IWR School

Applications in theory and beyond

Veronica Sanz (IFIC-UValencia/CSIC)

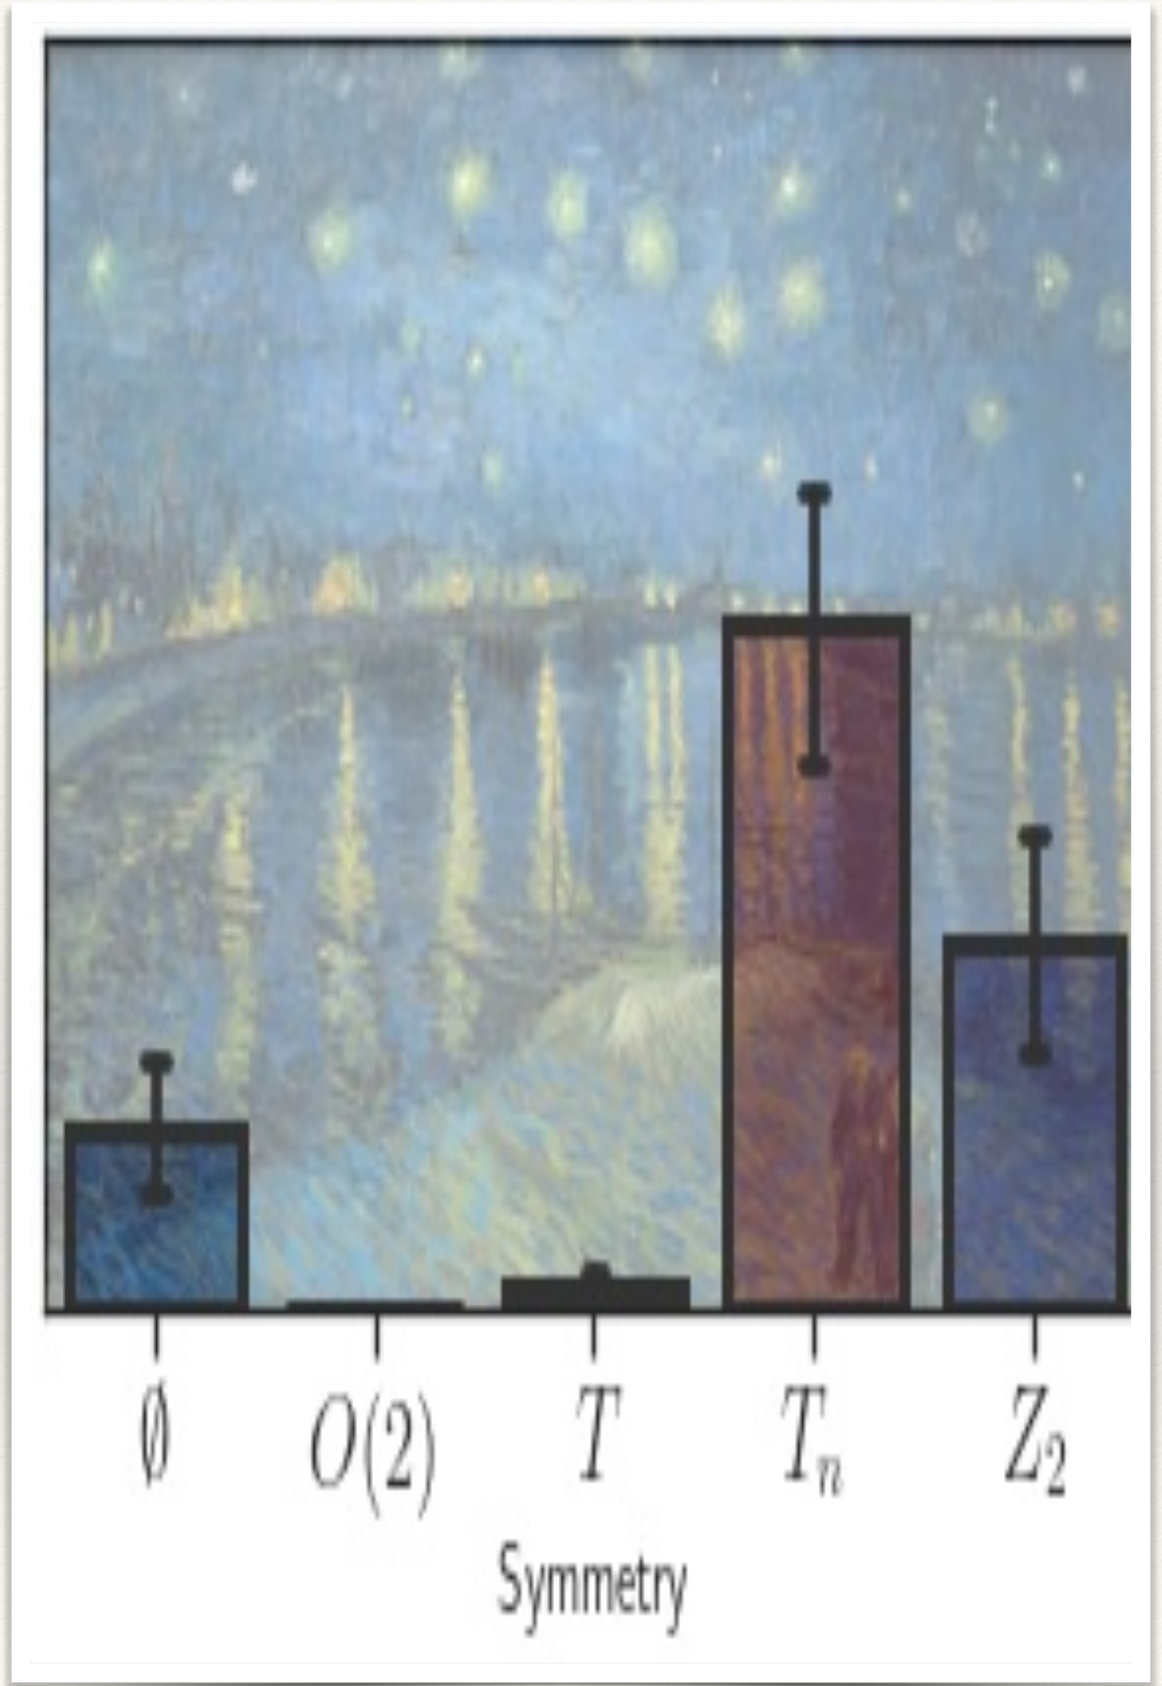

In these lectures, we will talk about

1. AI-assisted Science
2. Human surrender?
3. Looking under the hood
4. Variational Autoencoders
5. Symmetries

*classical mechanics

*particle collisions

*art

6. Beyond physics with AI

*Ecology and restoration

*Musical concepts

*Political adherence

*Traffic and pollution, Sun storms

AI-assisted Science

AI vs ML vs DL

Today we live in an era where the edge is Deep Learning

with new opportunities and new challenges

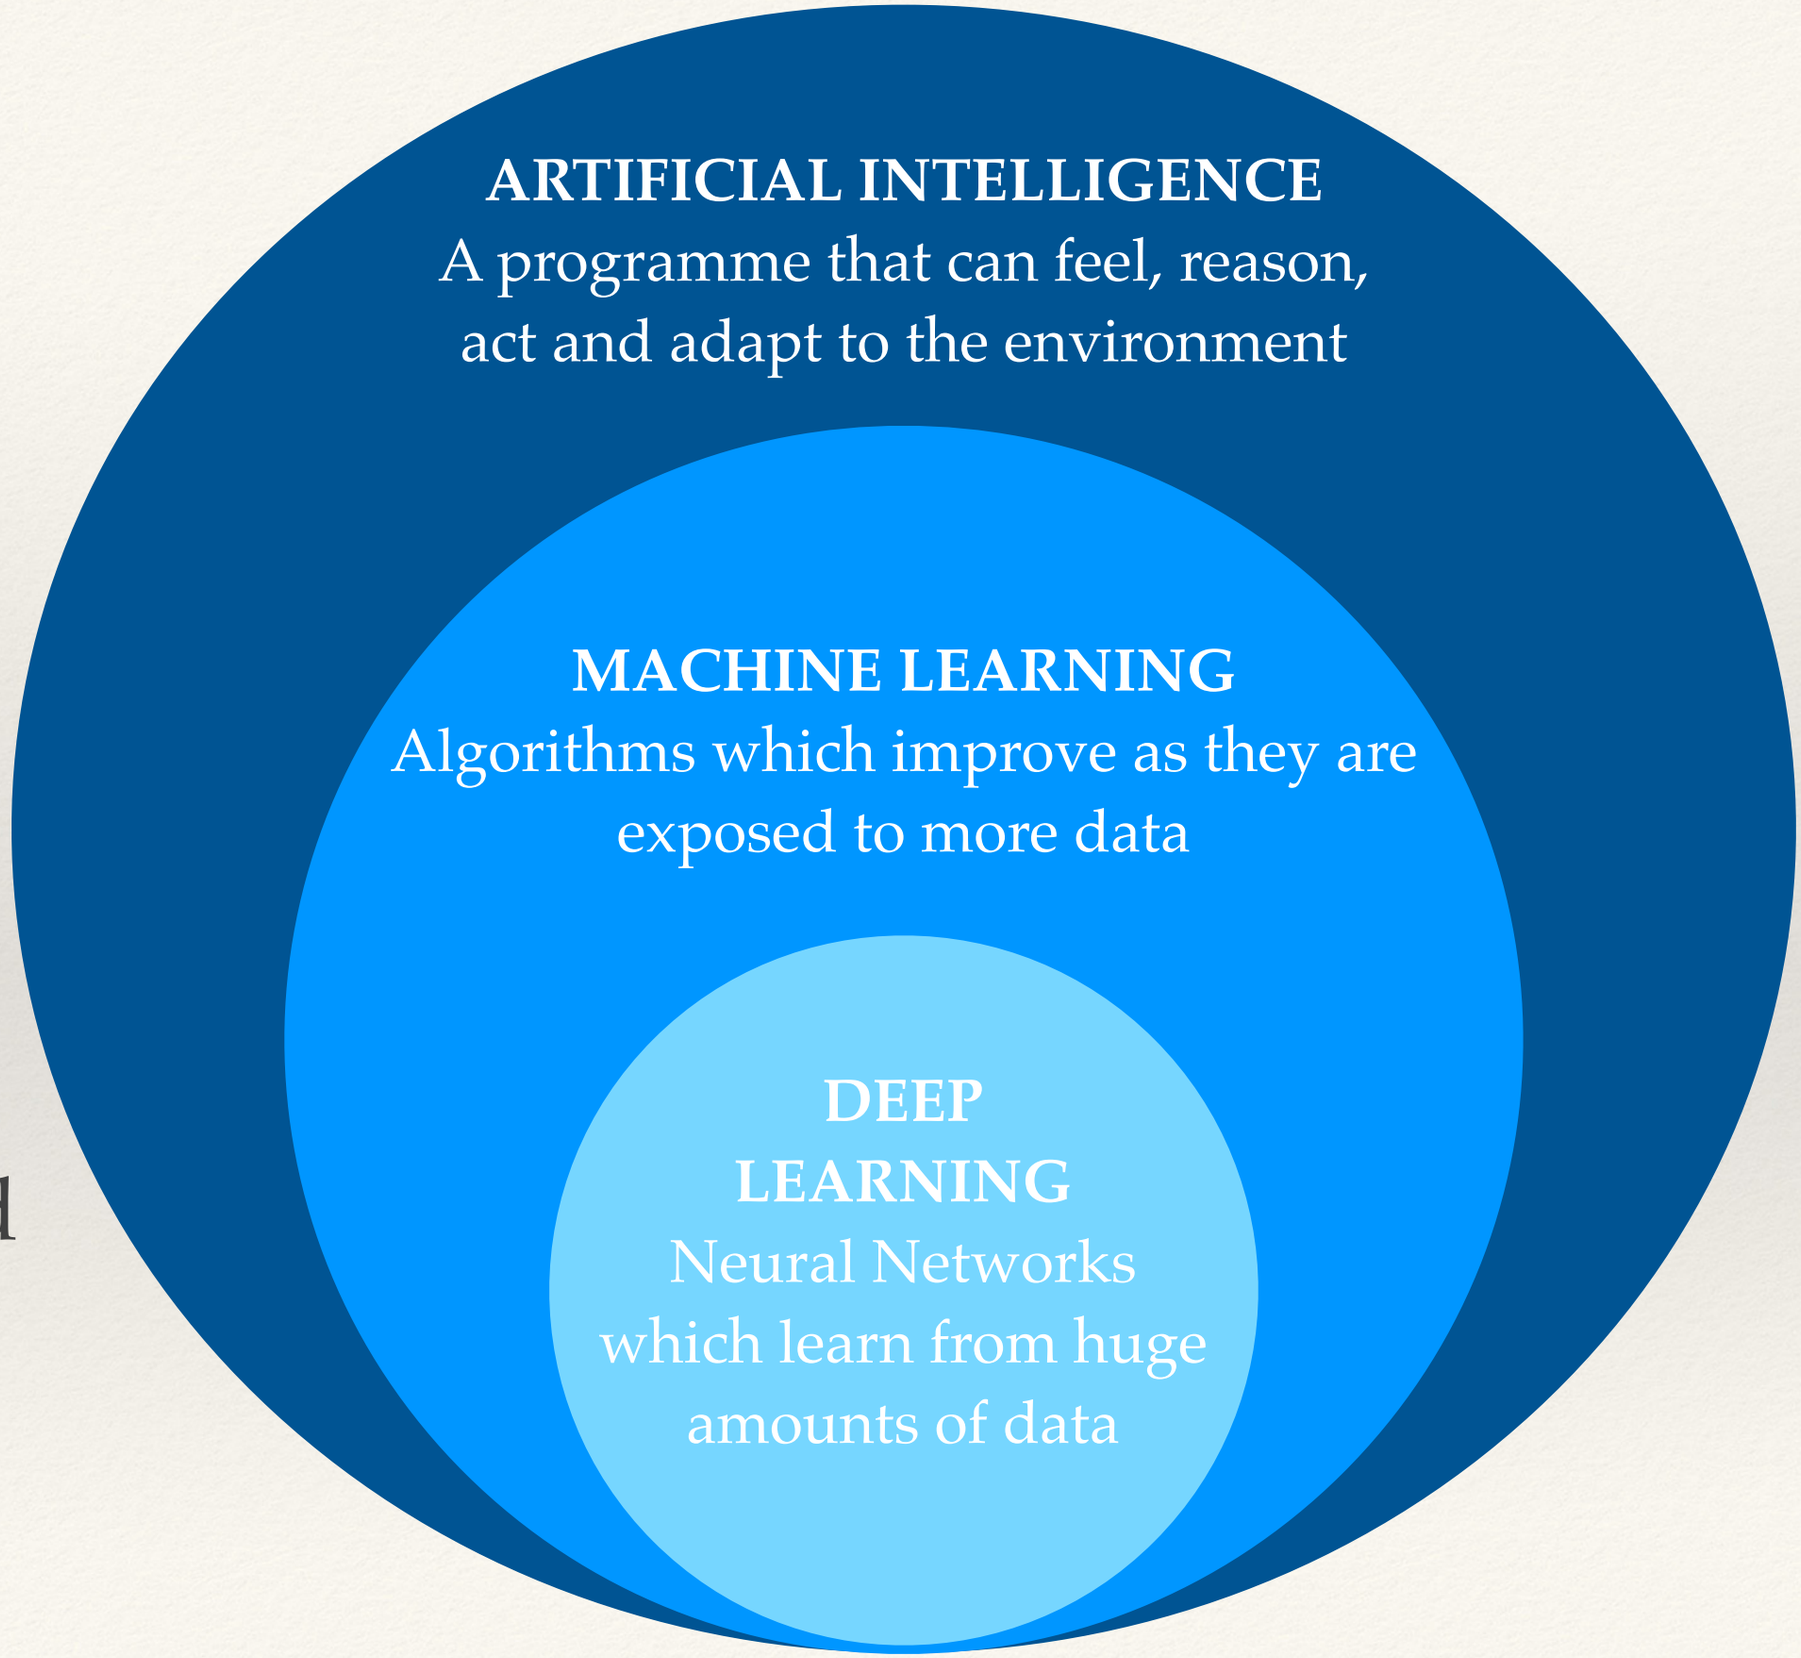

In particle, astro, cosmo & nuclear physics,
Artificial Intelligence techniques
are nowadays commonplace

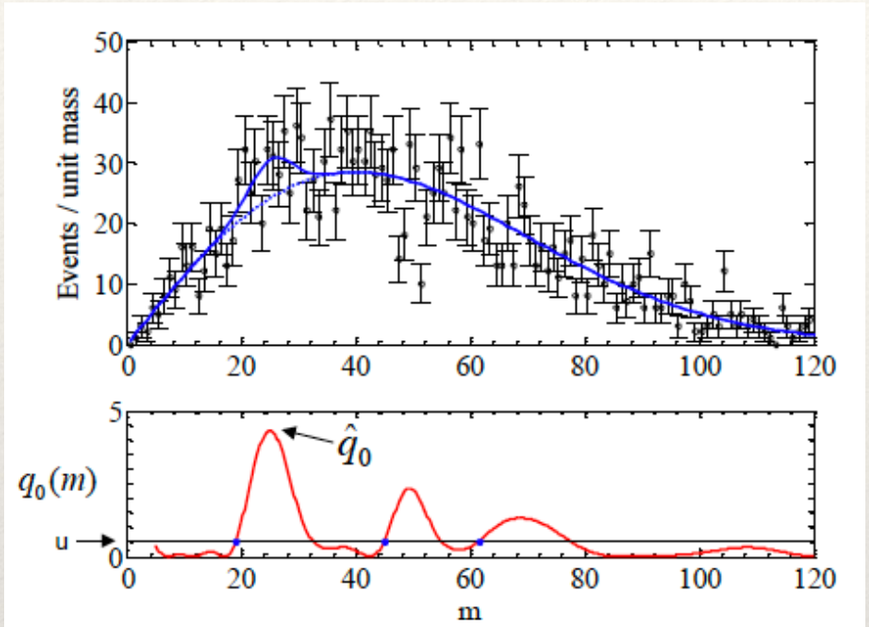

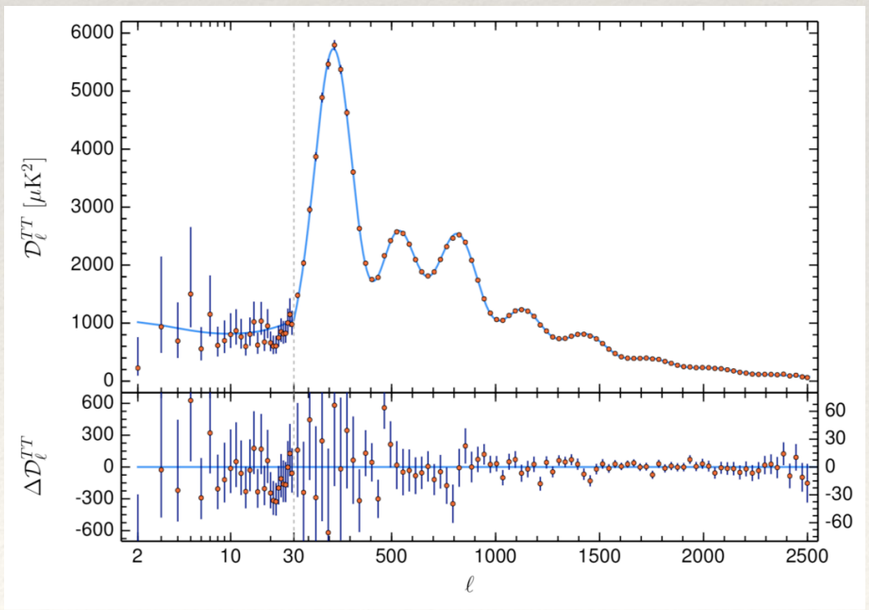

In fact,
we have always been early adopters and developers of sophisticated statistical techniques

why? the stakes are high, the community is large and connected, we usually have a good physical understanding and can face increasingly complex questions

Exps. have been using ML for a very long time
esp in channels where discovery = squeezing a bit more
(eg. TeVatron's monotop)

Exps. have been using ML for a very long time esp in channels where discovery = squeezing a bit more (eg. TeVatron's monotop)

e.g. ATLAS HDBS-2018-34

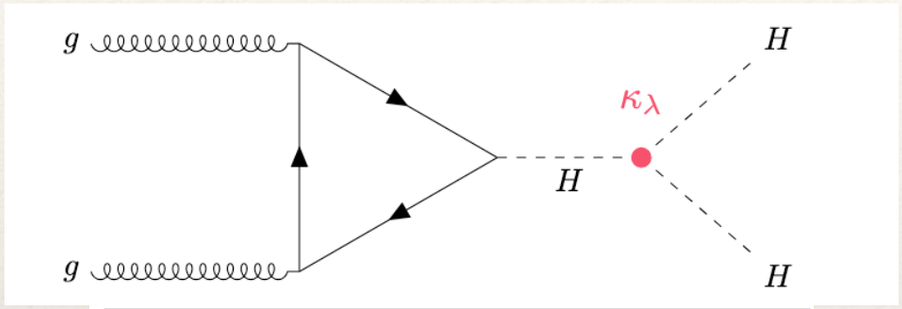

Photon-related kinematic variables	
$p_T/m_{\gamma\gamma}$	Transverse momentum of each of the two photons divided by the diphoton invariant mass $m_{\gamma\gamma}$
η and ϕ	Pseudorapidity and azimuthal angle of the leading and subleading photon
Jet-related kinematic variables	
b-tag status	Tightest fixed b-tag working point (60%, 70%, or 77%) that the jet passes
p_T, η and ϕ	Transverse momentum, pseudorapidity and azimuthal angle of the two jets with the highest b-tagging score
$p_T^{b\bar{b}}$, $\eta_{b\bar{b}}$ and $\phi_{b\bar{b}}$	Transverse momentum, pseudorapidity and azimuthal angle of the b-tagged jets system
$m_{b\bar{b}}$	Invariant mass of the two jets with the highest b-tagging score
H_T	Scalar sum of the p_T of the jets in the event
Single topness	For the definition, see Eq. (**??**)
Missing transverse momentum variables	
E_T^{miss} and ϕ^{miss}	Missing transverse momentum and its azimuthal angle

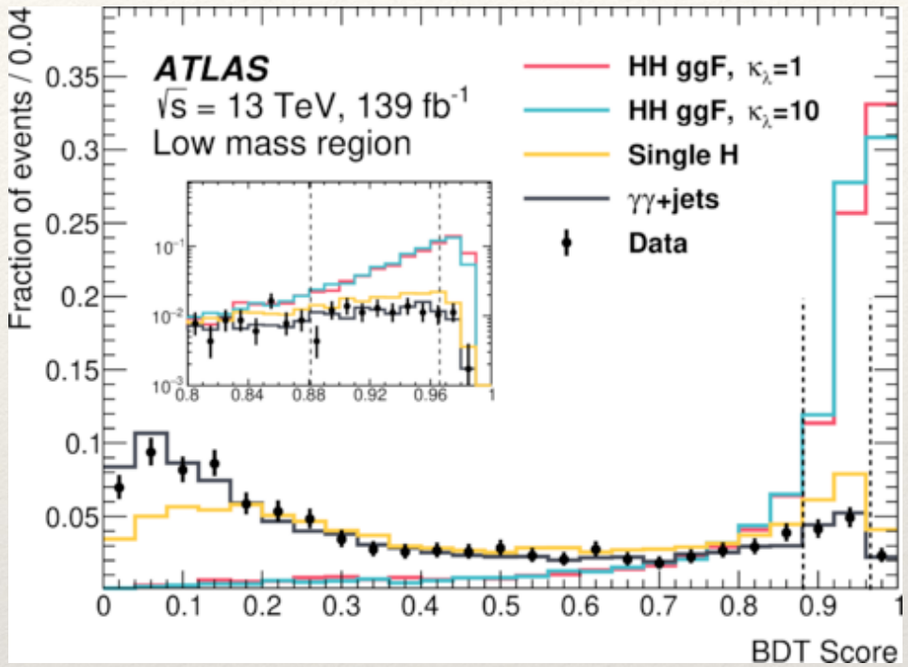

Exps. have been using ML for a very long time esp in channels where discovery = squeezing a bit more (eg. TeVatron's monotop)

e.g. ATLAS HDBS-2018-34

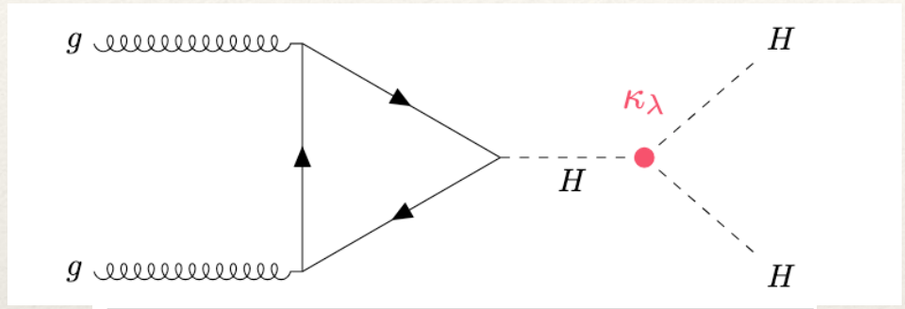

Photon-related kinematic variables	
$p_T/m_{\gamma\gamma}$	Transverse momentum of each of the two photons divided by the diphoton invariant mass $m_{\gamma\gamma}$
η and ϕ	Pseudorapidity and azimuthal angle of the leading and subleading photon
Jet-related kinematic variables	
b-tag status	Tightest fixed b-tag working point (60%, 70%, or 77%) that the jet passes
p_T, η and ϕ	Transverse momentum, pseudorapidity and azimuthal angle of the two jets with the highest b-tagging score
$p_T^{b\bar{b}}$, $\eta_{b\bar{b}}$ and $\phi_{b\bar{b}}$	Transverse momentum, pseudorapidity and azimuthal angle of the b-tagged jets system
$m_{b\bar{b}}$	Invariant mass of the two jets with the highest b-tagging score
H_T	Scalar sum of the p_T of the jets in the event
Single topness	For the definition, see Eq. (**??**)
Missing transverse momentum variables	
E_T^{miss} and ϕ^{miss}	Missing transverse momentum and its azimuthal angle

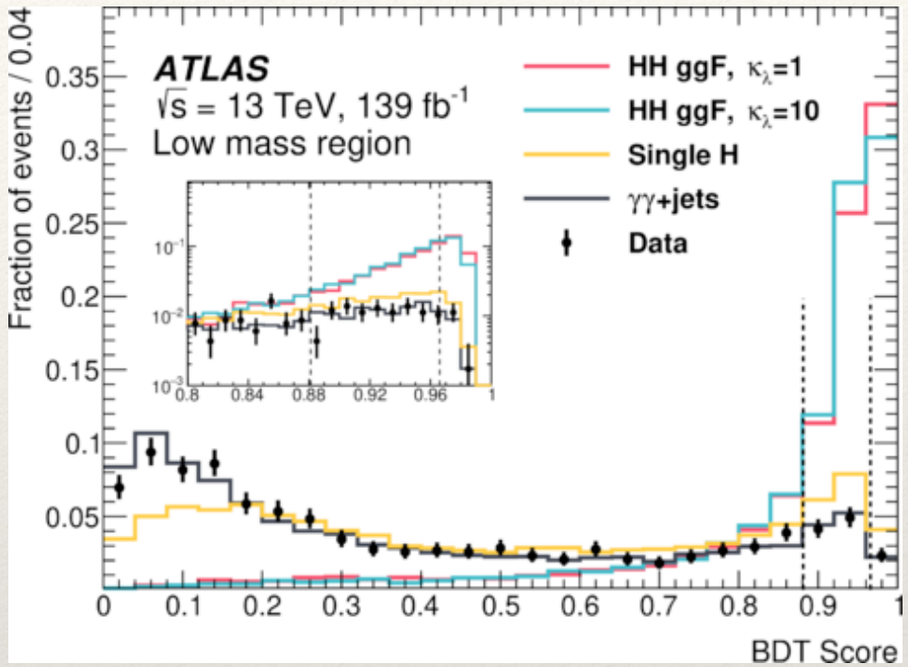

BDT $(p_T^\gamma/m_{\gamma\gamma}\,, p_T^{b\bar{b}}\ldots)$

Where this function is highly nonlinear and can express **complex relationships**

CAVEATs: interpretability, communication to theorists, reuse of data… all the stuff that happens when we move towards more complex/aggressive techniques

A lot of ML in Particle Physics is answering YES/NO questions

Is it a W?

