Beyond the Known: Probabilistic Inference for the AI Scientist

Marcin Sendera
PhD candidate at Jagiellonian University /
ex-research intern at Mila – AI Quebec Institute

Witold Lipski Award Winner's Talk, ML in PL 2025

The Grand Challenge: From Imitation to Discovery

LLMs (+ RL)?

HIC SVNT LEONES/DRACONES

The Real Challenge: Fantastic Beasts and **Where/How to Find Them**

LLMs (+ RL)?

HIC SVNT LEONES/DRACONES

The Principled Path is Probabilistic

$$P(Hypothesis \mid Data) = \frac{P(Data \mid Hypothesis) \cdot P(Hypothesis)}{P(Data)}$$

Only three components:

$P(Data \mid Hypothesis)$ - likelihood

$P(Hypothesis)$ - prior

$P(Data)$ - evidence

As Scientists We are All Bayesian, Even a Little Bit

Or at least we should be…

Example: Will we achieve AGI by 2035?

prior:

Pessimist

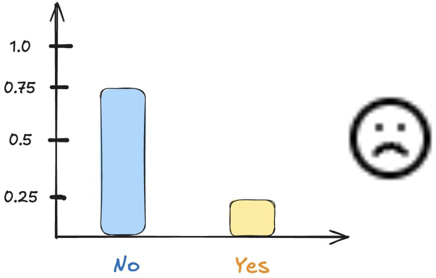

Neutral

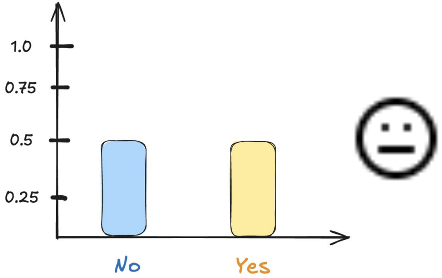

Optimist

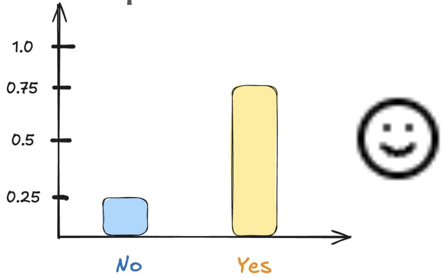

New evidence:

Google DeepMind @GoogleDeepMind · Jul 21

An advanced version of Gemini with Deep Think has officially achieved gold medal-level performance at the International Mathematical Olympiad. 🥇

It solved 5️⃣ out of 6️⃣ exceptionally difficult problems, involving algebra, combinatorics, geometry and number theory. Here's how 🧵

Example: Will we achieve AGI by 2035?

prior:

Pessimist

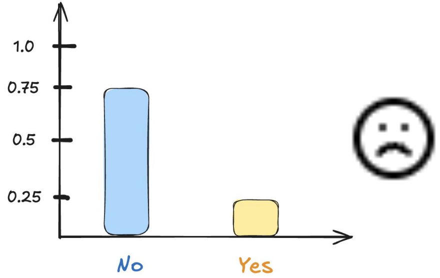

Neutral

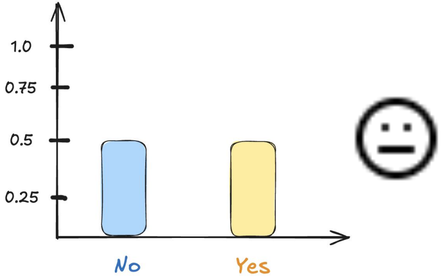

Optimist

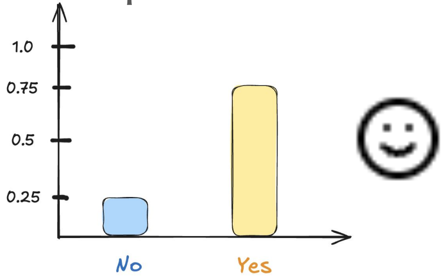

New evidence:

Google DeepMind @GoogleDeepMind · Jul 21

An advanced version of Gemini with Deep Think has officially achieved gold medal-level performance at the International Mathematical Olympiad. 🥇

It solved 5 out of 6 exceptionally difficult problems, involving algebra, combinatorics, geometry and number theory. Here's how 🧵

update:

