Scientific discovery in the era of AI

- a celebration of *RODEM* -

Tobias Golling,
University of Geneva

Sinergia

RODEM

Robust Deep Density Models for High-Energy Physics and Solar Physics

A Sinergia research project funded by the Swiss National Science Foundation SNSF 2021-2024

RODEM kick-off workshop in Kandersteg 2021

Prof. FLEURET

Prof. VOLOSHYNOVSKIY

UniGE
Département d'Informatique
Project coordinator
francois.fleuret(at)unige.ch
@francois.fleuret
UNIVERSITÉ DE GENÈVE

UniGE
Département de Physique Nucléaire et Corpusculaire
tobias.golling(at)unige.ch
@TGolling

UNIVERSITÉ DE GENÈVE

UniGE
Computer Vision and Multimedia Laboratory
slava.voloshynovskiy(at)unige.ch
@voloshynovskiy
UNIVERSITÉ DE GENÈVE

FHNW
School of Engineering
Institute for Data Science
andre.csillaghy(at)fhnw.ch
@FHNW_astro

n|w Fachhochschule Nordwestschweiz

ATLAS EXPERIMENT

syn·er·gy | 'sinərjē |

UNIVERSITÉ DE GENÈVE
FACULTY OF SCIENCE

PIs


TG


Slava Voloshynovskiy


François Fleuret

PhD students


Tomke Schröer


Malte Algren


Jona Ackerschott


Matthew Leigh


Debajyoti Sengupta


Sam Klein


Stephen Mulligan


Guillaume Quétant


Mariia Drozdova


Ivan Oleksiyuk


Bálint Máté


Atul Kumar Sinha


Daniele Paliotta

Postdocs


Kinga Wozniak


Johnny Raine replacement starting soon

This could be you !

Master


Franck Rothen

Alumni


Lukas Ehrke


Knut Zoch


Manuel Guth


Matthias Schlaffer


Sebastian Pina-Otey

Full spectrum: physics - CS

FONDATION ERNEST BONINCHI

Alexander von Humboldt Stiftung/Foundation

European Commission

Swiss National Science Foundation

Sinergia

Swiss Government Excellence Scholarships
STUDY IN SWITZERLAND

Publications (with code)

2022

- FETA
- Dequantisation
- Flows for Flows
- ν-Flows
- Flowification
- CURTAINs
- SUPA

2021

- Funnels
- Turbo-Sim

2023

- Drapes
- TURBO
- EPiC-ly fast
- Flows for flows
- Interplay of ML based resonant anomalies
- PC-Droid
- OT Decorrelation
- ν^2-Flows
- CURTAINs Flows for Flows
- Flow away your differences
- Topographs
- PC-JeDi

2024

- SkyCURTAINs
- Cluster Scanning
- Masked particle modelling

Lots in the pipeline…

https://github.com/rodem-hep

INSTITUTE OF ADVANCED STUDY

@ Villa Boninchi,
Sep 25 – Oct 6, 2023

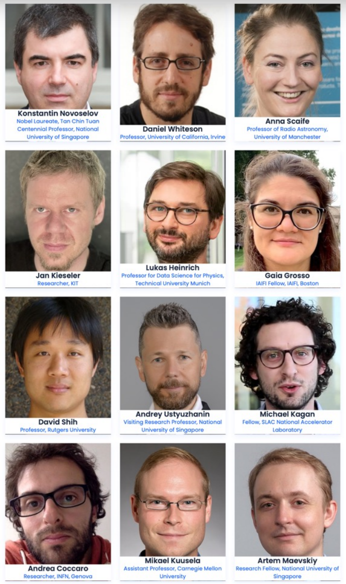

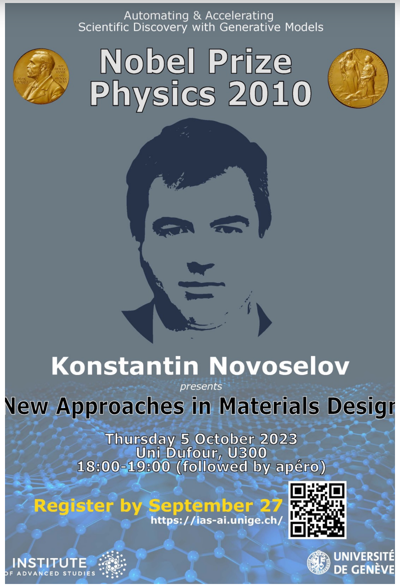

@ CSF, Ascona, Oct 29 – Nov 3, 2023

Many more events planned in the coming years…

Machine learning

- Statistical algorithms to **model** data & perform **tasks** without explicit instructions
- Thrives on **big data**
- **Generalizes** to unseen examples

The rise of AI/ML in science

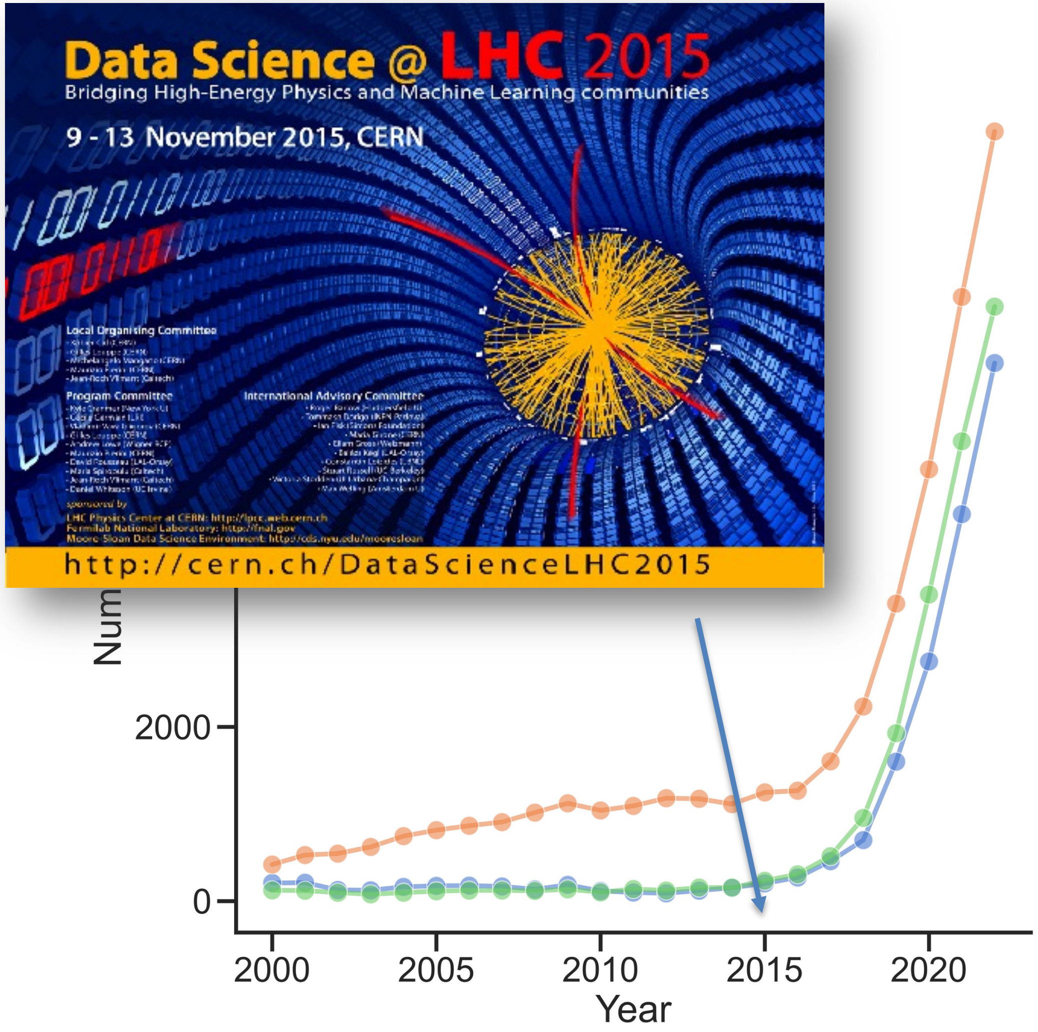

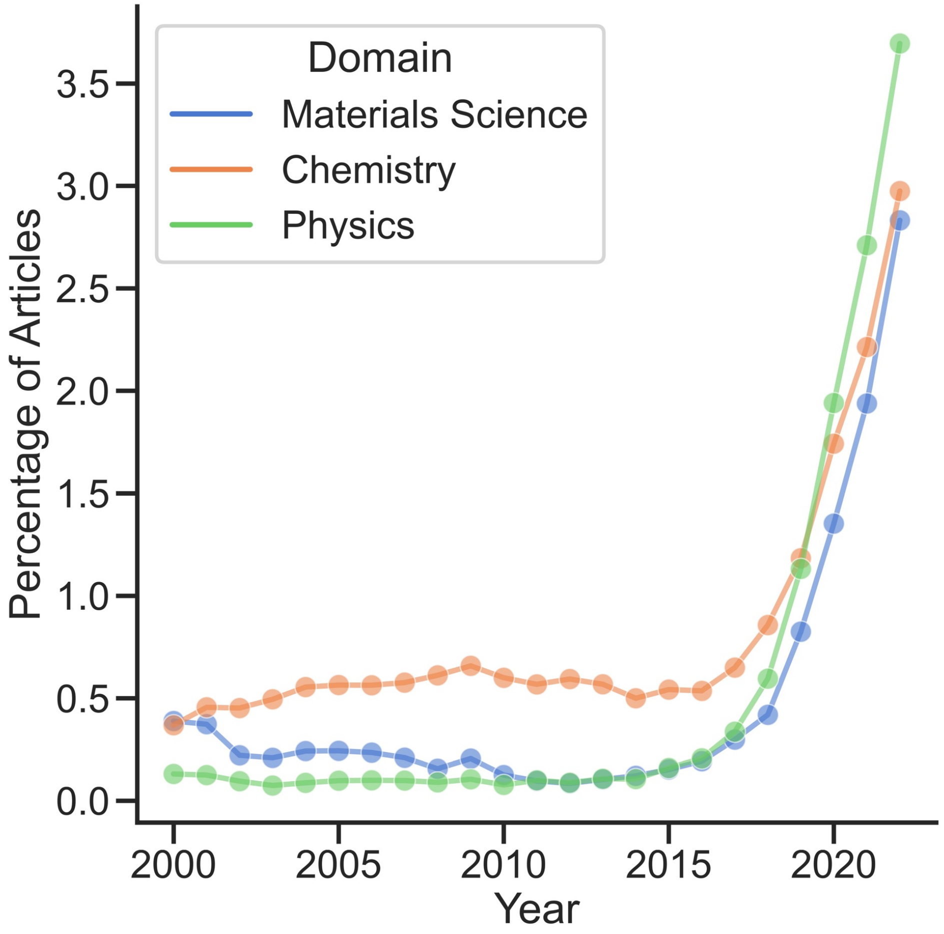

[Ben Blaiszik, “2021 AI/ML Publication Statistics and Charts”. Zenodo, Sep. 07, 2022. doi: 10.5281/zenodo.7057437.]

Part I – AI today

Science as usual – with an AI afterburner

State-of-the-art in ML@HEP

40k ML papers in hep-ex:

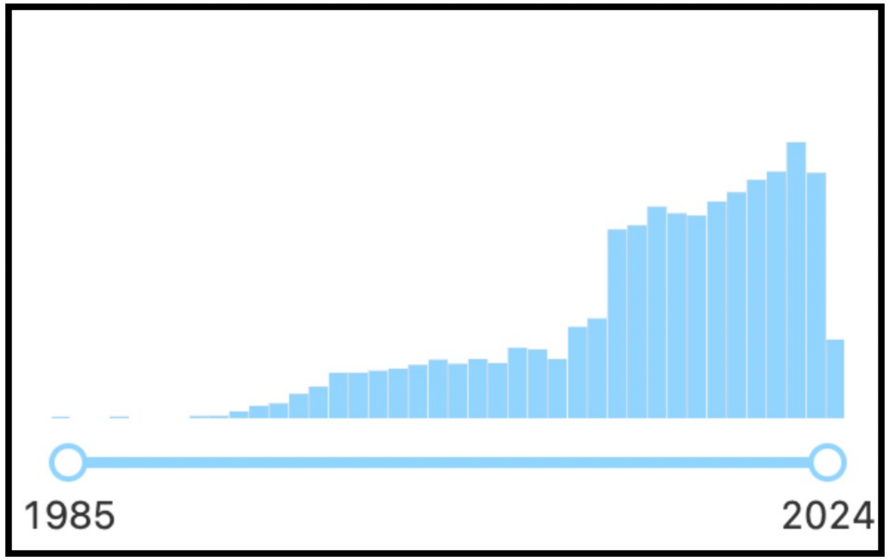

Very active ML@HEP community

Diverse *R&D* concept papers

+deployment in experiment
[90% of the work]

⇒ Time to get organized

AI / ML **everywhere** in our workflow

Optimal design

Classification

Search for unknown

Calibration

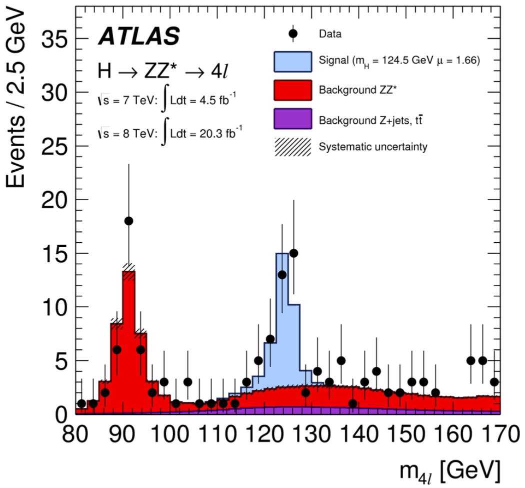

Classification without labels

Reconstruction

Fast simulators

In-situ background estimates

Decorrelate background

The *frontier* of classification

Top tagger comparison: R_{30} = BG rejection for 30% efficiency vs. #parameters

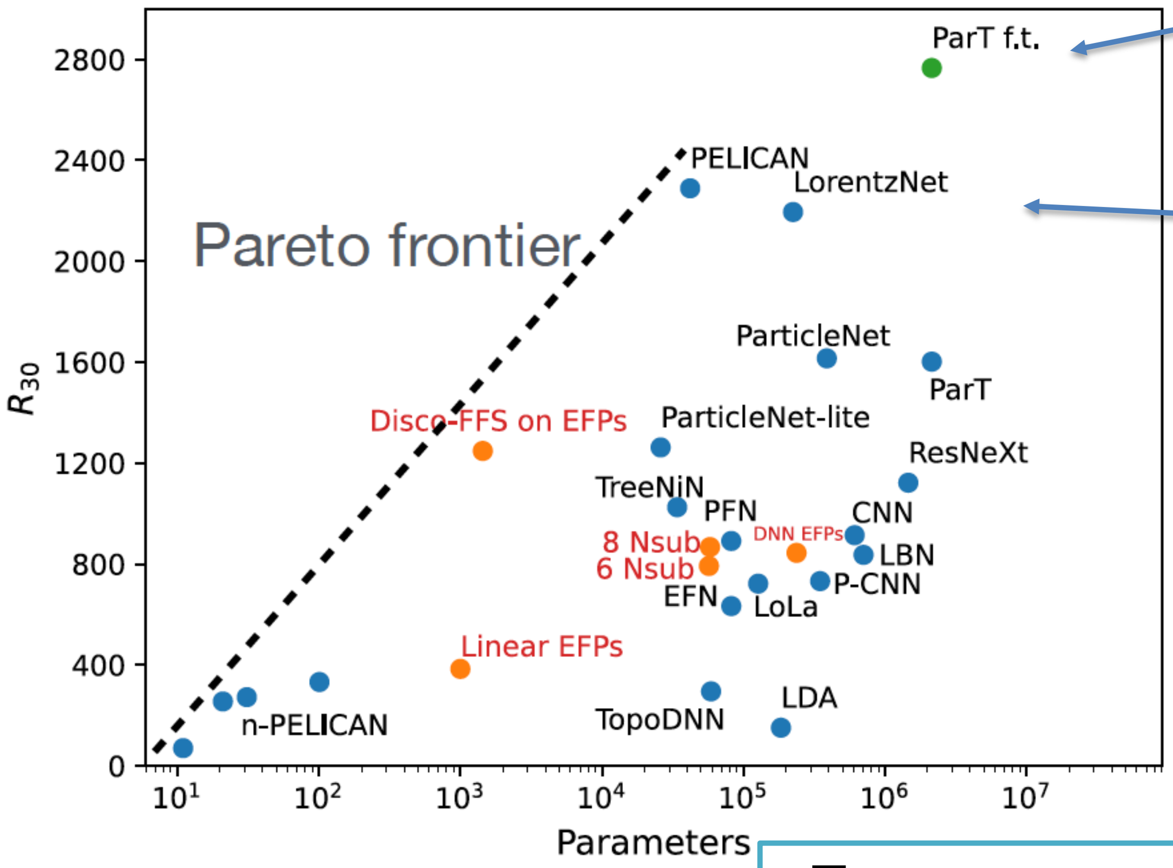

[G. Kasieczka, EuCAIF 2024]

Transformers rule the world

Inductive bias (Lorentz invariance, symmetries,…):
- More parameter efficient
- **BUT** less performant

The bitter lesson vs. heroic domain-specific modeling efforts

Still not all hope is lost: might be a sweet spot - we are far from *infinity*

Exact symmetries in latent space – hard to learn
Only approximate symmetries in data space

DL1, GN2, GN3, SALT

Enhance sensitivity

Plateau ?