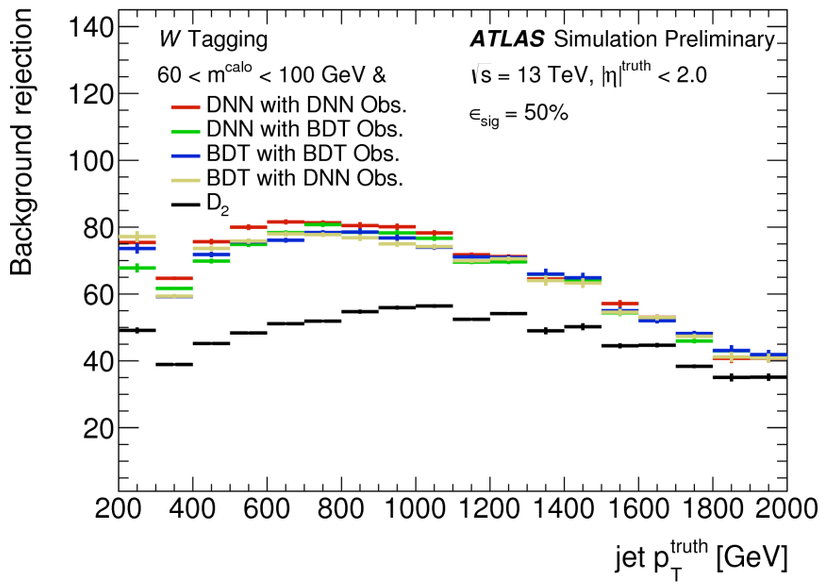

Is it a Higgs?

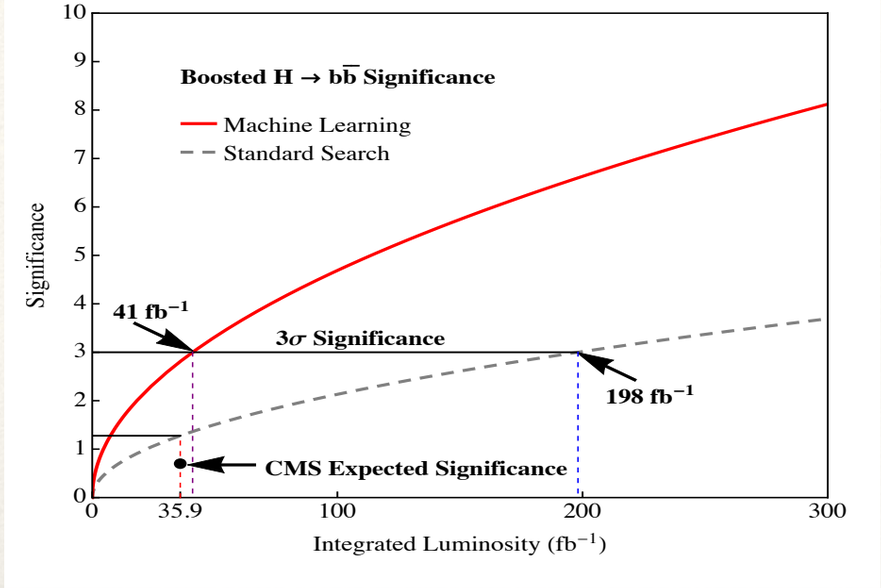

Is it DM?

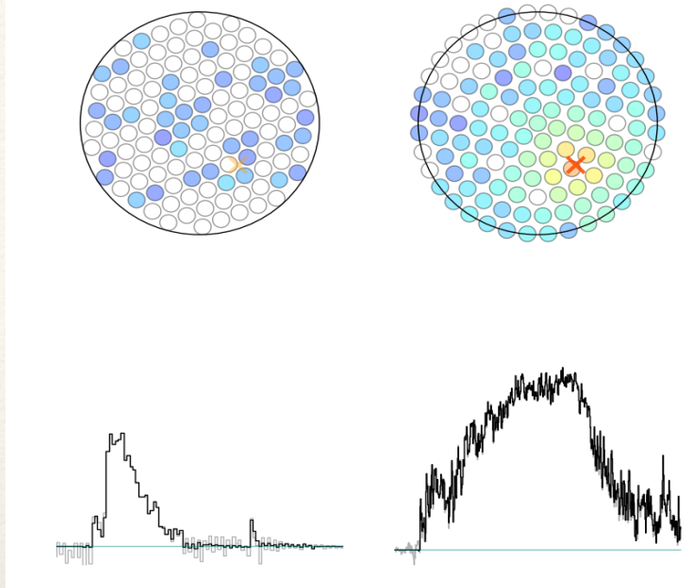

The gains in ID-ing phenomena are typically in the range of 5%-30%

for tricky environments: difference between **discovery or not**

Thanks to B Nachmann

Apart from better ID'ing, AI helps speeding up simulations

Here the gains respect to traditional techniques can be **huge**

Example: "*Up to two billion times acceleration of scientific simulations with deep neural architecture search*" (2001.08055)

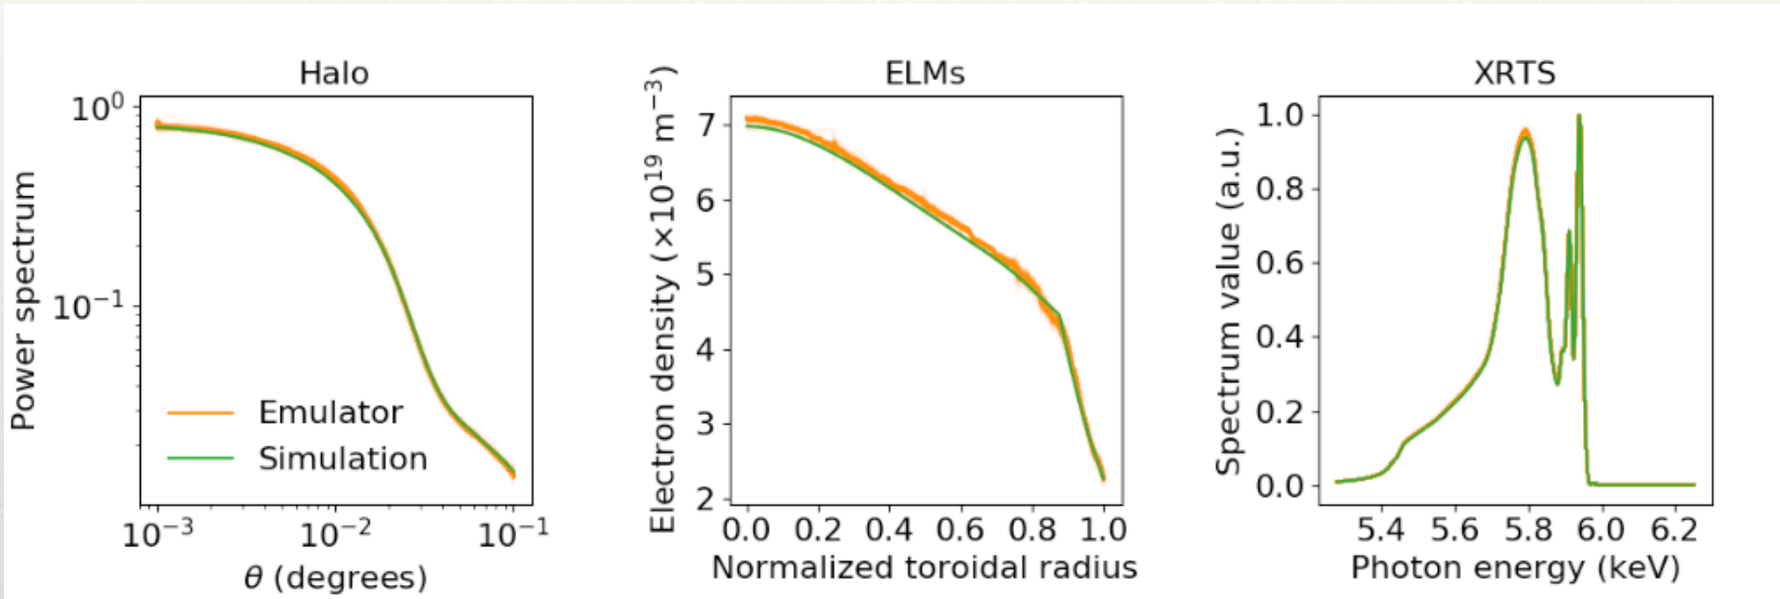

Galaxy halo | **Tokamak fusion** | **Inelastic XRay Thomson scattering**

Method	Platform	Time/shower (msec)	Speedup
Monte Carlo	2S Intel Xeon Platinum 8180	17000	1.0
3DGAN CPU		16	1063
3DGAN GPU	GTX 1080	4	4250

And we are getting **even better** thanks to Generative AI

e.g. 2109.07388

SEE MANUEL's LECTURES

Why are NNs so good at learning?

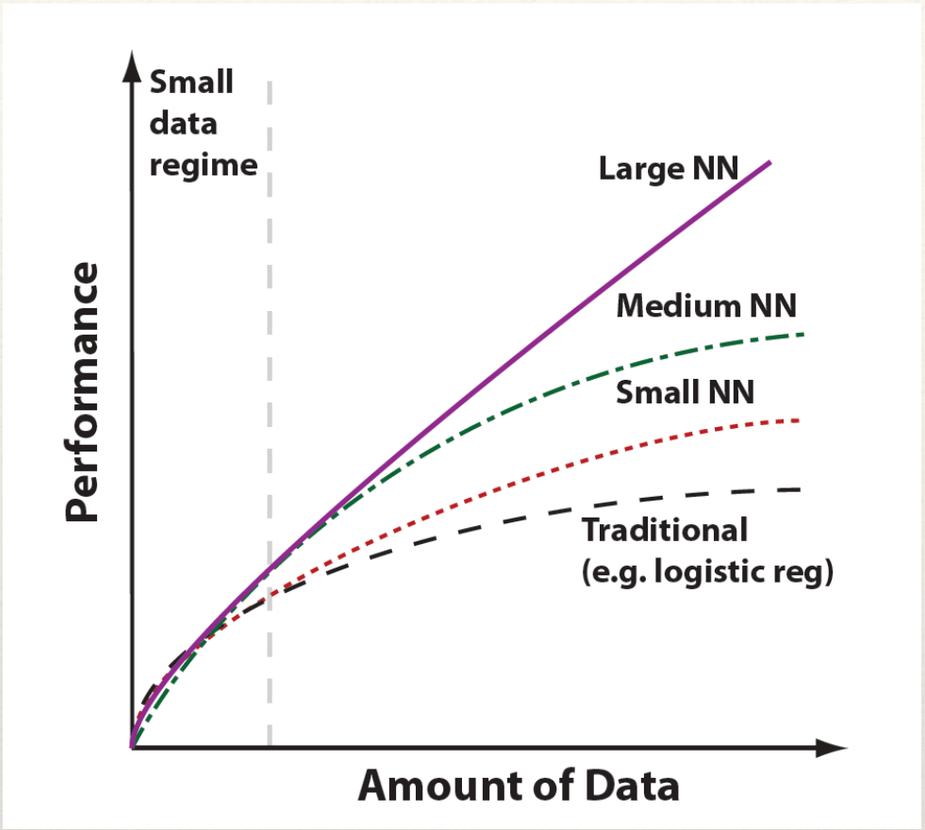

High-bias low-variance, 1803.08823

Good at handling large amounts of data: needle in a haystack

The NN structure (layers, 0/1 gates) allows a high representation power with moderate computational demands, e.g. allows parallelisation, use of GPUs…

It scales better than other learning methods (like SVMs)

Good at learning: ability to learn with little *domain knowledge*

NNs are able to take large streams of data and learn features with little guidance, work like *black boxes*

In both situations, ID and simulation,
AI comes to help in doing **better** what we
already know how to do

AI is *assisting*, enhancing the task
is not doing something radically new

Characteristic of fundamental physics:
we often have a way to tackle problems,
in our approximation to complex phenomena
we are strong followers of *reductionism*

Yet, Deep Learning seems to learn
beyond our naive expectations

Human surrender?

What's wrong with blackboxes?

Only open in case of a disaster

If it works, why fix it?
DL is very powerful, in a way that can be quantified and tensioned against human performance or other techniques

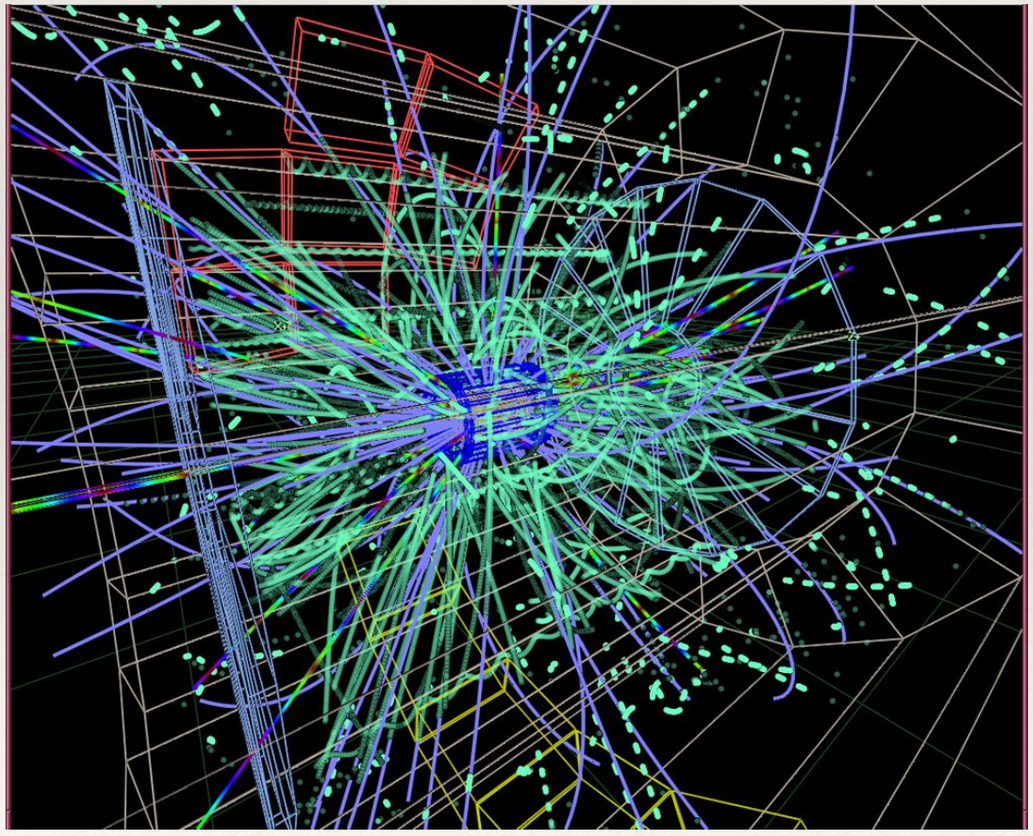

If the blackbox can help ID'ing this event faster and better than a traditional algorithm, who cares?

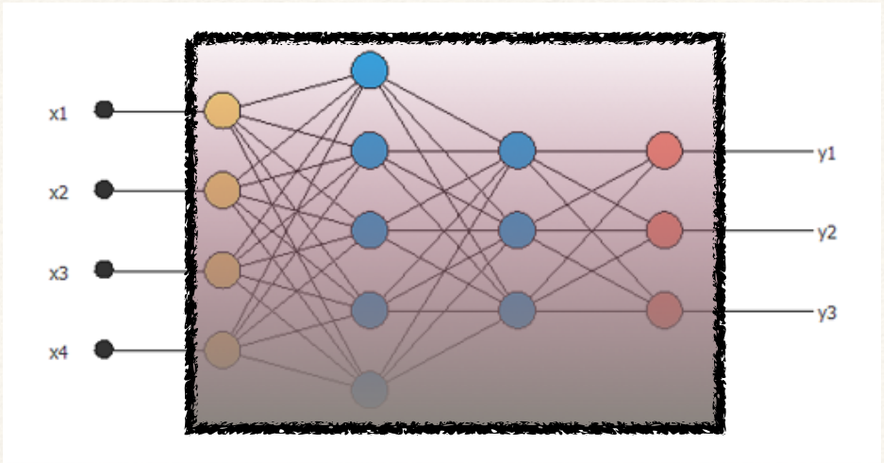

What's wrong with blackboxes?

If they do work, and help solve problems?

The lack of understanding hurts our *pride* as scientists

our job is to understand as much as we humanly can

"If you think you understand quantum mechanics, you don't understand quantum mechanics" R. Feynman, *The Character of Physical Law*

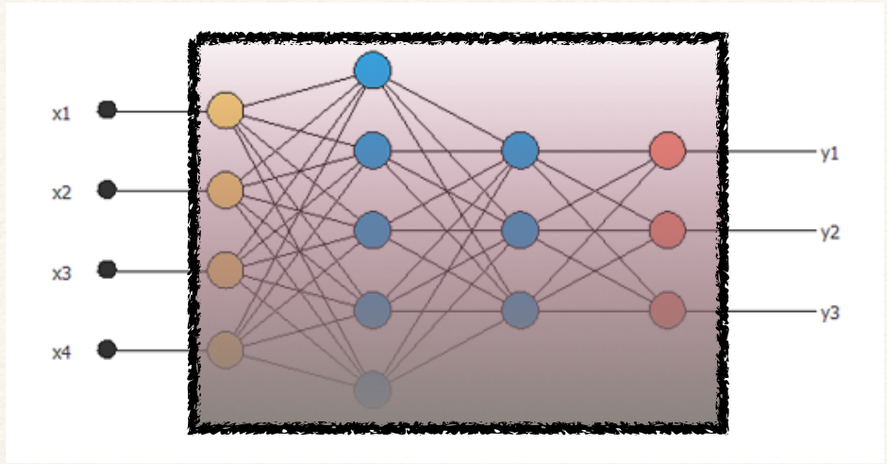

What's wrong with blackboxes?

If they do work, and help solve problems?

The lack of understanding hurts our *pride* as scientists
our job is to understand as much as we humanly can
"If you think you understand quantum mechanics, you don't understand quantum mechanics" R. Feynman, *The Character of Physical Law*

Any efforts we do to express the workings of NNs from different viewpoints may lead to *new ideas for machine learning*

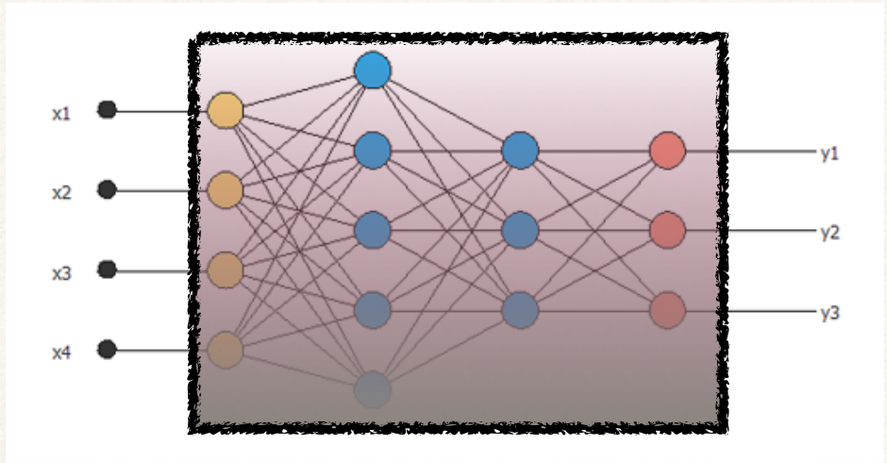

What's wrong with blackboxes?

If they do work, and help solve problems?

The lack of understanding hurts our *pride* as scientists
our job is to understand as much as we humanly can
"If you think you understand quantum mechanics, you don't understand quantum mechanics" R. Feynman, *The Character of Physical Law*

Any efforts we do to express the workings of NNs from different viewpoints may lead to *new ideas for machine learning*

The depth and reach of AI in *decision making* is growing very fast
we should be concerned about our lack of control over this
e.g. see EU efforts to regulate AI
XAI, Ethical AI… all these require a better understanding of DNNs

Looking under the hood

What if we didn't know what we were looking for?

what would **you** do?
as a physicist, you would start thinking on
possible physical relations, plotting things,
trying to obtain the best data representation
the representation which manifests a behaviour
Could an AI guide you?

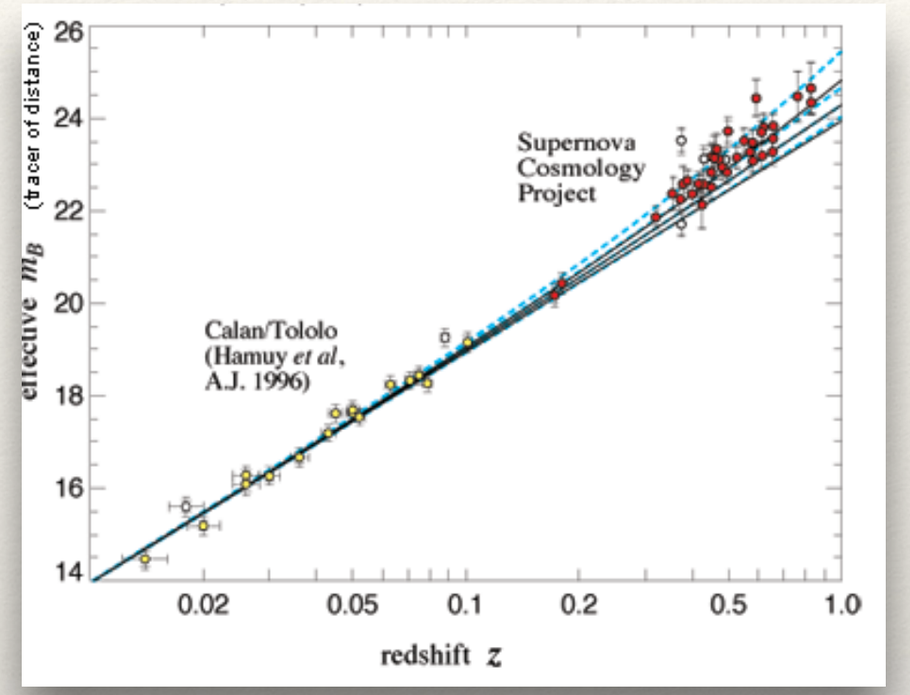

What if we didn't know what we were looking for?

what would **you** do?
as a physicist, you would start thinking on
possible physical relations, plotting things,
trying to obtain the best data representation
the representation which manifests a behaviour
Could an AI guide you?

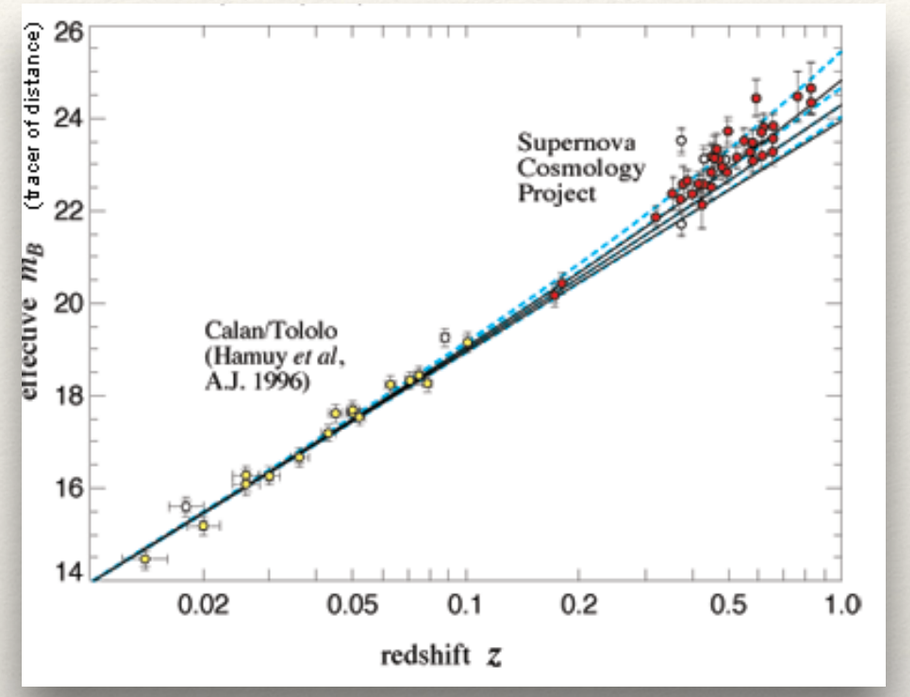

If I showed you many
examples of fluid
behaviour

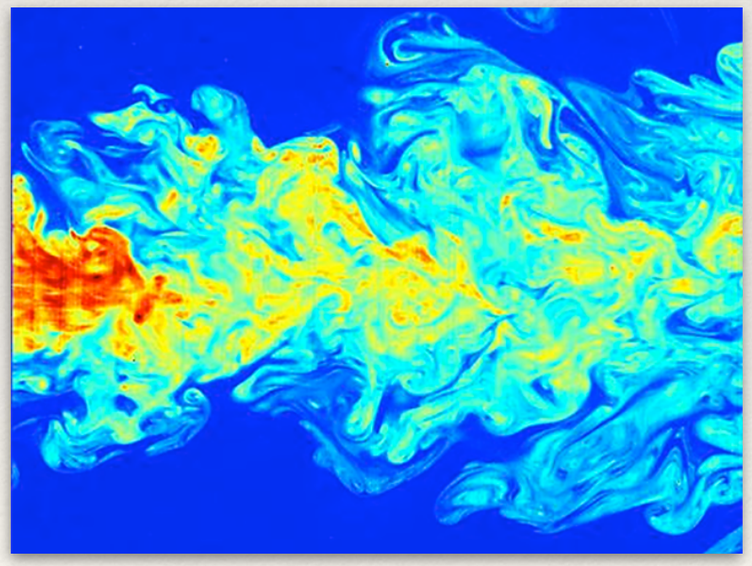

would you learn the
Navier-Stokes equation?

What if we didn't know what we were looking for?

what would **you** do?
as a physicist, you would start thinking on
possible physical relations, plotting things,
trying to obtain the best data representation
the representation which manifests a behaviour
Could an AI guide you?

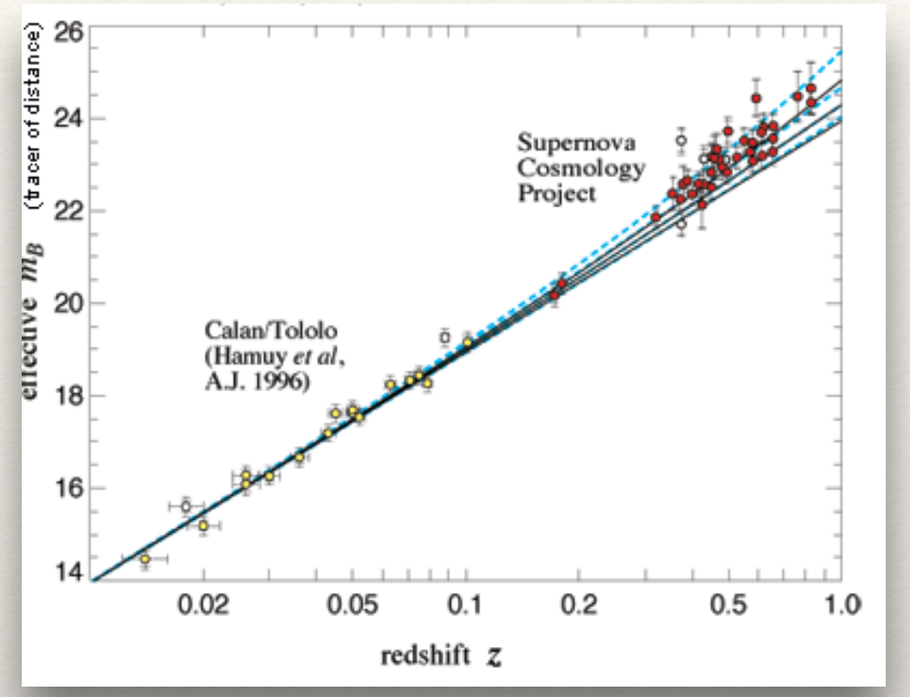

If I showed you many
examples of fluid
behaviour

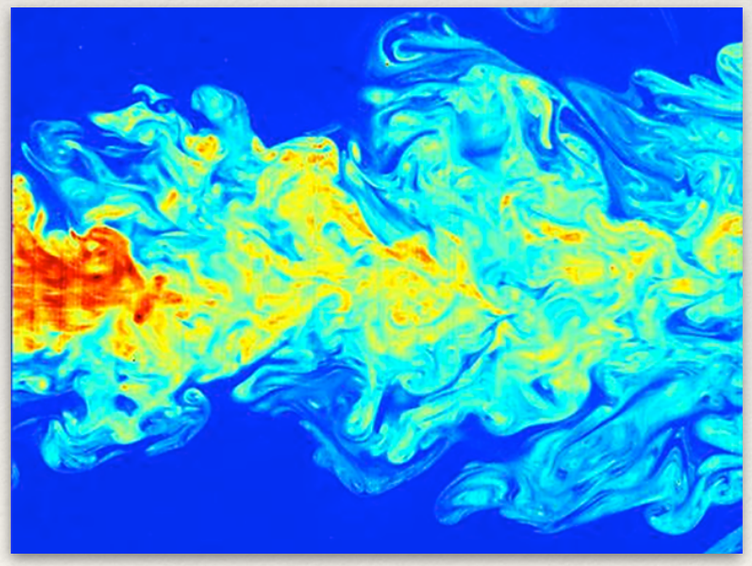

would you learn the
Navier-Stokes equation?