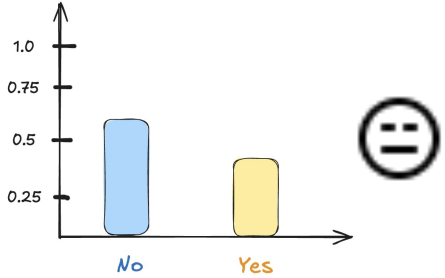

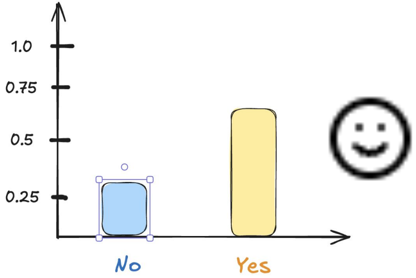

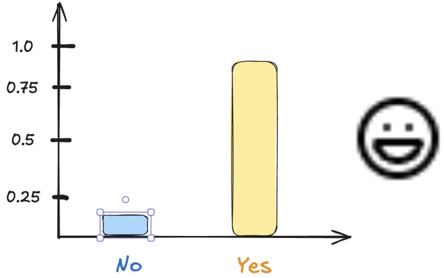

The Intractability Wall

$$P(Hypothesis \mid Data) = \frac{P(Data \mid Hypothesis) \cdot P(Hypothesis)}{P(Data)}$$

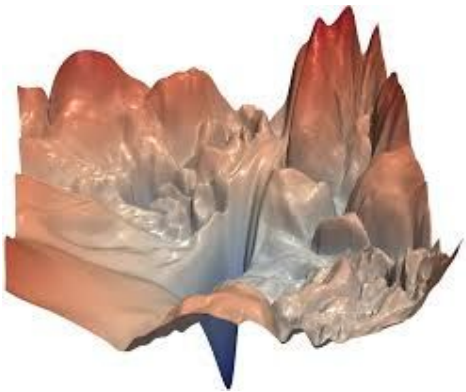

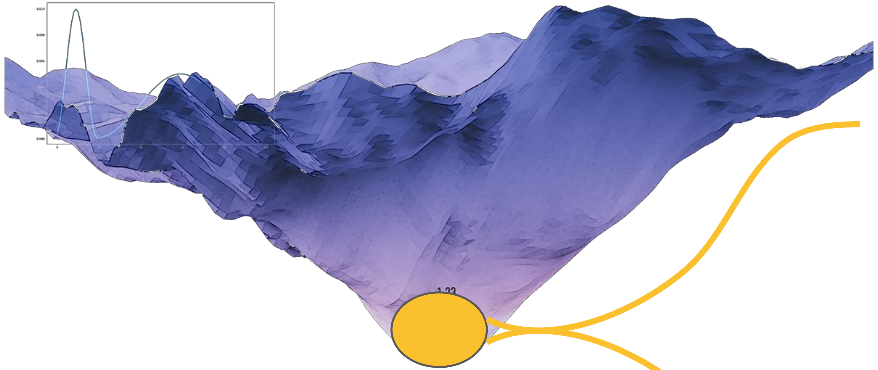

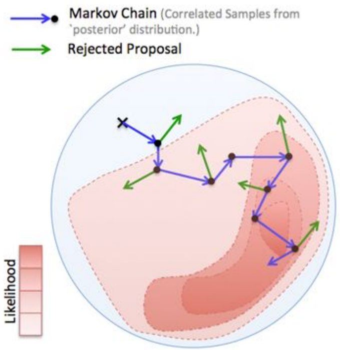

Wandering around the space is not the best idea!

slow MCMC sampler

A Better Strategy - Learn the Map!