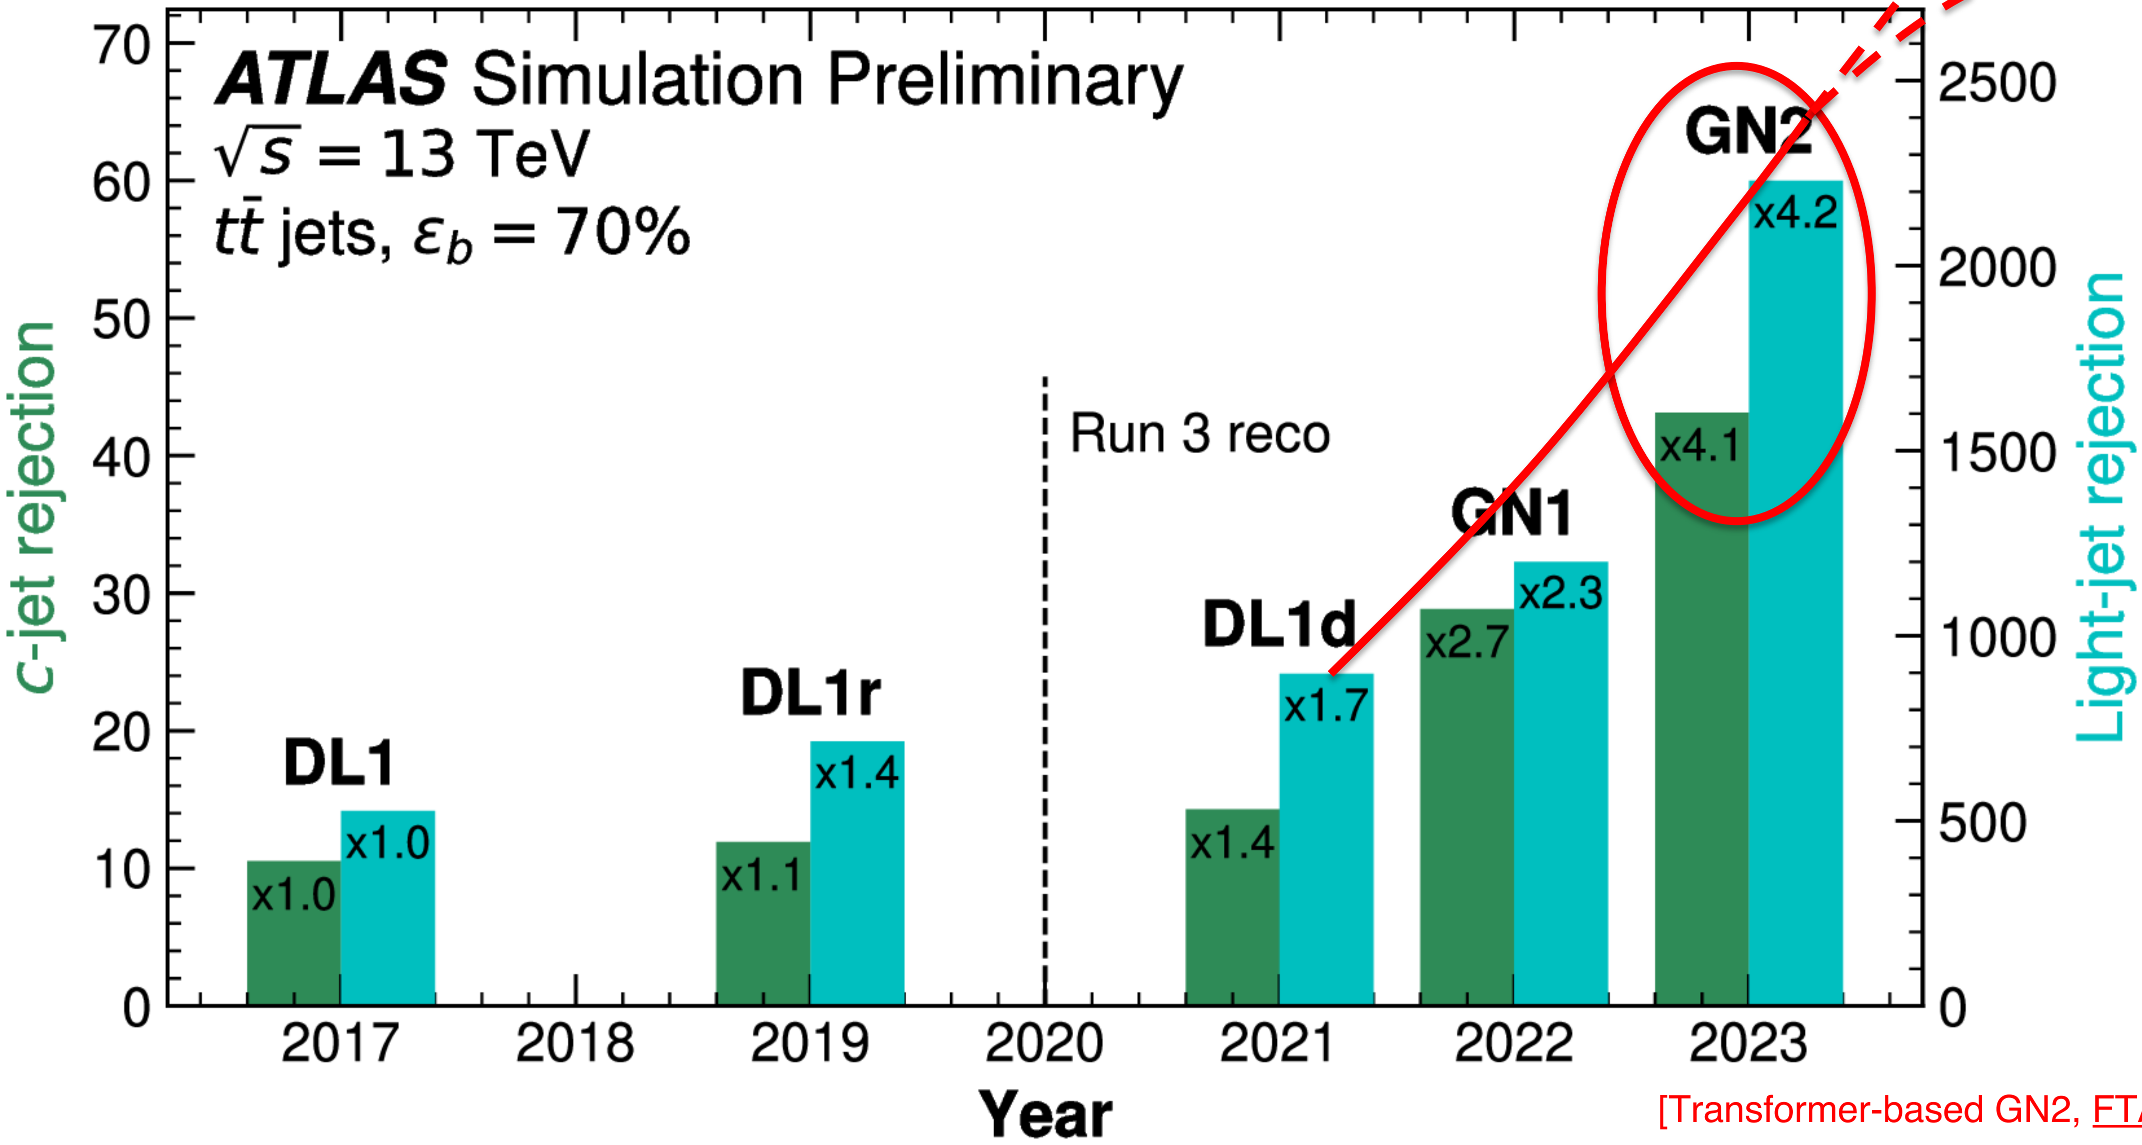

[Transformer-based GN2, FTAG-2023-01]

The ML toolbox: generative models

Fast surrogate model* which maps random numbers to structure

$$p_{\mathrm{model}} \approx p_{\mathrm{data}}$$

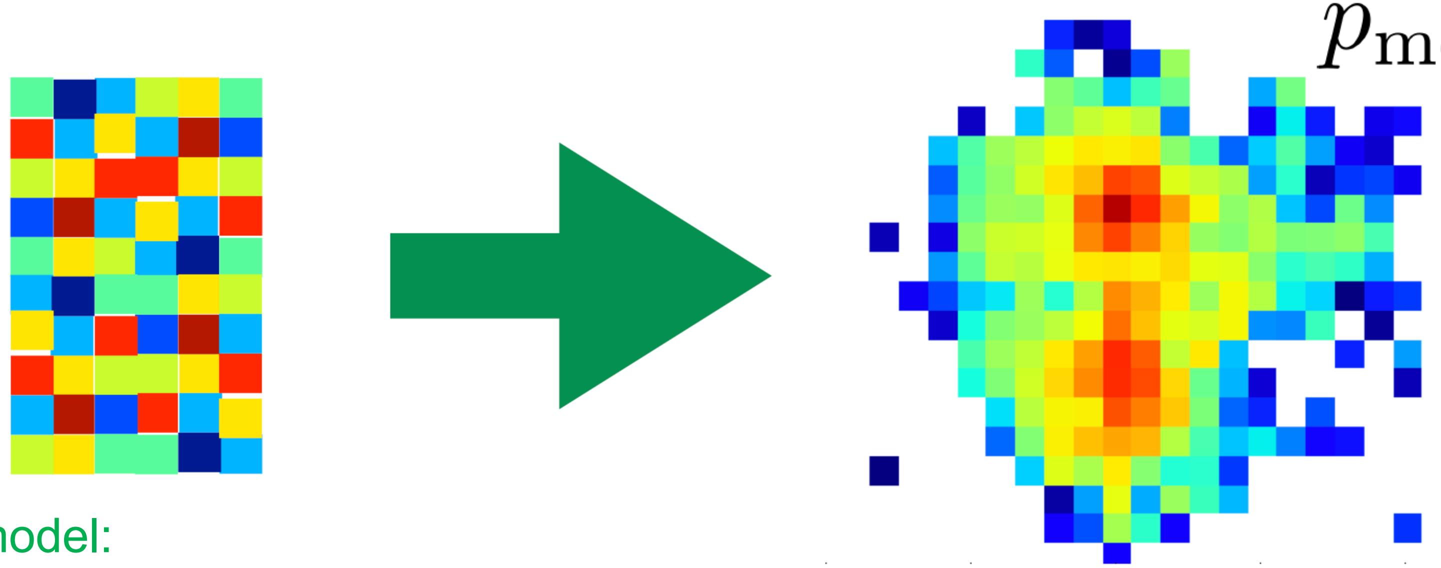

[image credit: Ben Nachman]

*Deep generative NN model:

- Variational Autoencoders (VAEs)
- Generative Adversarial Network (GANs)
- Normalizing Flows (NFs)
- Diffusion models

→ See my course 14P053 “Physics applications of AI”

Why generative models?

Density estimation & sampling

Outlier detection

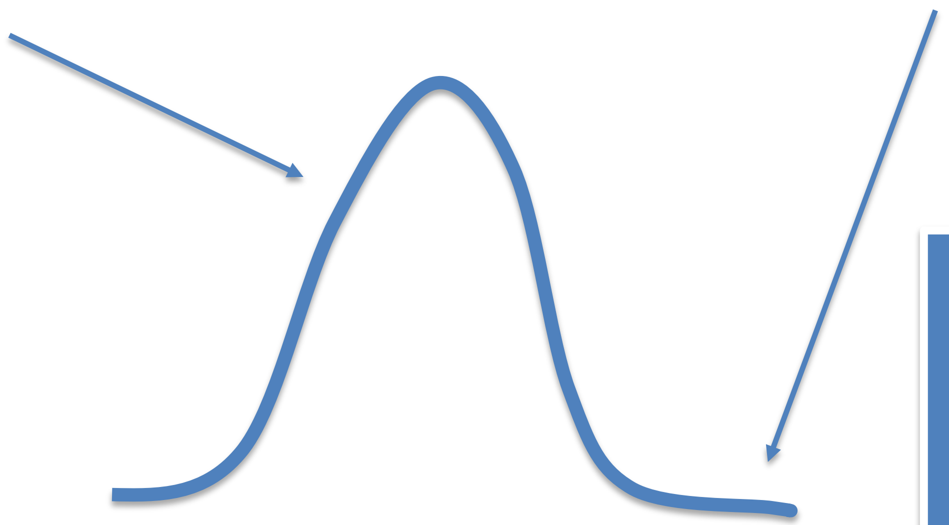

Representation learning
Understanding the data
…

Data compression

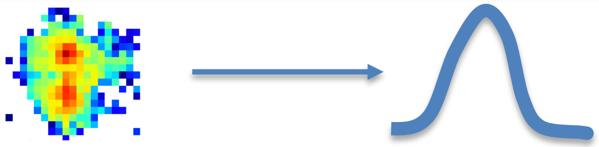

Density mapping

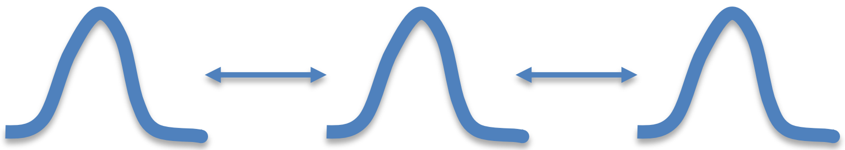

Evaluation of generative models

- Comparing high-dim joint distributions is **hard**
- No *best* GoF test
  - Need to know *relevant* alternative hypothesis

$$p_{\text{model}} \overset{?}{\approx} p_{\text{data}}$$

- *Pragmatic* tests to establish *trust*
  - *"Good enough* for task at hand"

[2211.10295]

Domain adaptation: calibrate synthetic to real data

1. *Scale factors*
<u>Issues</u>: support & dimensionality

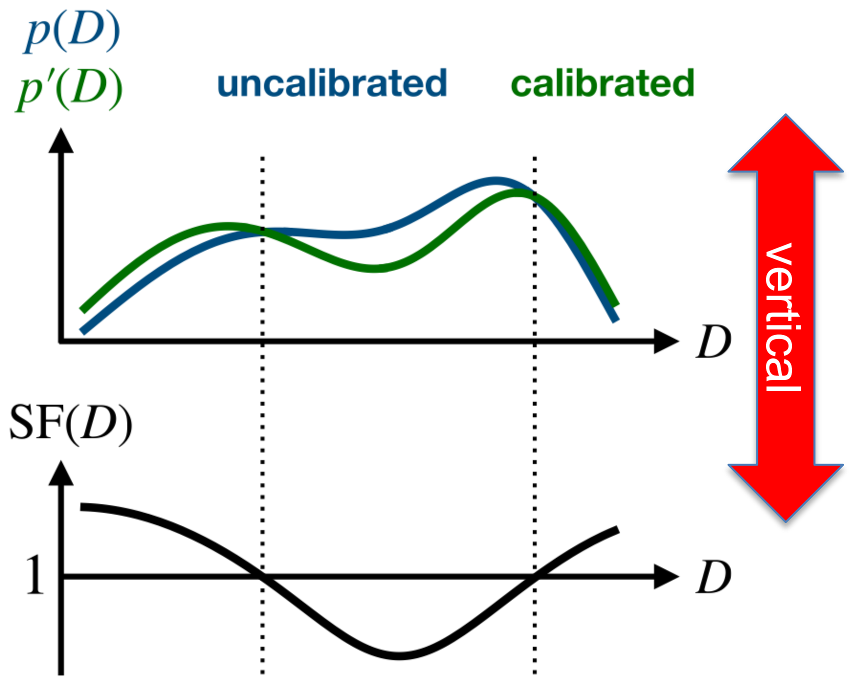

(b)

[innovation e.g. 1506.02169]

2. "Transport or flow your problems away"

ATLAS EXPERIMENT Jet flavor calibration

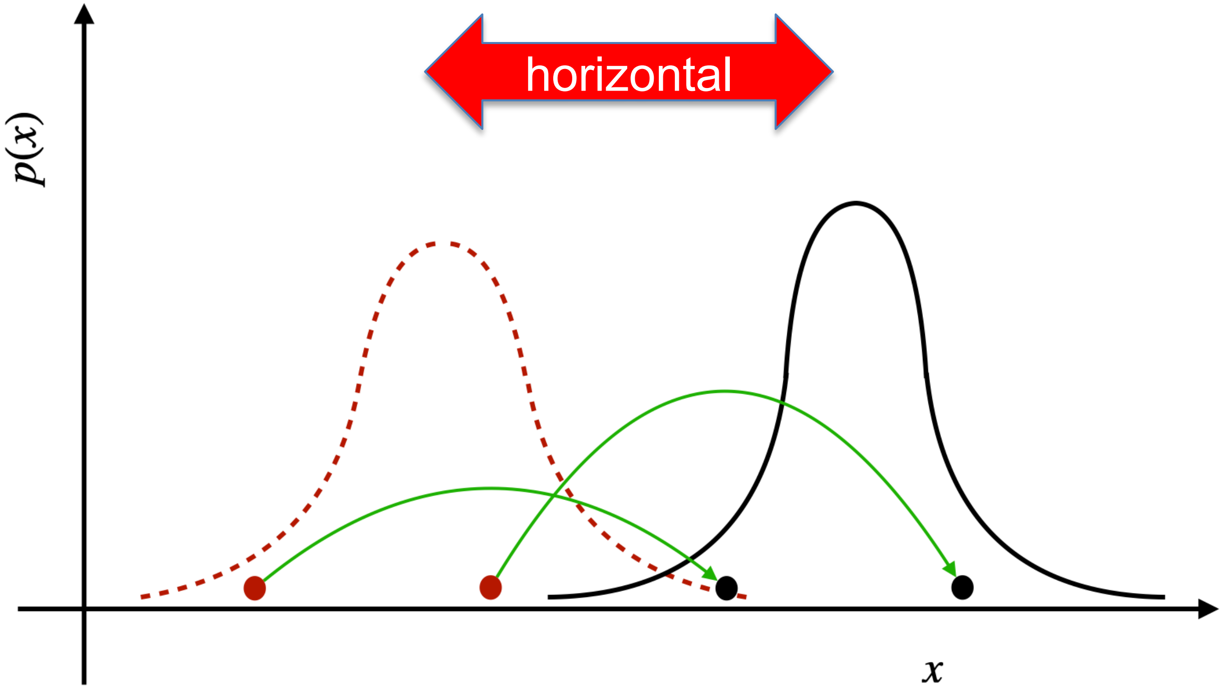

[2107.08648, 2304.14963]

Conditional neutrino regression with flows

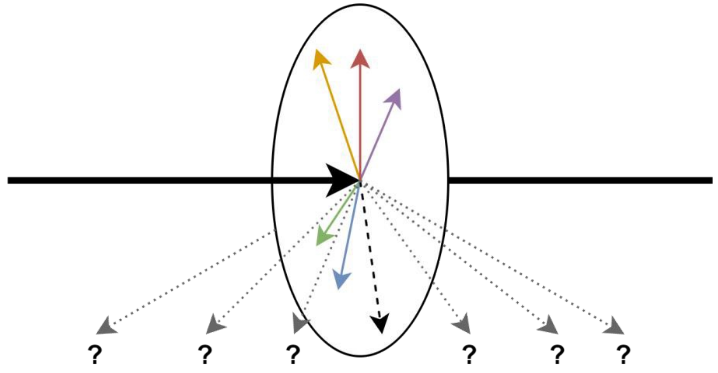

Improve over traditional method:

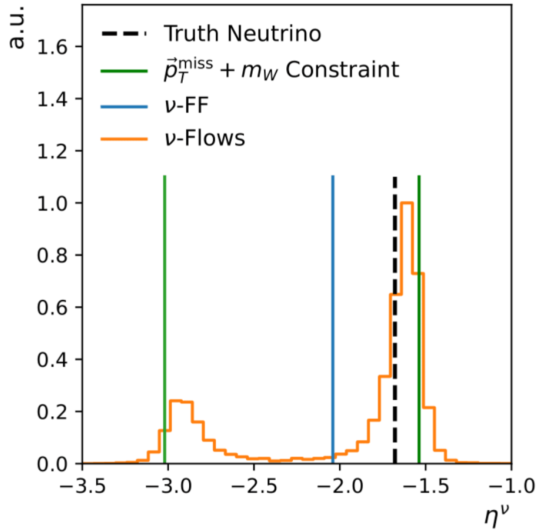

Transformer !