AI can help with that, e.g. **Symbolic Regression** (aka Eq Discovery)
or **symmetry**-ID

Variational Autoencoders

I am sure you have learned and used VAEs already, but do keep listening

VAEs: a type of GenAI architecture

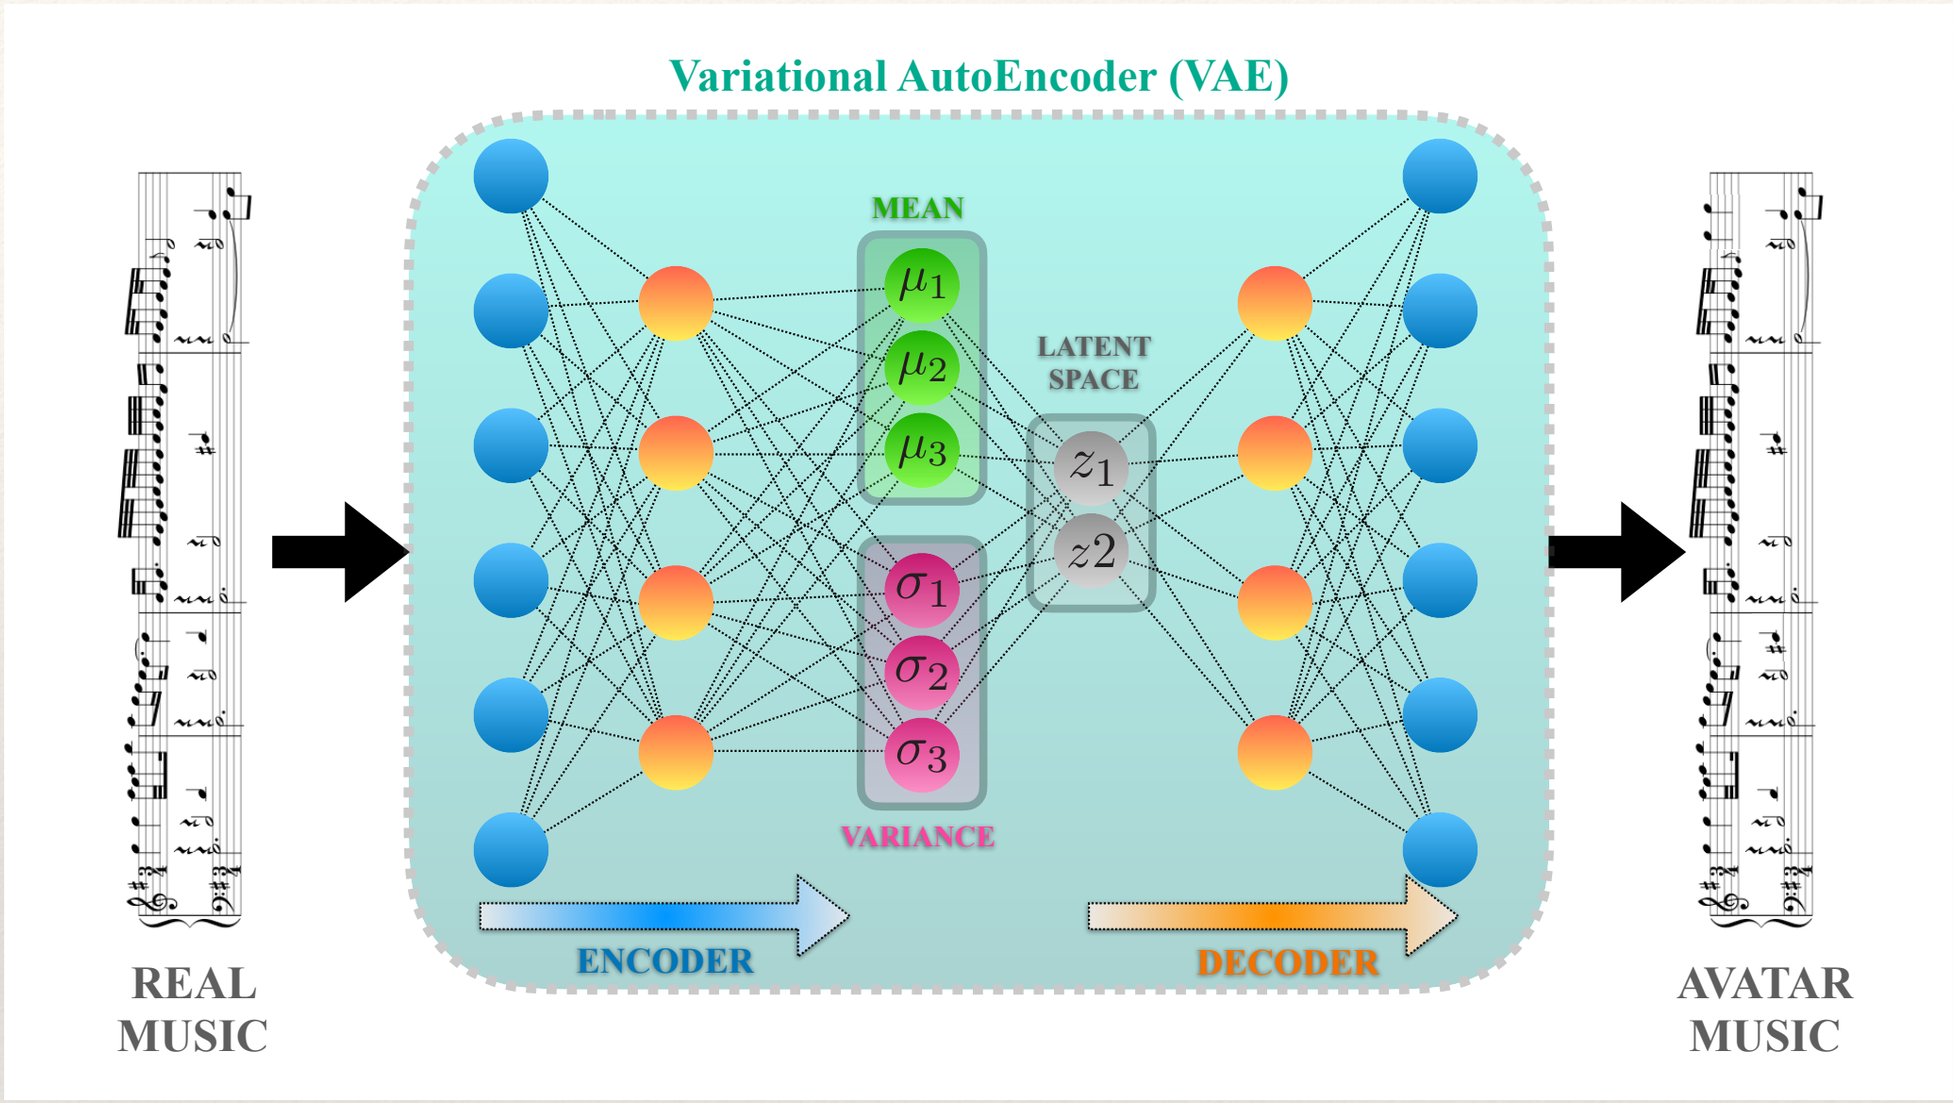

ENCODER: compresses the data, forcing higher levels of abstraction

DECODER: decompresses, creates new samples (Generative part)

Max level of compression (more abstraction)= bottleneck **LATENT SPACE**

Let's discuss how VAEs are trained

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

STEP 1 - 'LEARN' what is a human face

Take face images: x

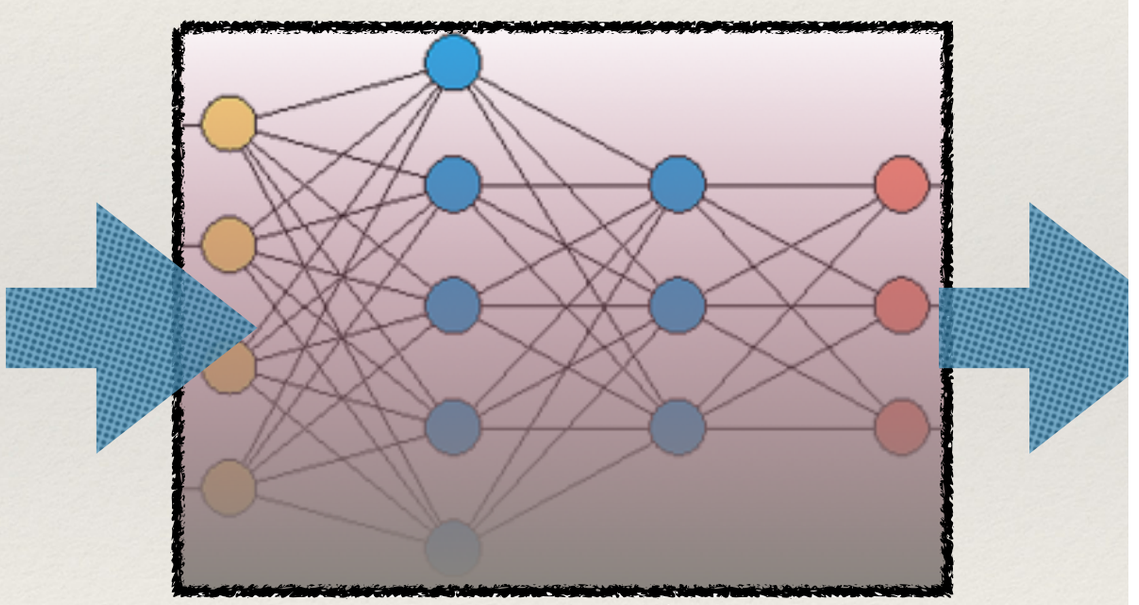

Transform them
in complicated ways

Create an avatar: x'

Doing this many times, while the DISCRIMINATOR says:

'You are going in the right direction', 'You are completely lost!'

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

STEP 2- AFTER MANY ITERATIONS…

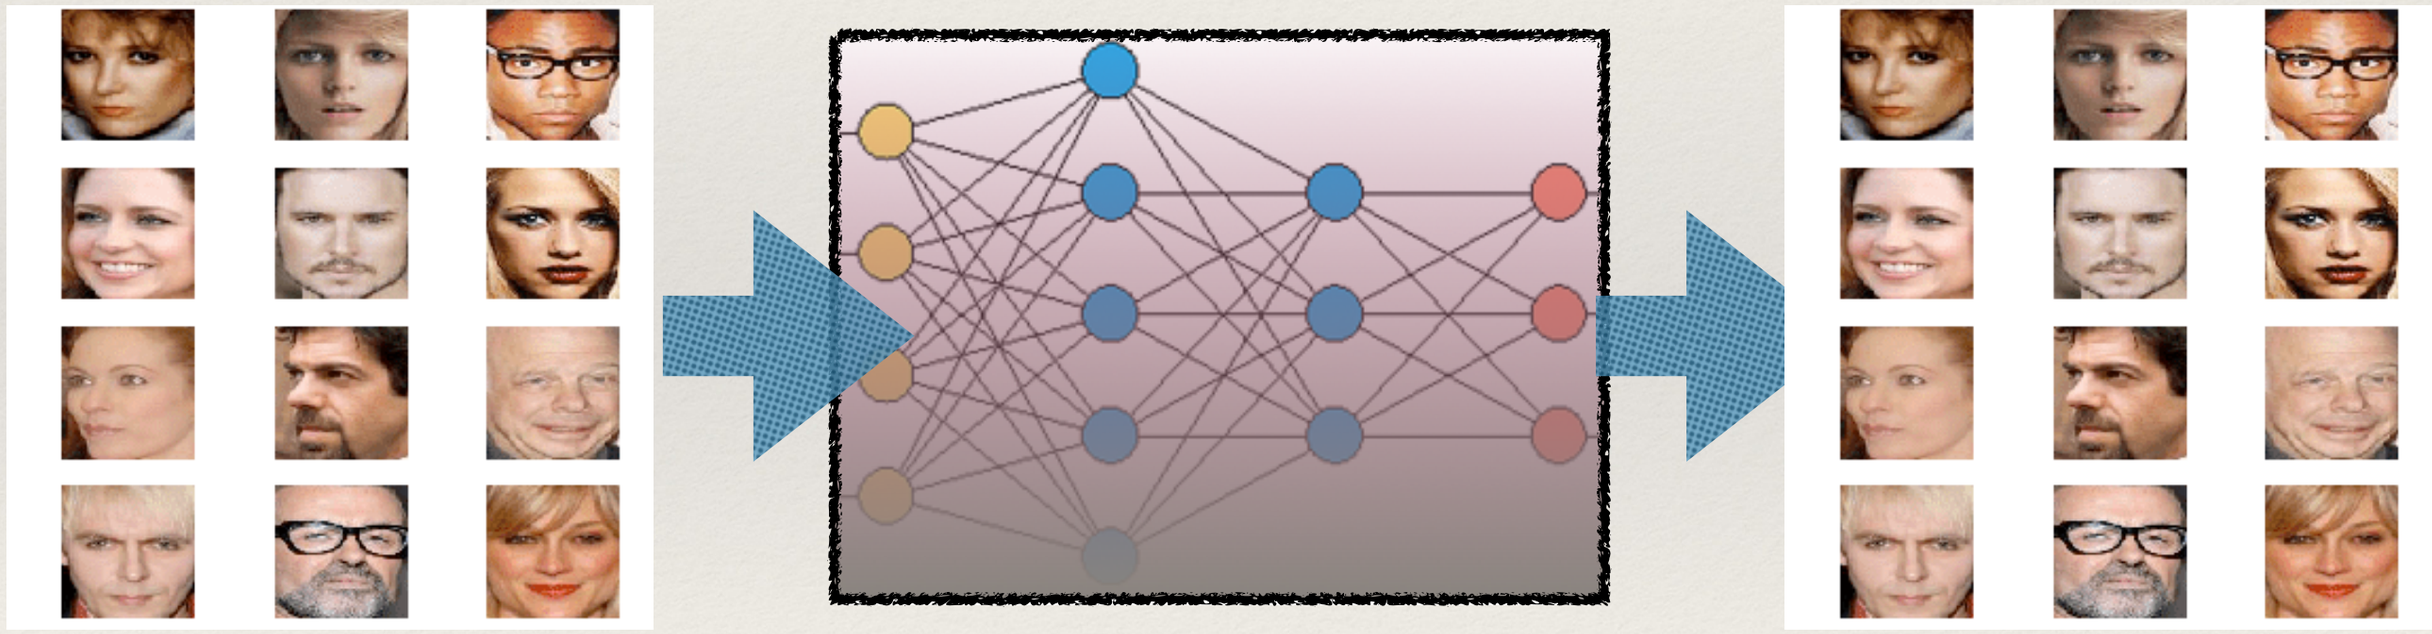

When the avatars are indistinguishable to the DISCRIMINATOR, **game is over**

Wait a minute!

Aren’t we just programming the identity transformation?

Wait a minute!

Aren’t we just programming the identity transformation?

*NNs have a huge number of parameters (**expressivity**)*
*so they could be just memorising all the faces (**overfitting**)*
Yet, we did control for that during the training

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

STEP 3- CREATE NEW POSSIBILITIES

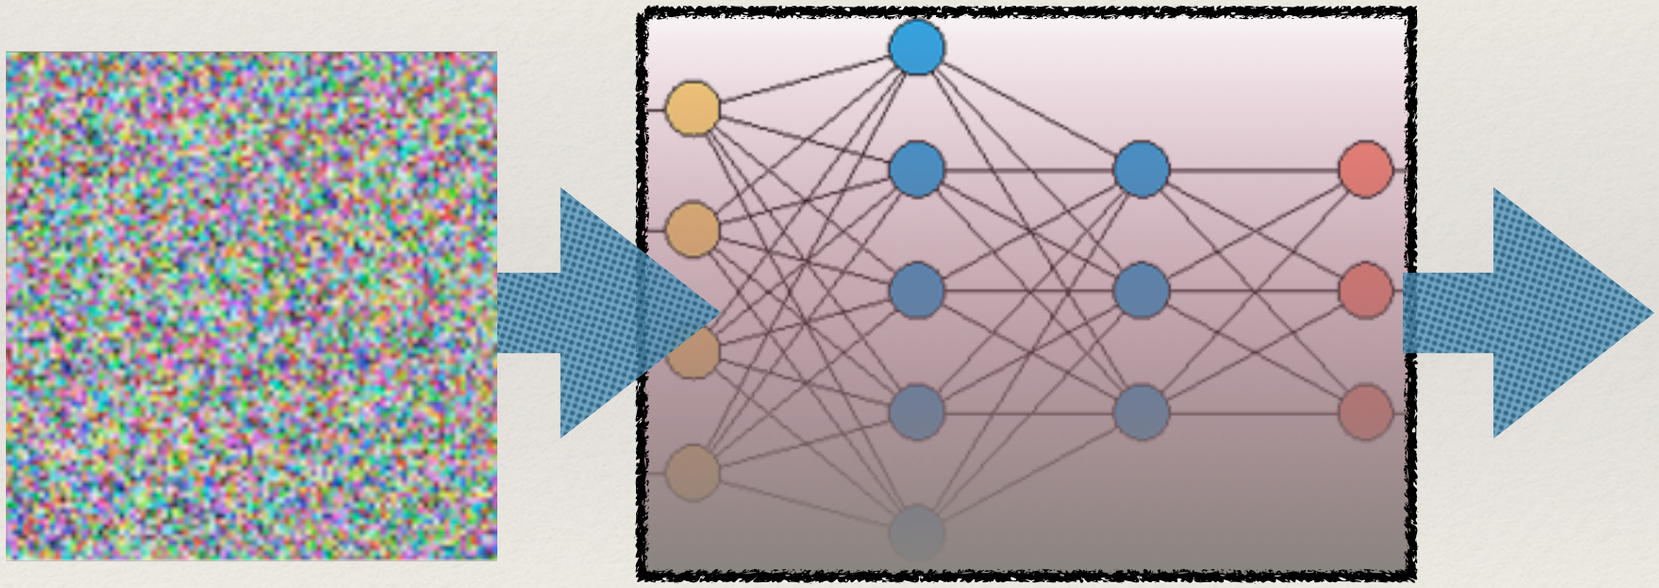

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

STEP 3- CREATE NEW POSSIBILITIES

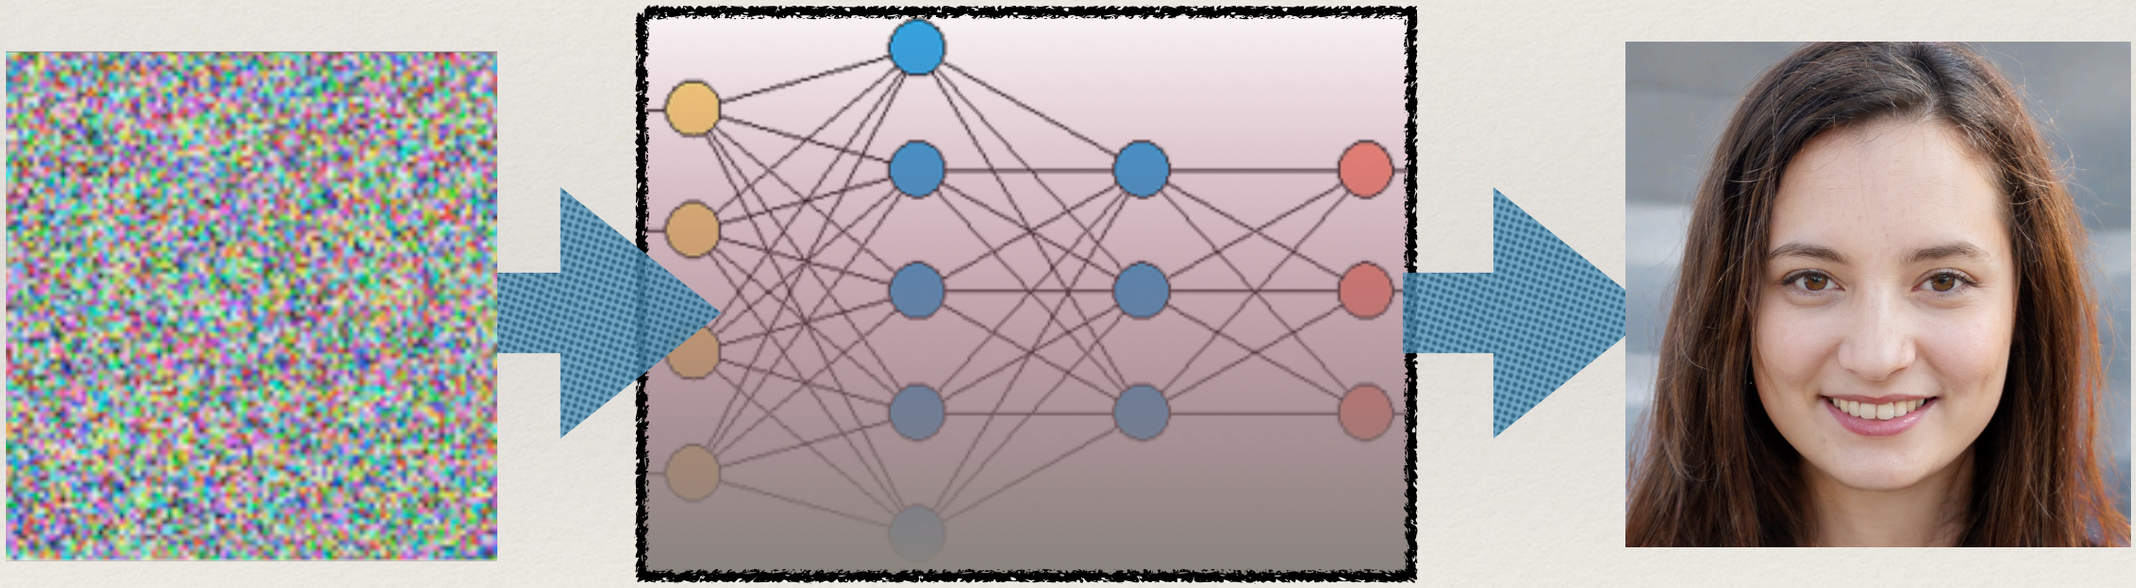

This woman **does not exist**. It has been generated from noise.

The NN has learnt the *concept* of 'human face' and now can create human faces from noise

What if we didn't ask for an outcome?

Supervised learning input-> predict output

what if we just asked 'look at this!' with no determined output?

GANs (Generative Adversarial Networks)

and **VAEs (Variational AutoEncoders)**

In CNNs, benchmarks were cats/dogs and hand-written digits (MNIST)

Here, human faces

STEP 3- CREATE NEW POSSIBILITIES

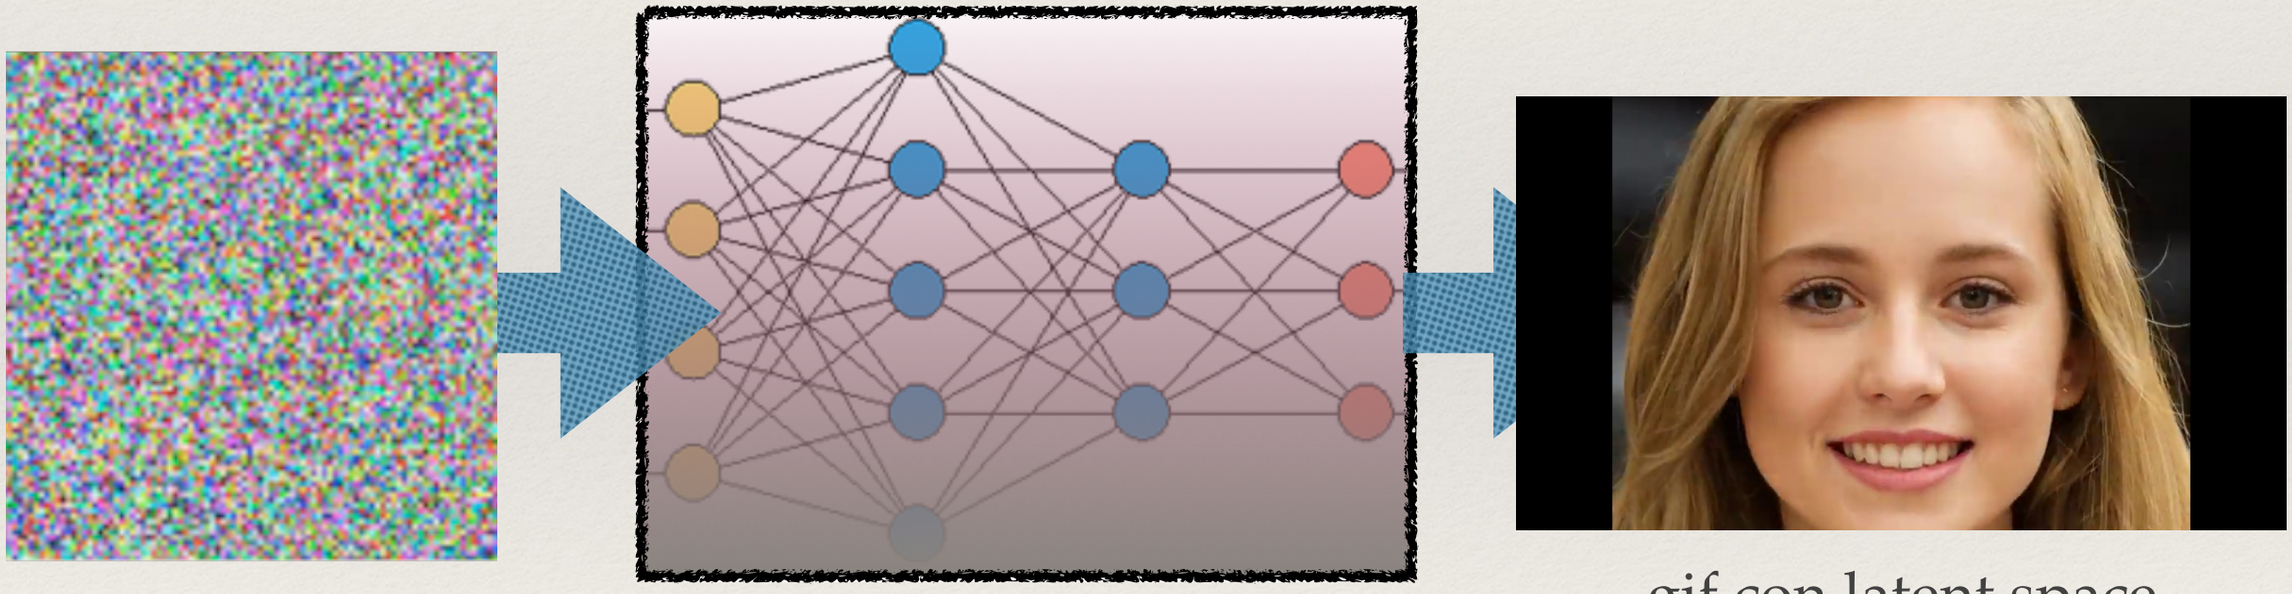

gif con latent space

Random noise generate deformations in the output,

leading to new *people*

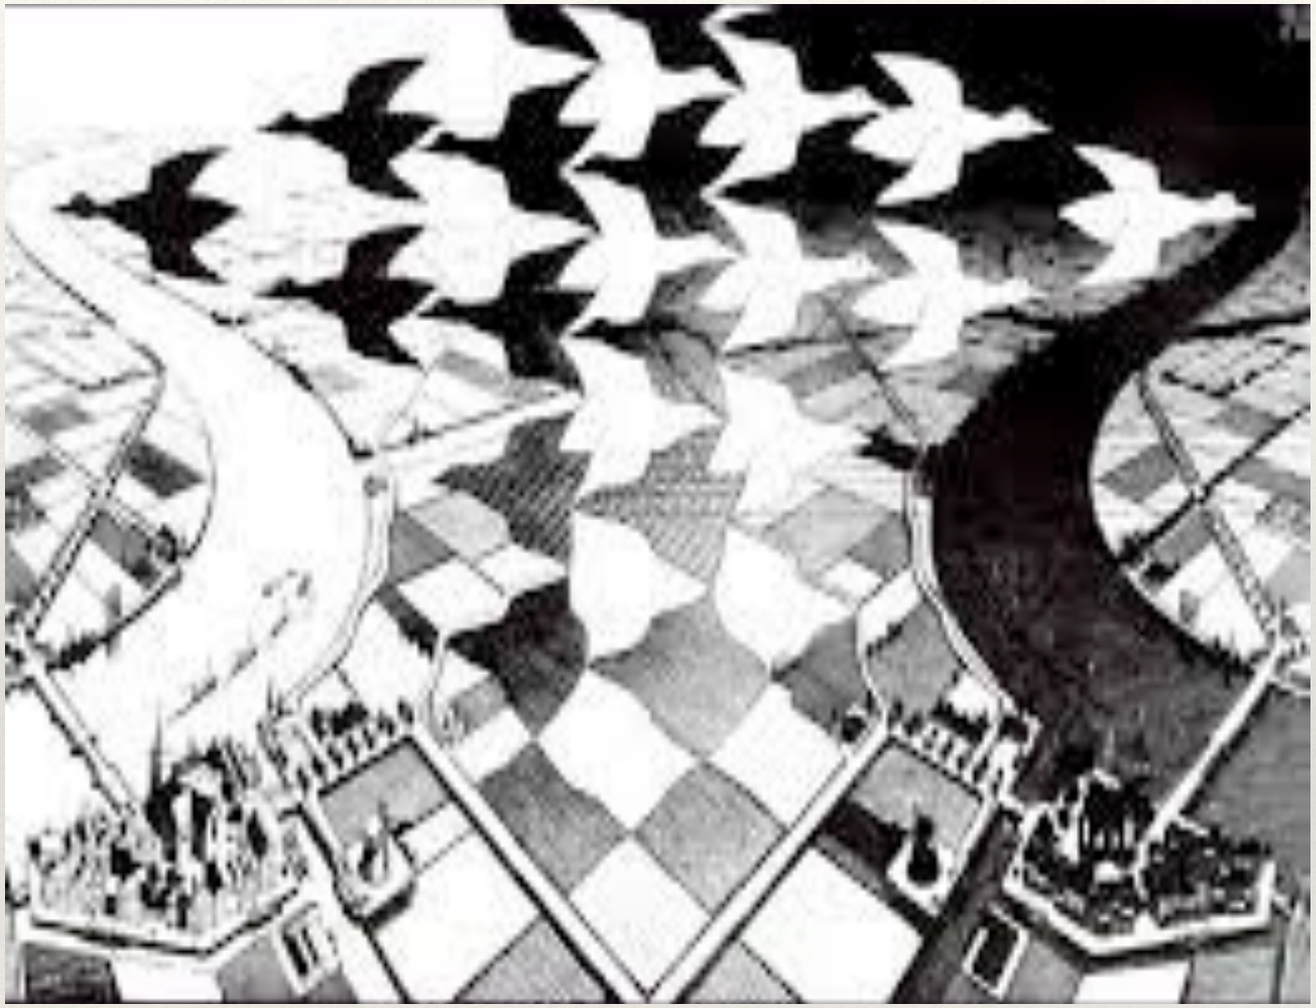

Symmetries

Can NNs learn **concepts** like symmetry?
We just saw how it *learned* the concept of human face