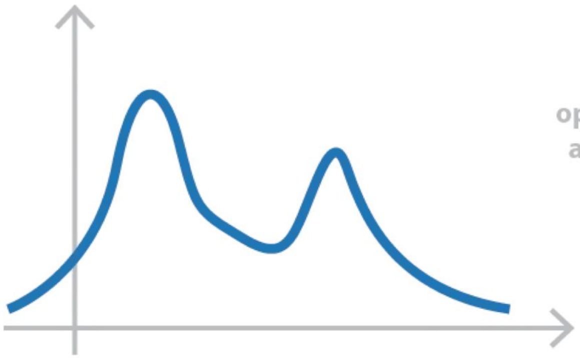

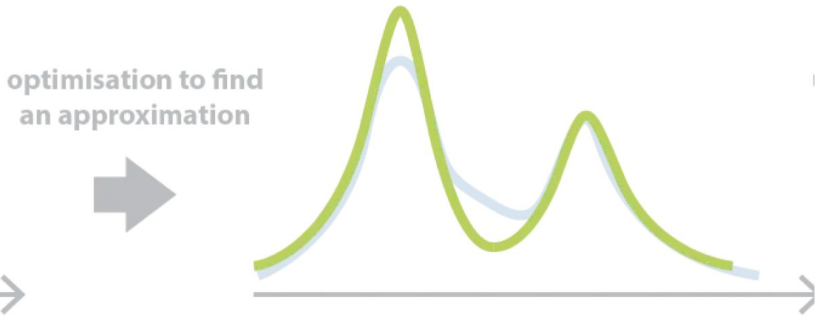

True posterior - intractable!

hard!

Variational Map - tractable!

simple!

The idea is simple but incredibly powerful:

learn a simpler proxy map of the territory

A New Paradigm: From Sequential to Joint

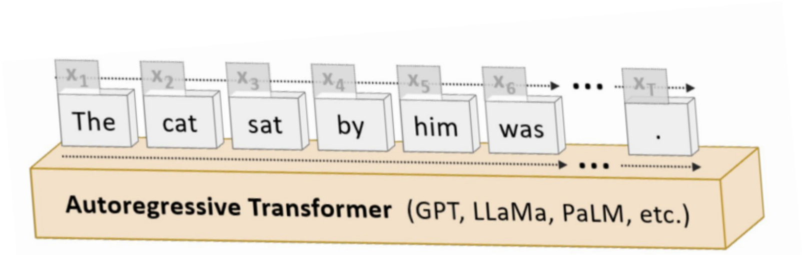

Sequential (Autoregressive) Models

$$p(x_i | x_{<i})$$

Joint (Holistic) Models

$$p(x)$$

How Diffusion Models Learn the Map?

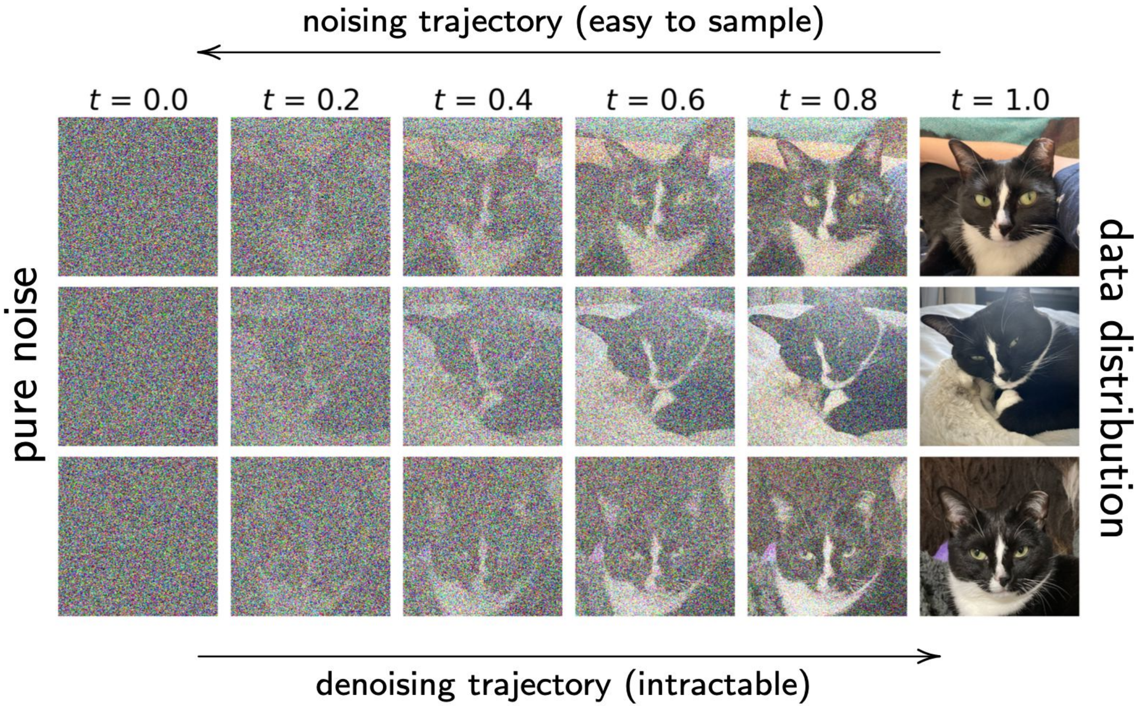

1. Noising trajectory - systematically adding noise
2. Learning how to reverse the noise

A diffusion model is...

