s)$ is the flow of all trajectories passing through the state s .
- Edge Flow $F(s \rightarrow s')$ is the flow through a particular edge $s \rightarrow s'$.
- Forward Policy $P_F: P_F(s' | s) = \frac{F(s \rightarrow s')}{F(s)}$
- Backward Policy $P_B: P_B(s | s') = \frac{F(s \rightarrow s')}{F(s')}$



Principle of conservation as a training objective


Consistent Flow: Flow F satisfies the *flow consistency equation*

$$\sum_{s' \in \text{Parents}(s)} F_{\theta}(s' \rightarrow s) = \sum_{s' \in \text{Children}(s)} F_{\theta}(s \rightarrow s')$$

Theorem: For a consistent flow F with terminal flow set as the reward $F(x \rightarrow s_f) = R(x)$, the forward policy samples x proportionally to $R(x)$.

$$\pi(x) \propto R(x)$$

Corollary: The flow at s_0 , $F(s_0)$ is the partition function Z !



Principle of conservation as a training objective

$$\sum_{s' \in \text{Parent}(s)} F_{\theta}(s' \rightarrow s) = \sum_{s'' \in \text{Child}(s)} F_{\theta}(s \rightarrow s')$$

- **Flow Matching Objective:**

$$L_{FM}(s; \theta) = \left(\log \frac{\sum_{s' \in \text{Parent}(s)} F_{\theta}(s' \rightarrow s)}{\sum_{s'' \in \text{Child}(s)} F_{\theta}(s \rightarrow s'')} \right)^2$$

Results

Tetris GFlowNets

If the model is sufficiently trained, the sampling policy $\pi(x)$ should be proportional to the reward $R(x)$:
 $\pi(x) \propto R(x)$



$\pi(x) =$
8.12 %



$\pi(x) =$
8.96 %



$\pi(x) =$
8.61 %



$\pi(x) =$
9.16 %



$\pi(x) =$
8.39 %

After training, GFlowNet samples multiple (diverse) modes with high probability.

The energy function $\varepsilon(x)$ is the fraction of the board occupied by pieces and the reward function is $R(X) = \varepsilon(x)^4$ to disproportionately favour the discovery of modes.

Multi-fidelity active learning with GFlowNets

Part 3



Why multi-fidelity?

In many areas of scientific applications we have access to multiple approximations of the objective function.

For example, for material discovery:

- Synthesis of a material and characterisation of a property in the lab
- Density Functional Theory (DFT)
- An ensemble of large graph neural networks trained on DFT data
- An efficient, smaller neural network

Why multi-fidelity?


In many areas of scientific applications we have access to multiple approximations of the objective function.

For example, for material discovery:

- Synthesis of a material and characterisation of a property in the lab
- Density Functional Theory (DFT)
- An ensemble of large graph neural networks trained on DFT data
- An efficient, smaller neural network

However, current multi-fidelity methods struggle with structured, large, high-dimensional search spaces and lack **diversity**.

Multi-fidelity active learning with GFlowNets



Multi-fidelity surrogate models

- Small (synthetic) tasks: exact Gaussian Processes
- Larger-scale, benchmark tasks: Deep Kernel Learning with stochastic variational Gaussian processes

Multi-fidelity kernel learning:

$$K(x, \tilde{x}, m, \tilde{m}) = K_1(x, \tilde{x}) \times K_2(m, \tilde{m})$$

- K_1 : RBF kernel
- K_2 : Downsampling kernel

- Wilson, Hu et al. Deep Kernel Learning, AISTATS, 2016.
- Wu et al. Practical multi-fidelity Bayesian optimization for hyperparameter tuning , UAI, 2019.

Multi-fidelity acquisition function

Maximum Entropy Search (MES)

MES it aims to maximise the mutual information between the value of the objective function f when choosing point x and the maximum of the objective function, f^* (instead of considering the $\arg \max$).

The multi-fidelity variant is designed to select the candidate x and the fidelity m that maximise the mutual information between f_M^* and the oracle at fidelity m , f_m , weighted by the cost of the oracle λ_m .

$$\alpha(x, m) = \frac{1}{\lambda_m} I(f_M^*; f_m \mid \mathcal{D})$$

- Moss et al. GIBBON: General-purpose Information-Based Bayesian Optimisation, JMLR, 2021.

Multi-fidelity GFlowNets (MF-GFN)

Given a baseline GFlowNet with state space S and action space A , we augment the state space with a new dimension for the fidelity $M' = 0, 1, 2, \dots, M$ (including $m = 0$, which corresponds to unset fidelity).