A symmetry is a redundancy, so the AI should learn it
to become more efficient at describing the data

The initial question is **how** do we interrogate a
blackbox and learn if it has learned the symmetry

Then we can continue by testing if it can learn
symmetries even if **approximate**

Ultimately, the idea would be to let the AI find
principles for us

Let's start with something very simple

Discovering physical concepts with NNs, 1807.10300 by Iten et al

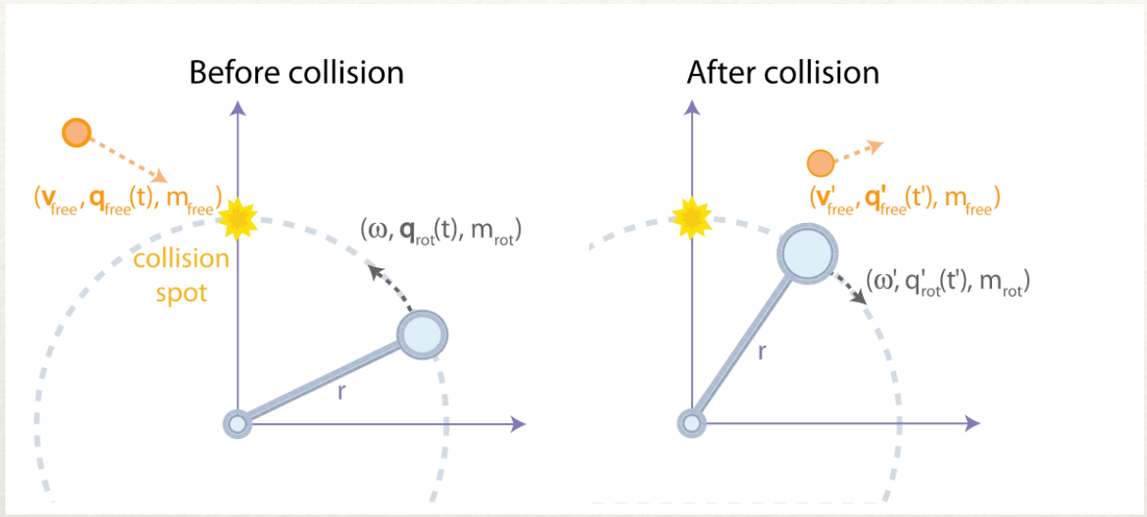

Classical mechanics:

Create a dataset of the characteristics of this type of collision

Train a VAE with the data

Reduce the latent space to a single neuron

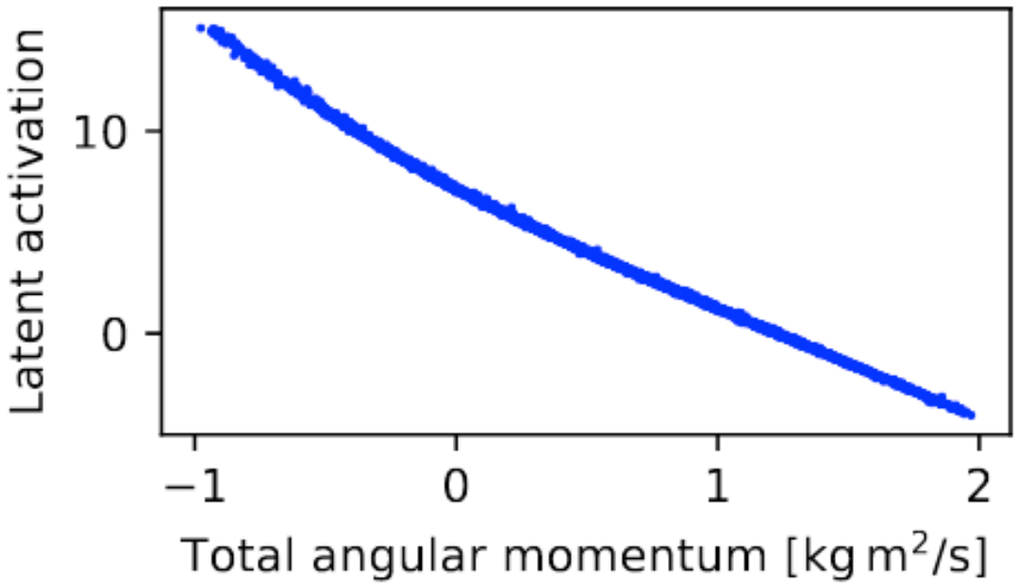

Interrogate the VAE= find a strong correlation between the neuron activation and the system

They showed that NNs were storing information of the angular momentum

The size of the latent activation was related to total angular momentum

Let's do something a bit more complicated

Learning Symmetries in datasets, 2504.05174 by VS

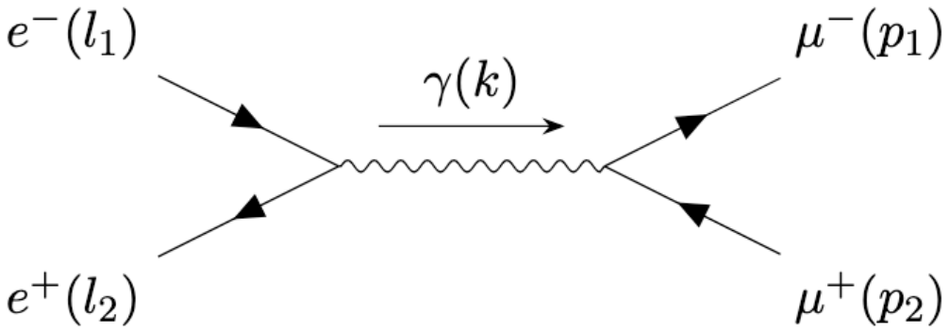

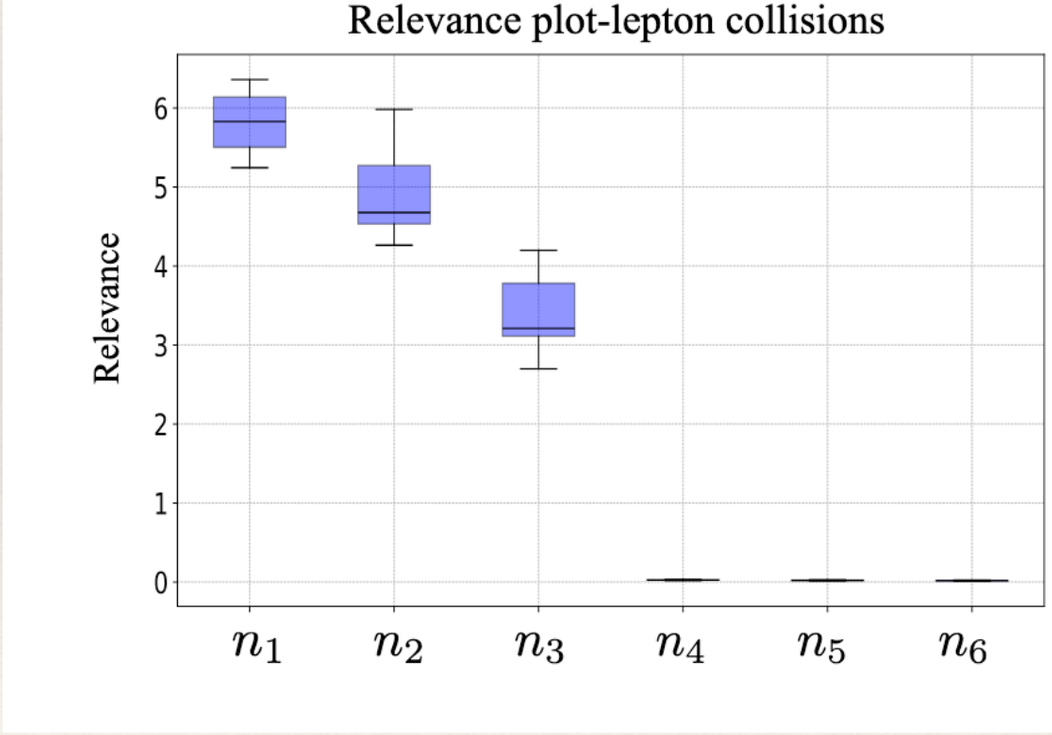

QED:

Create a dataset with the muons 4vectors

Train a VAE with the data

Interrogate the VAE

A bit more complicated than before

1.) What is the perceived dimensionality of the system?

$$\text{Relevance}_j = \frac{\text{std}(\langle z_j \rangle)}{\text{mean}(\sigma_j)} \ ,$$

From *Exploring how a GenAI understands music—> Tomorrow*

Input dim is 8, VAE sees 3 =>

there must be redundancies/symmetries

Let's do something a bit more complicated

Learning Symmetries in datasets, 2504.05174 by VS

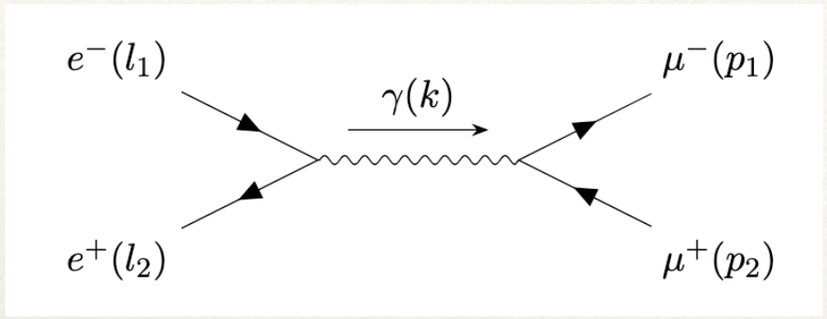

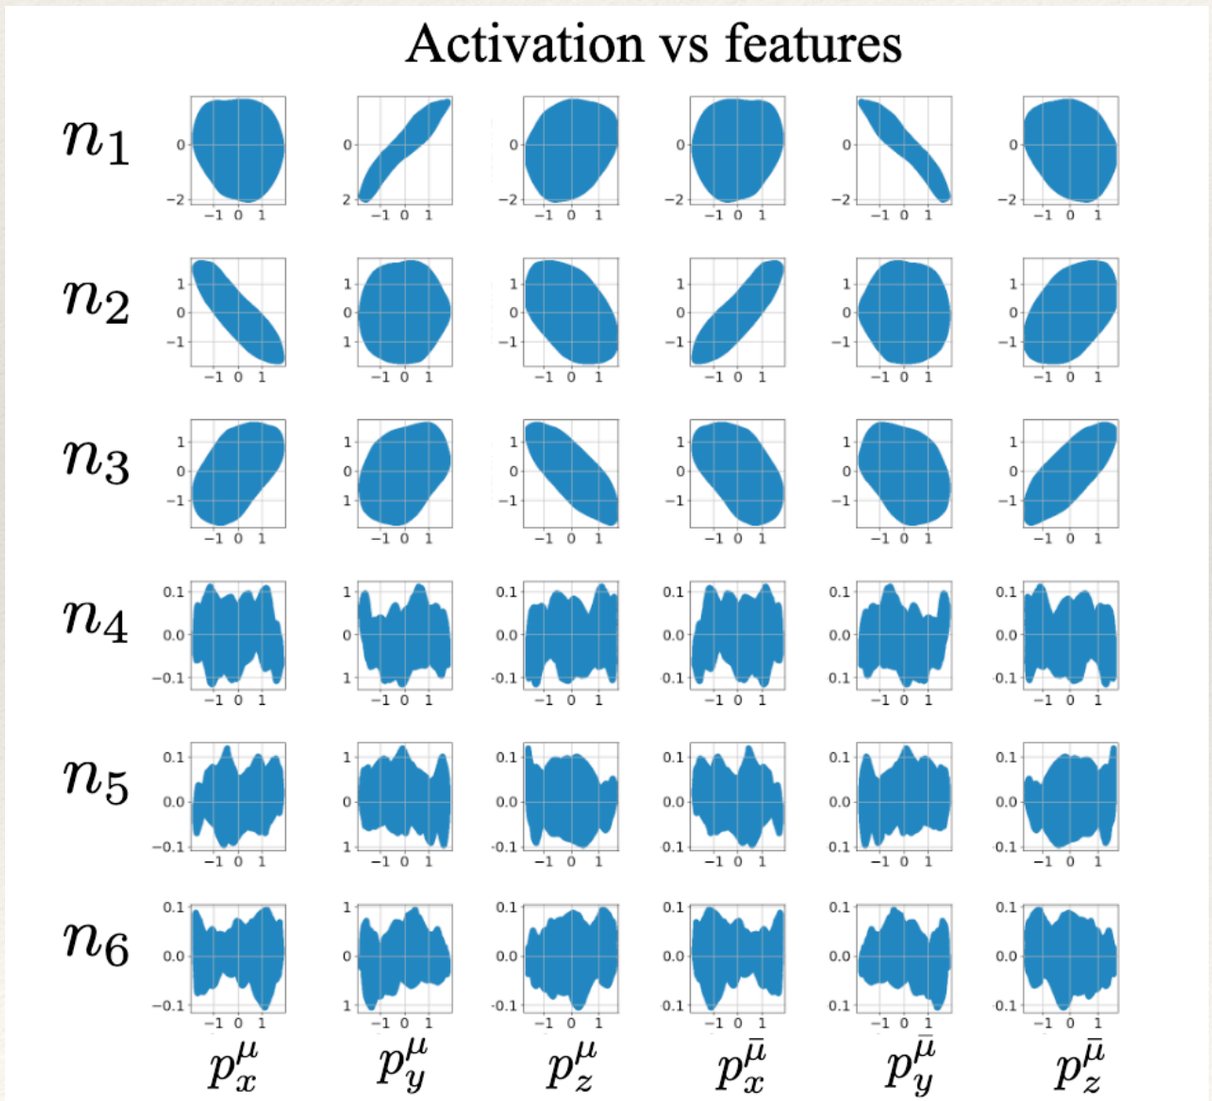

2.) What did it learn?

In these collisions, energy-momentum is conserved
Energy fixed, 3D momentum conservation
So we go from **8** to 8-2-3=**3** quantities

Momentum conservation

$$p_x^\mu + p_x^{\bar{\mu}} = 0 \ ,$$
$$p_y^\mu + p_y^{\bar{\mu}} = 0 \ ,$$
$$p_z^\mu + p_z^{\bar{\mu}} = 0 \ .$$

Have been **discovered** and **encoded** in the three relevant neurons, leaving the rest as **generative**

$$\langle z_2 \rangle \sim p_x^\mu - p_x^{\bar{\mu}} \ ,$$
$$\langle z_1 \rangle \sim p_y^\mu - p_y^{\bar{\mu}} \ ,$$
$$\langle z_3 \rangle \sim p_z^\mu - p_z^{\bar{\mu}} \ .$$

Let's see if this stands with approx symms

Symmetry meets AI, SciPost Phys (2021) by Barenboim, Hirn, VS

1.) Create a dataset of 2D images of potentials with symmetries (or not)

For example

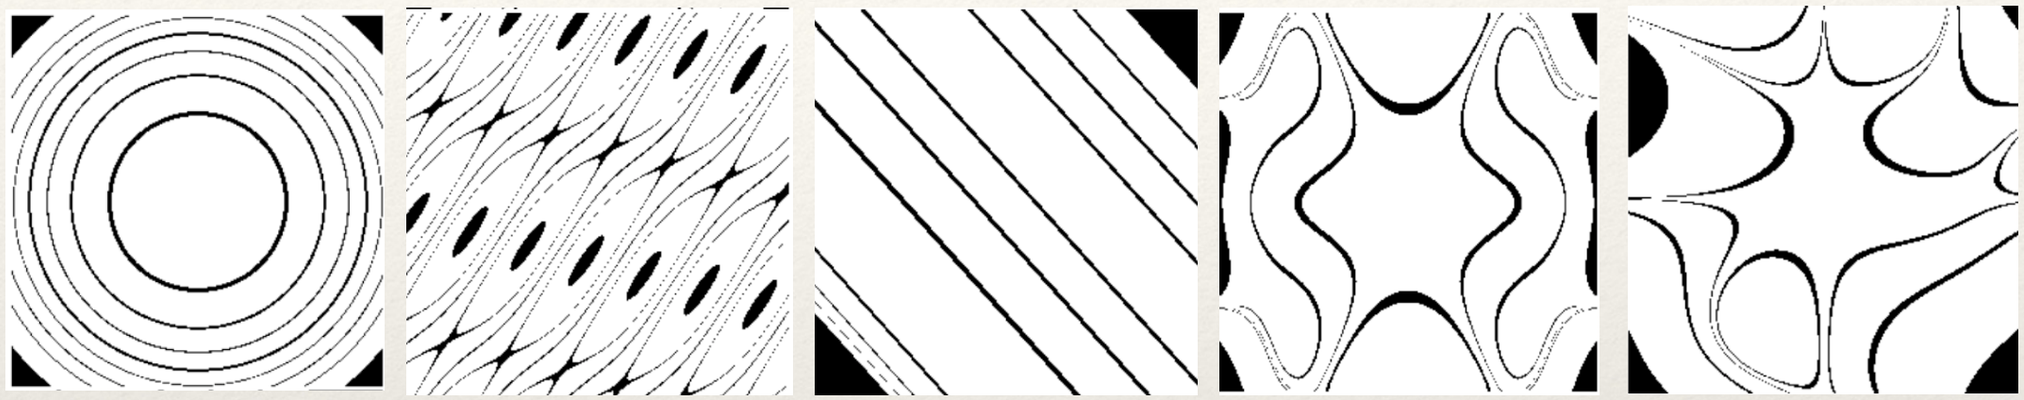

Rotation	Discrete translation	Translation	Reflection	Nothing
$O(2)$	T_n	T	Z_2	$\emptyset$

2.) Train this set in a decoy task (again like an identity) and inspect the last hidden layer

3.) With the true labels, train a **symmetry detector**

4.) Apply it to art, see if it can handle approx symmetries

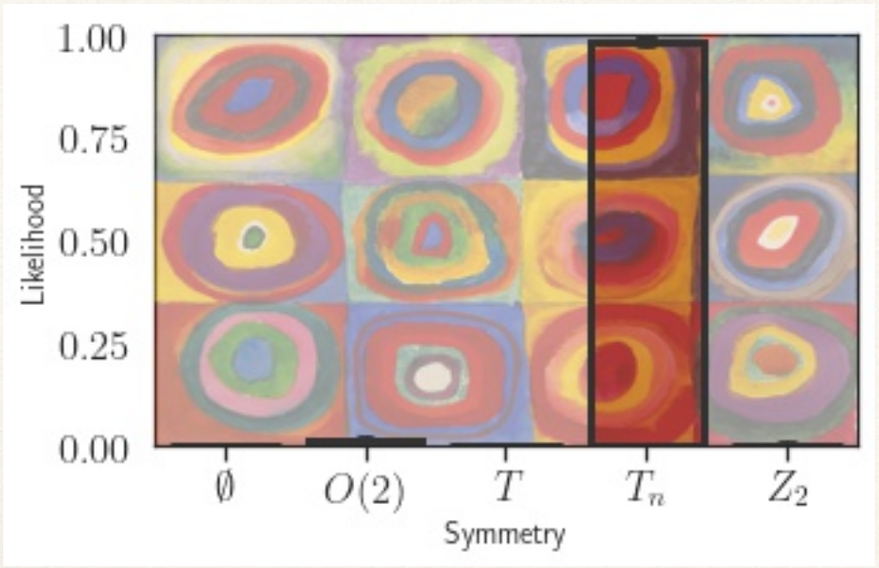

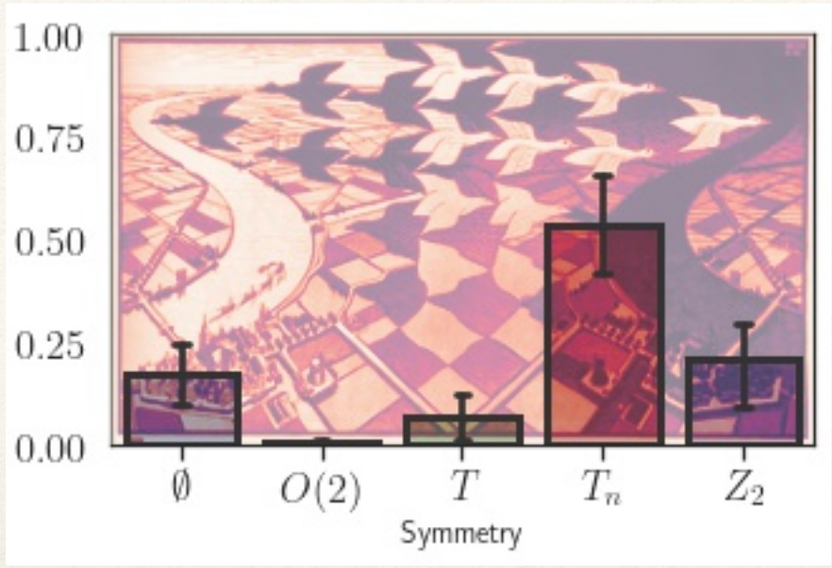

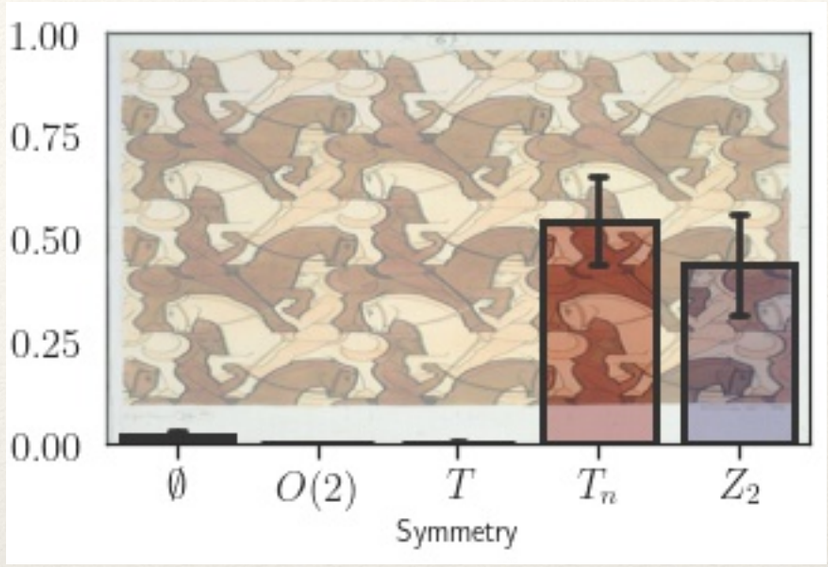

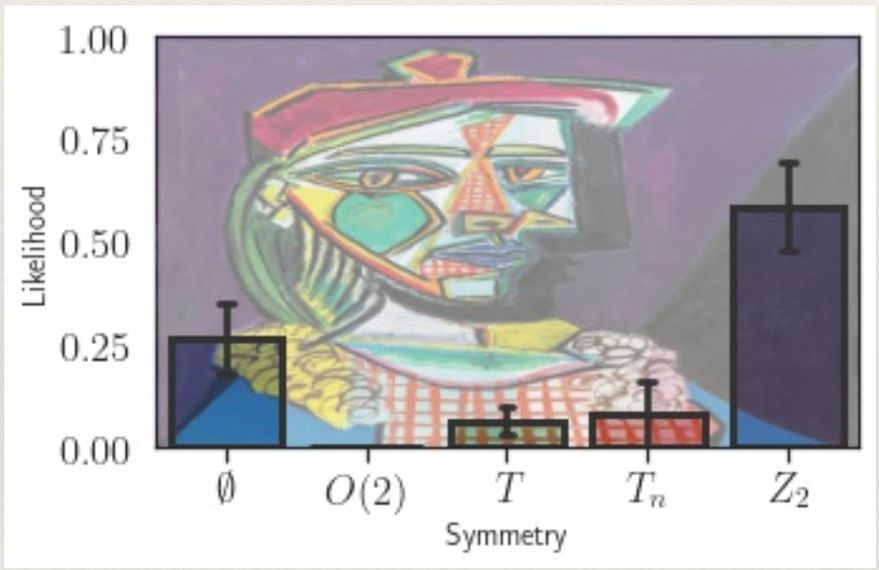

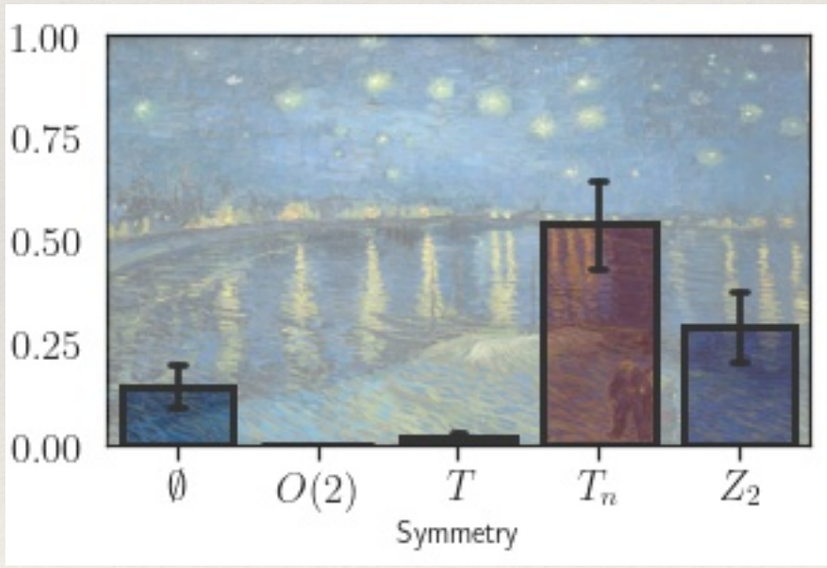

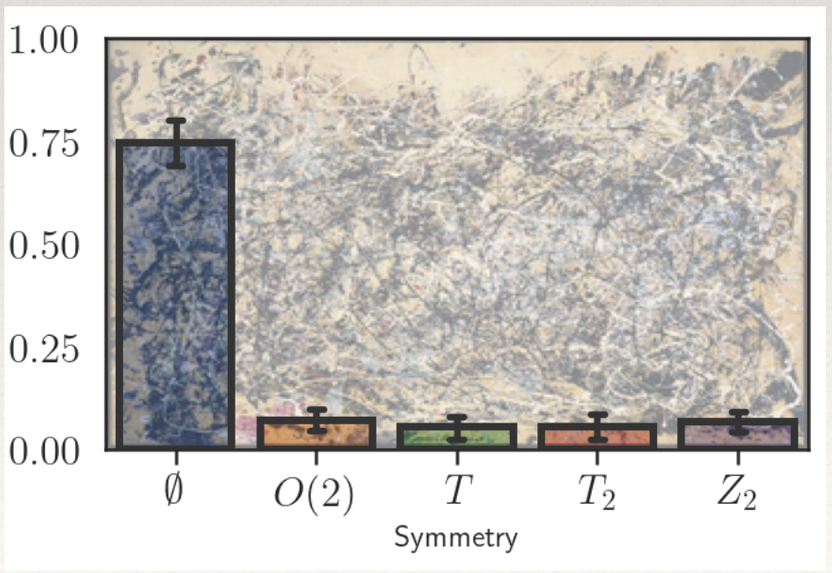

and many more, including children's drawings, fractals, photographs etc

Beyond fundamental physics

Understanding, R+D, precision

In fundamental physics
we ALWAYS have a firm ground to step from
A theoretical framework: SM, LambdaCDM, mean field, perturbation theory in QFT…