1. A stochastic differential equation (that can be discretised)
2. A deep hierarchical model
3. **A policy in a Markov decision process (taking Gaussian steps)!**

Diffusion + (entropy-regularised) RL superpowers

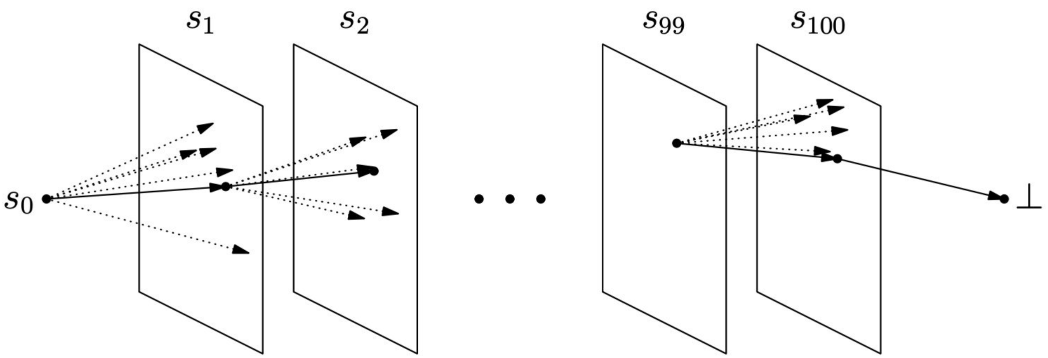

If a think of a diffusion model as a policy, we can:

- Learn to sample a **target density** (no data)
- Use **non-Gaussian policies** if desired
- Train using **off-policy methods** $\Rightarrow$ flexible exploration **(but how?)**
- **Asymptotically equivalent to neural SDEs!**

Problems in samplers: finding all types of good solutions (mode-seeking; **finding!**) and giving them appropriate approximate probabilities (mode-covering; **mapping!**)

E.g.: finding all types of "nice" chemical reactions vs. explore each of them

Exploration in RL-based Diffusion

Improving explorations by:

1. Adding Langevin dynamics (better finding)
2. Local search with prioritised replay buffer (better mapping)
3. And many others

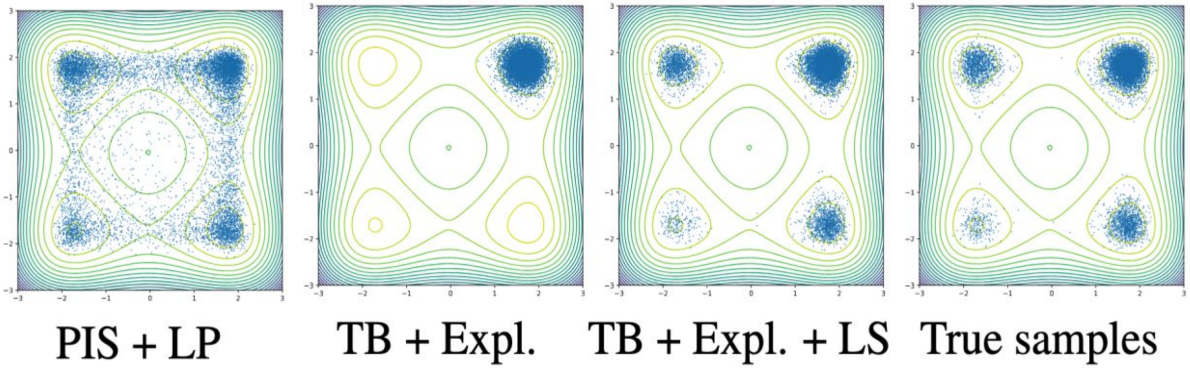

Sendera, Marcin, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. "Improved off-policy training of diffusion samplers." *Advances in Neural Information Processing Systems* 37 (2024): 81016-81045.