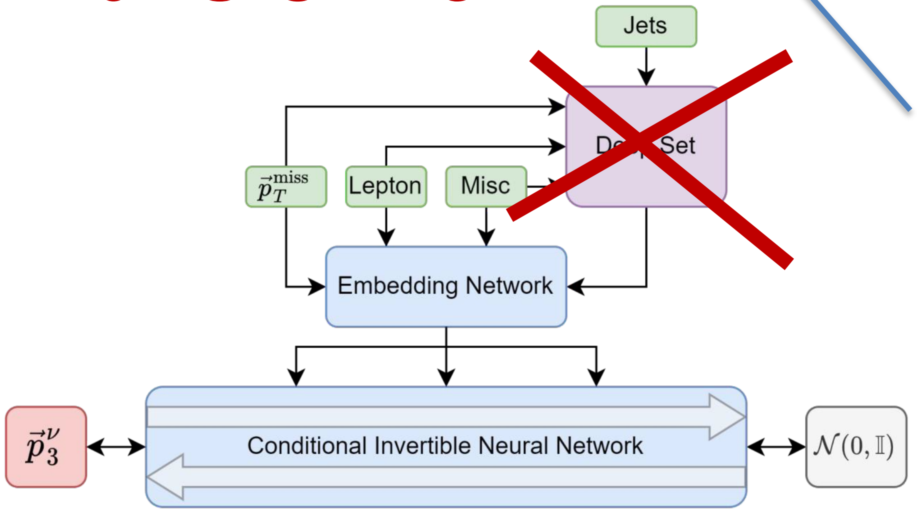

Learn conditional likelihood over neutrino 3-momenta assuming an underlying process (inductive bias)

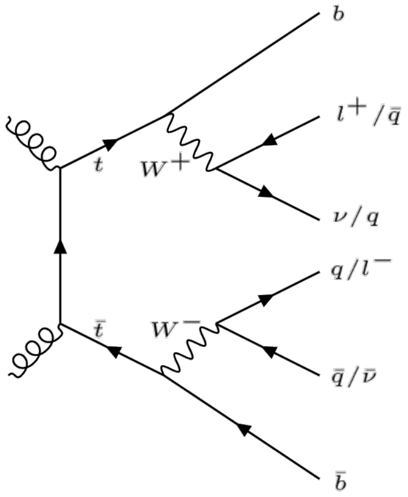

[2207.00664, 2-nu: 2307.02405 & @ event-level 2303.13937]

Simulate faces or… …detector images

[Karras et al., 2018]

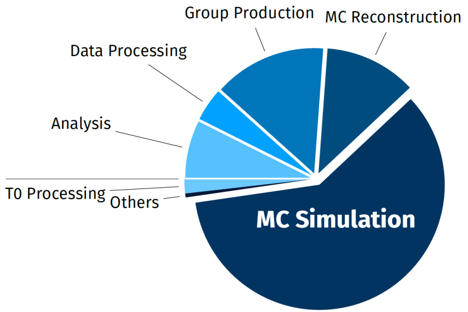

Fast surrogate simulator

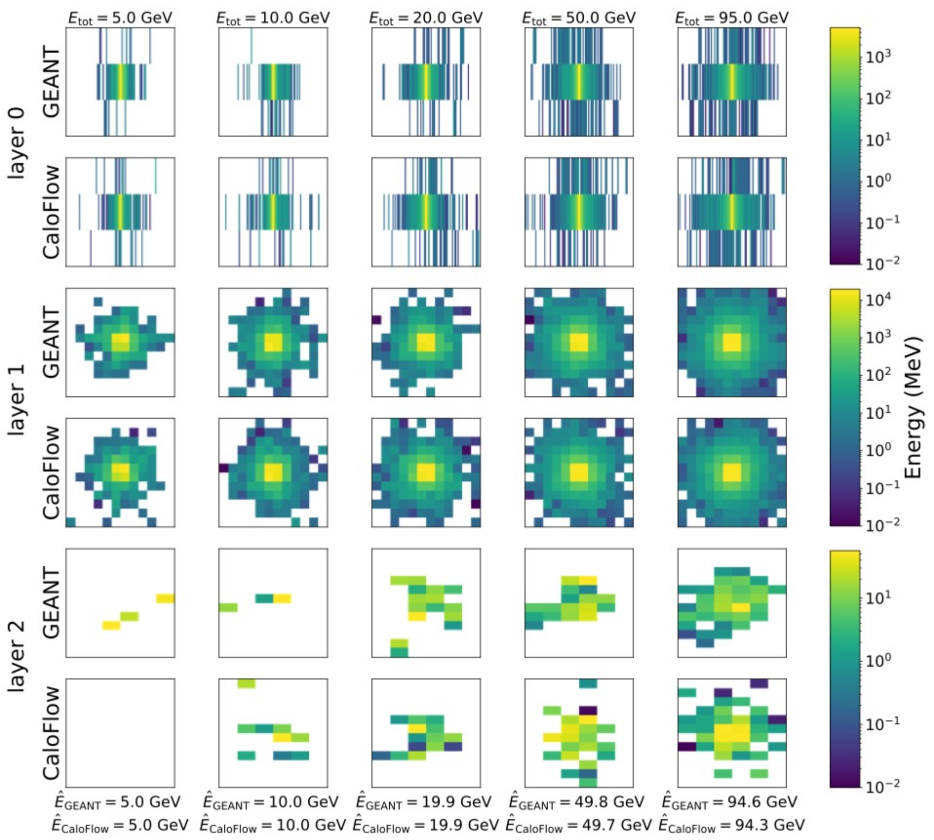

[CaloFlow,…]

ATLAS EXPERIMENT

[2210.06204, SUPA]

Image $\rightarrow$ Point cloud

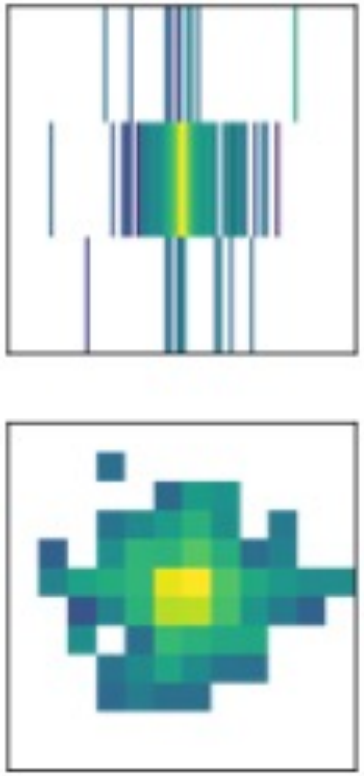

Addresses sparsity issue

Promotes portable solutions:
decouples modeling from detector geometry

Point cloud diffusion

[source]

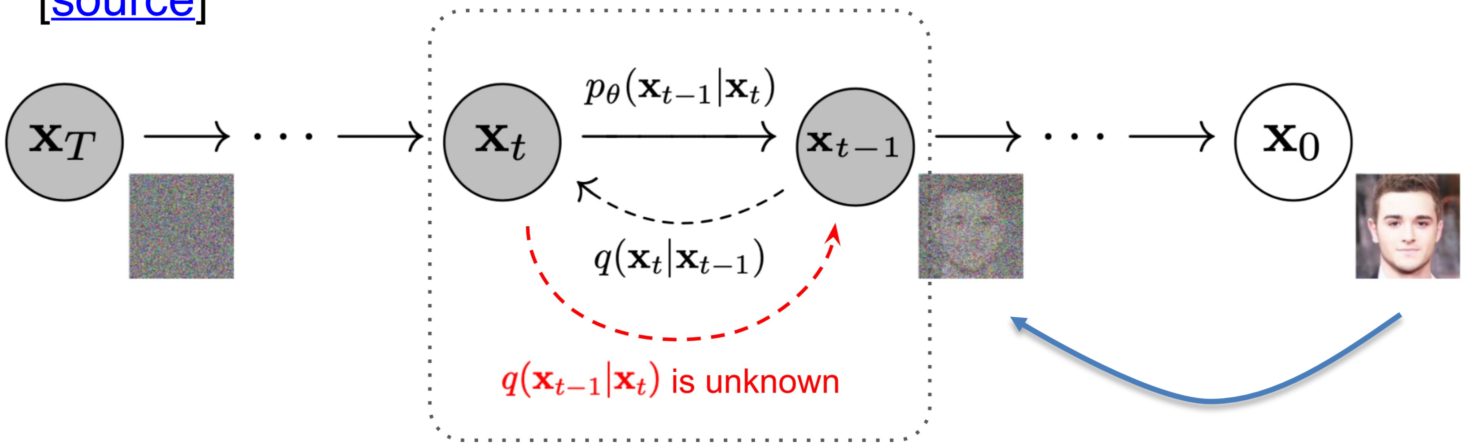

"Learn the noise"

Gradually add Gaussian noise

Images → Point cloud

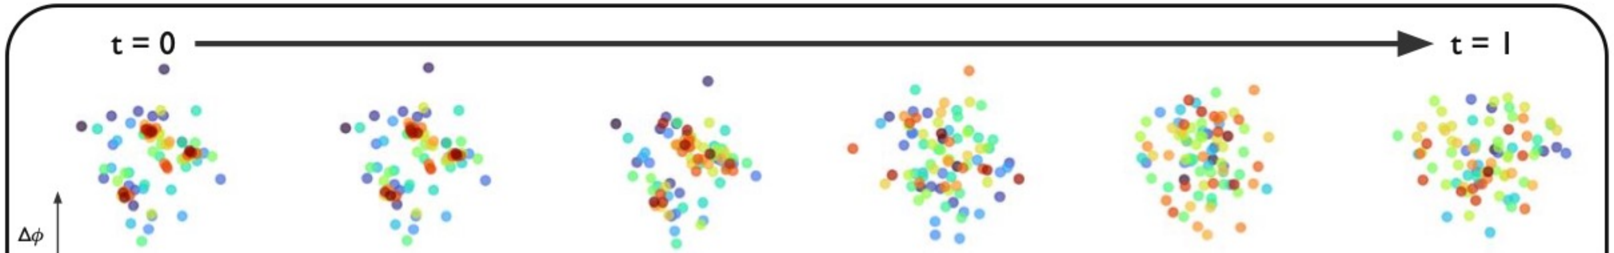

[JetNet data]

Transformer Encoder (TE) Block

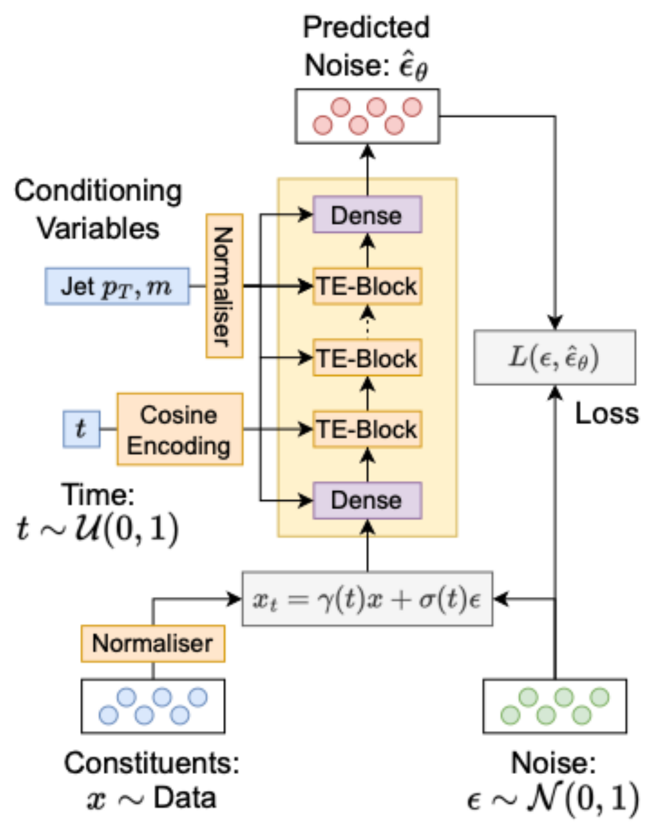

[2303.05376 & faster: 2307.06836 & 2310.00049]

Issue: background sculpting for bump hunting

Signal

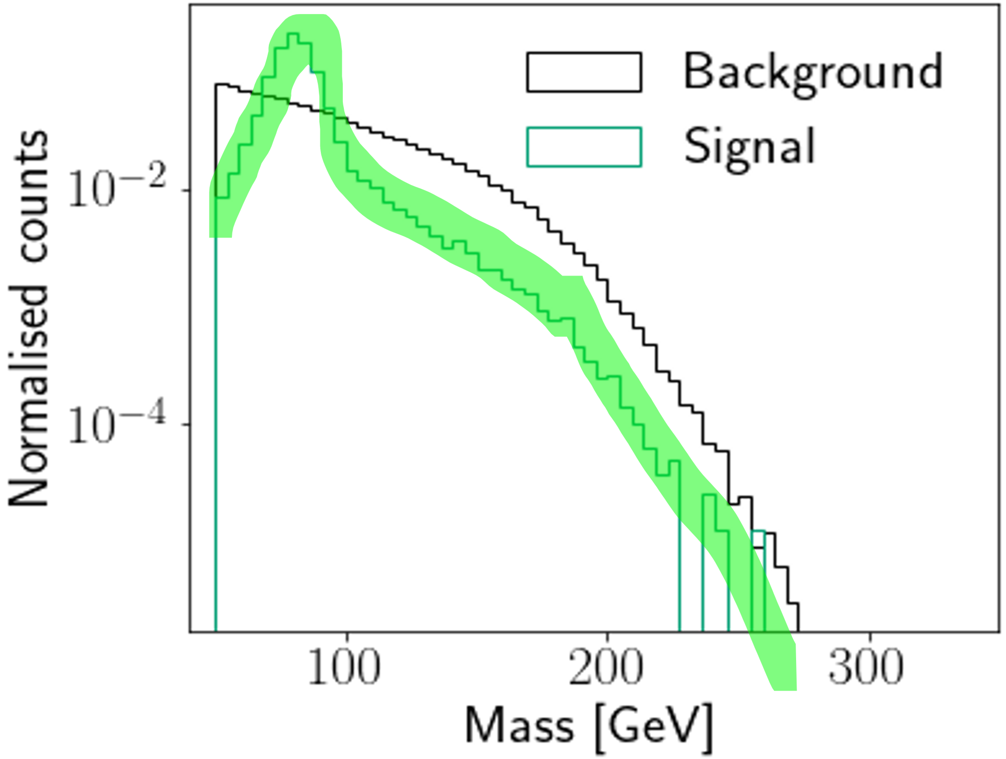

Background after cut on classifier

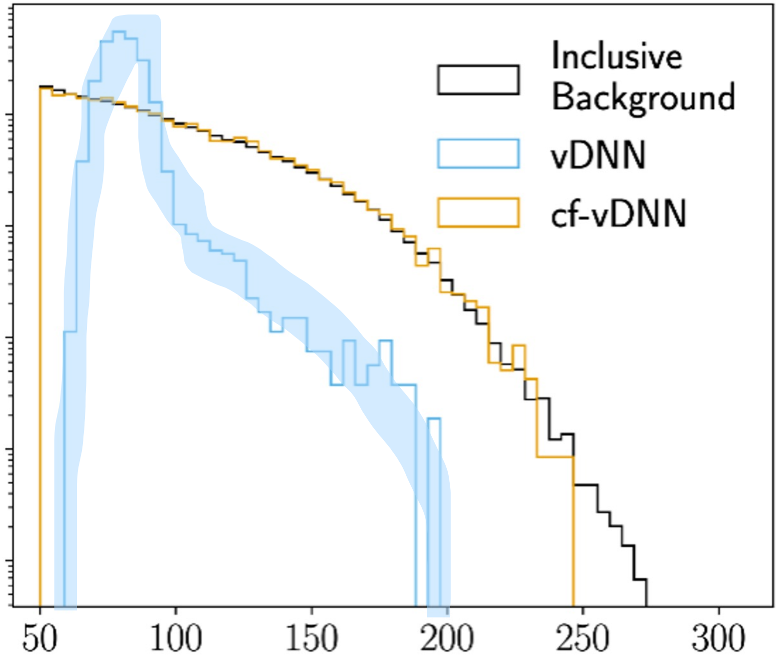

Goal: decorrelate background from mass

Decorrelation with normalizing flows

Flow = map between distributions

<u>Invertible</u>:

no change in separation power

Can be made conditional

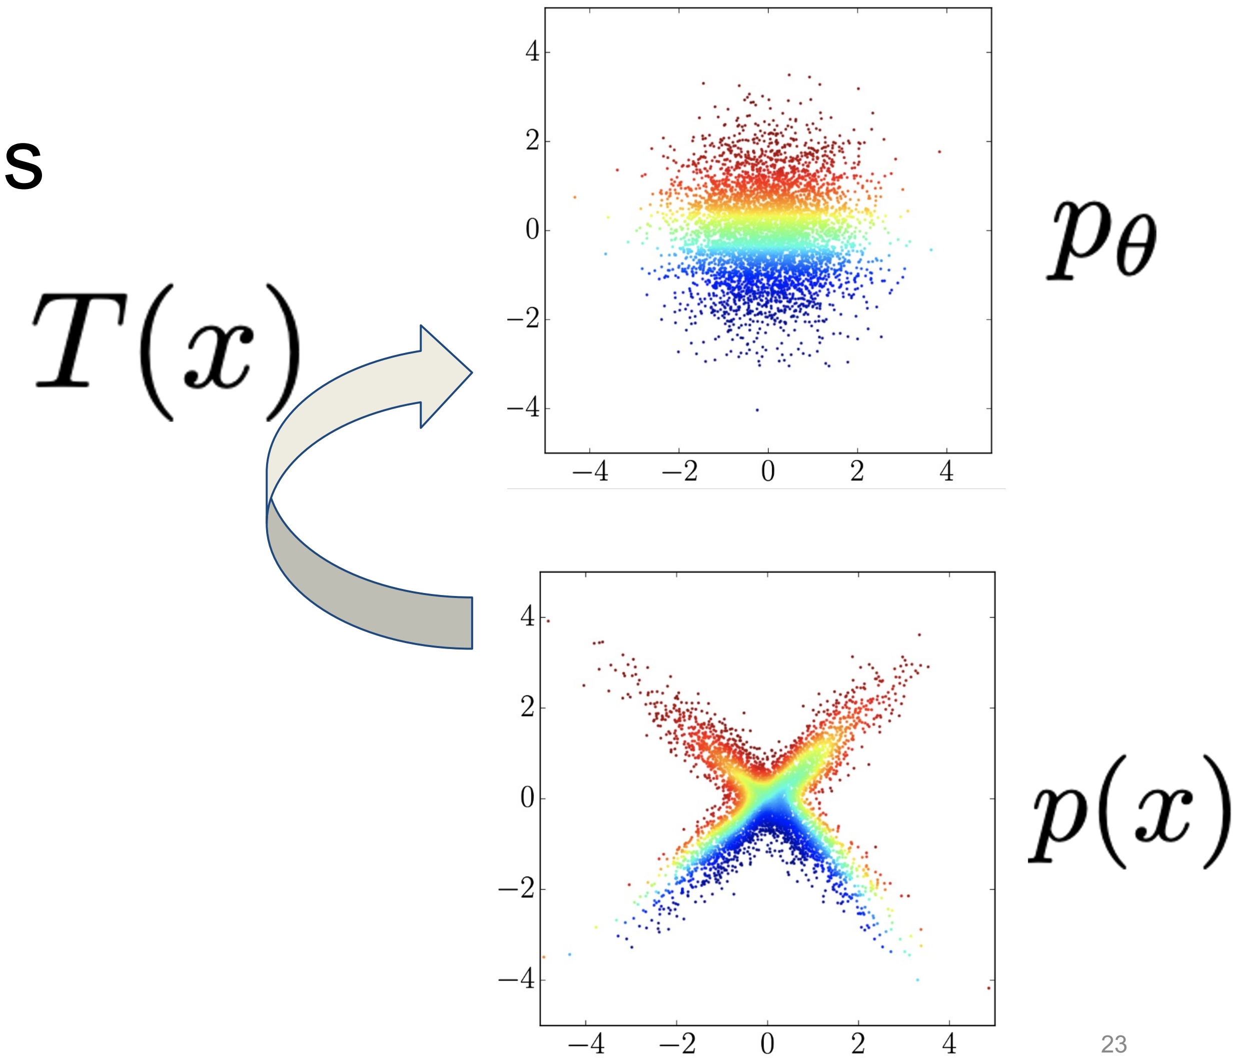

Train a flow to learn p(vDNN | m)

Decorrelated background

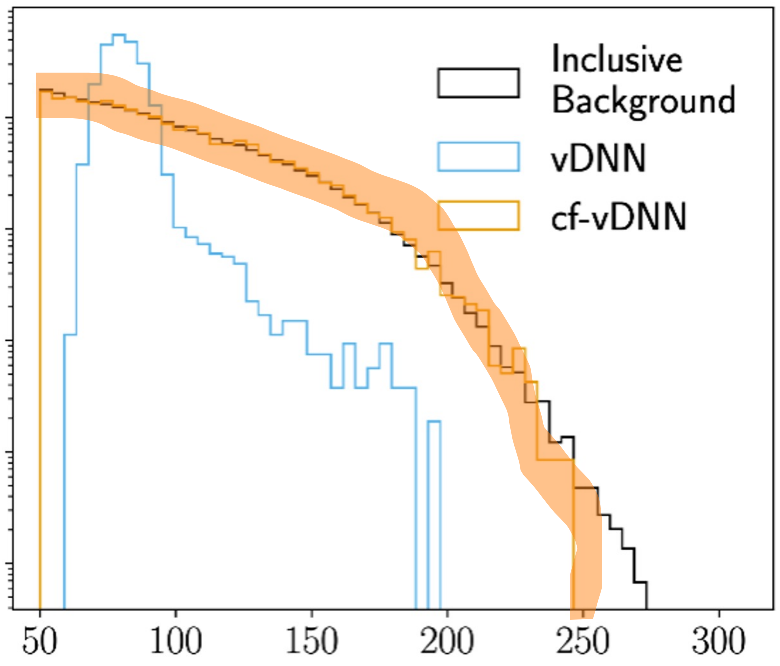

$T(x)$

Background:

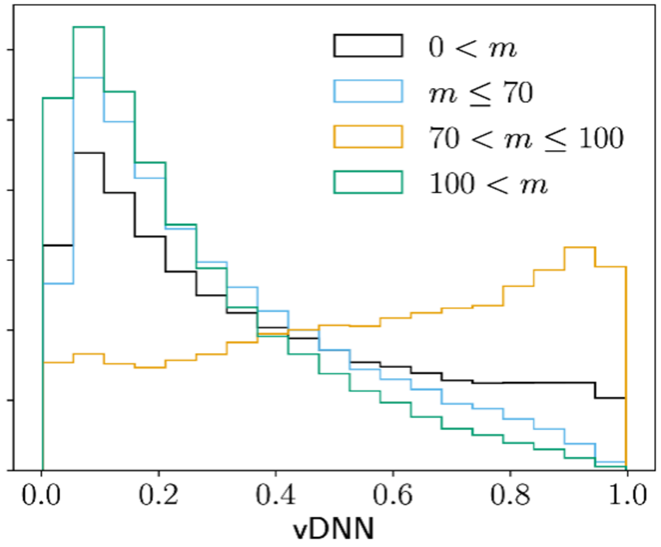

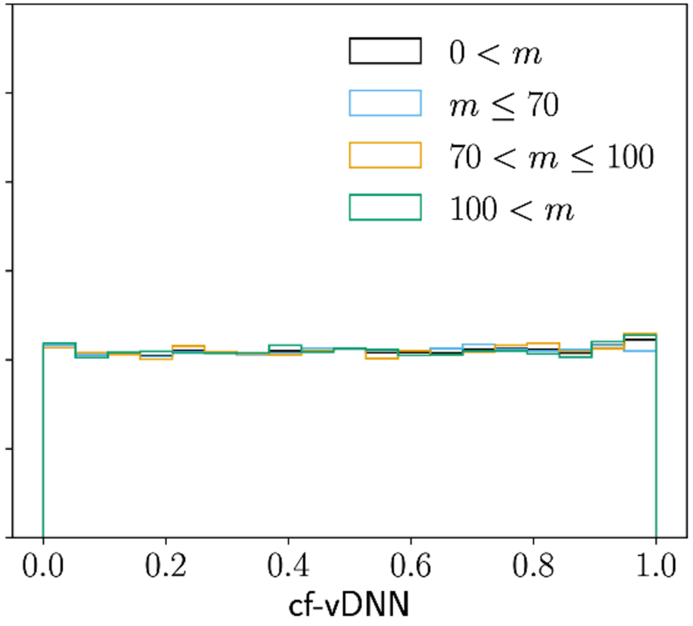

ATLAS EXPERIMENT

Flavor tagging mass decor.