We say *terra incognita* or similar,
but we actually have a **range** of possibilities in mind

Understanding, R+D, precision

In fundamental physics
we ALWAYS have a firm ground to step from
A theoretical framework: SM, LambdaCDM, mean field, perturbation theory in QFT…

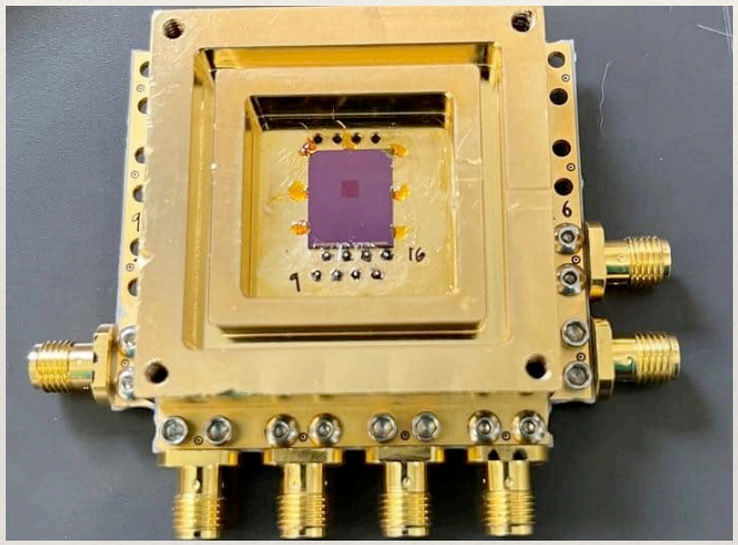

Single photon det. SMSPD. Fermilab

+some physical understanding of the observable environment: we build our purpose-made instruments, experiments are reproducible, measurements can be correlated…

Understanding, R+D, precision

In fundamental physics
we ALWAYS have a firm ground to step from
A theoretical framework: SM, LambdaCDM, mean field, perturbation theory in QFT…

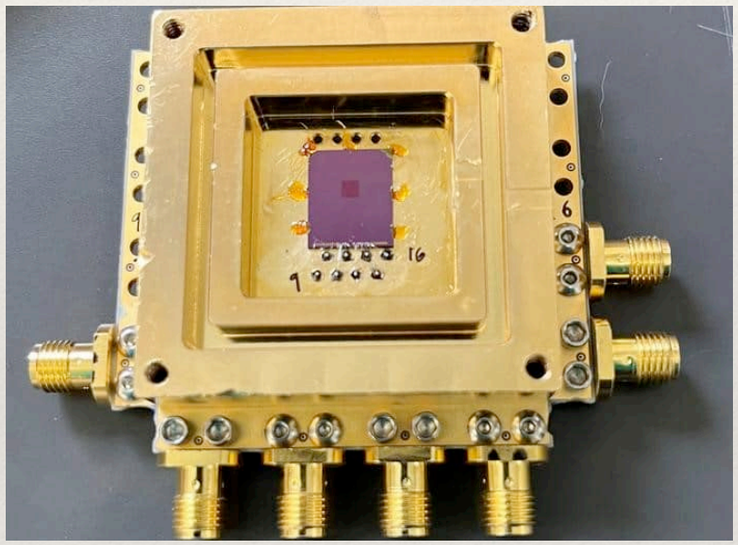

Single photon det. SMSPD. Fermilab

+some physical understanding of the observable environment: we build our purpose-made instruments, experiments are reproducible, measurements can be correlated…

Very often cross-checks are possible, limiting cases are known
Improvements can be done incrementally
Precision is an achievable goal

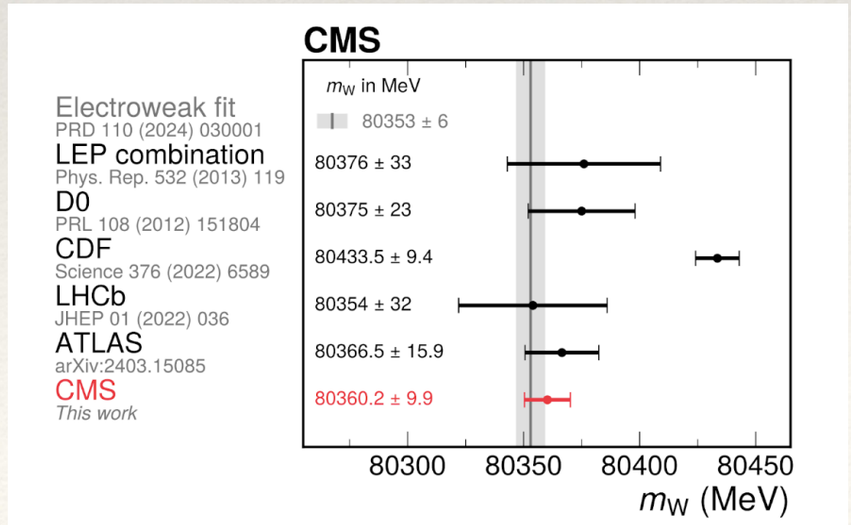

Understanding, R+D, precision

When attacking a problem

CMS's W to lepton-neutrino, electron tracking

we develop frameworks that allow us to correlate with others

EWPTs global fit, GEANT4 electron's behaviour in a material

progress in one corner of the data often illuminates other areas

Understanding, R+D, precision

When attacking a problem

CMS's W to lepton-neutrino, electron tracking

we develop frameworks that allow us to correlate with others

EWPTs global fit, GEANT4 electron's behaviour in a material

progress in one corner of the data often illuminates other areas

Robust tools for DAQ, simulation and theoretical interpretation

MC@NLO, GEANT4…

These capabilities are distinctive to us

However, when engaging with other scientific domains

fields equally rich in data and eager for deeper understanding

this mindset must evolve

Understanding, period

*Learn unknown patterns *e.g. beginning solar storm*

*Discover relations, equations

e.g. competition/cooperation among species

*Find more abstract, simplifying description of a complex system *e.g. find the political axes*

* Evolution of a situation *e.g. predict a traffic jam* & robustify your prediction *e.g. set an alarm threshold*

Understanding, period

*Learn unknown patterns *e.g. beginning solar storm*

*Discover relations, equations

e.g. competition/cooperation among species

*Find more abstract, simplifying description of a complex system *e.g. find the political axes*

* Evolution of a situation *e.g. predict a traffic jam* & robustify your prediction *e.g. set an alarm threshold*

But outside fundamental physics, *data* may be more sparse, w/ genuinely more variance, bunch of unknown unknowns, prone to glitches… data prep and viz are more important
Even the target variables may not be clear

Let's discuss examples of all this

Understanding ecological interactions with AI

Based on papers in
Methods in Ecology and Evolution (2022),
Applications of Generative AI (2024),
Ecological Informatics (2024),
Journal of Applied Ecology (2025)

The beginning

CIDE Centro de Investigaciones Sobre Desertificación

Discussing with colleagues in the Research Center for Desertification in Valencia possible multidisciplinary projects

Painfully collected a dataset of plant species in arid regions in the Mediterranean (typical landscape)

Done by an expert human, cost is huge

Hope was to use this dataset (from the natural environment) to help restoration efforts

Restoring the ecosystem in an arid region is complex, conditions are harsh and policy had been to plant pines for exploitation (bad idea)

Hypothesis:
phylogenetic closeness should impact restoration

The beginning - Petrer

The dataset was grouped as a patch (a small area on the ground) with annotations of which species where present, number of them, and which was the nurse

We spent a lot of time with cleaning and viz tasks

Rough ideas learn what makes species **cooperate** or **compete**, under what circumstances, what creates more **diversity**

The beginning - Petrer

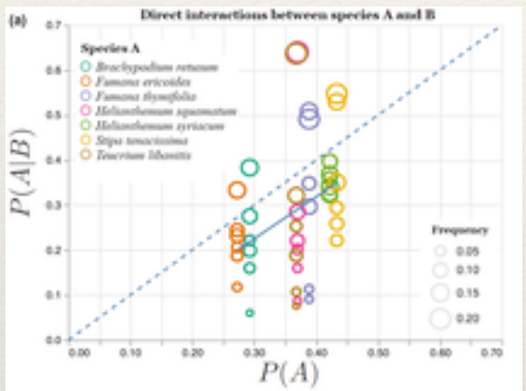

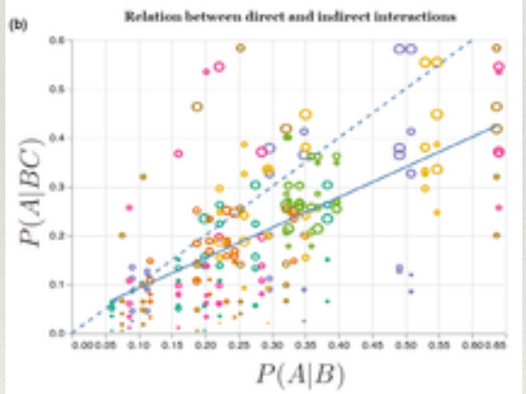

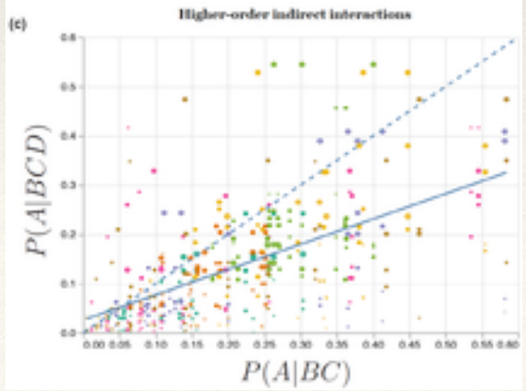

The dataset was grouped as a patch (a small area on the ground) with annotations of which species where present, number of them, and which was the nurse

We spent a lot of time with cleaning and viz tasks

Rough ideas learn what makes species **cooperate** or **compete**, under what circumstances, what creates more **diversity**

Obstacle was lack of stats

just few thousands of examples for dozens of species

limitation **intrinsic** to the study

We also knew that **interactions were complex**

We toyed with many statistical and AI techniques recommendation systems, LLMs, correlation+kurtosis etc searching for a clear question and a clear answer

Setting the problem

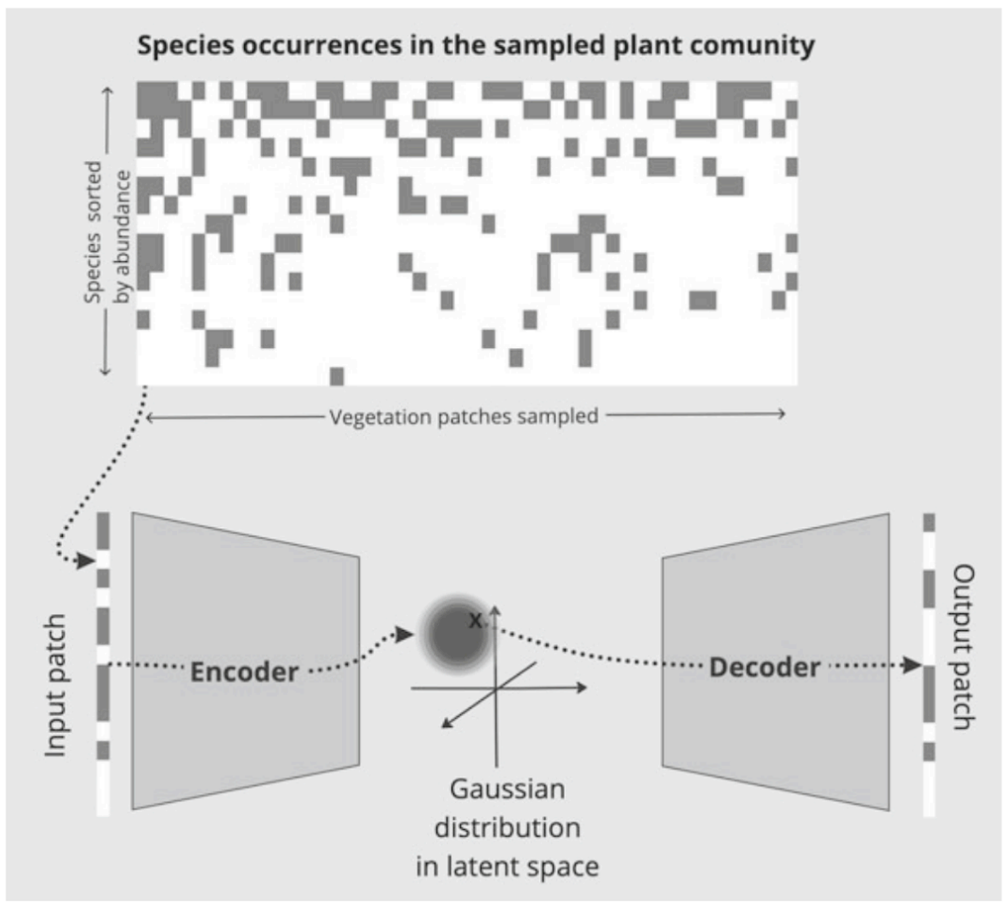

We settled to represent the data as absence/presence in a patch

We trained (successfully) a GAN and VAE with varying number of species 8-16-32 most common species

Aim to get a model of the species interactions that can interpolate in the regions with limited stats

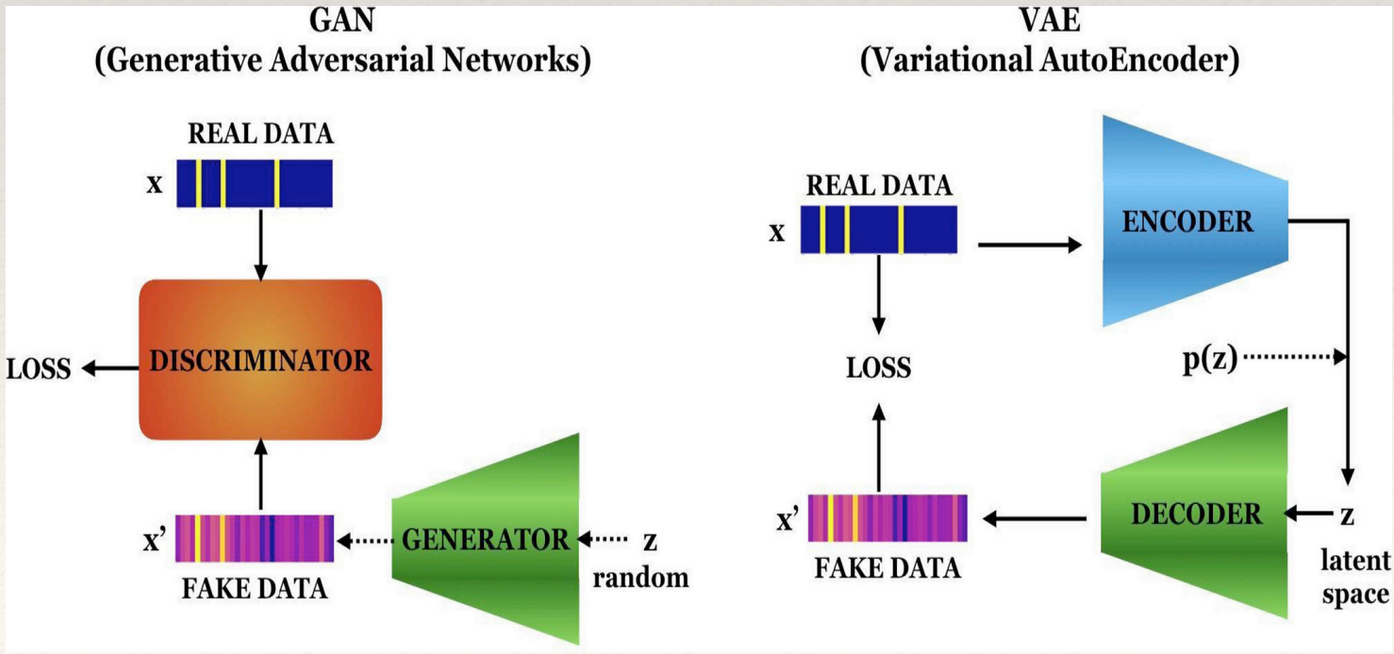

Results on Petrer

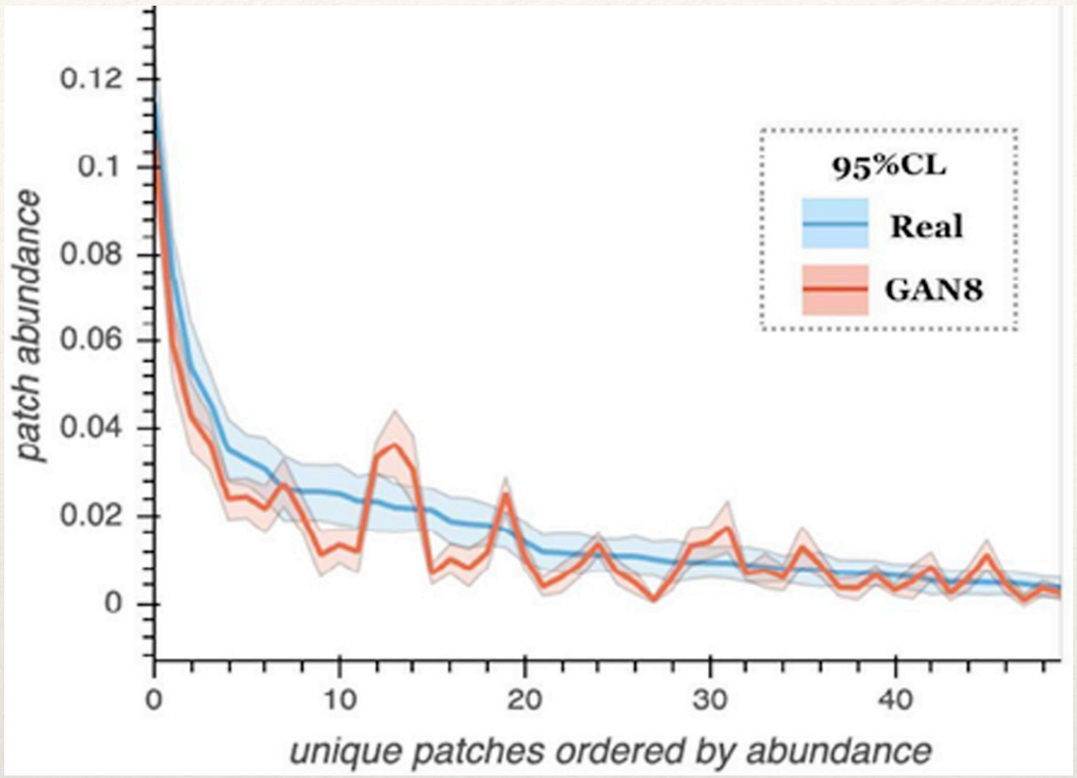

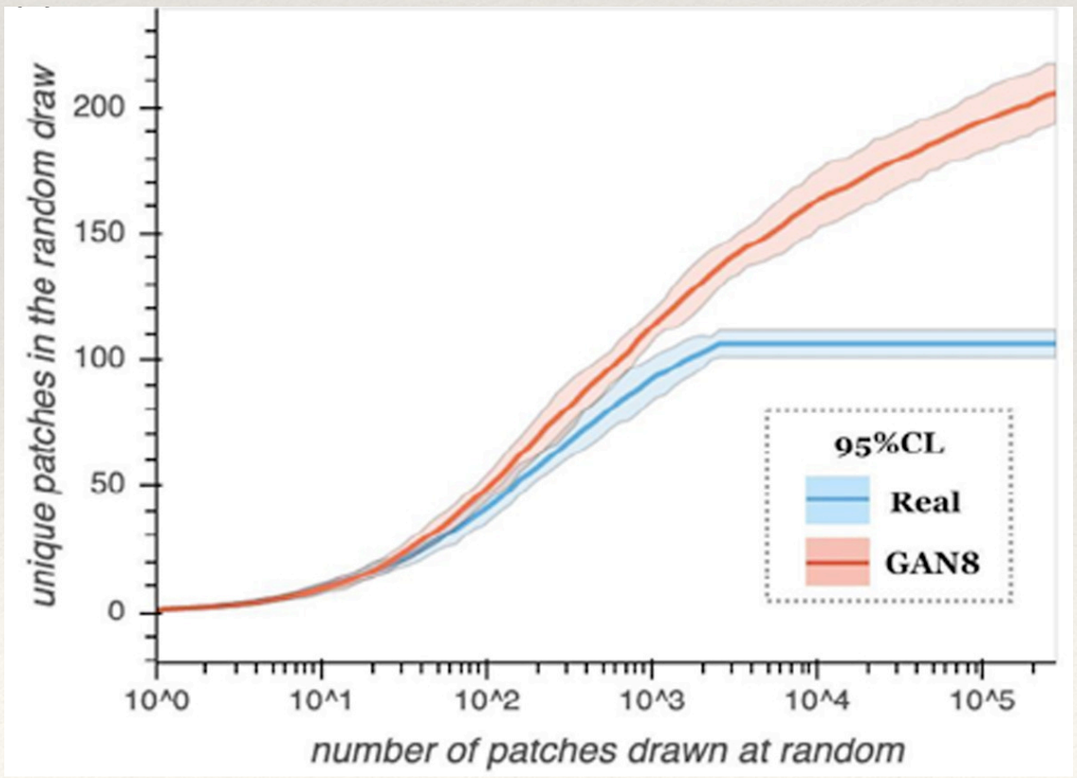

Once you have a GenAI model trained on data, you can start sampling from it and see how far it goes

The GenAI model was reproducing the typical distribution in the natural environment
e.g. had learned gypsum terrain
and we could refute the phylogenia hypothesis

And could be extrapolated to regions where there was no data in a continuous way

Results on Petrer

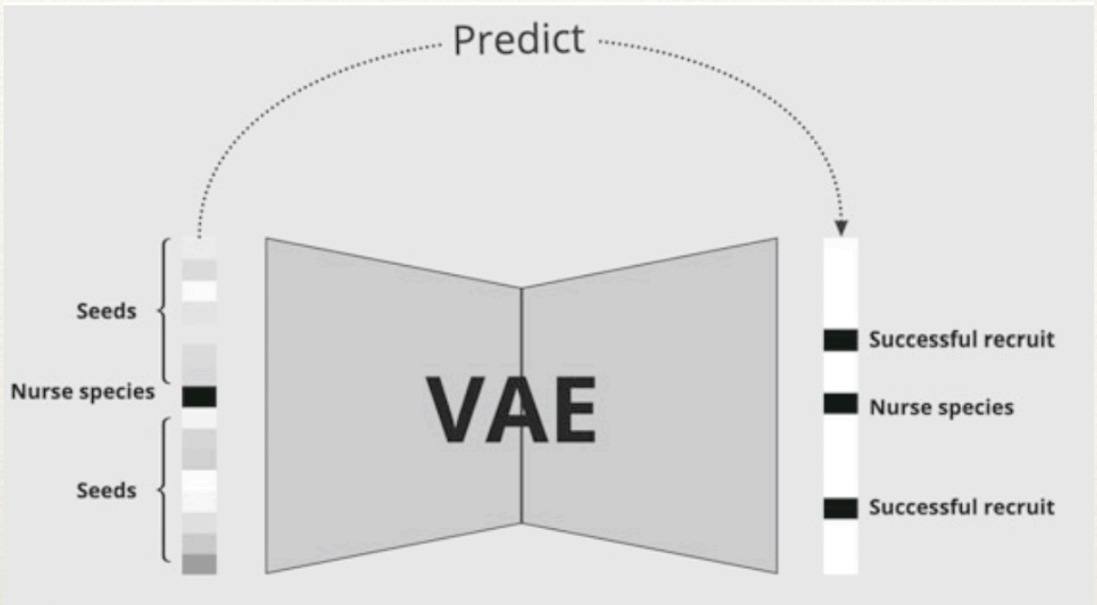

With the GenAI model we can ask restoration questions such as **What seed would produce more diverse patches?**

The VAE can be used to test pioneer species by introducing them into its latent space and analyzing the generated patch distributions

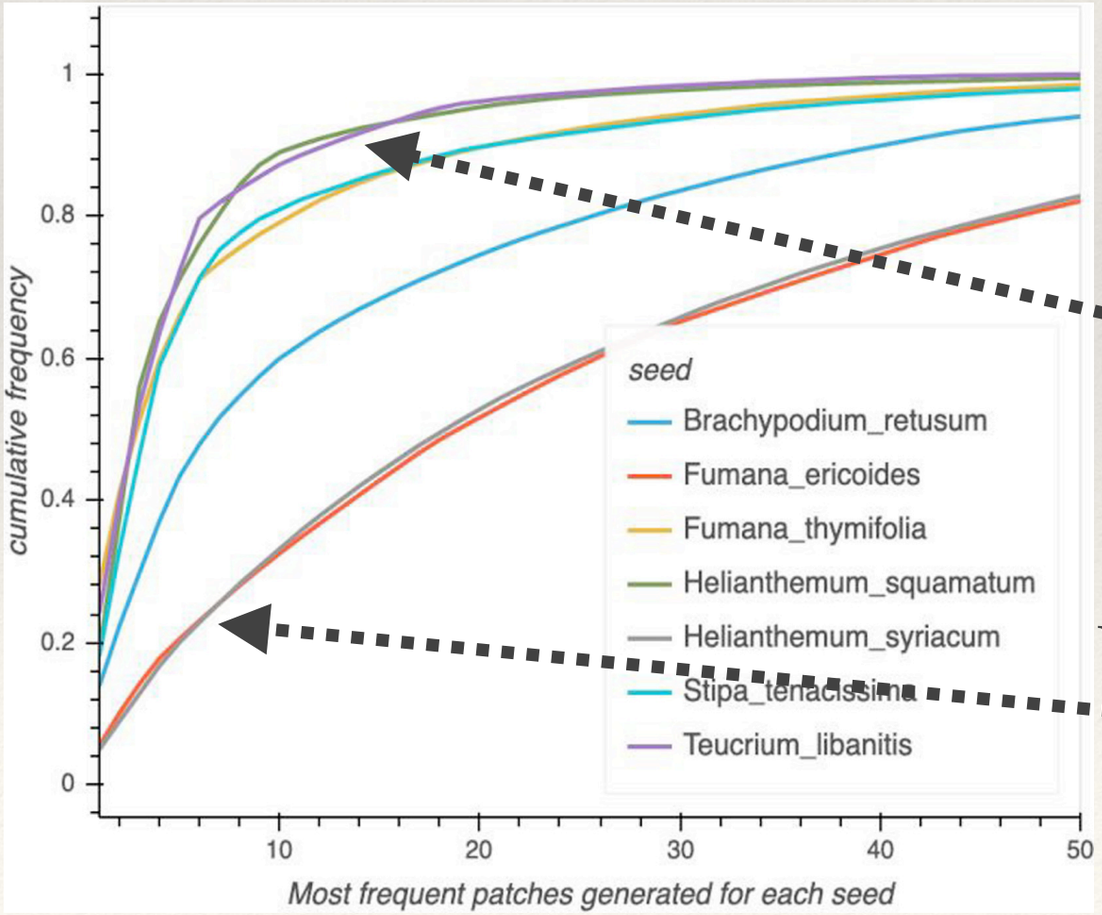

produce only a few dominant patch types with high probability, saturating quickly

generate a larger variety of distinct patches with low individual probabilities, suggesting they **promote greater biodiversity**

Transfer learning

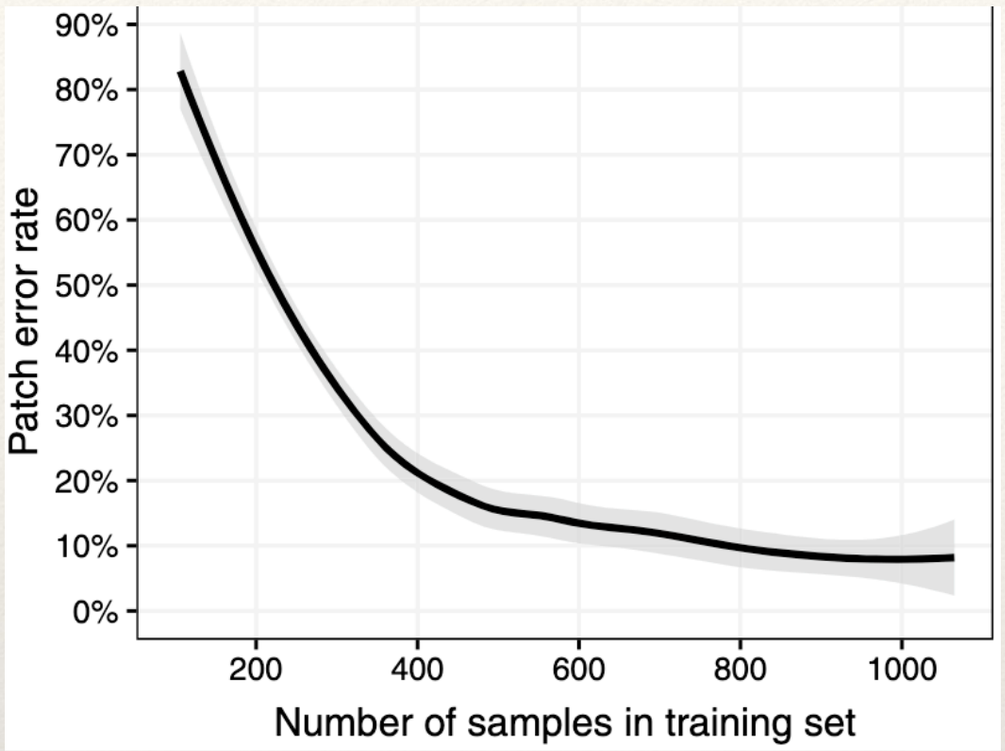

Unsurprisingly
the more data the better, to a point

But datasets in this area a really small (hundreds of samples) and scaling up is a costly field campaign

Also, maybe the rules of the game learned by the GenAI model are somewhat universal, albeit complex

Can we combine different datasets? How?