Sampling from the Posterior

What if we have a great prior and we want to sample from the posterior?

$$\mathbf{x} \sim p^{\text{post}}(\mathbf{x}) \propto p(\mathbf{x})r(\mathbf{x})$$

-> New loss function!

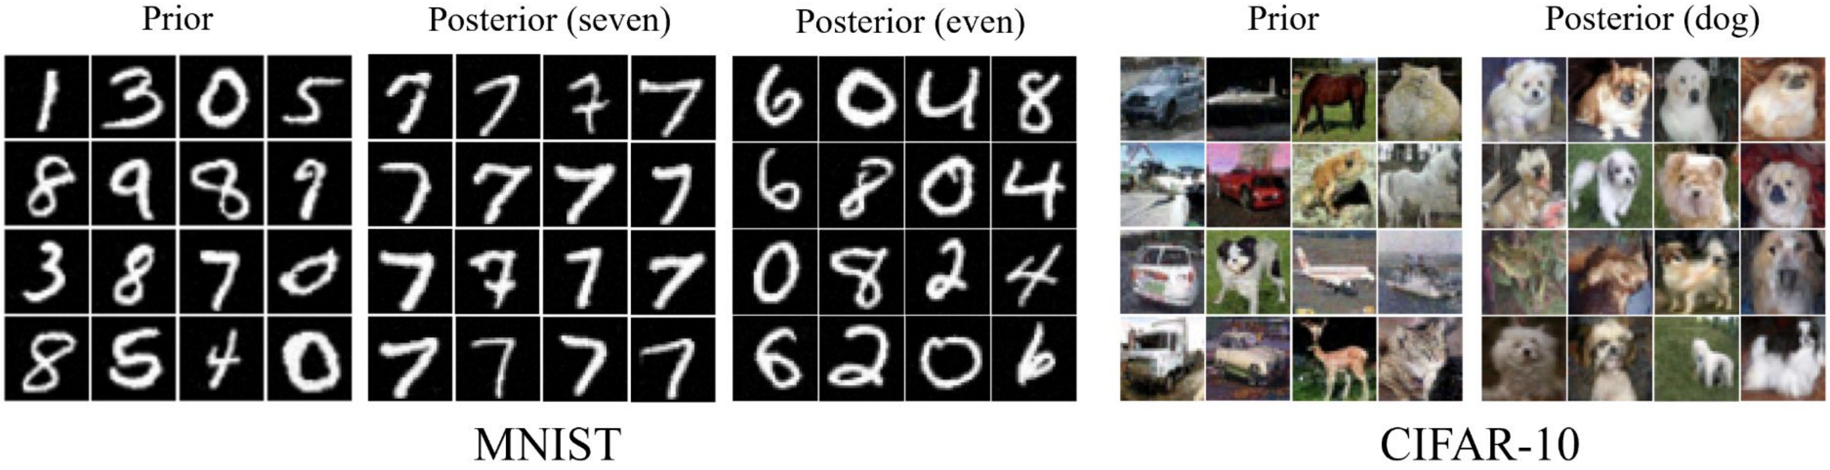

Venkatraman, Siddarth*, Moksh Jain*, Luca Scimeca*, Minsu Kim*, **Marcin Sendera***, Mohsin Hasan, Luke Rowe et al. "Amortizing intractable inference in diffusion models for vision, language, and control." *Advances in neural information processing systems* 37 (2024): 76080-76114.