[2211.02486, 2307.05187]

Learning high-D background templates*

Learn from simulation	Learn from data sidebands (SB)
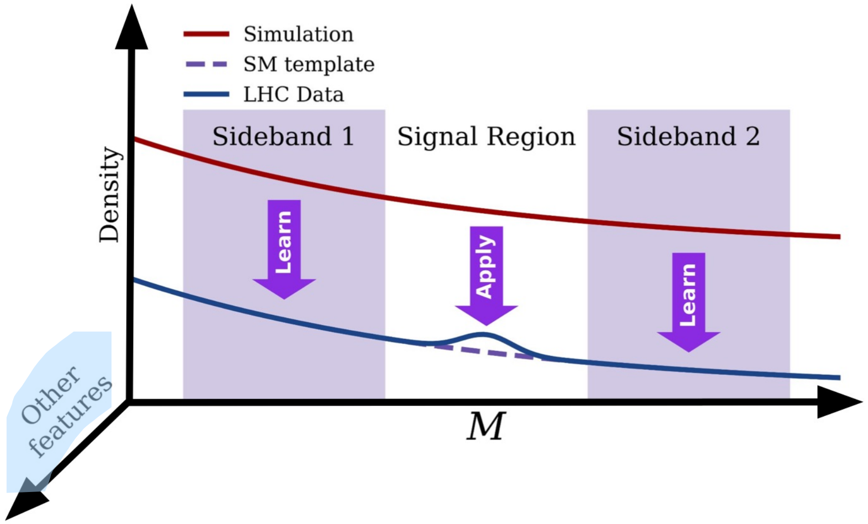	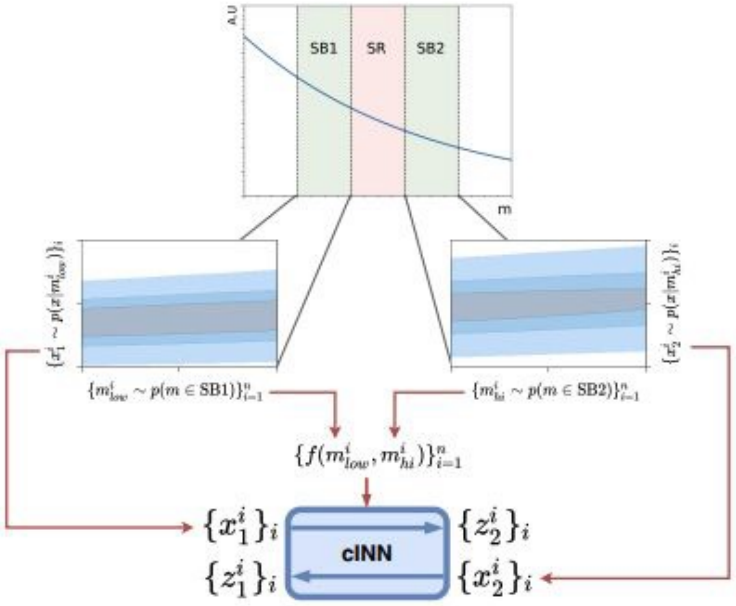

2212.11285 & comparison: 2307.11157

2203.09470 & 2211.02487 & 2305.04646 & 2309.06472 & using diffusion: 2312.10130 & applied to DM search in stellar streams: 2405.12131

ATLAS EXPERIMENT

Search

[*Fidelity of simulation alone insufficient]

Classification without labeling (CWoLa)

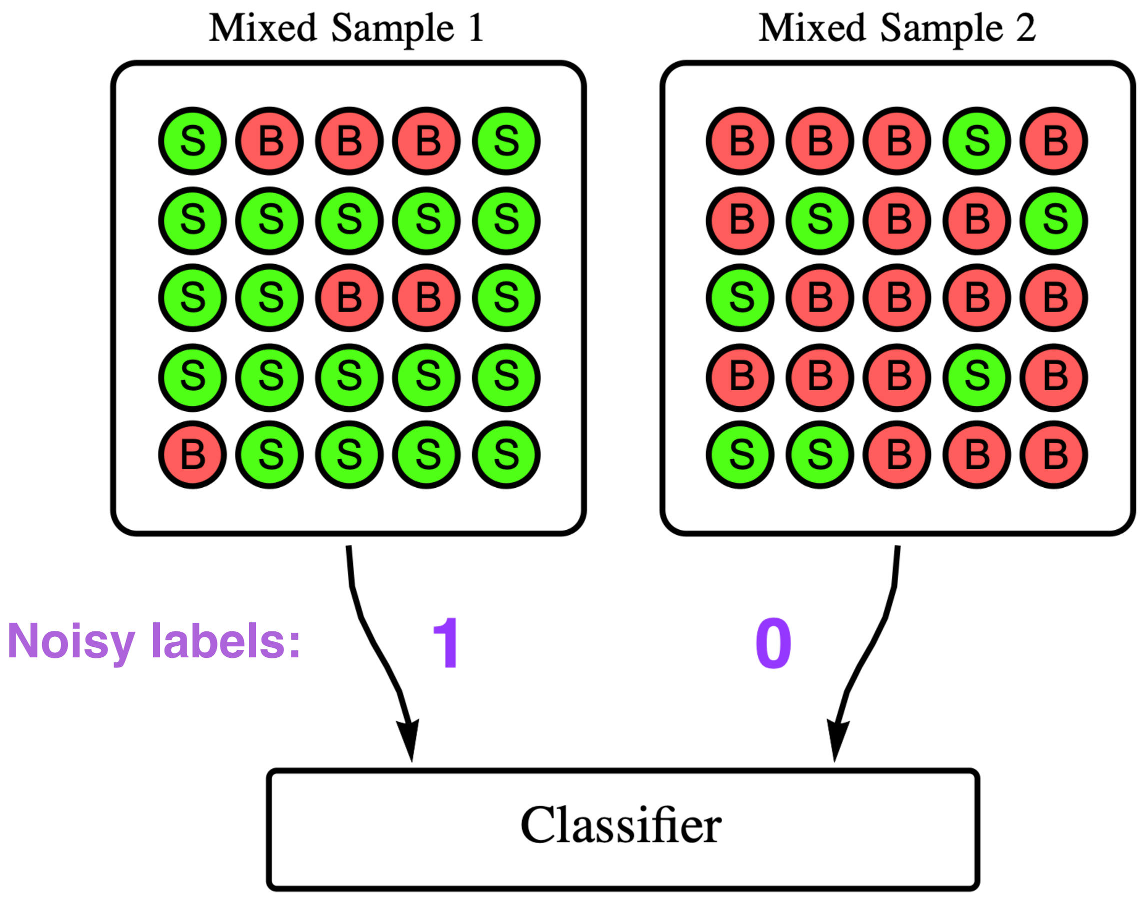

Maximize sensitivity to signal

Abandon notion of *event label*

Noisy labels to be S or B

Bump hunt [1902.02634]
ATLAS analysis [2005.02983]

Beyond resonances
e.g. symmetries [2203.07529]

ATLAS EXPERIMENT Search

Part II – AI tomorrow

Transformative science: automate & accelerate

Speculative, provocative, exploratory,…

Lots of open-ended questions

Humanity at the brink

Energy, climate, SDGs,…

Human history is a story of enabling technology

The AI box is opened
– *obligation* to see it through

Science = AI demonstrator

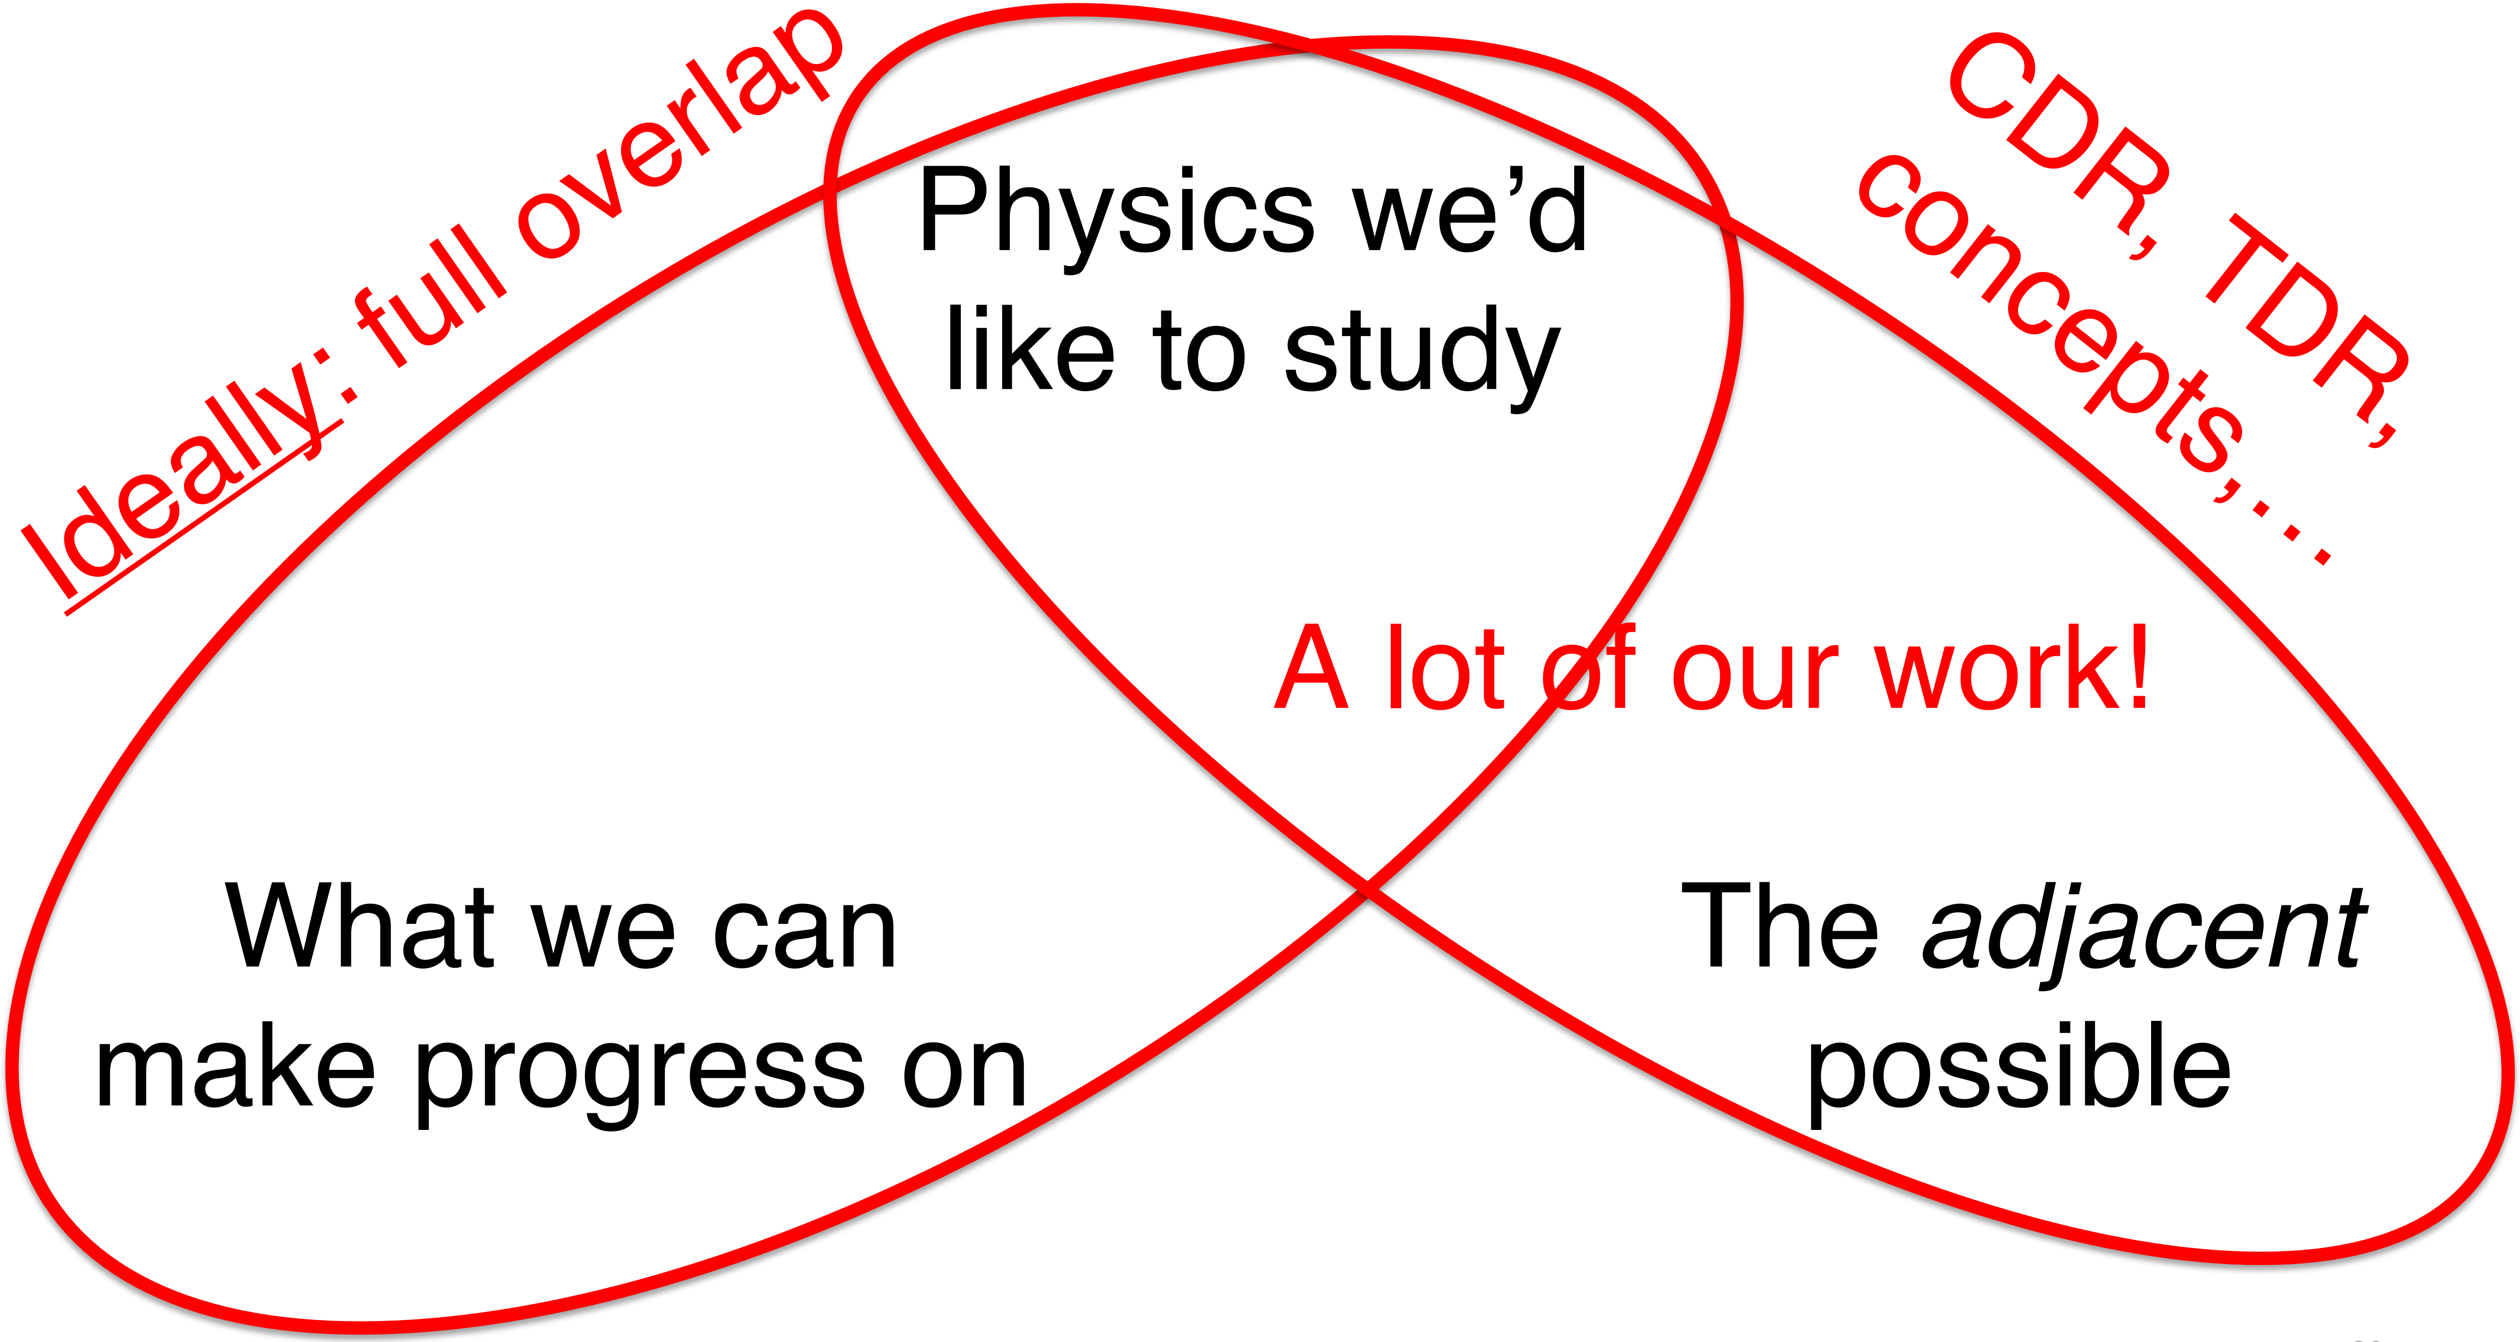
Ideally: full overlap
Physics we'd like to study
CDR, TDR, concepts,…
A lot of our work!
What we can make progress on
The *adjacent* possible