Transfer learning

Explored how Petrer could help understanding
other arid ecosystems

Those ecosystems contained different species:
there was **no direct translation between the systems**

We needed to create a **dictionary**
tested translations based on phenotypic, phylogenetic
and abundance information

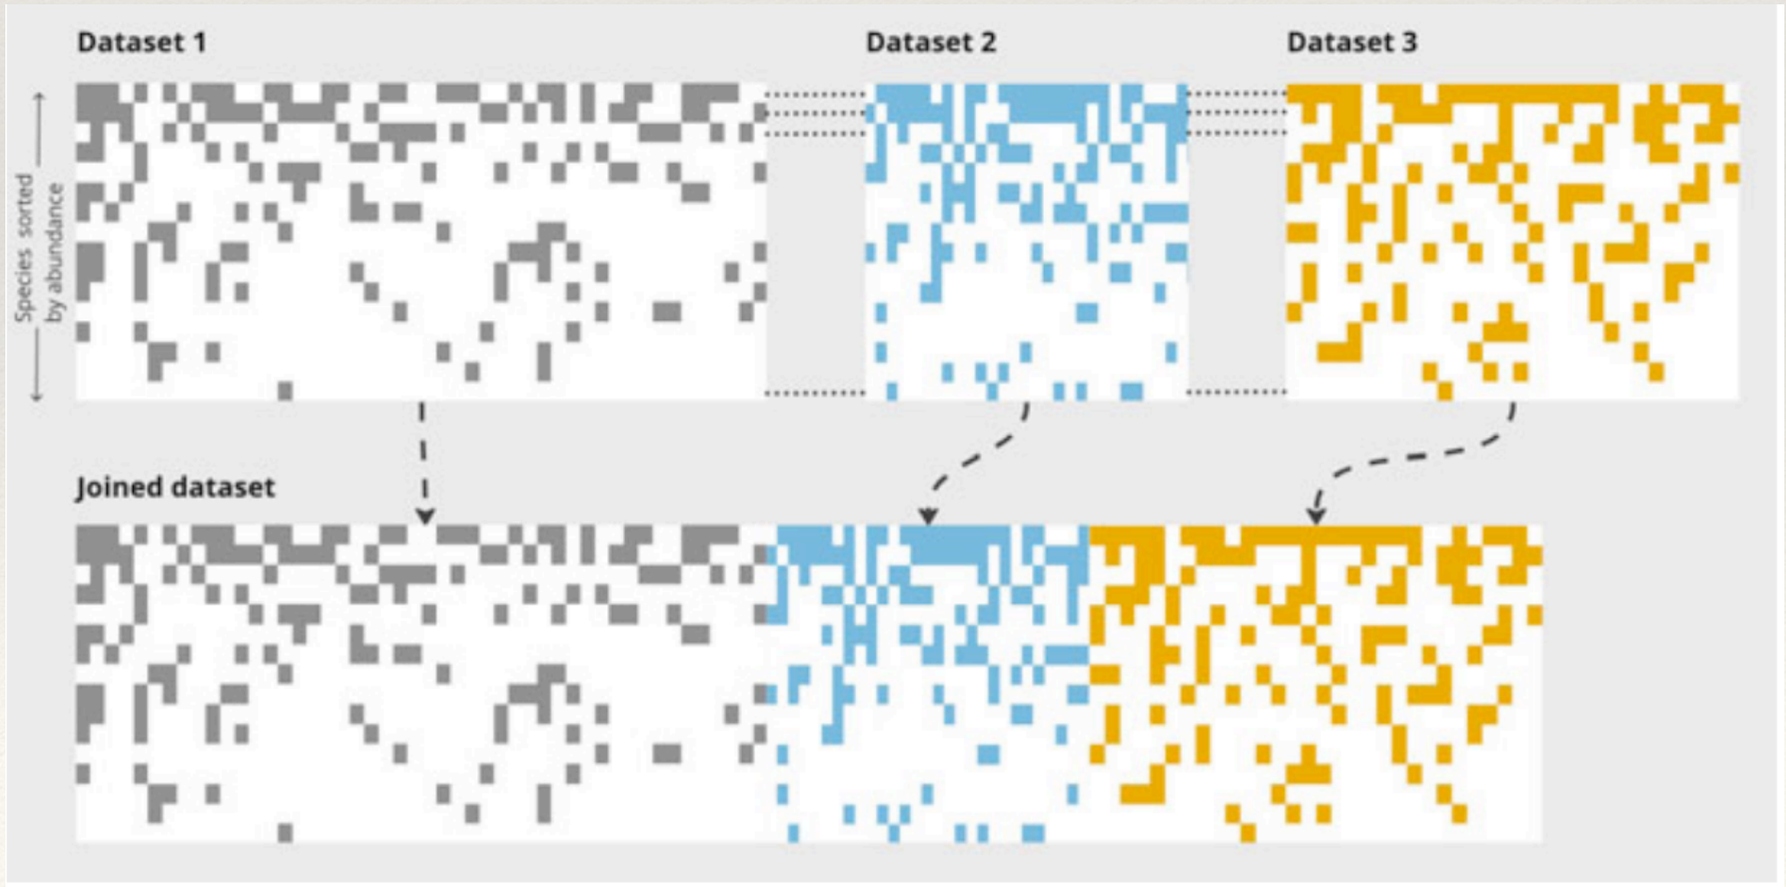

Transfer learning

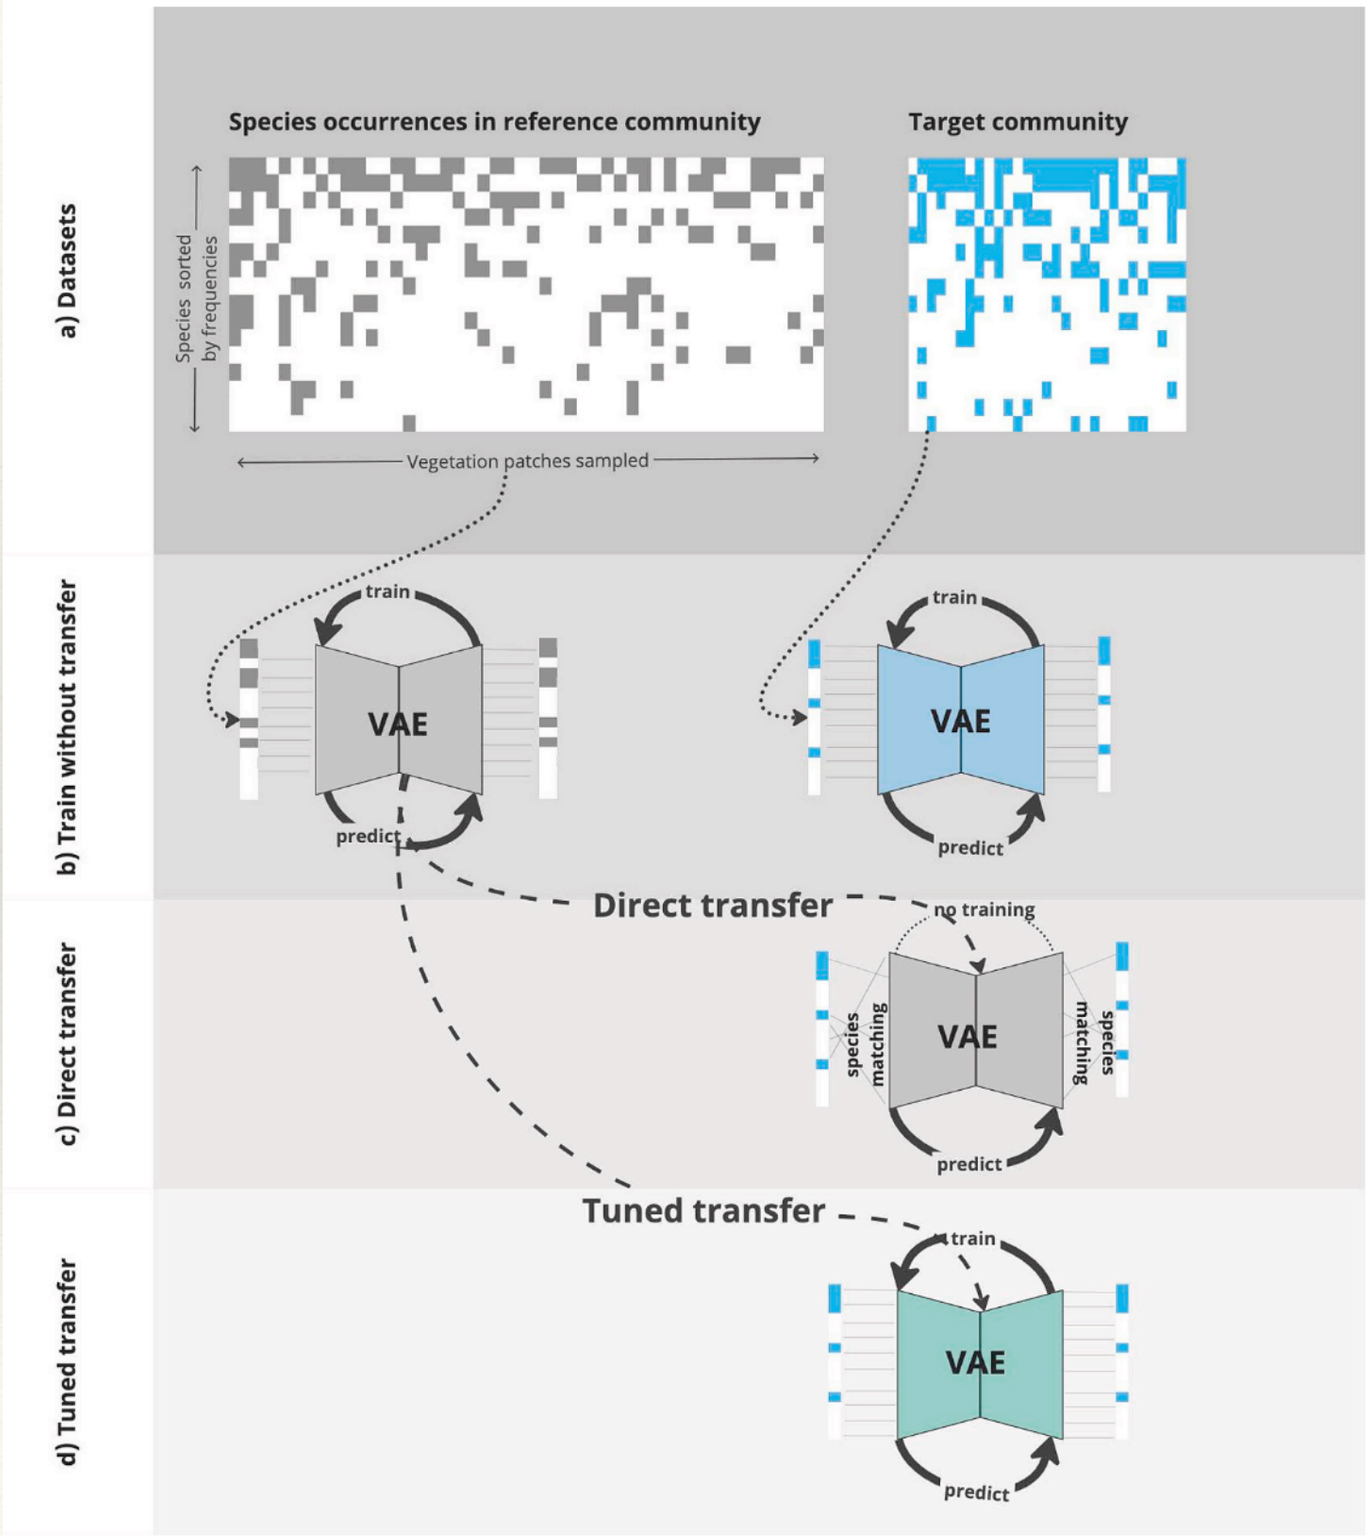

Compare

Train without transfer

Direct transfer w/ dictionary

Tuned transfer

Transfer learning

Compare
Train without transfer
Direct transfer w/ dictionary
Tuned transfer w/ dictionary

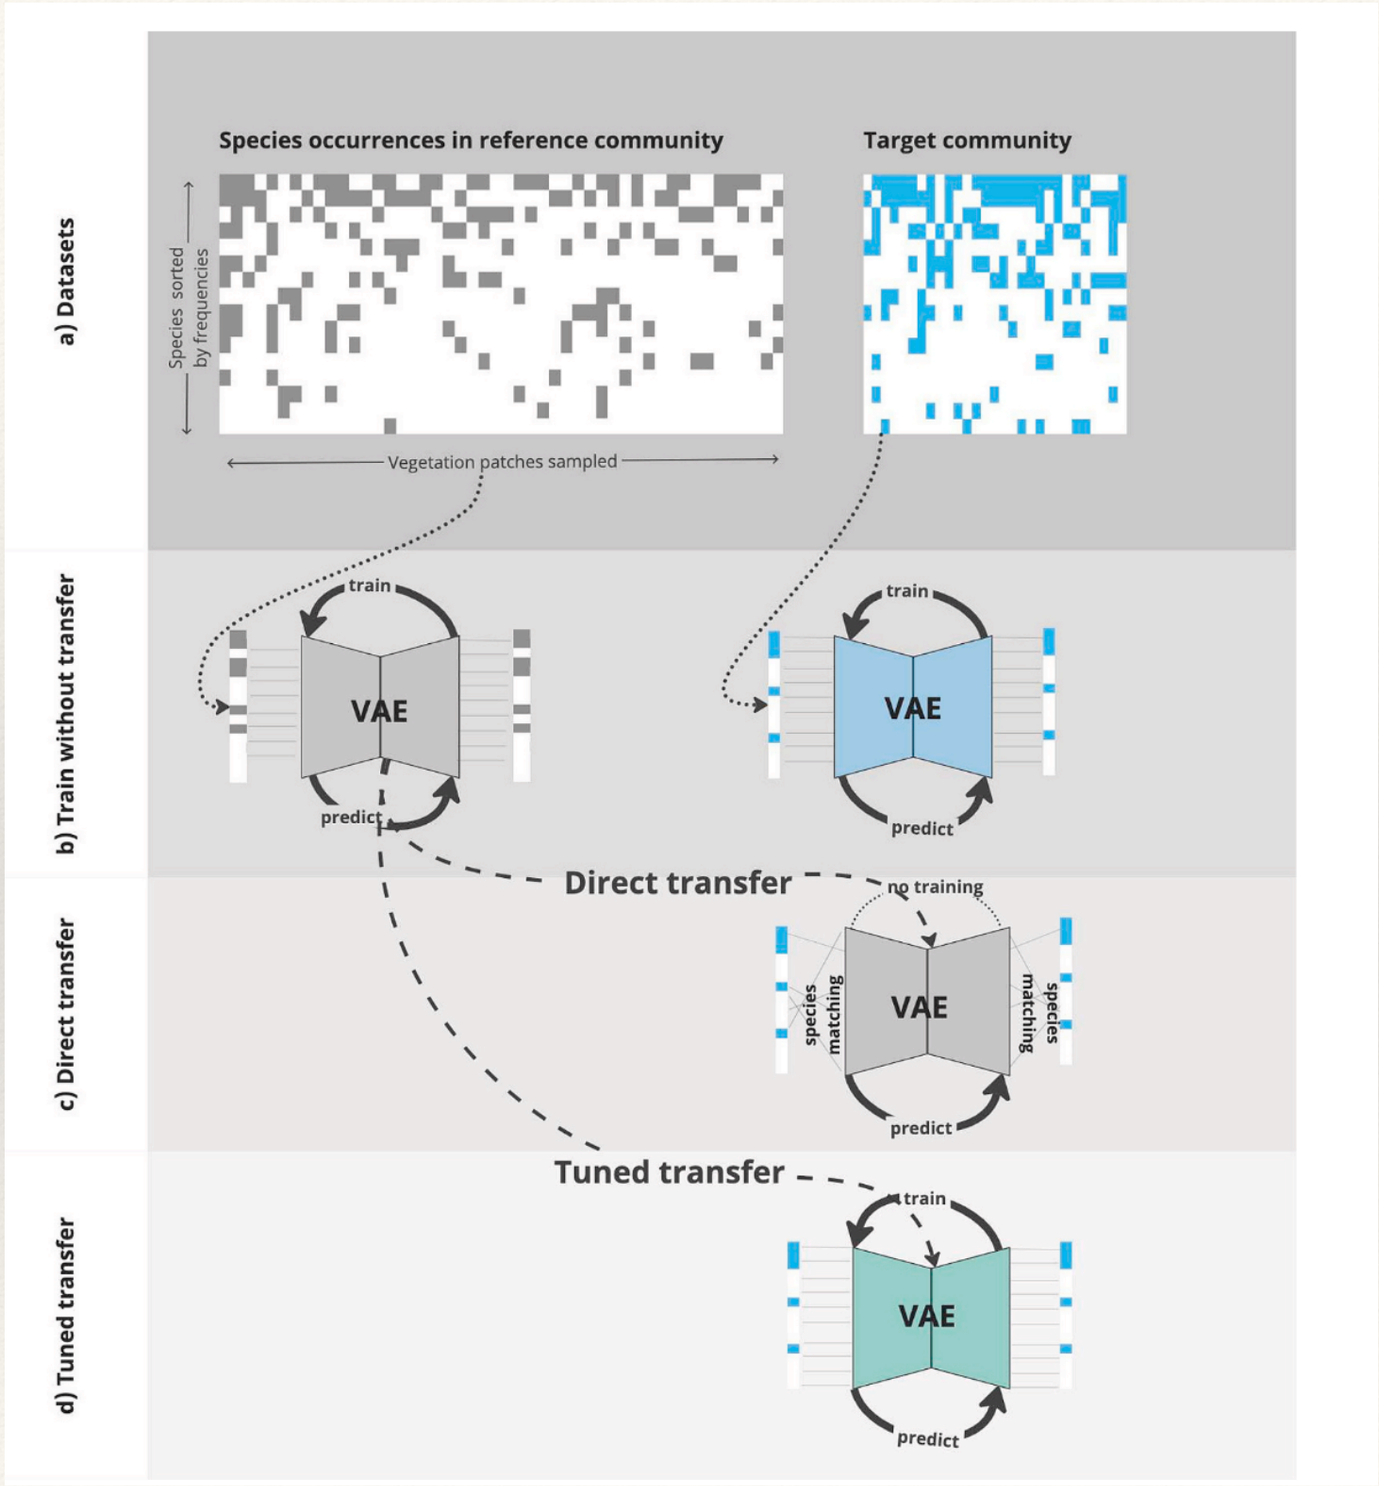

Transfer learning

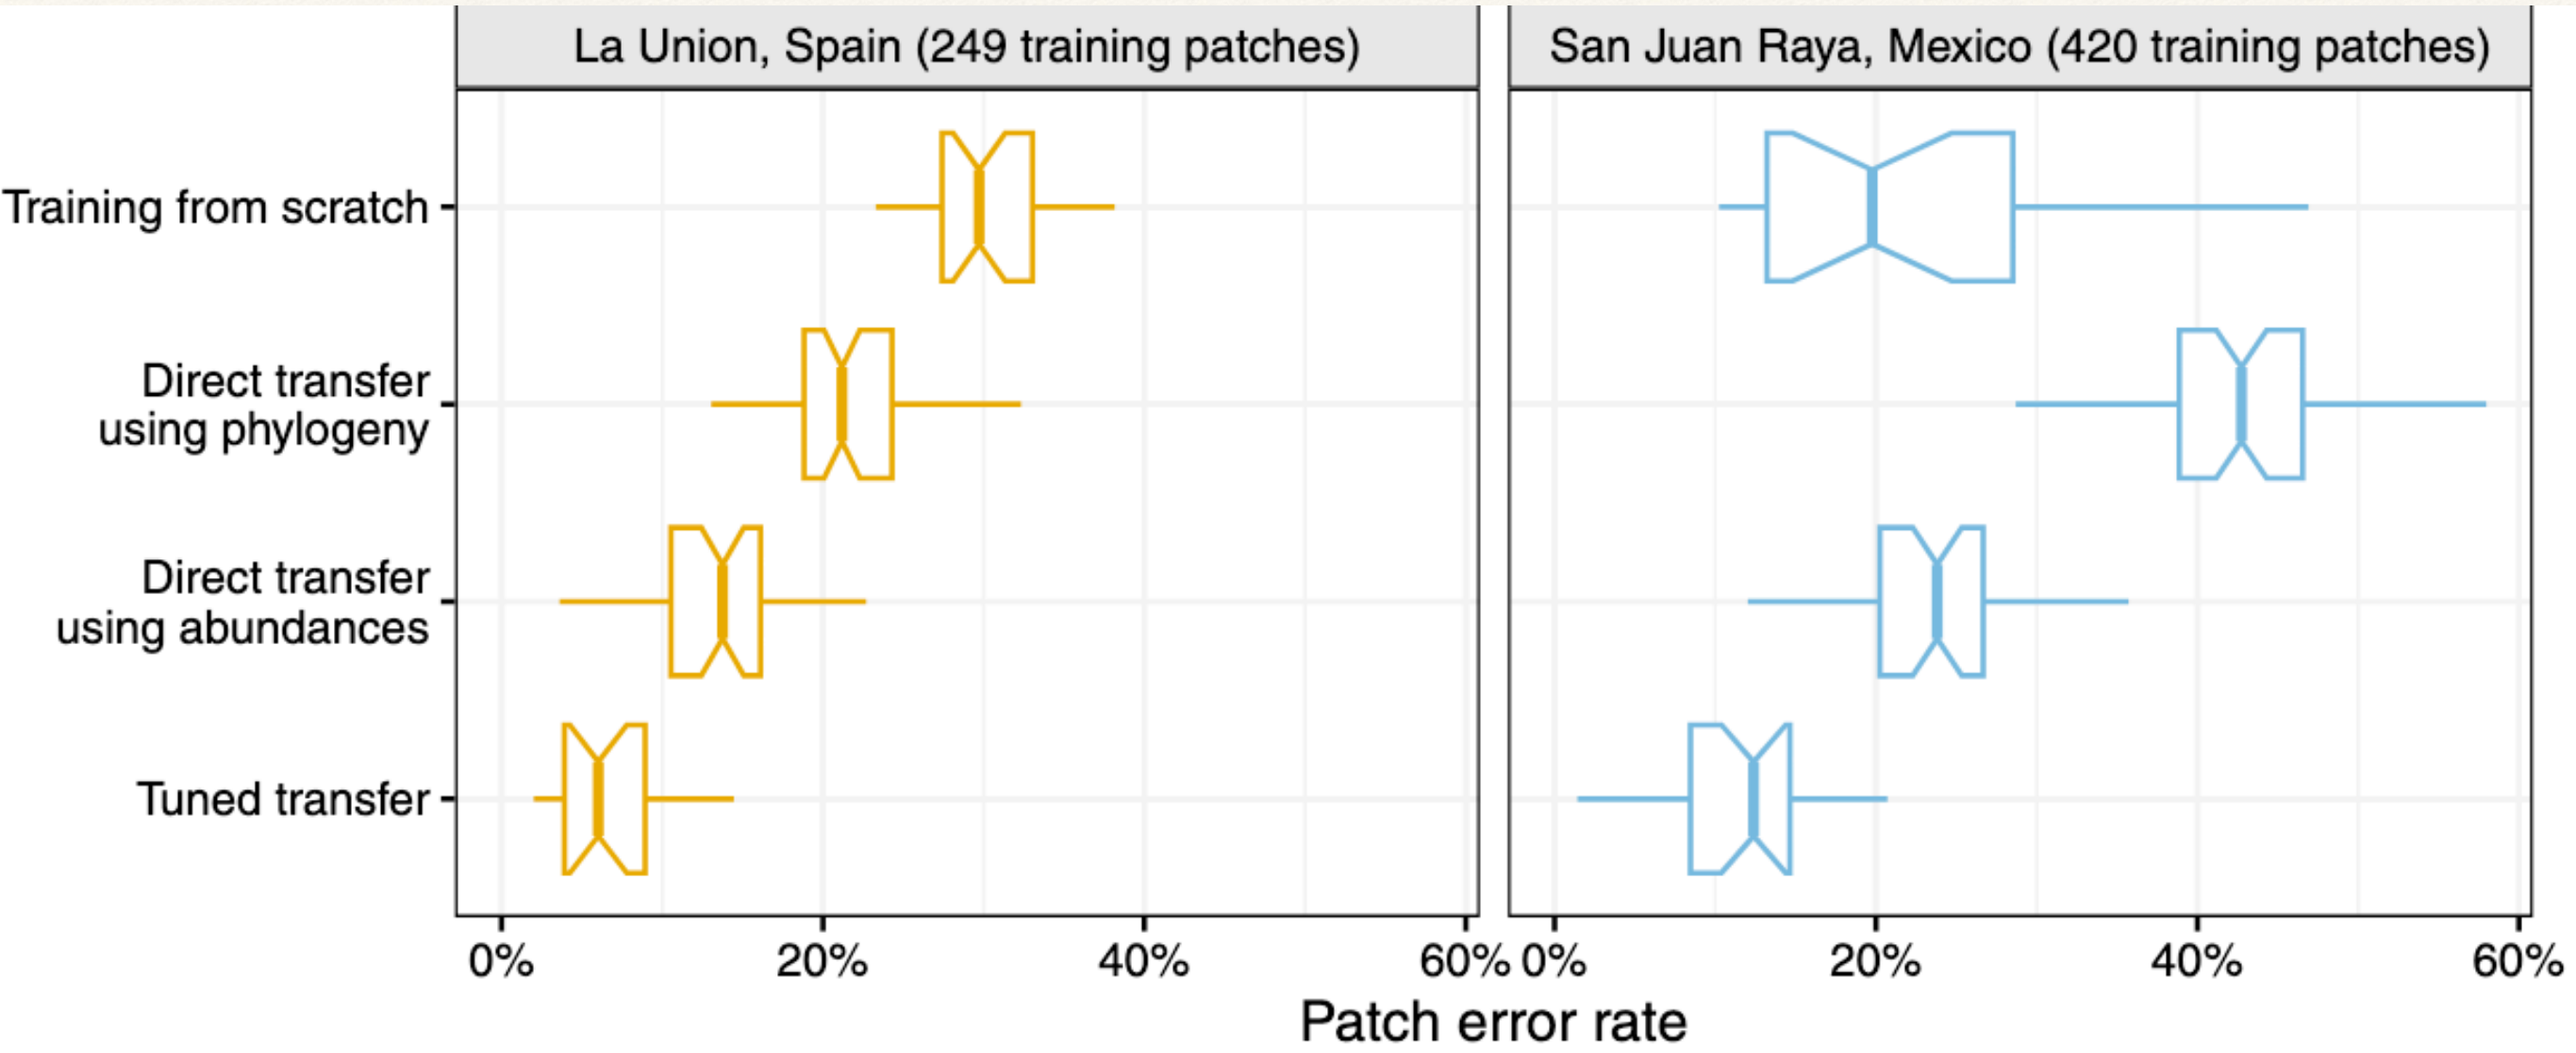

- Clear improvement
- Best dictionary based on abundance
- Settles a methodology to combine different datasets