Sampling from the Posterior

Usage in Science -> Lensing problem, working with astrophysicists

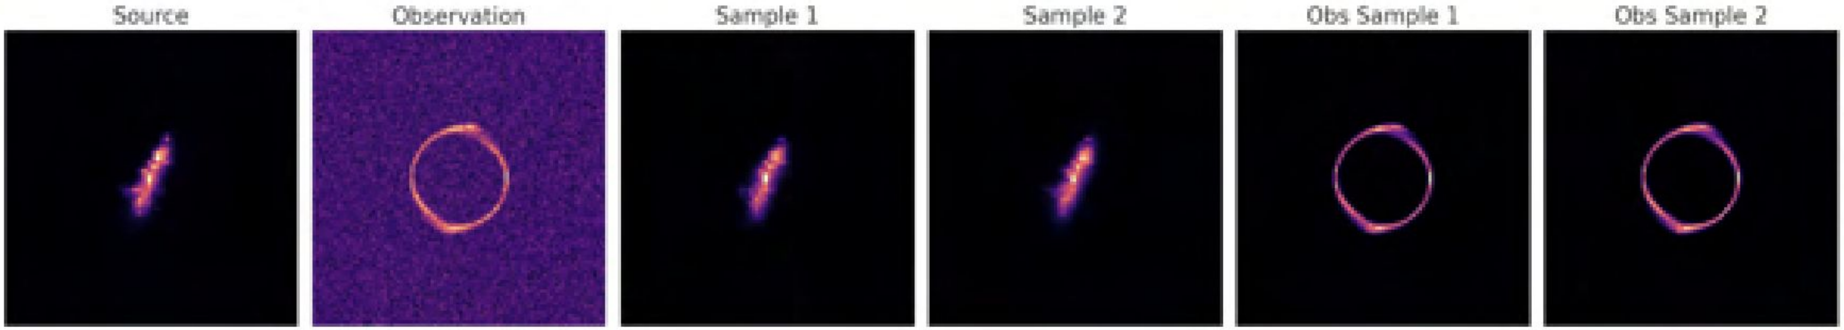

AI Scientist - exploring revisited

What we know:

1. Amortised variational inference gives us faster searching across the space, mode seeking and mode coverage than MCMC
2. Diffusion models are great for scientific discovery since they are considering the objects holistically
3. RL-based approaches allows for the search outside the known Science

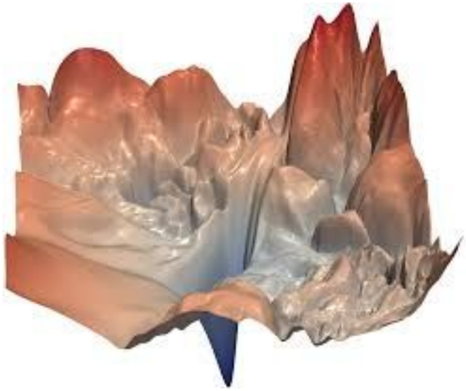

Well, can we really find new knowledge? Isn't that all only a sci-fi story?

Well, not entirely!

Wes Roth @WesRothMoney

AlphaEvolve Just Helped Prove New Theorems in Complexity Theory

Google DeepMind's AlphaEvolve just made real breakthroughs in theoretical computer science.

Instead of generating full proofs, it discovered new combinatorial structures that plug into existing proof frameworks, leading to verified, publishable theorems in complexity theory.

The team improved the inapproximability bound for MAX-4-CUT and found massive Ramanujan graphs never seen before, all with provable correctness.

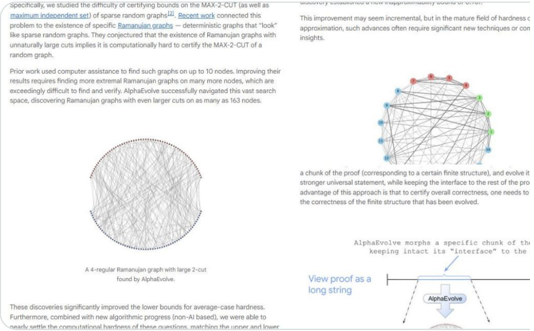

sakana.ai

The AI Scientist Generates its First Peer-Reviewed Scientific Publication

March 12, 2025

Math, Inc. @mathematics_inc

Today we're announcing Gauss, our first autoformalization agent that just completed Terry Tao & Alex Kontorovich's Strong Prime Number Theorem project in 3 weeks—an effort that took human experts 18+ months of partial progress.

7:38 PM · Sep 11, 2025 · **1.2M** Views

It's happening now!

Thank You & Acknowledgements

Mila & friends

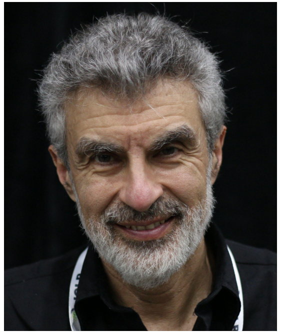




JAGIELLONIAN UNIVERSITY IN KRAKÓW & friends




A Final Thought: Investing in Science

Q&A

Contact info

e-mails: marcin.sendera@{gmail.com; uj.edu.pl; mila.quebec}