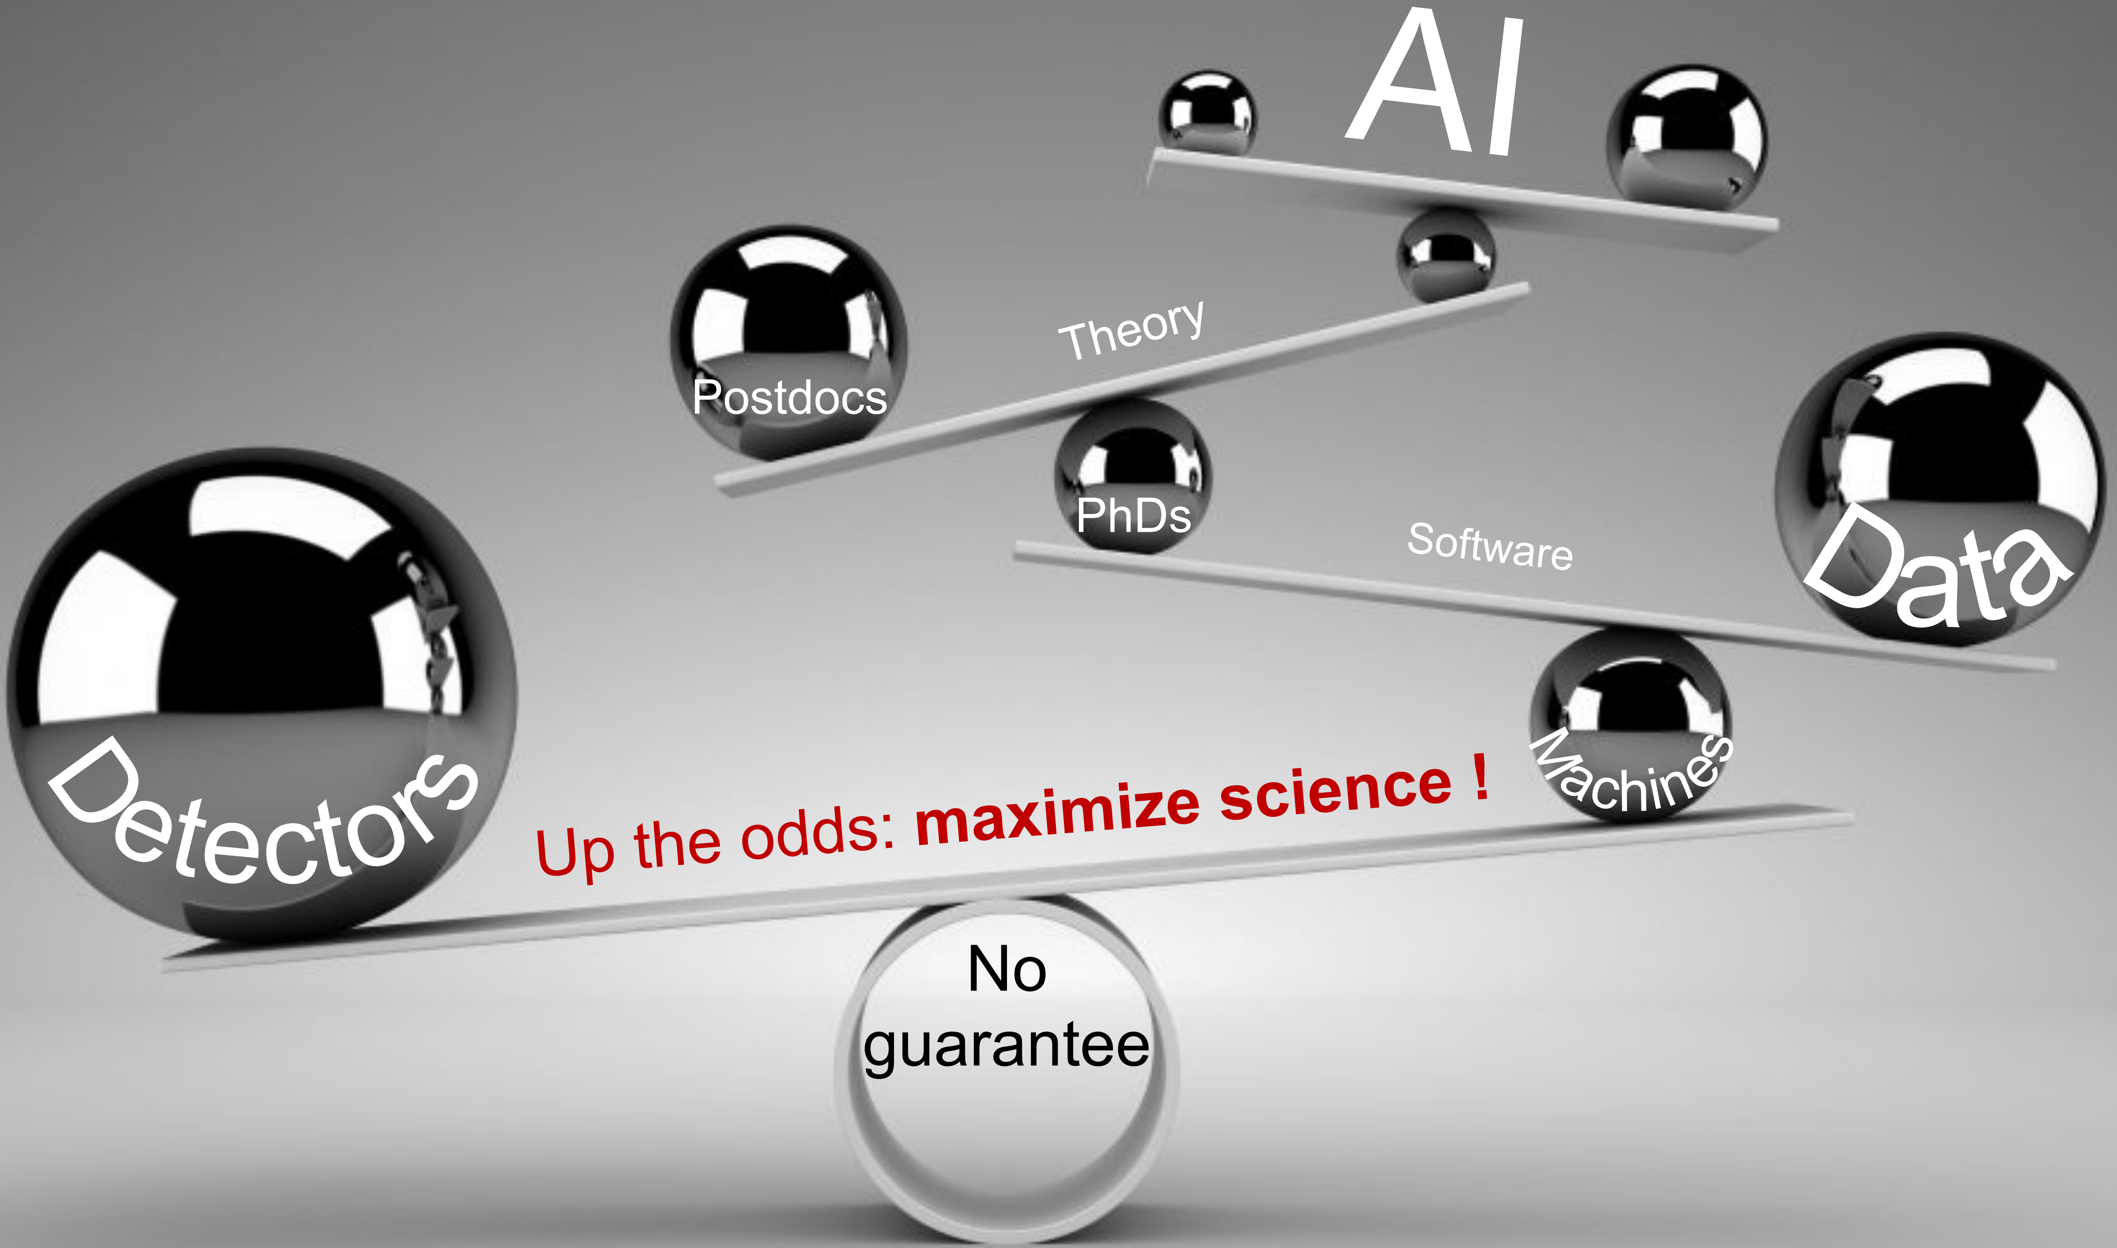

Research is exploration

AI to up the odds !

Known: SM

Unknown: BSM

$$\frac{knowledge\ \ gain}{resources\ \ used}$$

Human

Compute

Cost

Time

…

“New directions in science are launched by new tools much more often than by new concepts.”

- Freeman Dyson

"If your life's work can be accomplished in your lifetime, you're not thinking big enough."

— Wes Jackson

"There is no power for change greater than a community discovering what it cares about."

— Margaret Wheatley

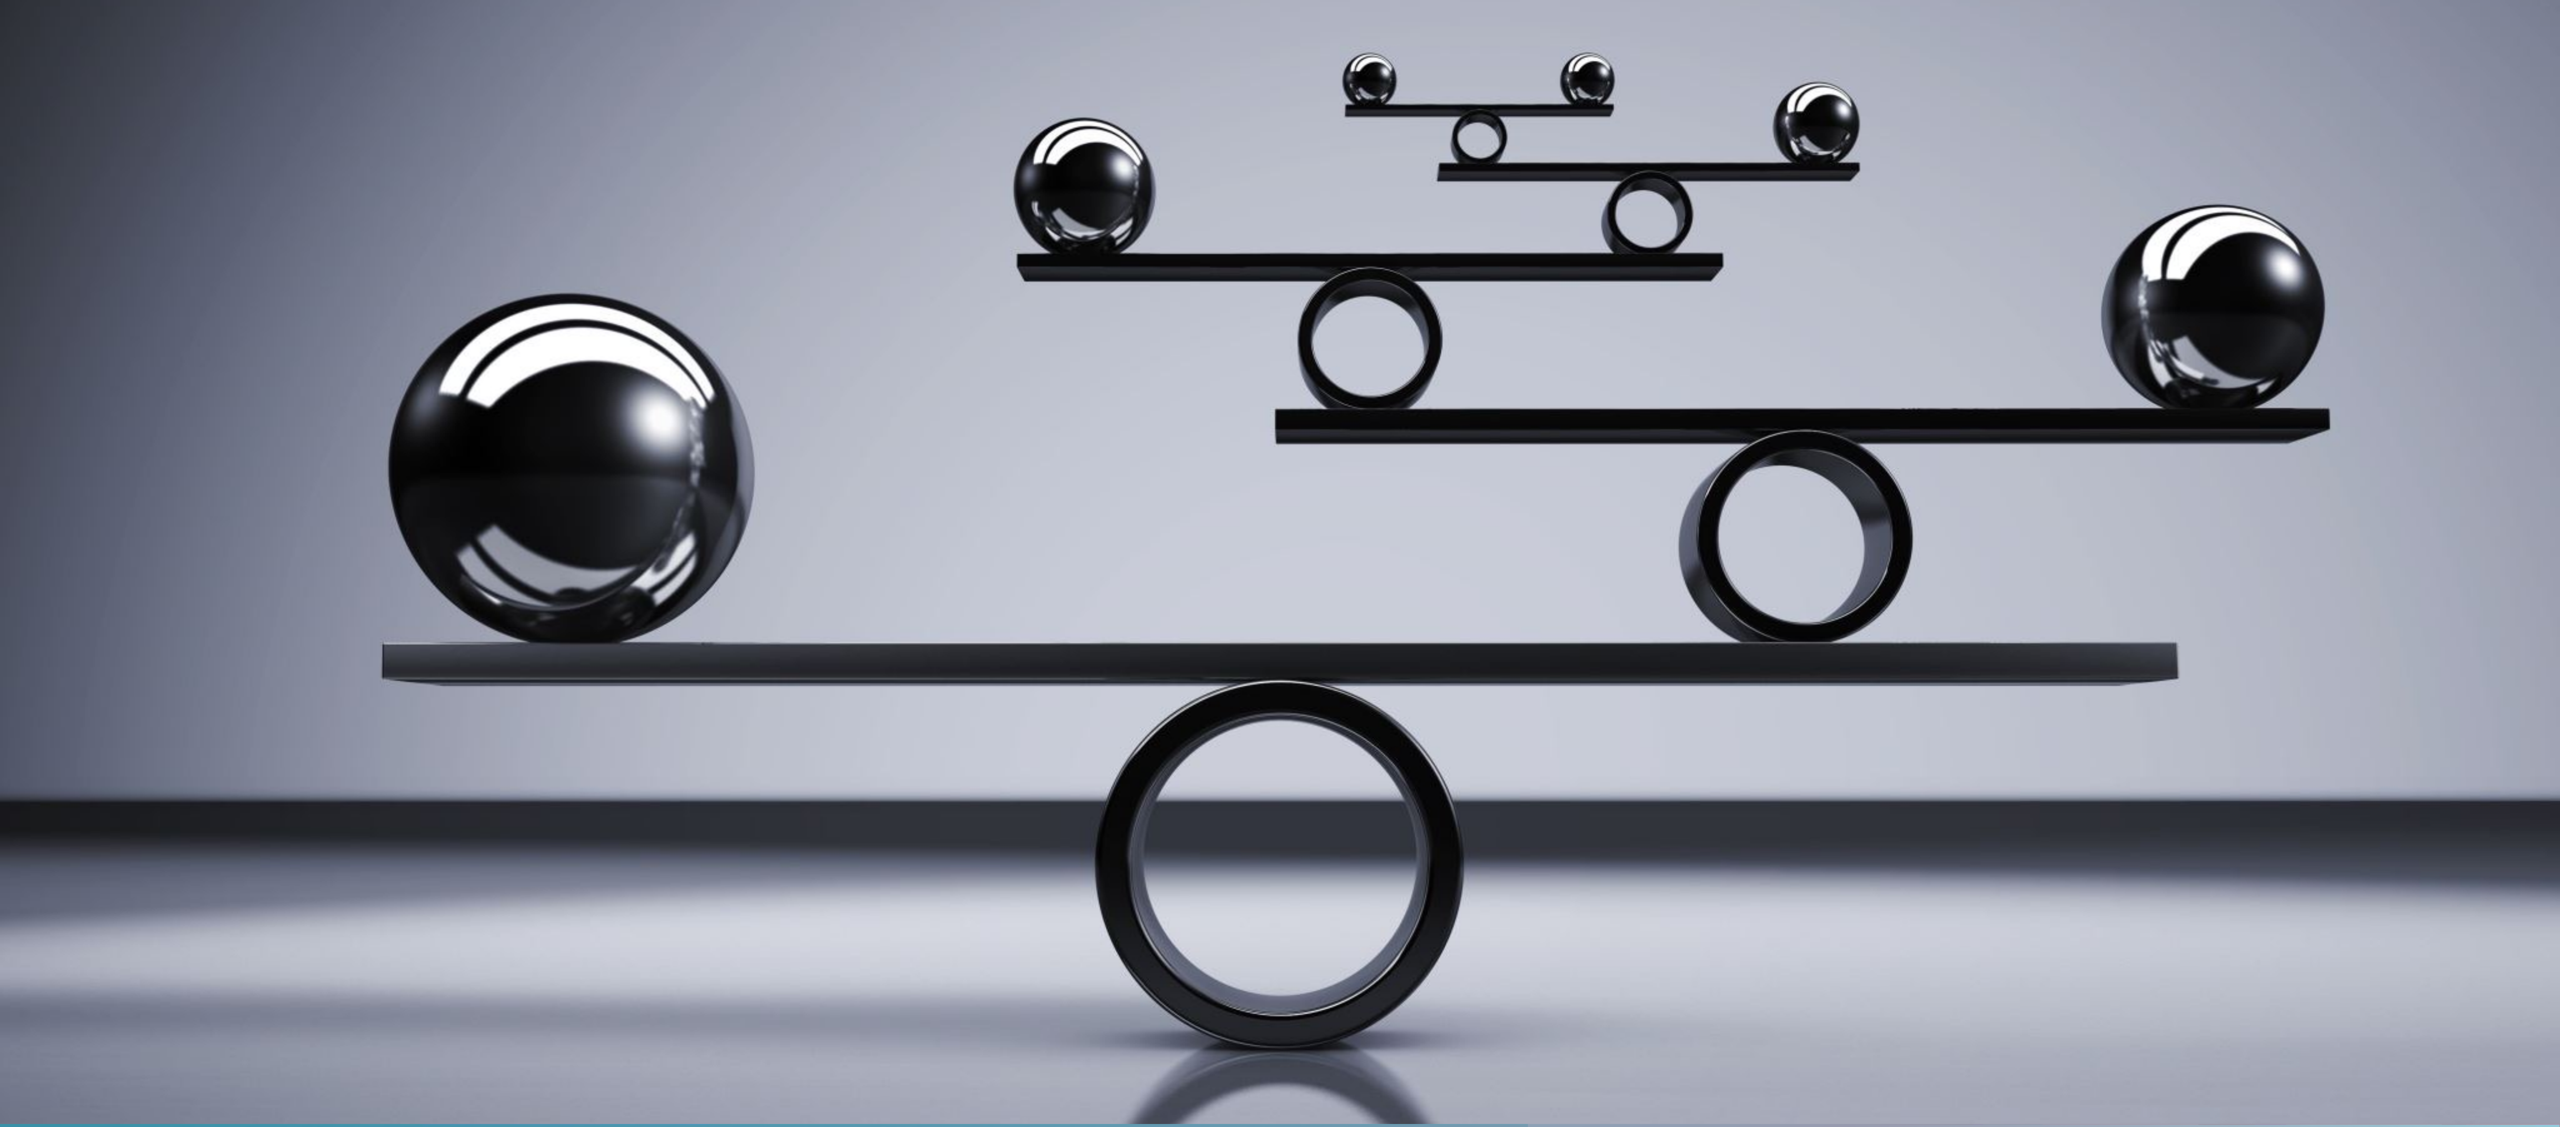

Three visions

Vision Nº 1

Foundation models

or

Legacy of ChatGPT

or

Grand ideas too beautiful to be missed

The *essence* of science…

Prediction machine

Finding new regularities

Learning saves computational resources

Reduce dimensionality of problem

<u>Example</u>: Kepler

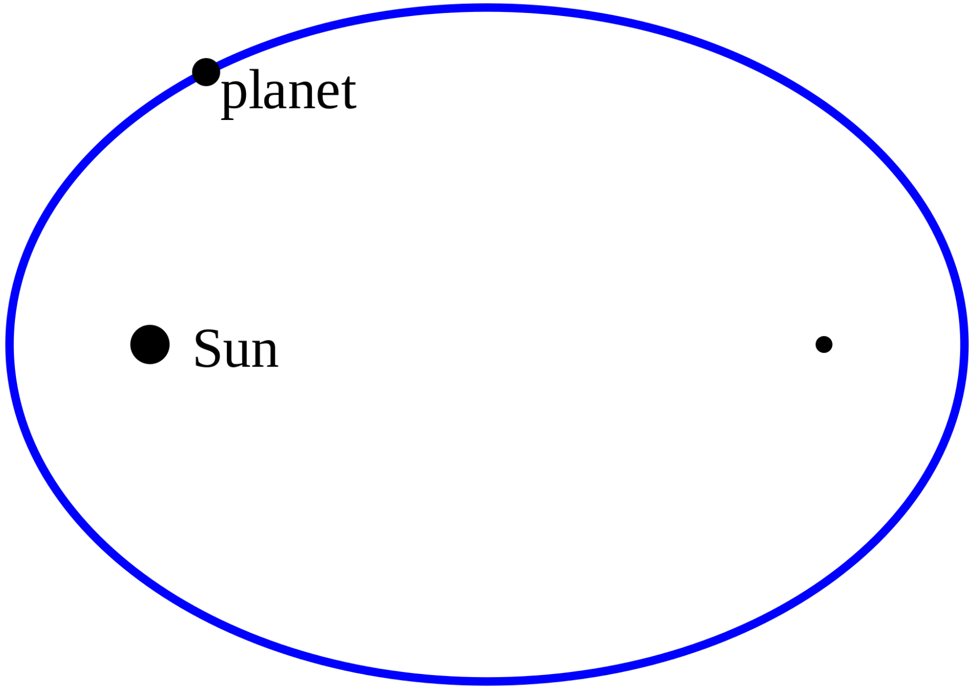

= the *essence* of ML

Scientists **model** the world

[Leo Breiman 2001 on statistical modeling: the two cultures]

X → Nature → Y

Data → Model → Evaluation: Generalization

Human ingenuity [e.g. SM]
↕
Data-learned [foundation model]

Recap: what is a generative model?