Restoration

The goal of the research group was to guide restoration efforts in damaged ecosystems
old mines, consolidated dumps, burned areas…

Those restored areas were seeded years ago and left to evolve

We did a field campaign in two **restoration** sites and their surrounding **natural** area

The data was very sparse but we checked that the GenAI model with the dictionary was good at predicting the restoration area

Restoration

The goal of the research group was to guide restoration efforts in damaged ecosystems
old mines, consolidated dumps, burned areas…

Those restored areas were seeded years ago and left to evolve

We did a field campaign in two **restoration** sites and their surrounding **natural** area

The data was very sparse but we checked that the GenAI model with the dictionary was good at predicting the restoration area

BUT restoration was often seeding of non-natural species
here we had to ask a different question
quantify the level of restoration achieved wrt natural area

Restoration

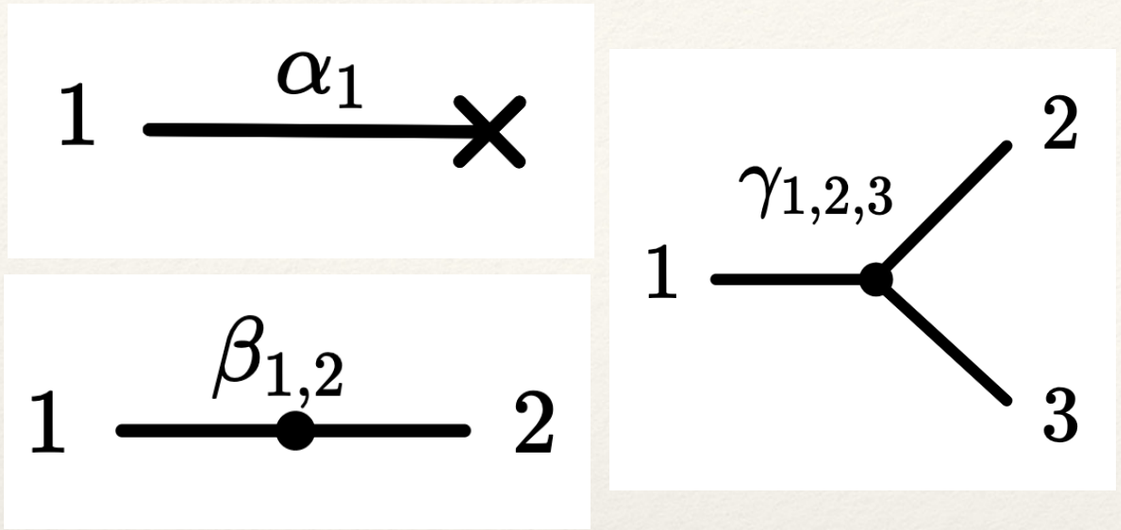

We moved to an interpretable method
Best fit of a kind of Lagrangian

The natural areas were **different**

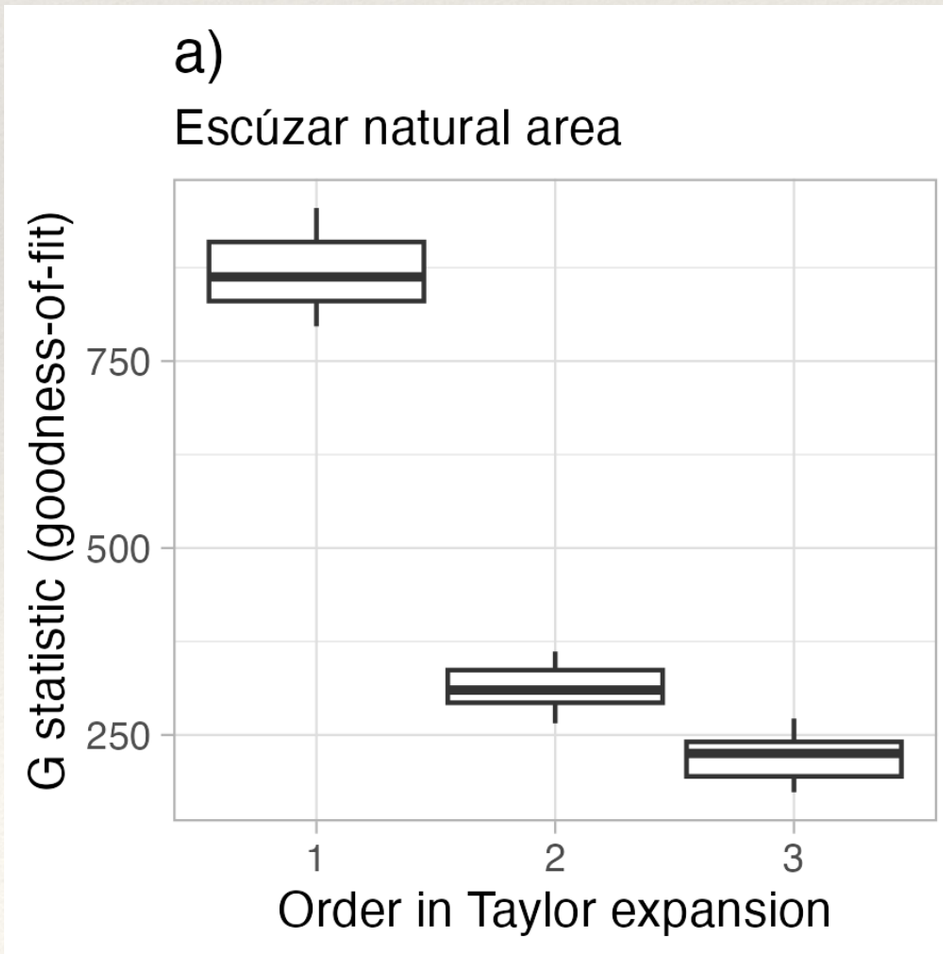

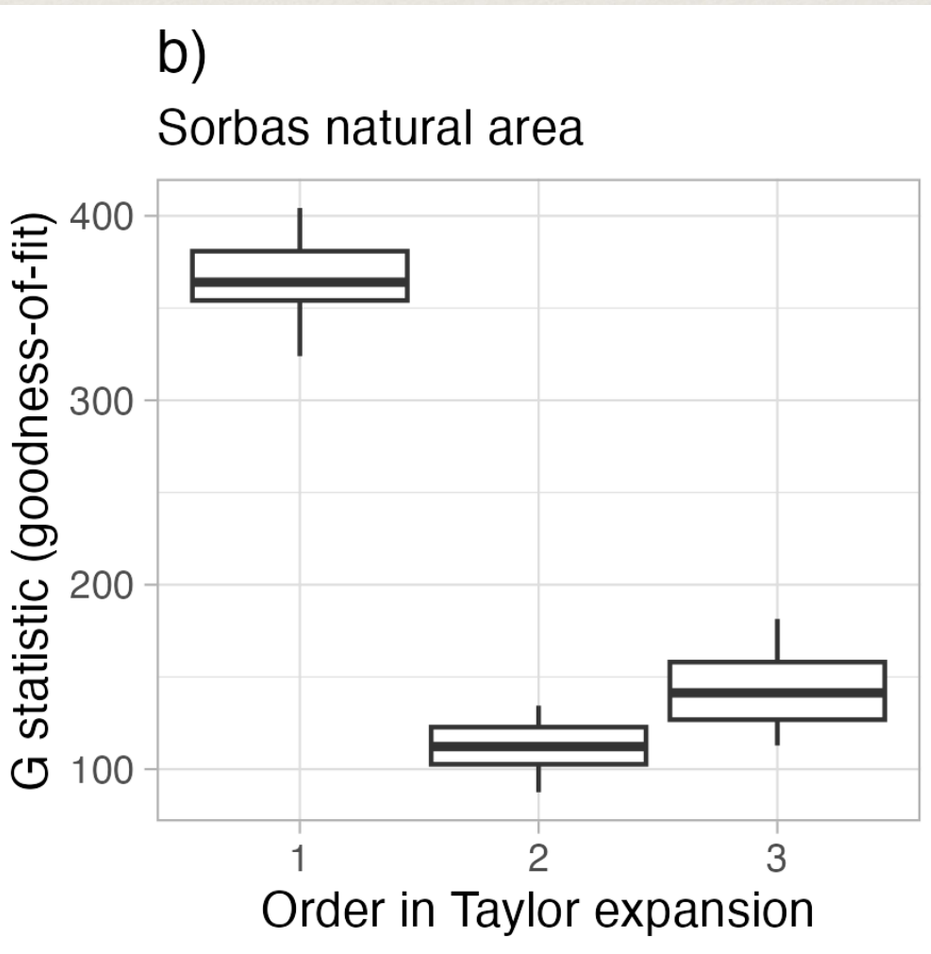

Restoration

Assess the level of compatibility between the natural and restored ecosystem on the species present in both

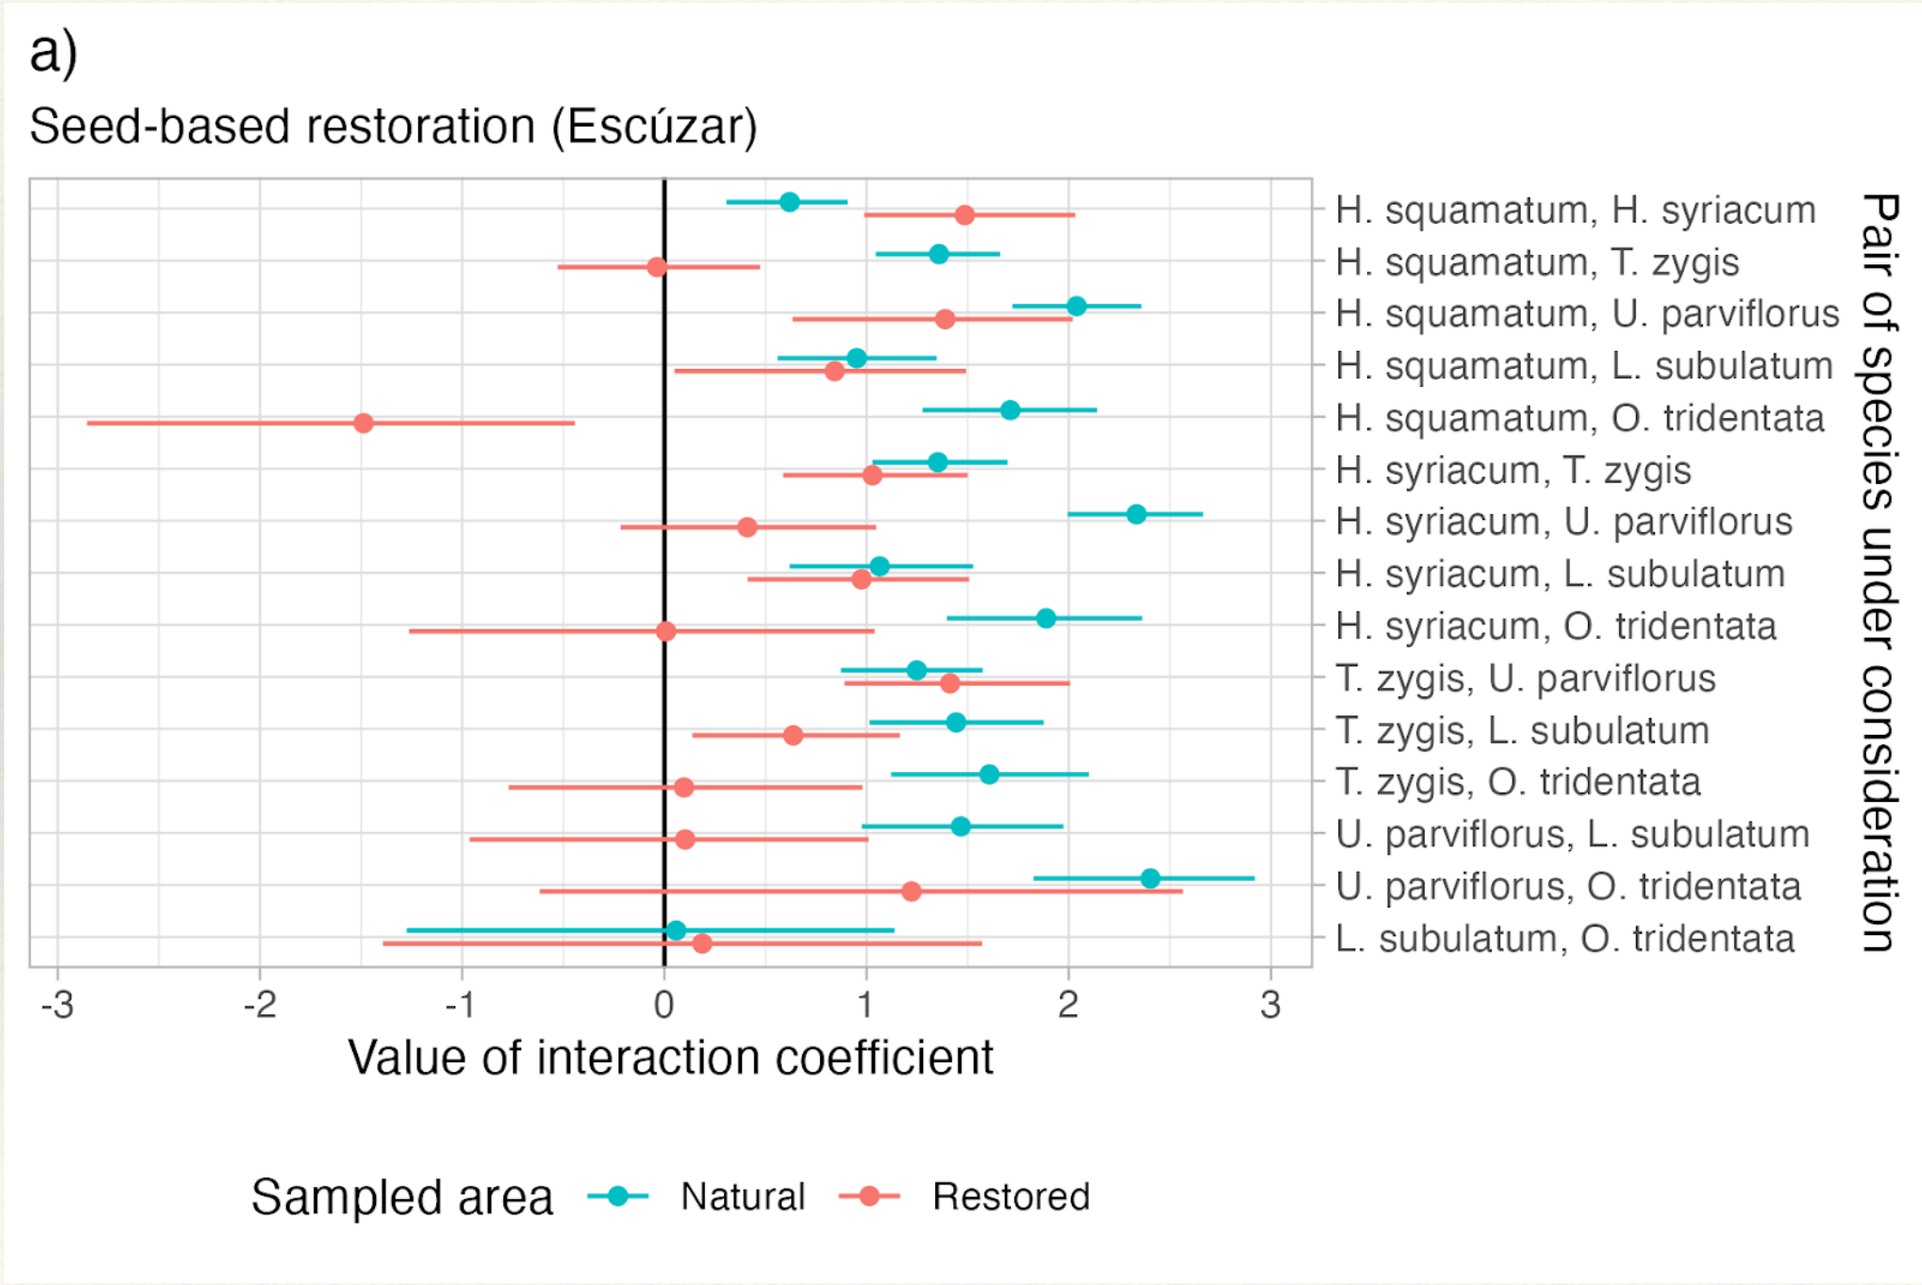

Music and AI

Based on
Exploring how a GenAI understands music.
Neural Computing and Applications (2024)

A VAE trained on music, does it *understand* music?

Google's MAGENTA is a music generator trained on a huge music dataset
The architecture is massive, latent space of 512 neurons

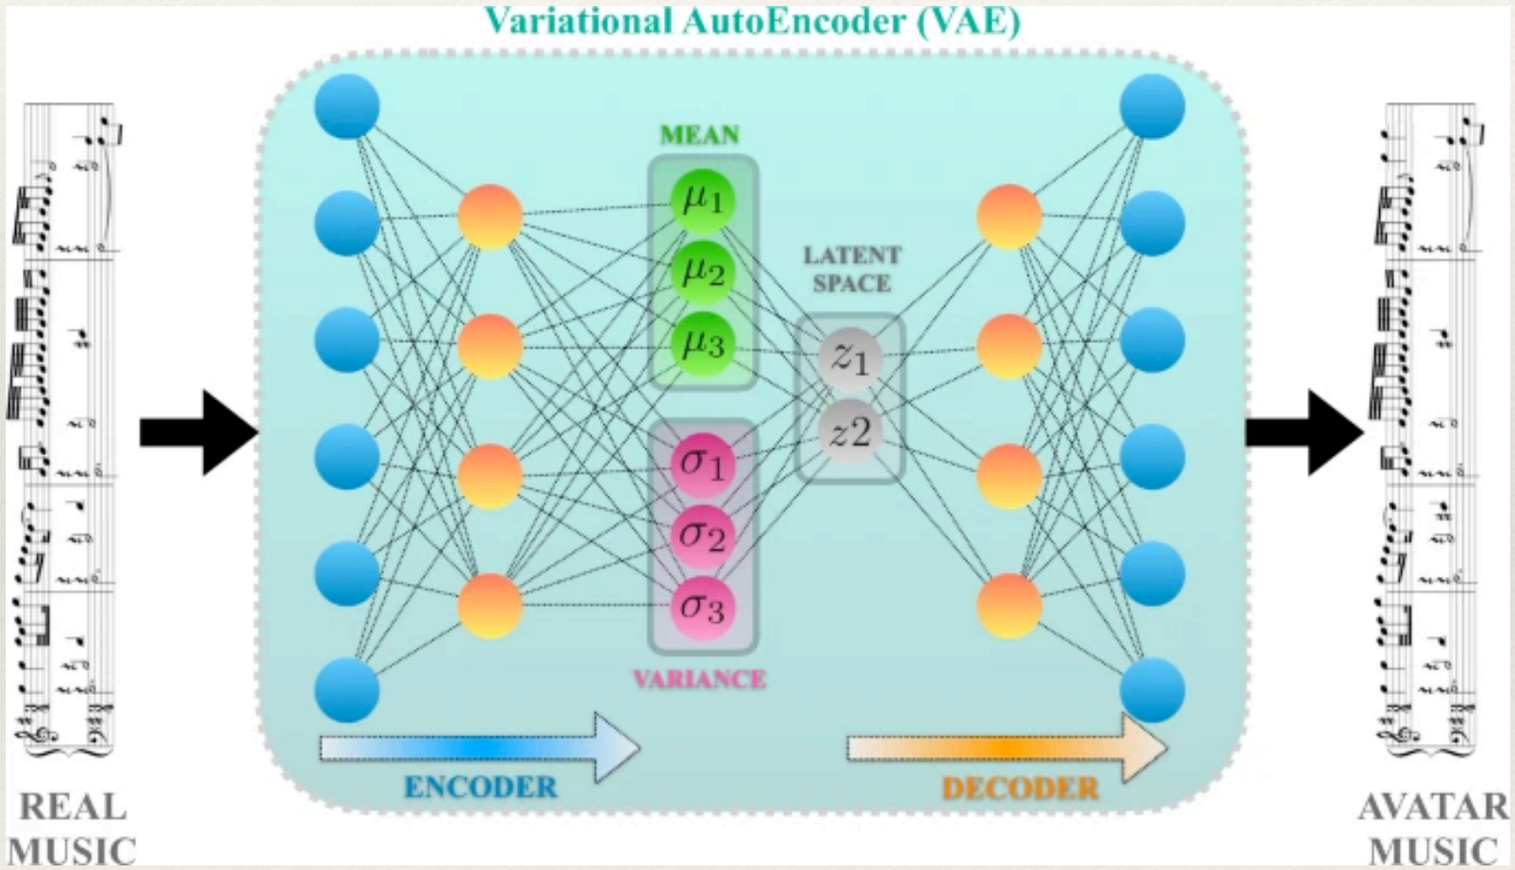

As in the symmetry examples I showed you,
did MAGENTA learn something about music? **What makes music, music?**

Self-organization

We can ask specific questions like

-Can latent dimensions be **ordered** by importance?

-**How many** latent dimensions are really necessary for the VAE to perform well, i.e., can some latent dimensions be classified as irrelevant?

-Do the relevant latent dimensions have an **interpretation**? correspond to some concepts of musicality as understood by humans?

Self-organization

-Can latent dimensions be **ordered** by relevance/importance?

Distribution in musical pieces

σ_i

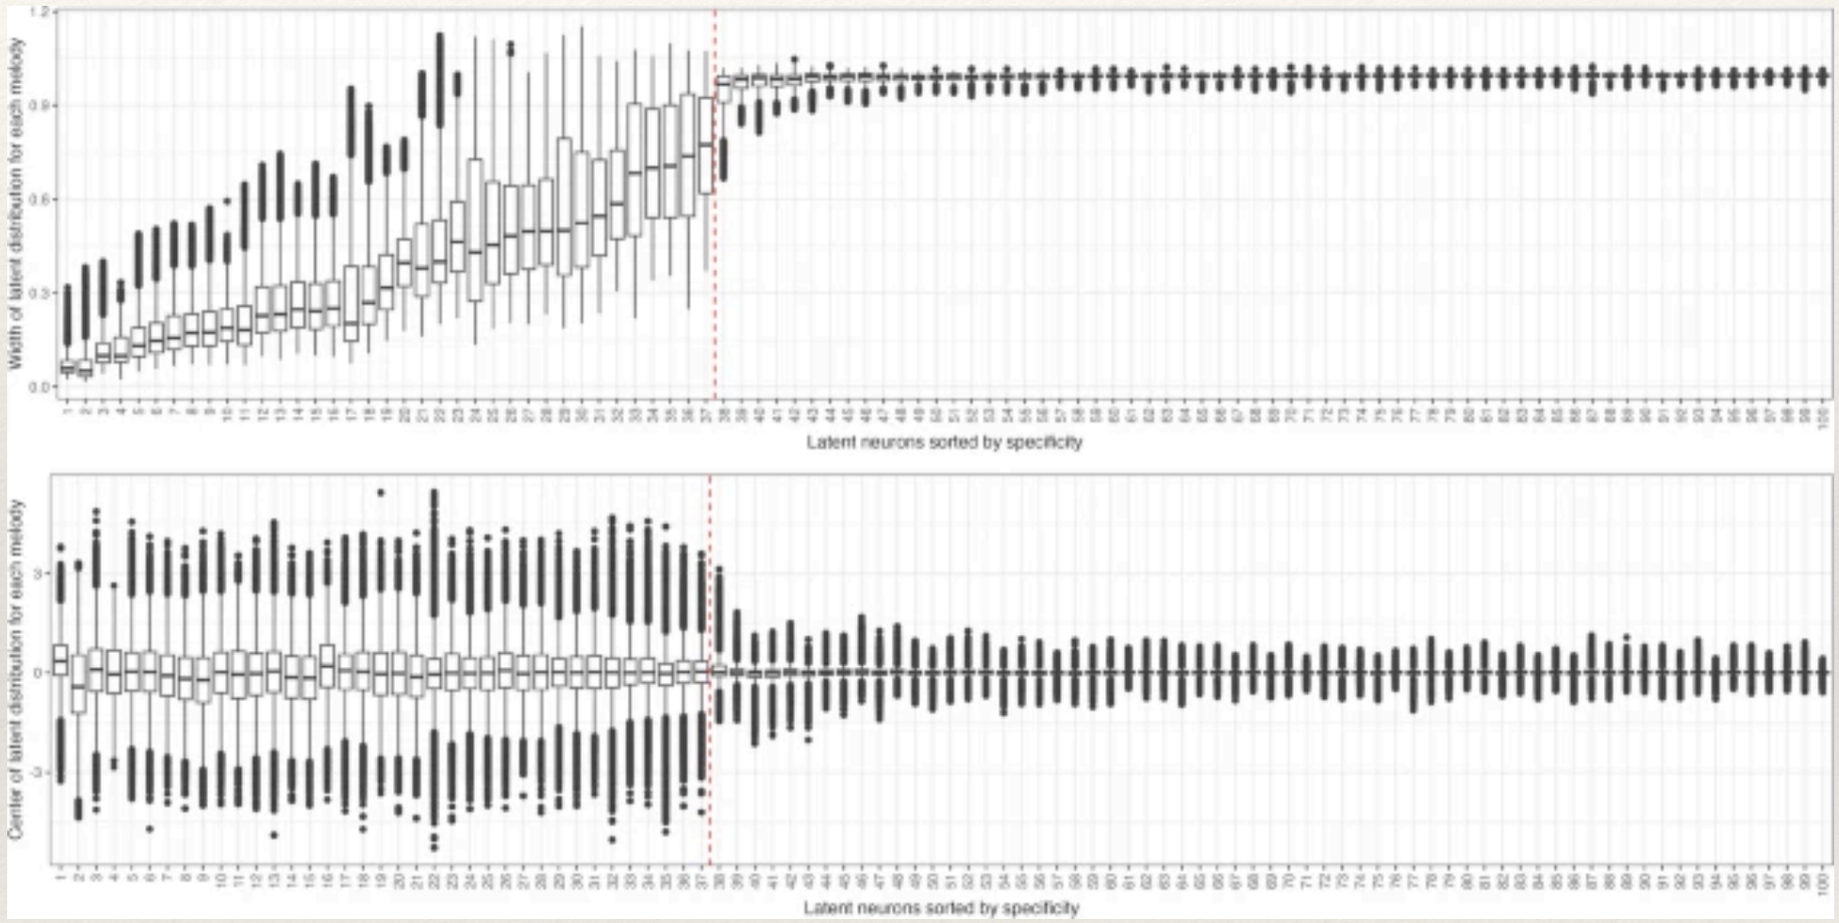

μ_i

Neurons (i)

There is indeed a **hierarchical behaviour**

Self-organization

$$\text{Relevance}_j = \frac{\text{std}(\langle z_j \rangle)}{\text{mean}(\sigma_j)} \ ,$$