The set of allowed transitions A_M is augmented such that a fidelity $m > 0$ of a trajectory must be selected once, and only once, from any intermediate state. This is meant to provide flexibility and improve generalisation.

Finished trajectories are the concatenation of an object x and the fidelity m .

GFlowNet is trained with the acquisition function $\alpha(x, m)$ as reward function.

Experiments


Baselines

- **SF-GFN**: GFlowNet with highest fidelity oracle to establish a benchmark for performance without considering the cost-accuracy trade-offs.
- **Random fid. GFN**: GFlowNet with random fidelities, that is a variant of SF-GFN where the candidates are generated with the GFlowNet but the fidelities are picked randomly and a multi-fidelity acquisition function is used, to investigate the benefit of deciding the fidelity with GFlowNets.
- **Random**: Quasi-random approach where the candidates and fidelities are picked randomly and the top (x, m) pairs scored by the acquisition function are queried.
- **MF-PPO**: Instantiation of multi-fidelity Bayesian optimisation where the acquisition function is optimised using proximal policy optimisation (reinforcement learning).


Synthetic tasks: Branin and Hartmann

Branin: 100×100 grid, 3 oracles (from the BO literature).

Hartmann: 6D grid of length 10, 3 oracles (from the BO literature).



Branin task




Hartmann task


DNA aptamers and antimicrobial peptides (AMP)

DNA: GFlowNet adds one nucleobase (A, T, C, G) at a time up to length 30. This yields a design space of size $|X| = 4^{30}$. The objective function is the free energy estimated by NUPACK. The (simulated) lower fidelity oracle is a transformer trained with 1 million sequences.

AMP: Protein sequences with variable length (max. 50). The oracles are 3 ML models trained with different subsets of data.




DNA task




AMP task

Small molecules

More realistic experiments, with oracles that correlate with experimental results as approximations of the scoring function. The costs reflect the computational demands of each oracle (1, 3, 7).



Ionisation potential task



Electron affinity task

Summary and conclusions



Summary and conclusions

- Tackling the most pressing problems for humanity, such as the climate crisis and the threat of global pandemics, requires **accelerating the pace of scientific discovery**.
 - Current AI tools are not enough to truly utilize all the information and resources at our disposal.
 - AI-driven scientific discovery demands learning methods that can **efficiently discover diverse candidates in very large, multi-modal search spaces**.
 - **GFlowNet** is a learning method for amortised inference that can sample proportionally to a reward function.
 - **Multi-fidelity active learning with GFlowNets** enables **cost-effective exploration** of large, high-dimensional and structured spaces, and discovers multiple, diverse modes of black-box score functions.
-
- Hernandez-Garcia, Saxena et al. Multi-fidelity active learning with GFlowNets. Soon on arXiv!
 - Jain et al. GFlowNets for AI-Driven Scientific Discovery. Digital Discovery, Royal Society of Chemistry, 2023.
 - Jain et al. Biological Sequence Design with GFlowNets, ICML, 2022.

Summary and conclusions

- Tackling the most pressing problems for humanity, such as the climate crisis and the threat of global pandemics, requires **accelerating the pace of scientific discovery**.
- Current AI tools are not enough to truly utilize all the information and resources at our disposal.
- AI-driven scientific discovery demands learning methods that can **efficiently discover diverse candidates in very large, multi-modal search spaces**.
- **GFlowNet** is a learning method for amortised inference that can sample proportionally to a reward function.
- **Multi-fidelity active learning with GFlowNets** enables **cost-effective exploration** of large, high-dimensional and structured spaces, and discovers multiple, diverse modes of black-box score functions.


Open source code: github.com/alexhernandezgarcia/gflownet

- Hernandez-Garcia, Saxena et al. Multi-fidelity active learning with GFlowNets. Soon on arXiv!
- Jain et al. [GFlowNets for AI-Driven Scientific Discovery](#). Digital Discovery, Royal Society of Chemistry, 2023.
- Jain et al. [Biological Sequence Design with GFlowNets](#), ICML, 2022.

Acknowledgements

- Nikita Saxena (Mila) - equivalent contribution
- Moksh Jain (Mila)
- Chenghao Liu (Mila)
- Yoshua Bengio (Mila)
- IVADO
- CIFAR

Thanks! Questions?



Alex Hernández-García (he/il/él)



Mila

Université
de Montréal



alexhernandezgarcia.github.io | alex.hernandez-garcia@mila.quebec

[@alexhg@scholar.social](https://scholar.social/@alexhg)  | [@alexhdezgcia](https://twitter.com/alexhdezgcia) 

Slides: alexhernandezgarcia.github.io/slides/msml23