An **implicit model**
that describes
how data was generated

[There is no model-less model]

[ChatGPT = implicit model of human language]

[DALL·E = implicit model of natural images]

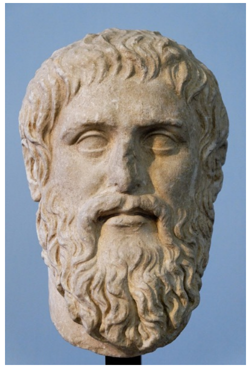

Models with *meaningful* **latent** representations

Plato: myth of the cave

<u>The quest of science</u>:

Learn true underlying objects (latent variables)

from observed data (shadows)

The promise of foundation models

The idea of a foundation model

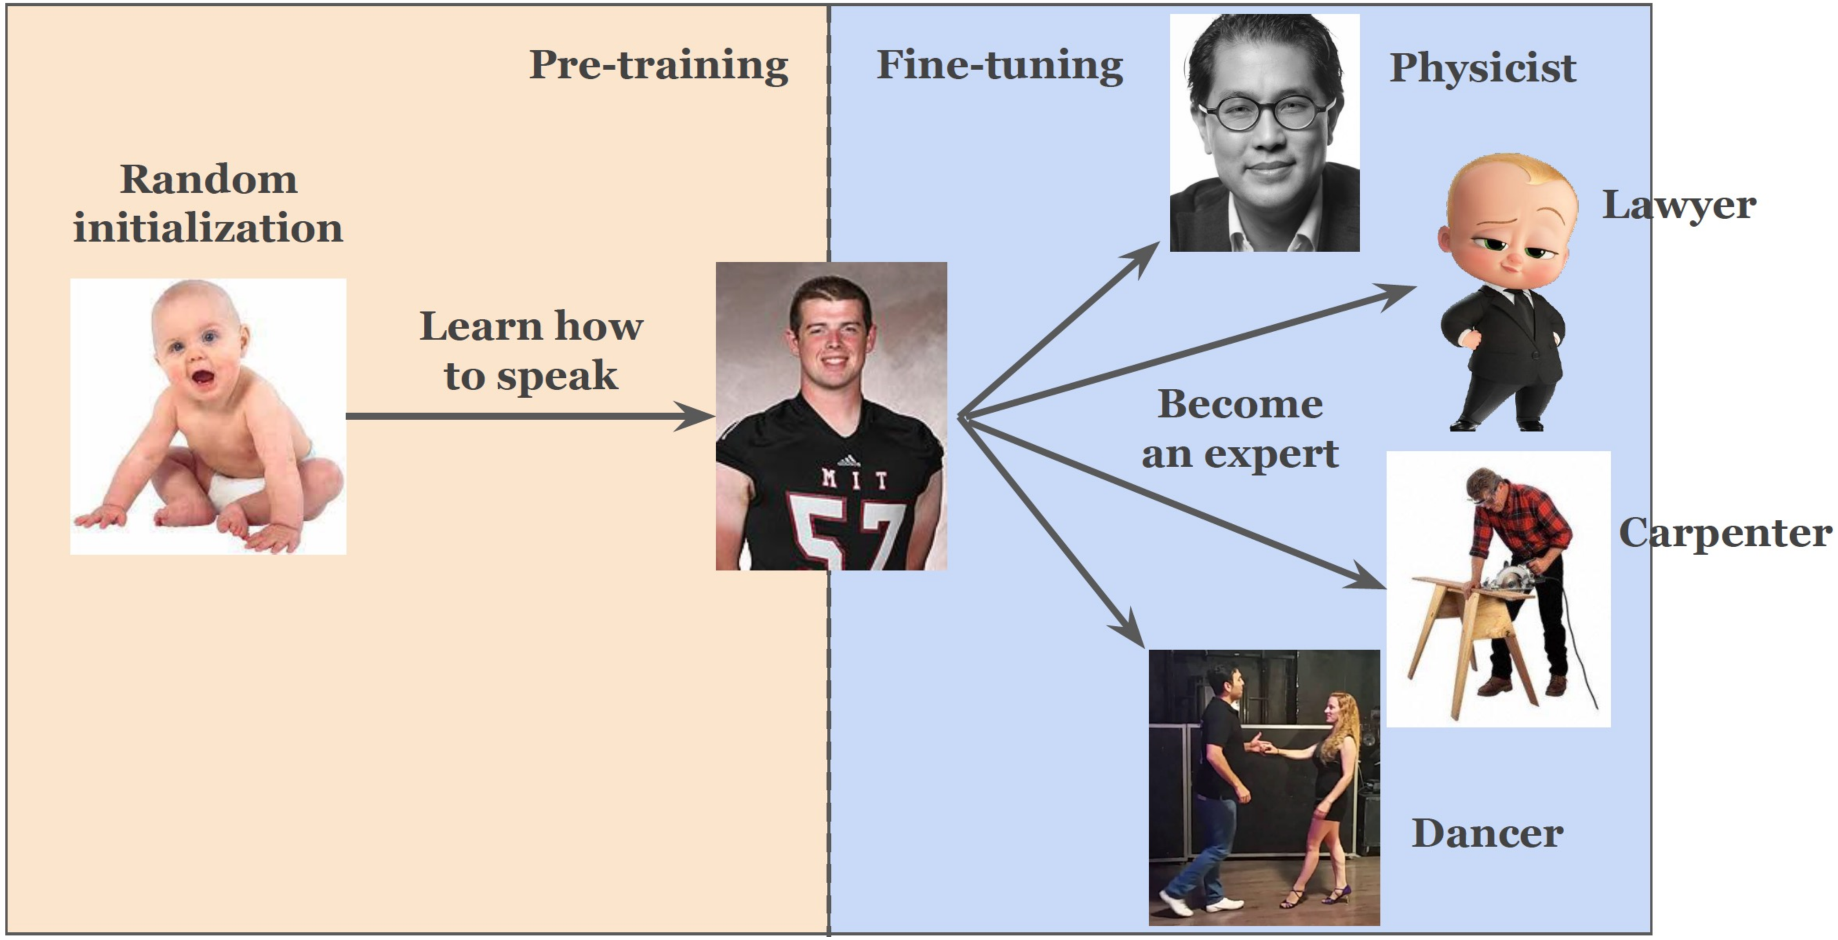

1. **Pre-train** on big unlabeled data

2. **Fine-tune** on labeled data + transfer learning

[Image credit: Kazuhiro Terao]

Characteristics of a foundation model

Pre-train using SSL* – no labels needed: can train on data

Learn meaningful data representation

Transfer & finetune: adopt to multiple downstream tasks

Multimodality: common embedding [e.g. text & images]

[*SSL = Self-supervised learning]

Pre-training

Augmentation [Re-sim]

Masking [next word prediction]

Novel physics-inspired training schemes?

Train using auxiliary tasks [e.g f-tag]

Encode physics to guide model

[2002.05709]

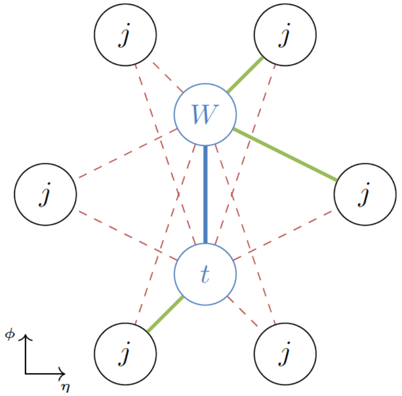

[2303.13937]

(b) Topograph

Example: masked particle modeling

Flavor tagging

Pre-training task:

Mask & predict constituents of a jet

Fine-tune for downstream tasks:

- Classification
- Weak supervision
- …

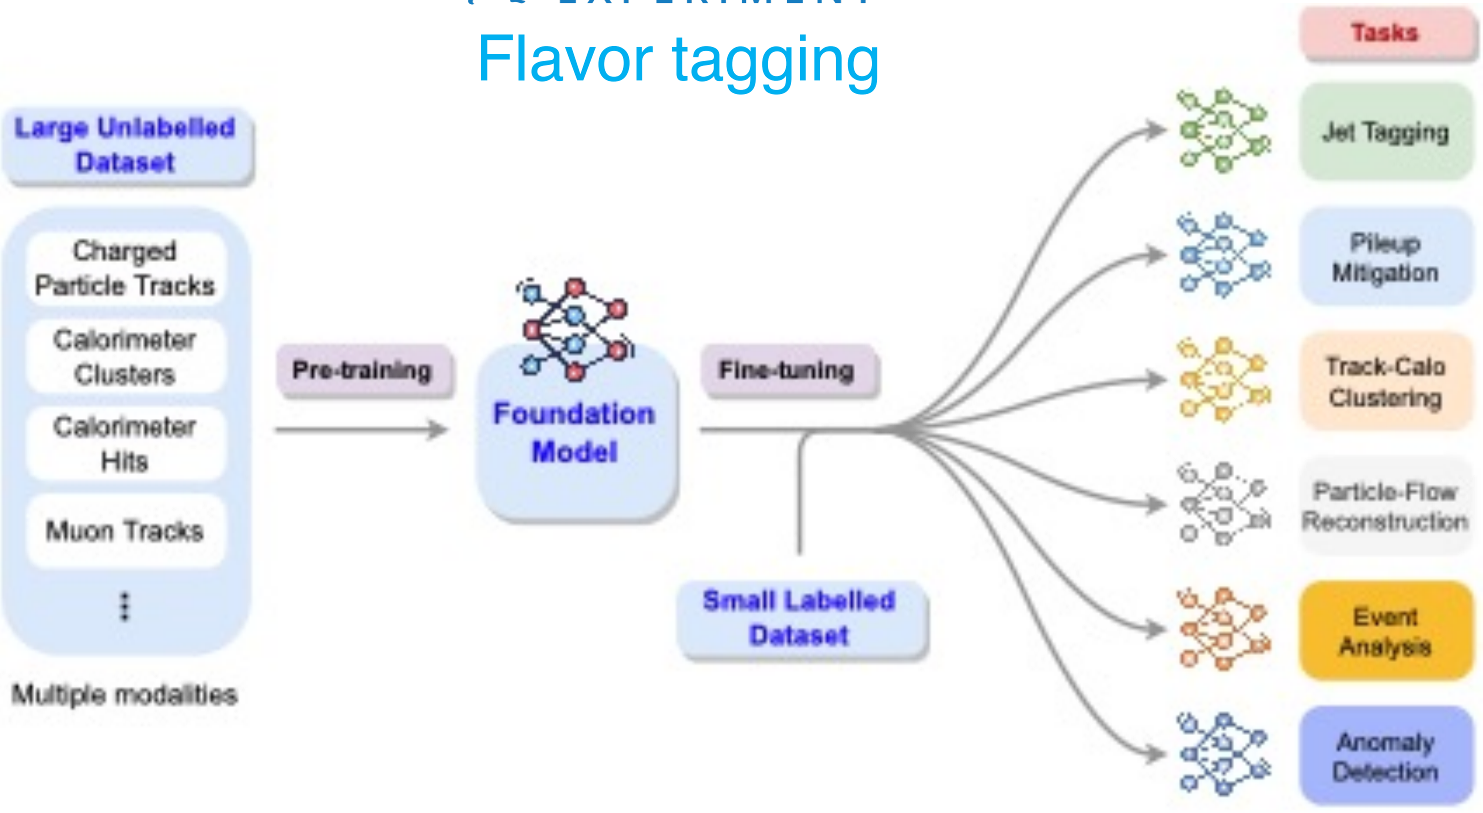

[2401.13537]

We have **our own** embedding spaces

Reconstruction & theory spaces

What do foundation models add to this?

- **End-to-end**
- **Differentiable**
- **Amortization & democratic** – reuse model
- **Multimodal** [importance of language?]
- 🐘 **Interpretability:** symbolic regression,…

A
U
T
O
M
A
T
A
B
L
E

Vision Nº 2

Push the frontier
of the unknown

or

How to optimize our
search strategy

or

Automation

BSM stubbornly resists discovery

ATLAS + CMS = O(1000) search papers

O(8’000) person years

- ~2 years per analysis
- Average of ~4 people

Best use of resources ?

[Shikma Bressler, Hammers & Nails ’22]

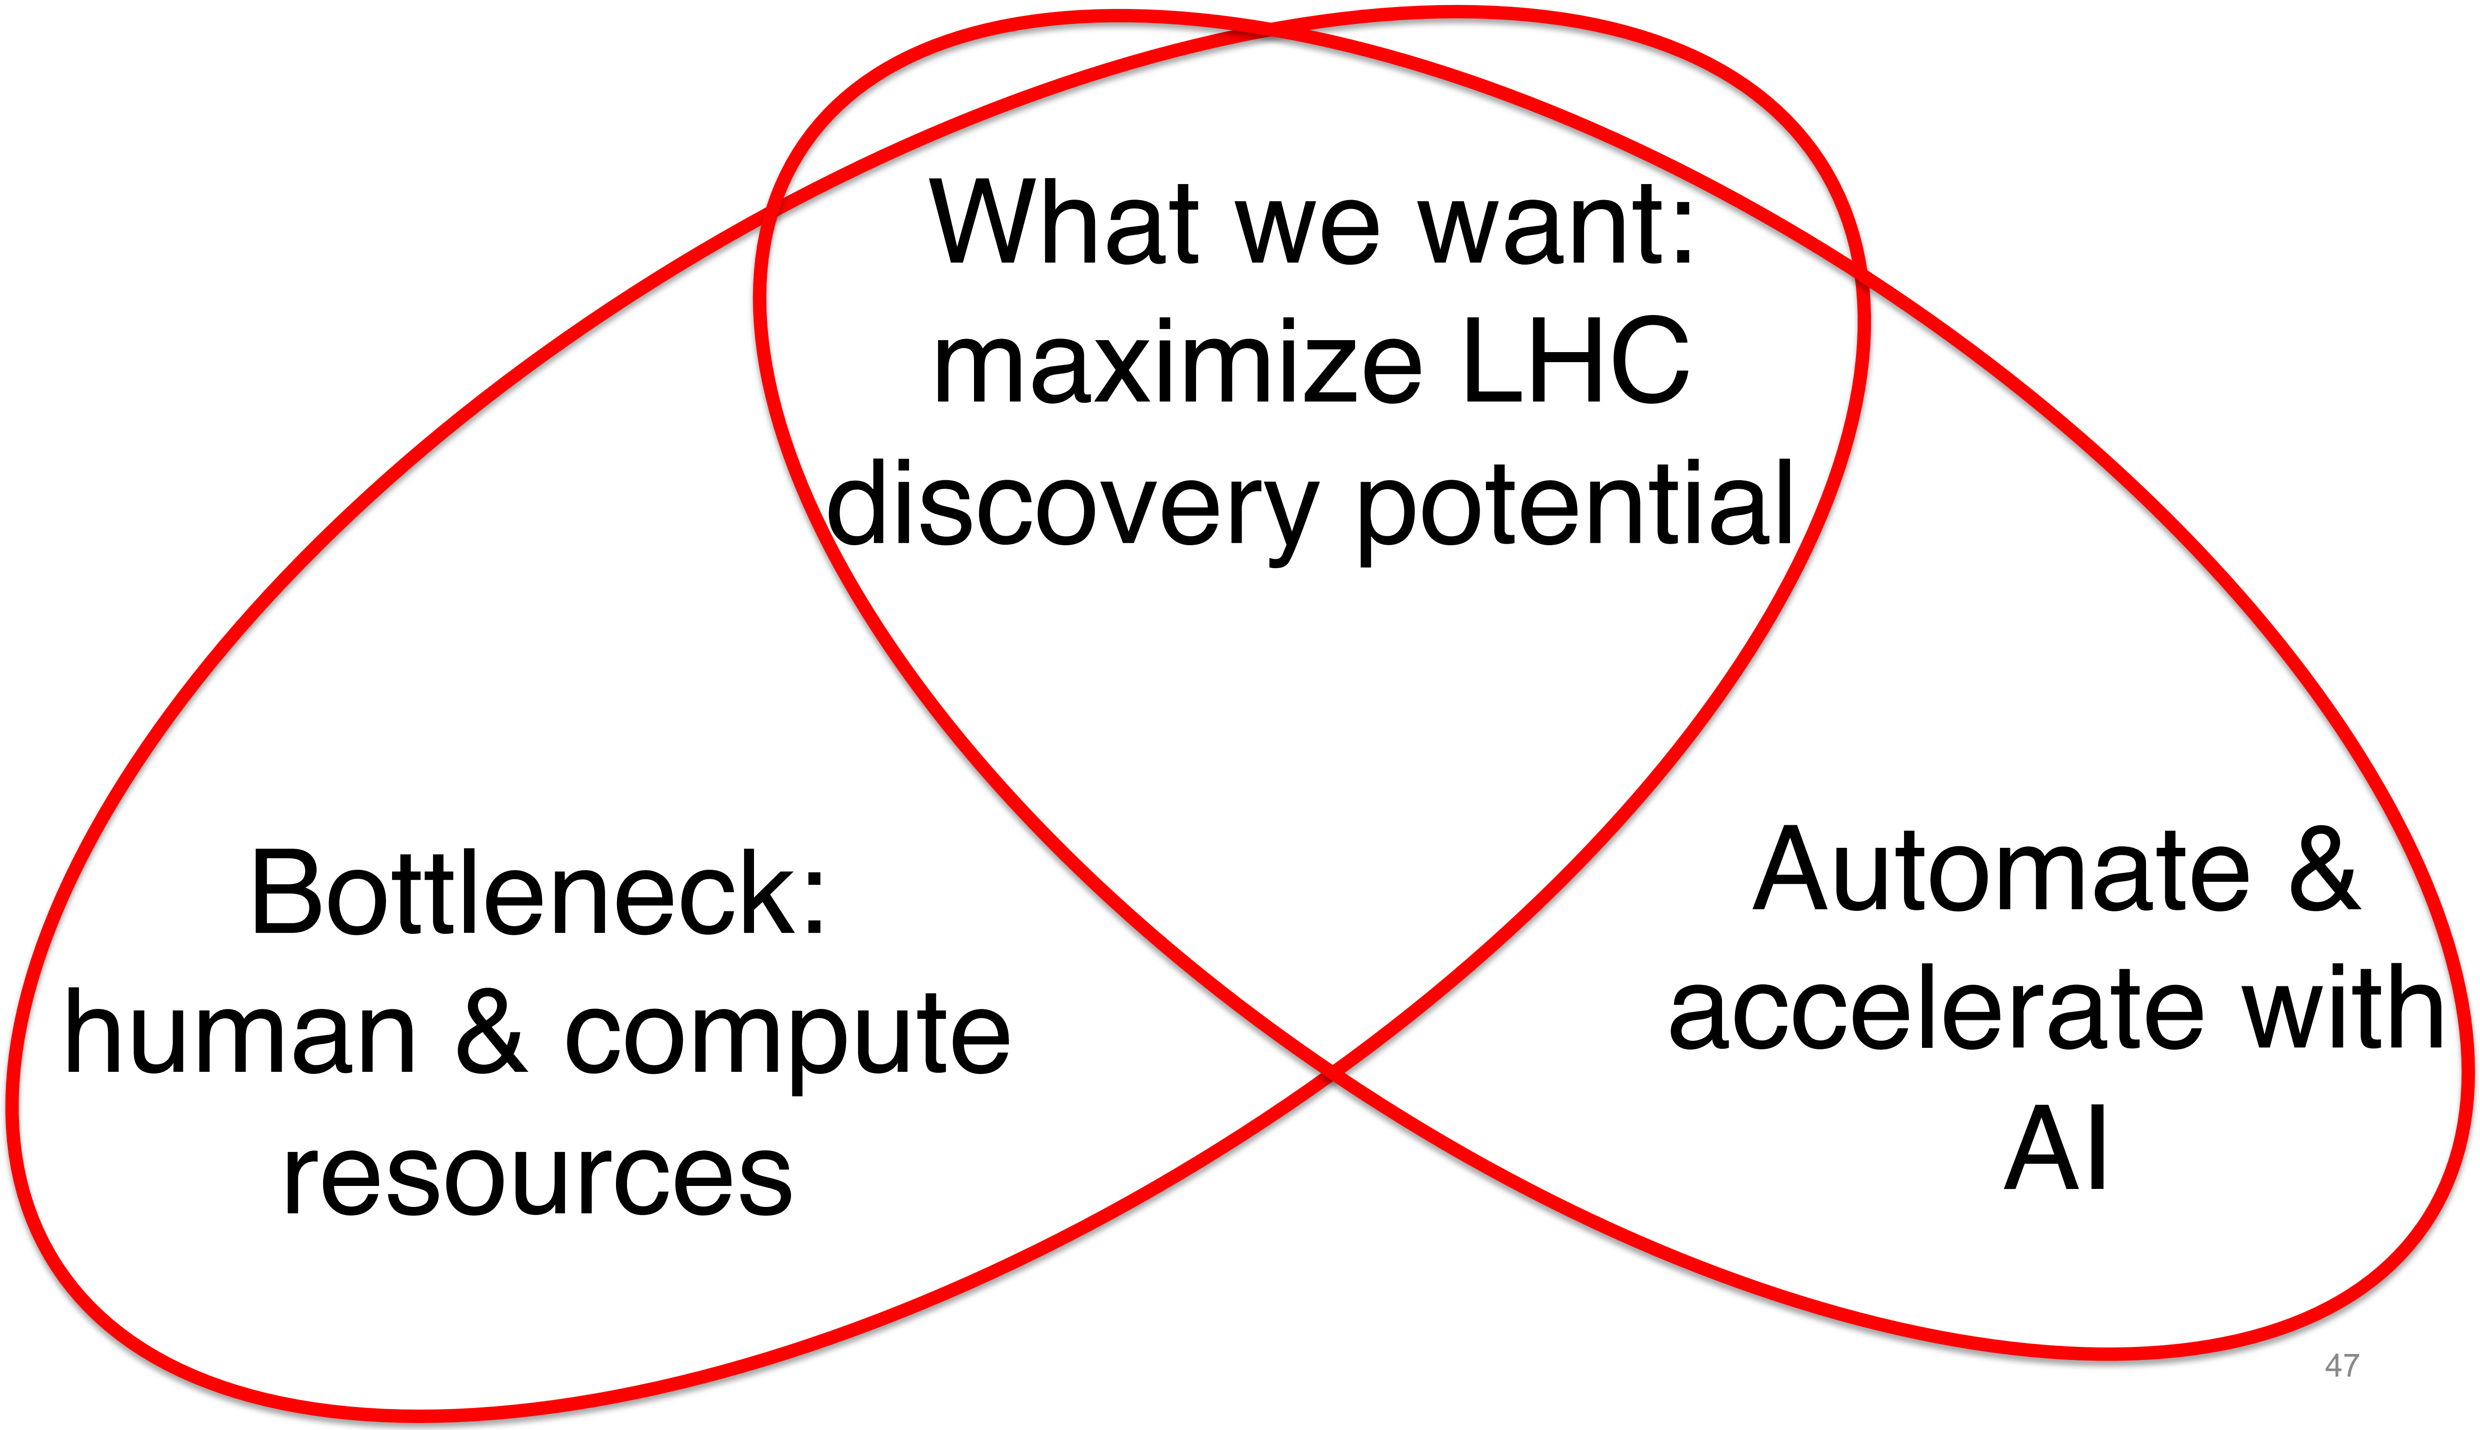
What we want: maximize LHC discovery potential
Bottleneck: human & compute resources
Automate & accelerate with AI

How much **signature** space have we explored?

	e	μ	τ	q/g	b	t	γ	Z/W	H	BSM $\to$ SM$_1$ $\times$ SM$_1$				BSM $\to$ SM$_1$ $\times$ SM$_2$				BSM $\to$ complex			
										q/g	γ/π^0's	b	$\cdots$	tZ/H	bH		$\cdots$	$\tau qq'$	eqq'	$\mu qq'$	$\cdots$
e	[37, 38]	[39, 40]	[39]	∅	∅	∅	[41]	[42]	∅	∅	∅	∅		∅	∅	∅		∅	[43, 44]	∅	
μ		[37, 38]	[39]	∅	∅	∅	[41]	[42]	∅	∅	∅	∅		∅	∅	∅		∅	∅	[43, 44]	
τ			[45, 46]	∅	[47]	∅	∅	∅	∅	∅	∅	∅		∅	∅	∅		[48, 49]	∅	∅	
q/g				[29, 30, 50, 51]	[52]	∅	[53, 54]	[55]	∅	∅	∅	∅		∅	∅	∅		∅	∅	∅	
b					[29, 52, 56]	[57]	[54]	[58]	[59]	∅	∅	∅		[60]	∅	∅		∅	∅	∅	
t						[61]	∅	[62]	[63]	∅	∅	∅		[64]	[60]	∅		∅	∅	∅	
γ							[65, 66]	[67–69]	[68, 70]	∅	∅	∅		∅	∅	∅		∅	∅	∅	
Z/W								[71]	[71]	∅	∅	∅		∅	∅	∅		∅	∅	∅	
H									[72, 73]	[74]	∅	∅		∅	∅	∅		∅	∅	∅	
BSM $\to$ SM$_1$ $\times$ SM$_1$: q/g										∅	∅	∅		∅	∅	∅		∅	∅	∅	
BSM $\to$ SM$_1$ $\times$ SM$_1$: γ/π^0's											[75]	∅		∅	∅	∅		∅	∅	∅	
BSM $\to$ SM$_1$ $\times$ SM$_1$: b												[76, 77]		∅	∅	∅		∅	∅	∅	
$\vdots$																					
$\vdots$																					

Vast signature space **un**explored

[1907.06659]

How to quantify *coverage*?

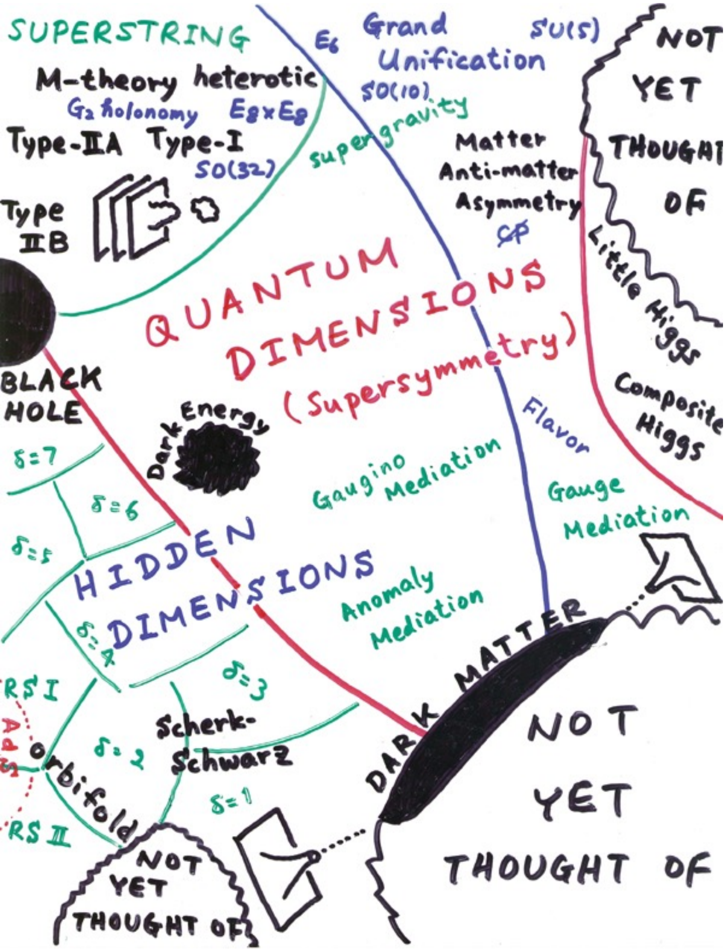

What theory prior? [Frequentist vs. Bayesian]

How to interpret "model-agnostic" null results?

Go beyond benchmarking (i.e. Frequentist)

Recastability!

Follow-up strategy after an "anomalous" signal?

Balance cost of follow-up against frequency alerts?

"What is the next best search given all existing search results?"

Our go-to method: **2**-hypothesis test*

Works great if you know what you’re looking for !

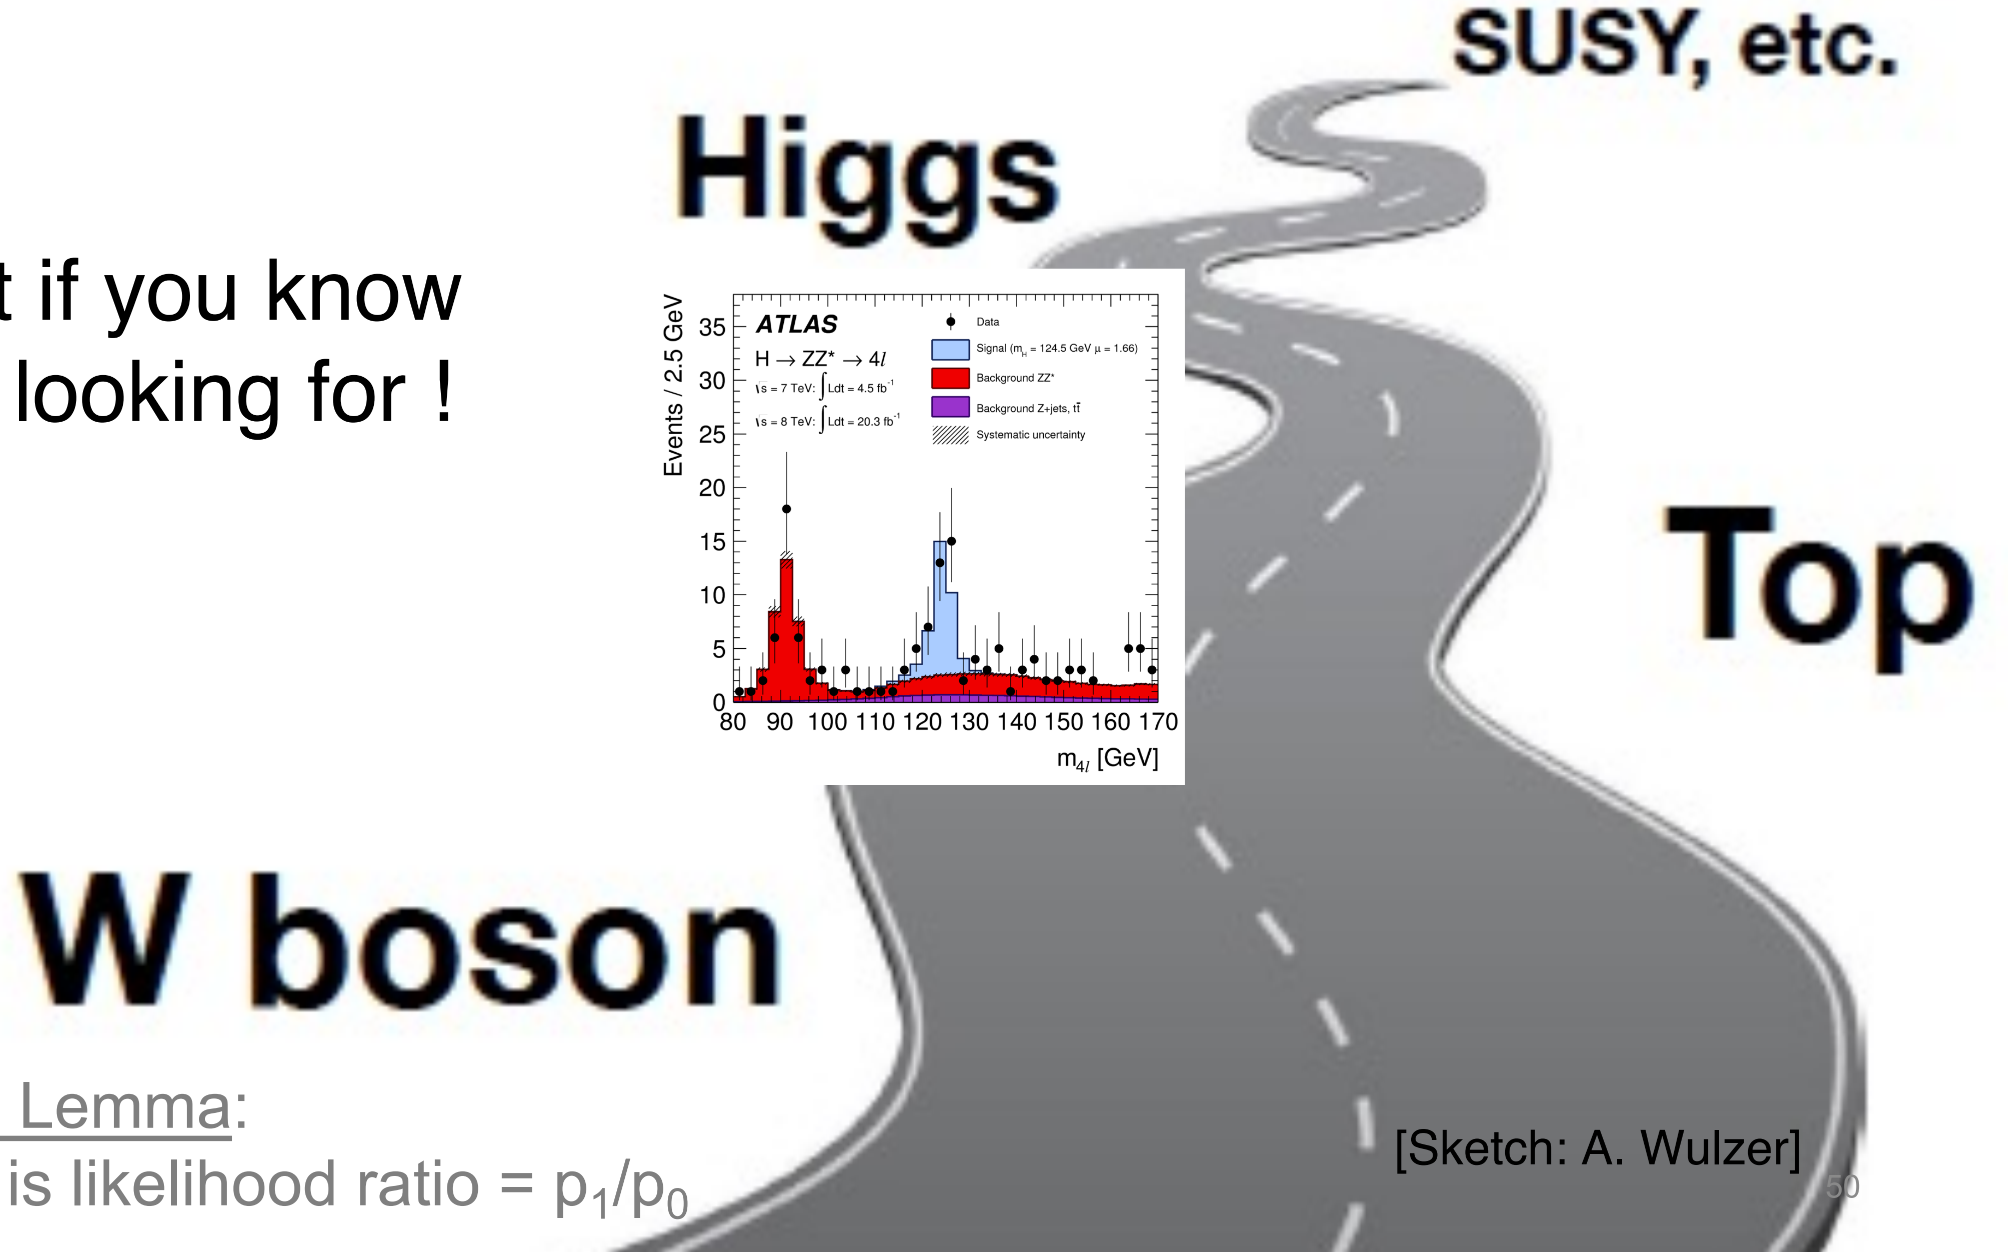

[Sketch: A. Wulzer]

*Neyman-Pearson Lemma:
Best test statistics is likelihood ratio = p_1/p_0

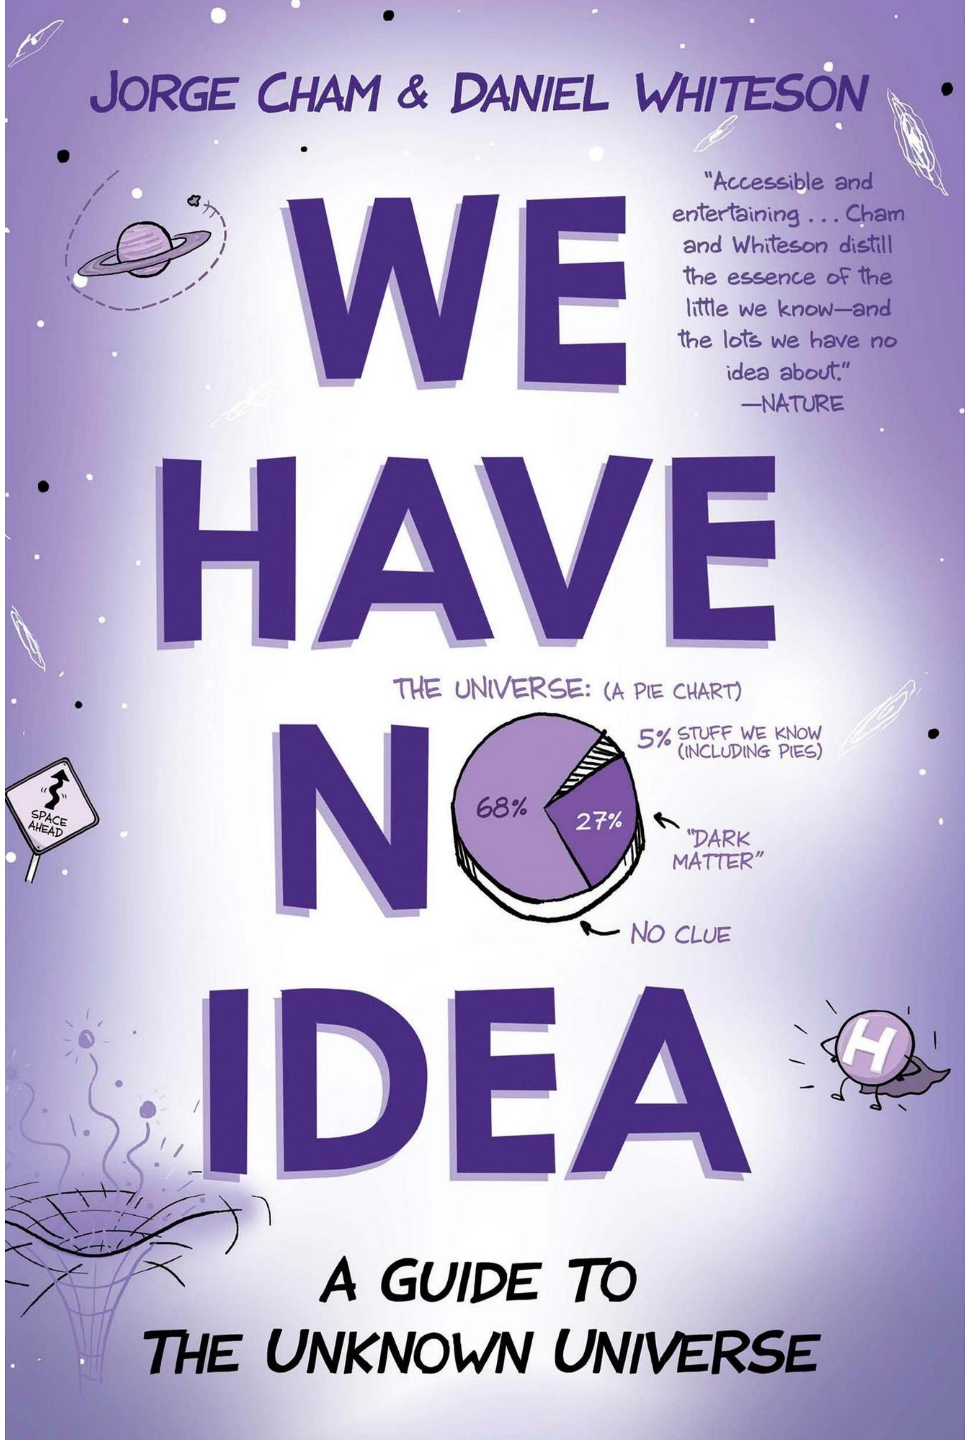

"No convincing theoretical guidance"

No *trust* in p_1 = playing the lottery!

p_0 = SM

p_1 = *everything else*

How to design *complementary* search strategy?

Theory guided $\leftrightarrow$ data-inspired

$$\frac{p_1}{p_0}$$

Foundation model:
discrete BSM $\rightarrow$ continuous embedding

Door to alternative metric: volume in *embedded space* [2208.05484]

$\rightarrow$ compare *reach* of colliders

MC, in-situ BG estimate,…

Becomes question of automation

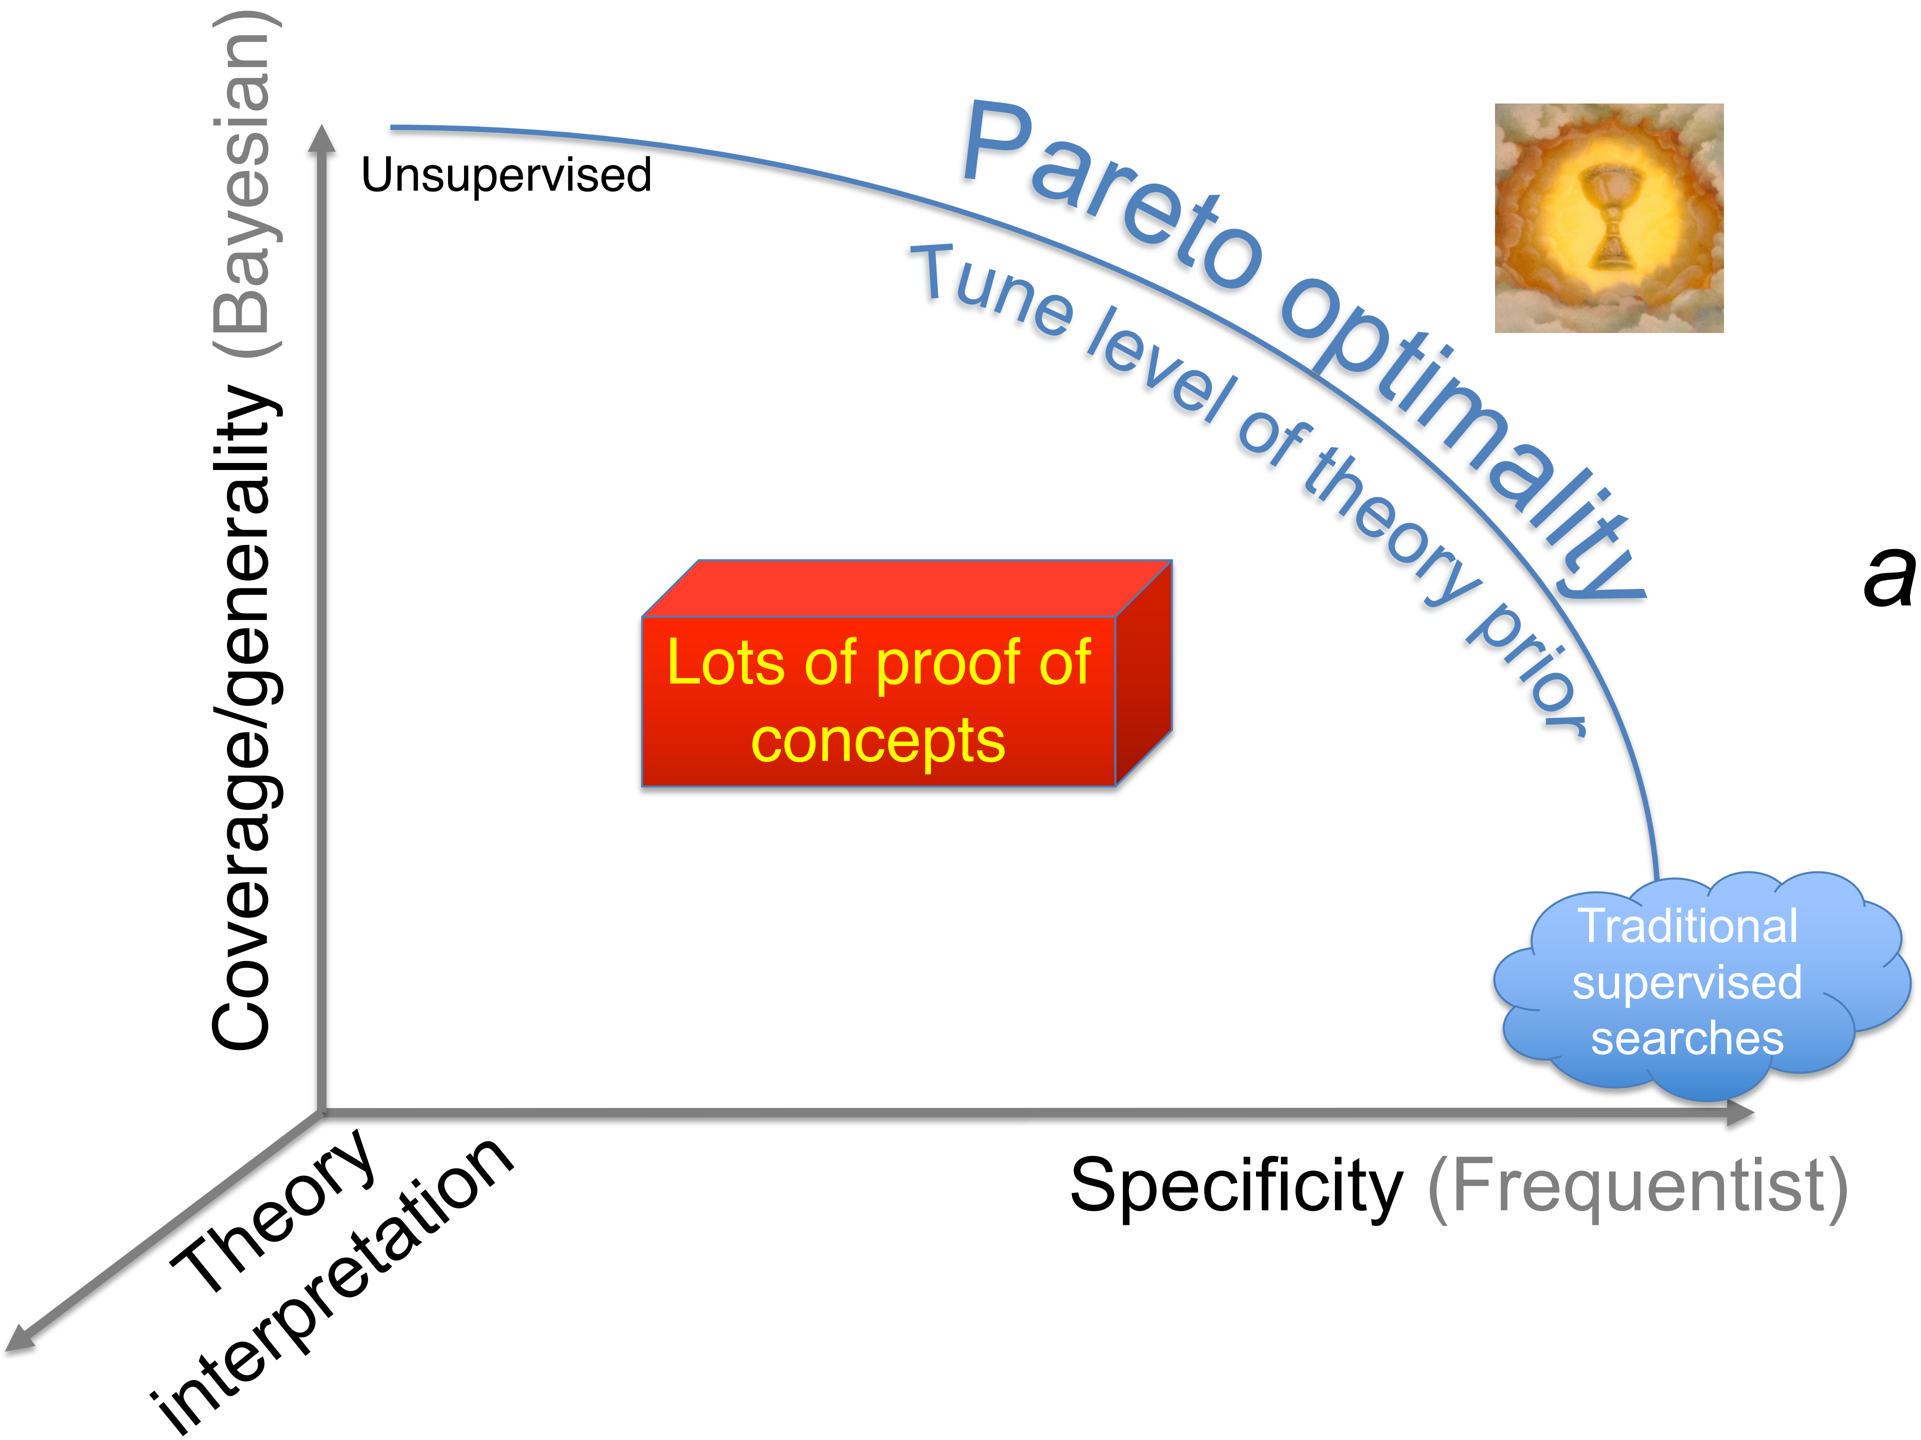

Diverse
&
automated
search
strategy

Vision Nº 3

The future of particle physics

or

What if secrets of nature are NOT in our current data?

Ultimate goal: learning about nature

Likelihood

$$p(\text{theory}\,|\,\text{data}) = \frac{p(\text{data}\,|\,\text{theory})\,p(\text{theory})}{p(\text{data})}$$

Posterior

Optimizing the science output

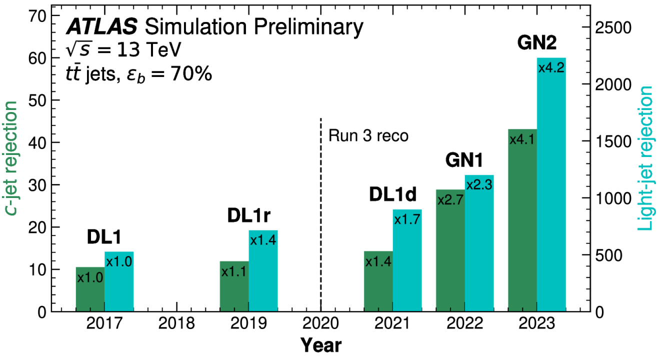

Optimal classification

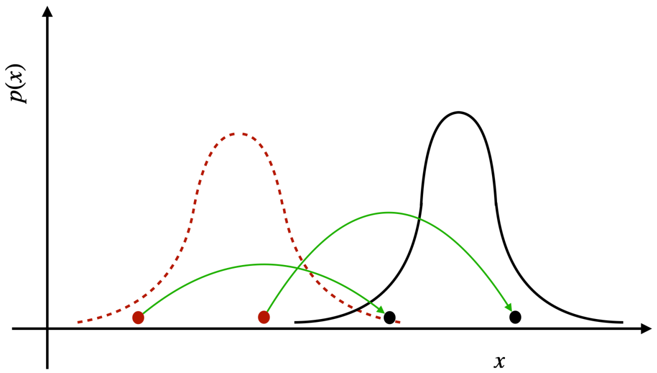

Optimal calibration

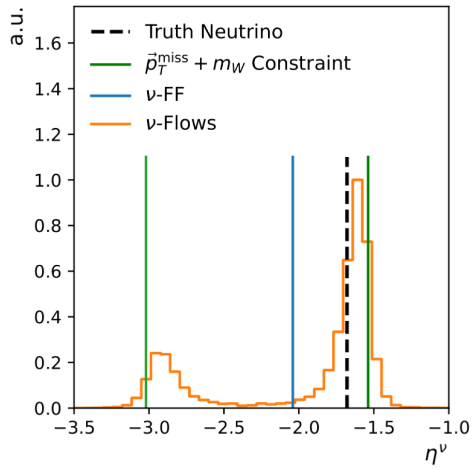

Optimal reconstruction

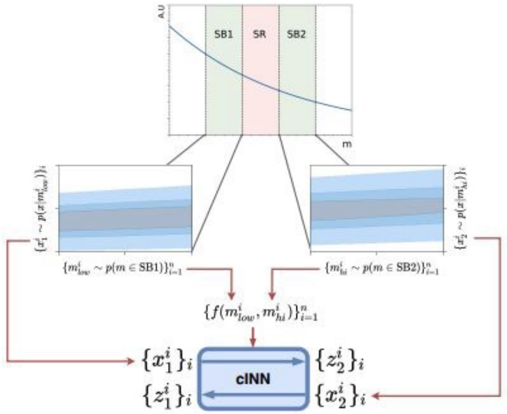

Optimal simulation

Natural limit: true posterior p(theory | data)

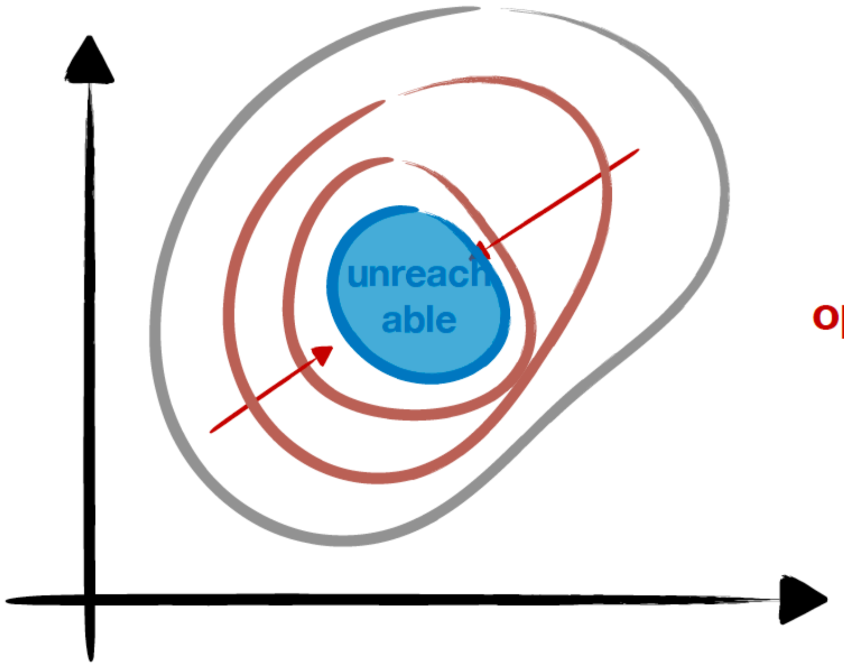

unoptimized

optimized (e.g. w/ ML)

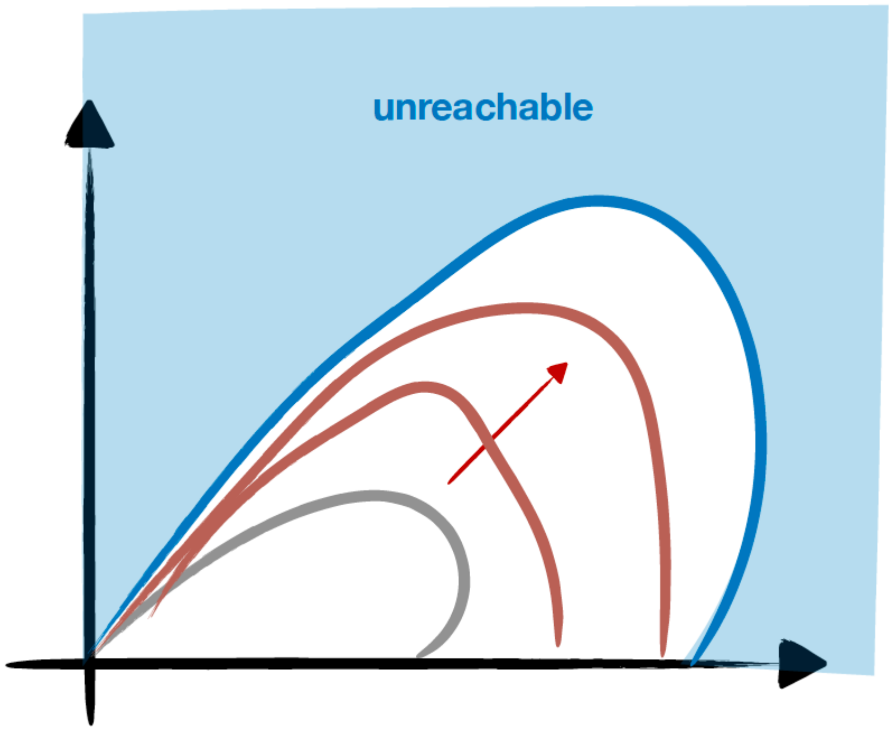

Measurements
(e.g. Higgs Couplings)

Searches
(e.g. Supersymmetry)

Need better data

[Lukas Heinrich - Detector design using differential programing]

Design new ***optimal*** experiment
to optimize p(theory | data)

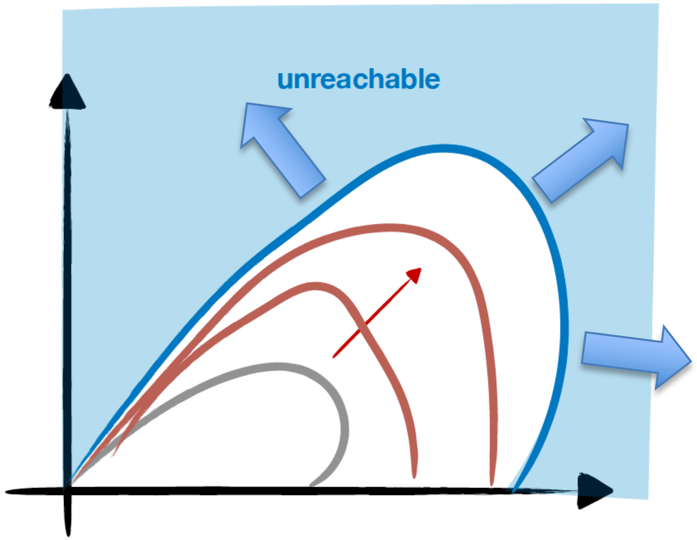

Inform & augment experts
→ collaboration!

[Lukas Heinrich - Detector design using differential programing]
[see also MODE]

COMMUNITY EFFORTS

Get organized !

European **C**oalition for **AI** in **F**undamental Physics

JENA Expressions of Interest

ECFA
European Committee for Future Accelerators

NuPECC

APPEC

EuCAIF mission: community consensus, provide structure & support

Topics of interest: [feel free to sign up]

- Foundation models for fundamental physics
- Optimal design
- …

eucaif

Concluding remarks

AI = enabling technology ⇒ time to harness

AI for Science & Science for AI

RODEM = enabler of my research

syn·er·gy | ˈsinərjē |

ATLAS EXPERIMENT

UNIVERSITÉ DE GENÈVE
FACULTY OF SCIENCE

PIs


TG


Slava Voloshynovskiy


François Fleuret

PhD students


Tomke Schröer


Malte Algren


Jona Ackerschott


Matthew Leigh


Debajyoti Sengupta


Sam Klein


Stephen Mulligan


Guillaume Quétant


Mariia Drozdova


Ivan Oleksiyuk


Bálint Máté


Atul Kumar Sinha


Daniele Paliotta

Postdocs


Kinga Wozniak


Johnny Raine replacement starting soon

This could be you !

Master


Franck Rothen

Alumni


Lukas Ehrke


Knut Zoch


Manuel Guth


Matthias Schlaffer


Sebastian Pina-Otey

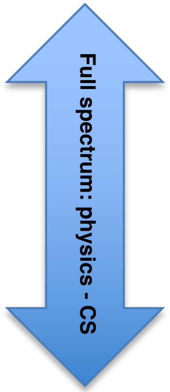

FONDATION ERNEST BONINCHI

Alexander von Humboldt Stiftung / Foundation

European Commission

Swiss National Science Foundation

Sinergia

Swiss Government Excellence Scholarships
STUDY IN SWITZERLAND