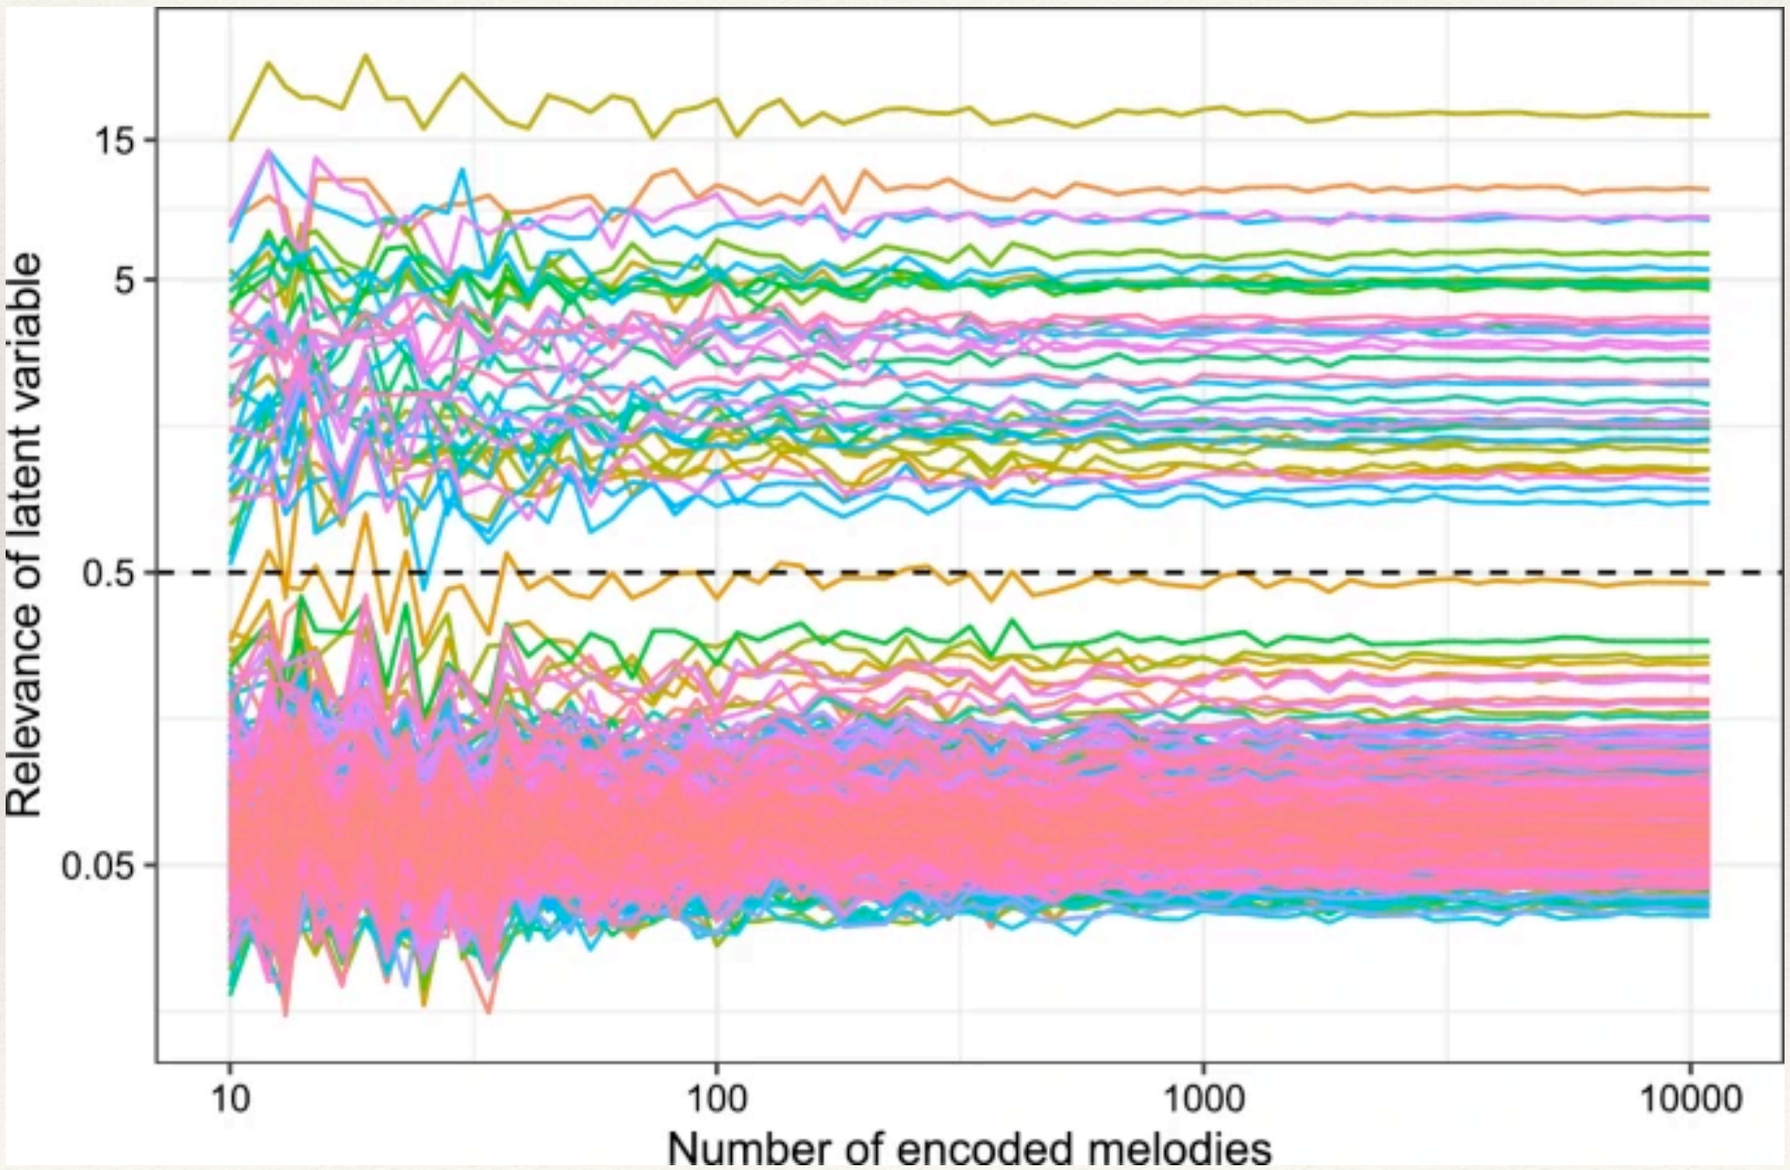

Self-organization

-**How many** latent dimensions are really necessary for the VAE to perform well, i.e., can some latent dimensions be classified as irrelevant?

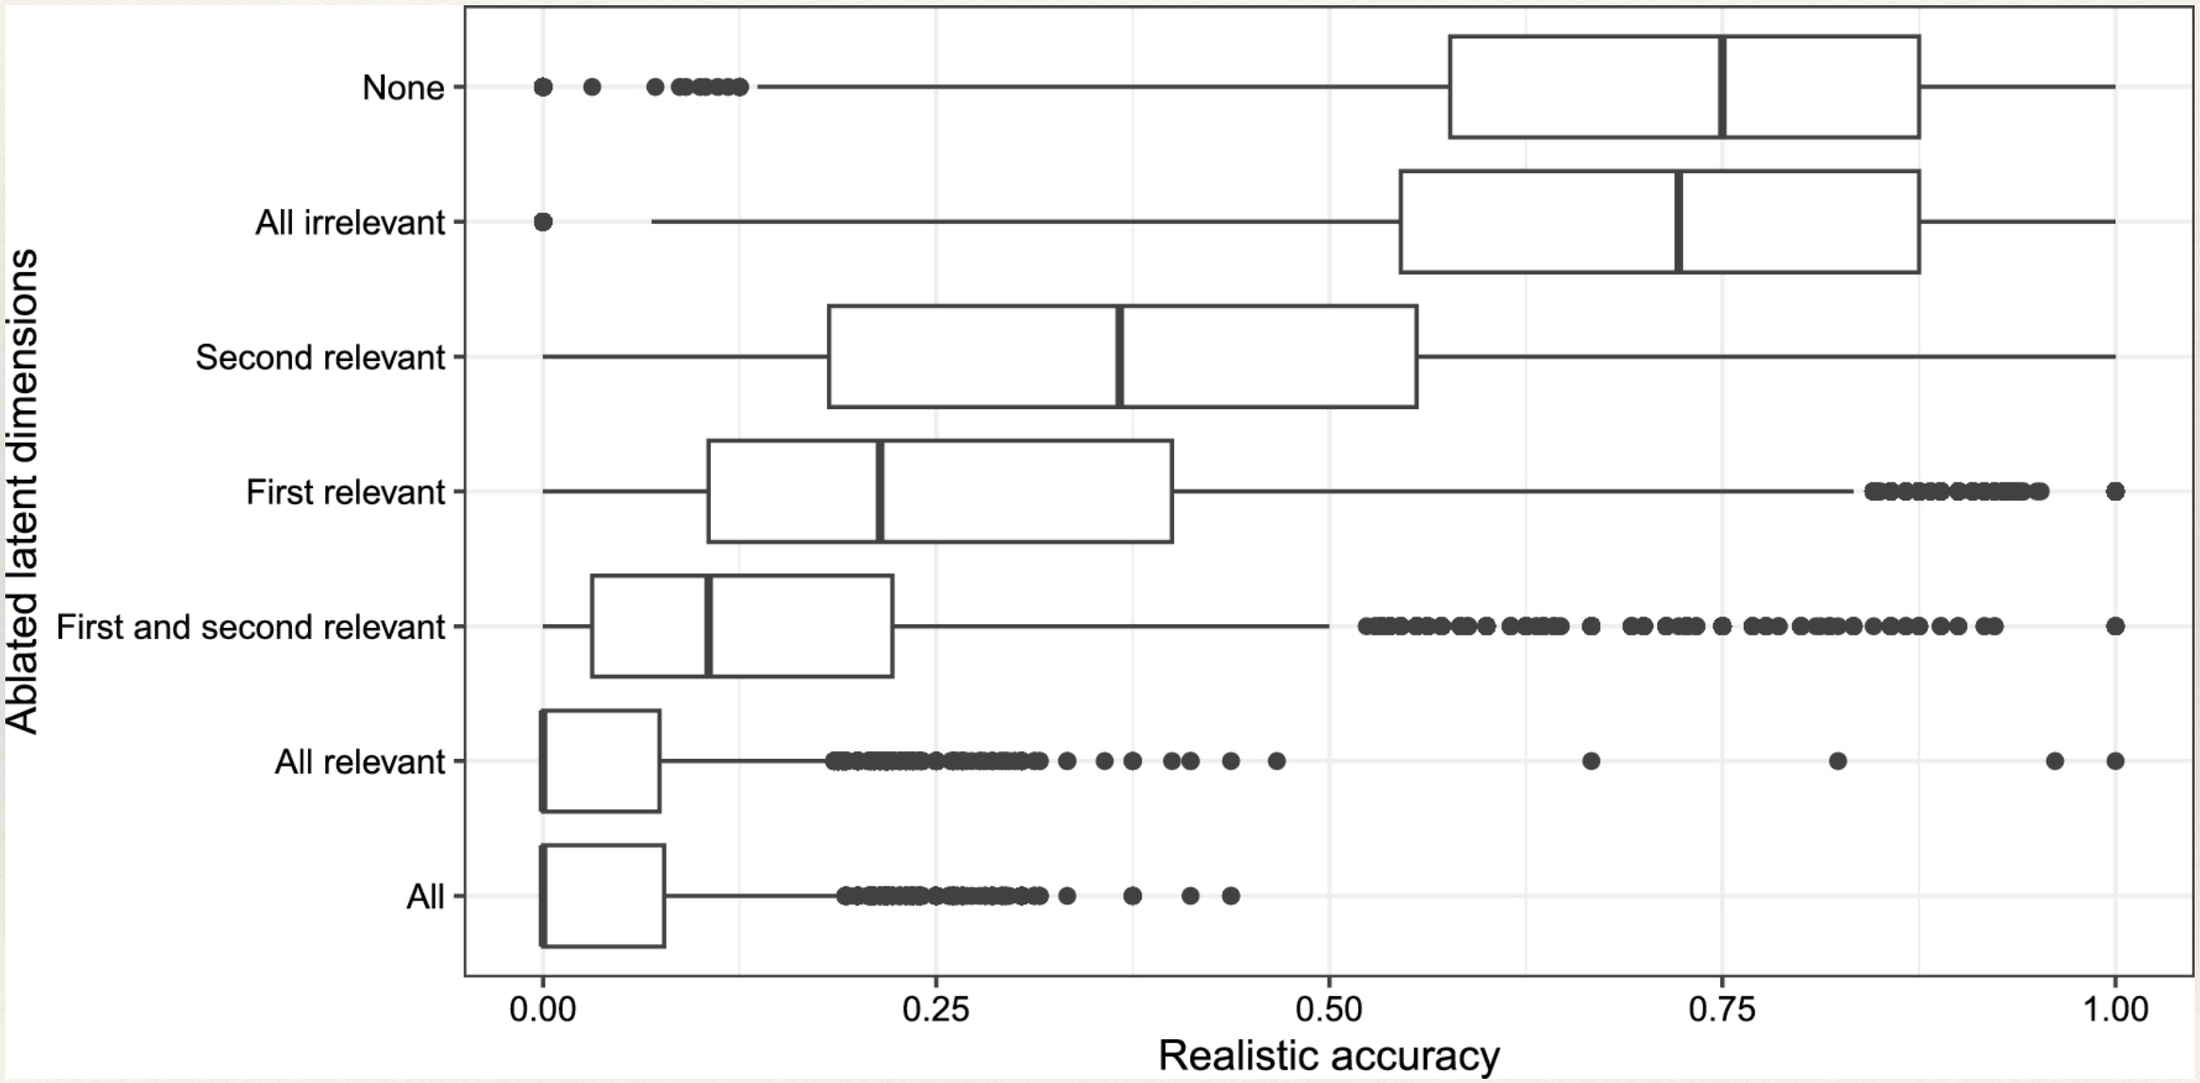

very few…

Interpretation

-Do the relevant latent dimensions have an **interpretation**? correspond to some concepts of musicality as understood by humans?

First two most relevant neurons are highly correlated with rythm, pitch and melody musical features

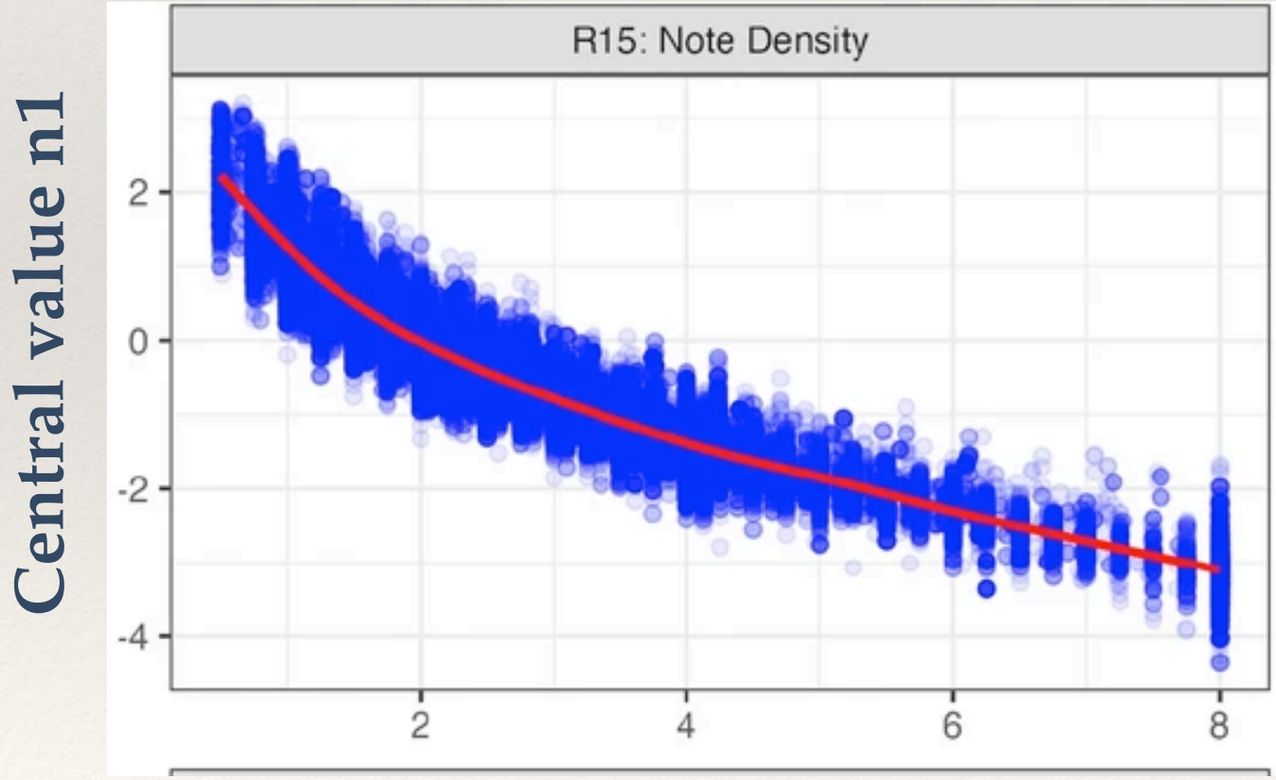

For example, n1 vs note density

All the irrelevant neurons do not show correlation
but they are key to produce new distinctive musical pieces

Politics

Based on
Exploring the political pulse of a country using DS.
Journal of Computational Social Science (2021)
and
Polarisation of political elites with ML techniques.
Journal of Computational Social Science (2025, to appear)

Evolution of the Spanish Parliament in the last years

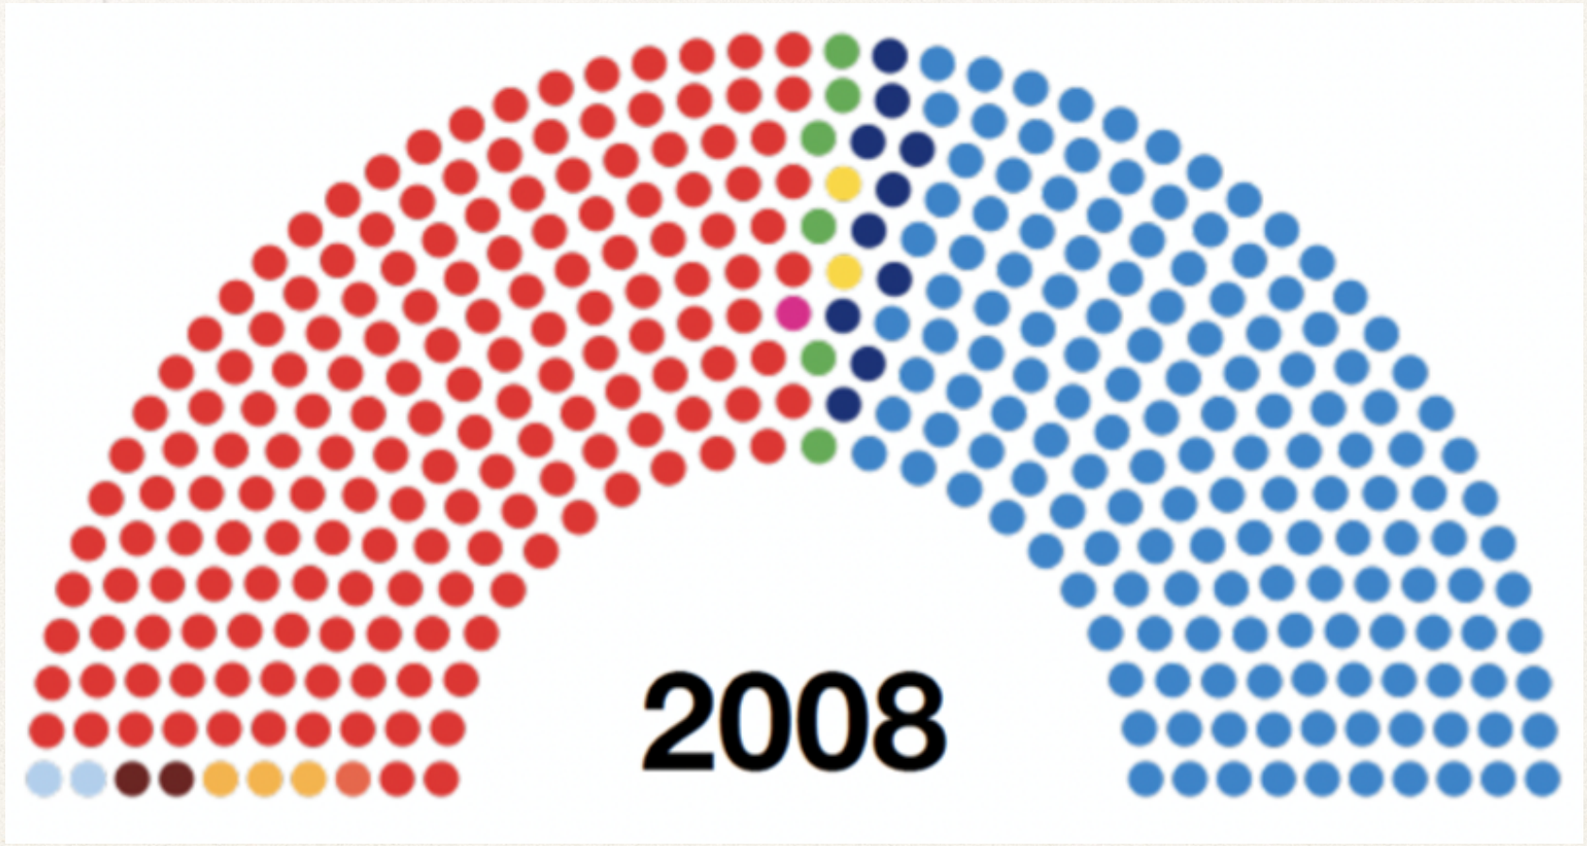

Highly Bi-partisan

Evolution of the Spanish Parliament in the last years

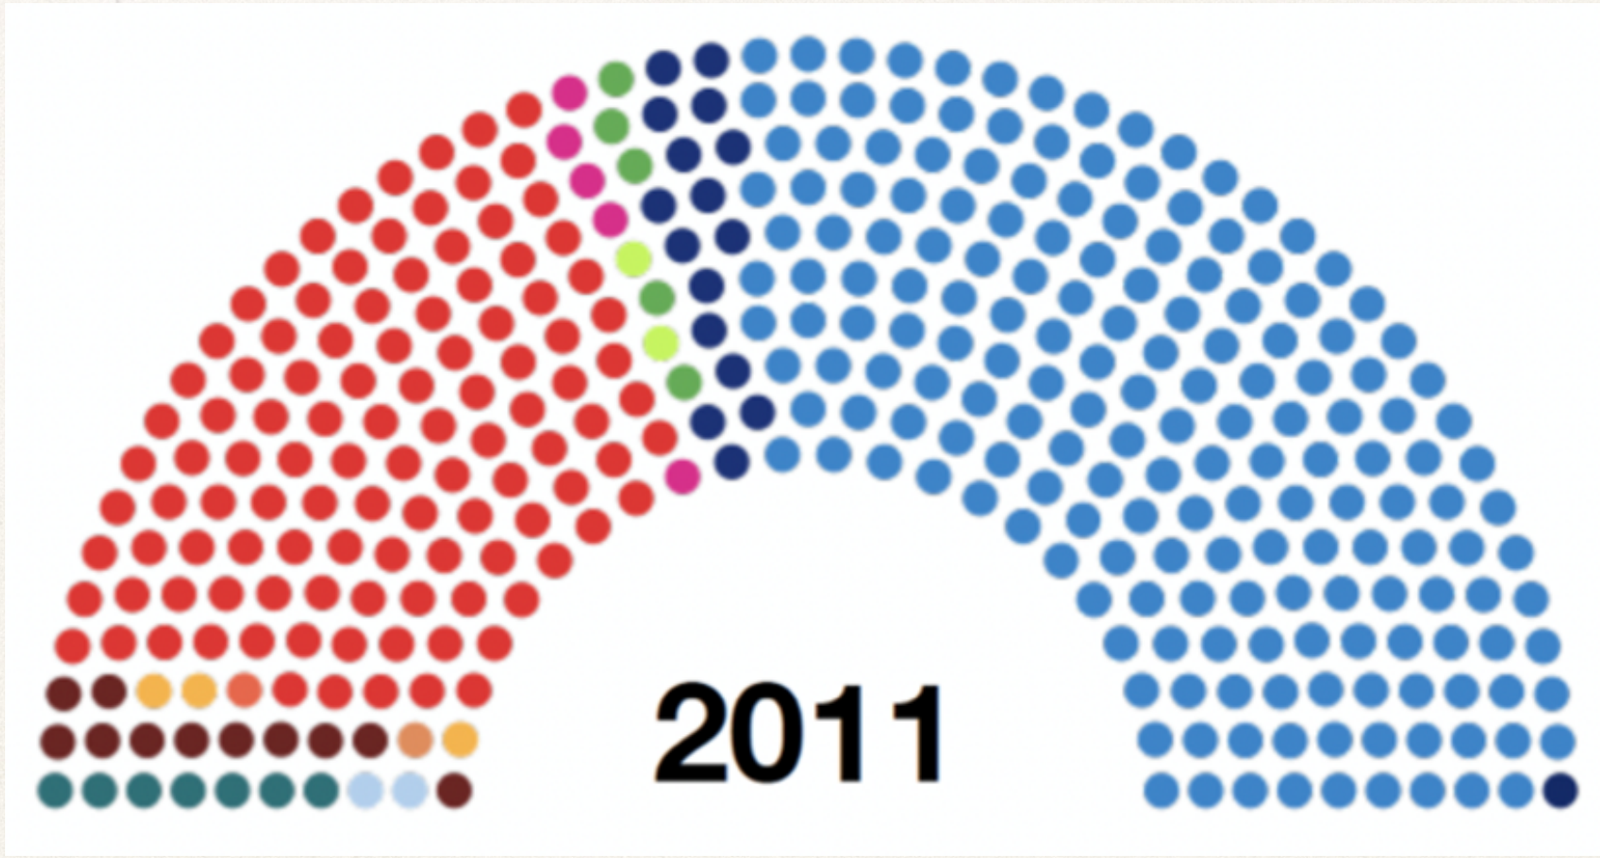

Highly Bi-partisan

Evolution of the Spanish Parliament in the last years

PSOE

C's

PP

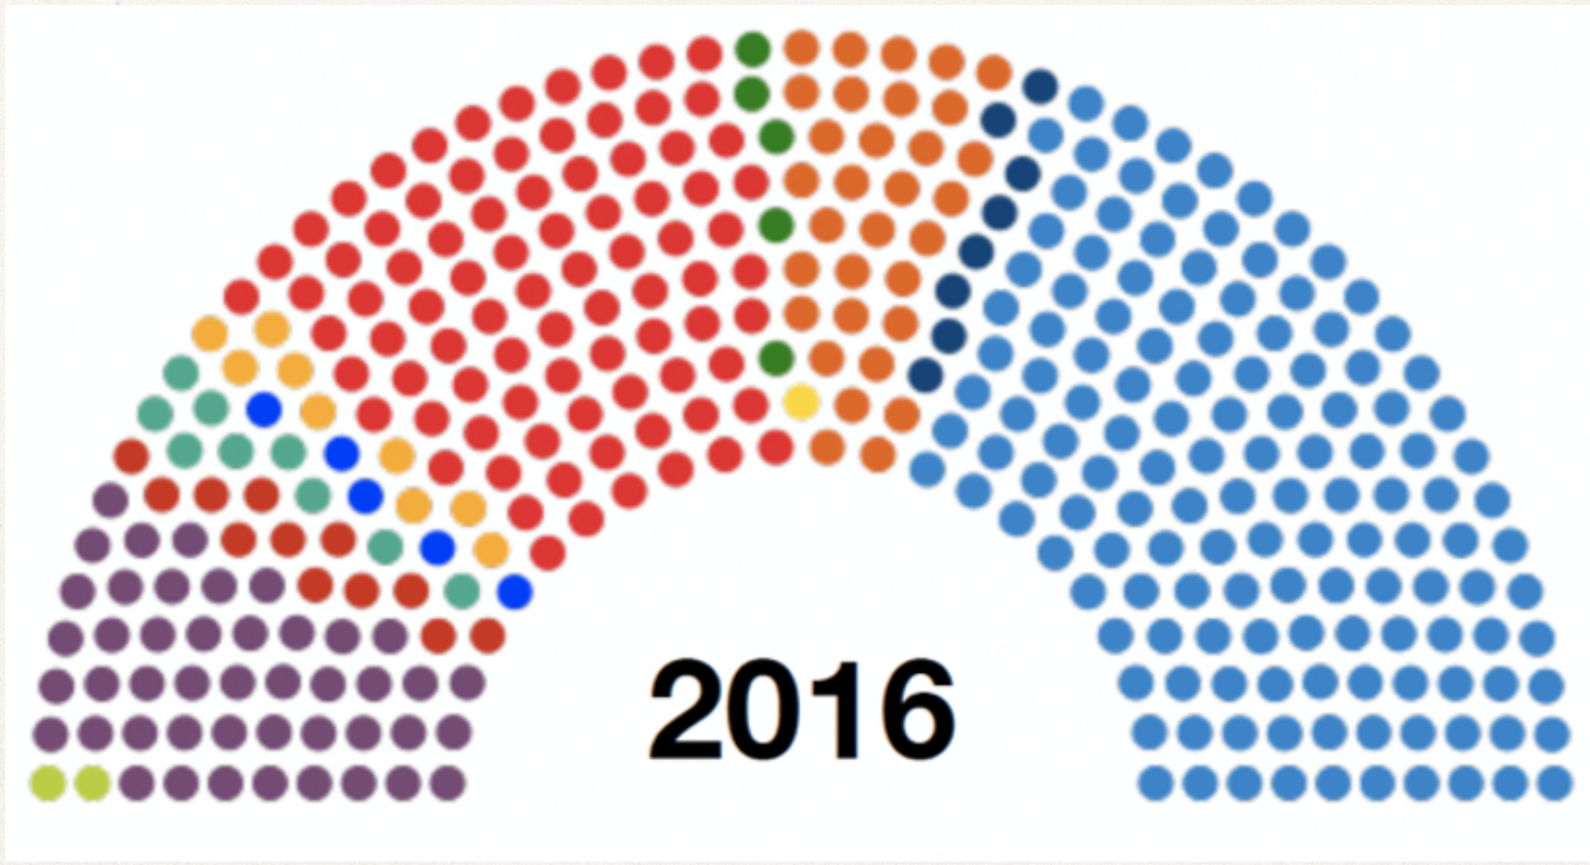

PODEMOS

Evolution of the Spanish Parliament in the last years

PSOE

C's

PP

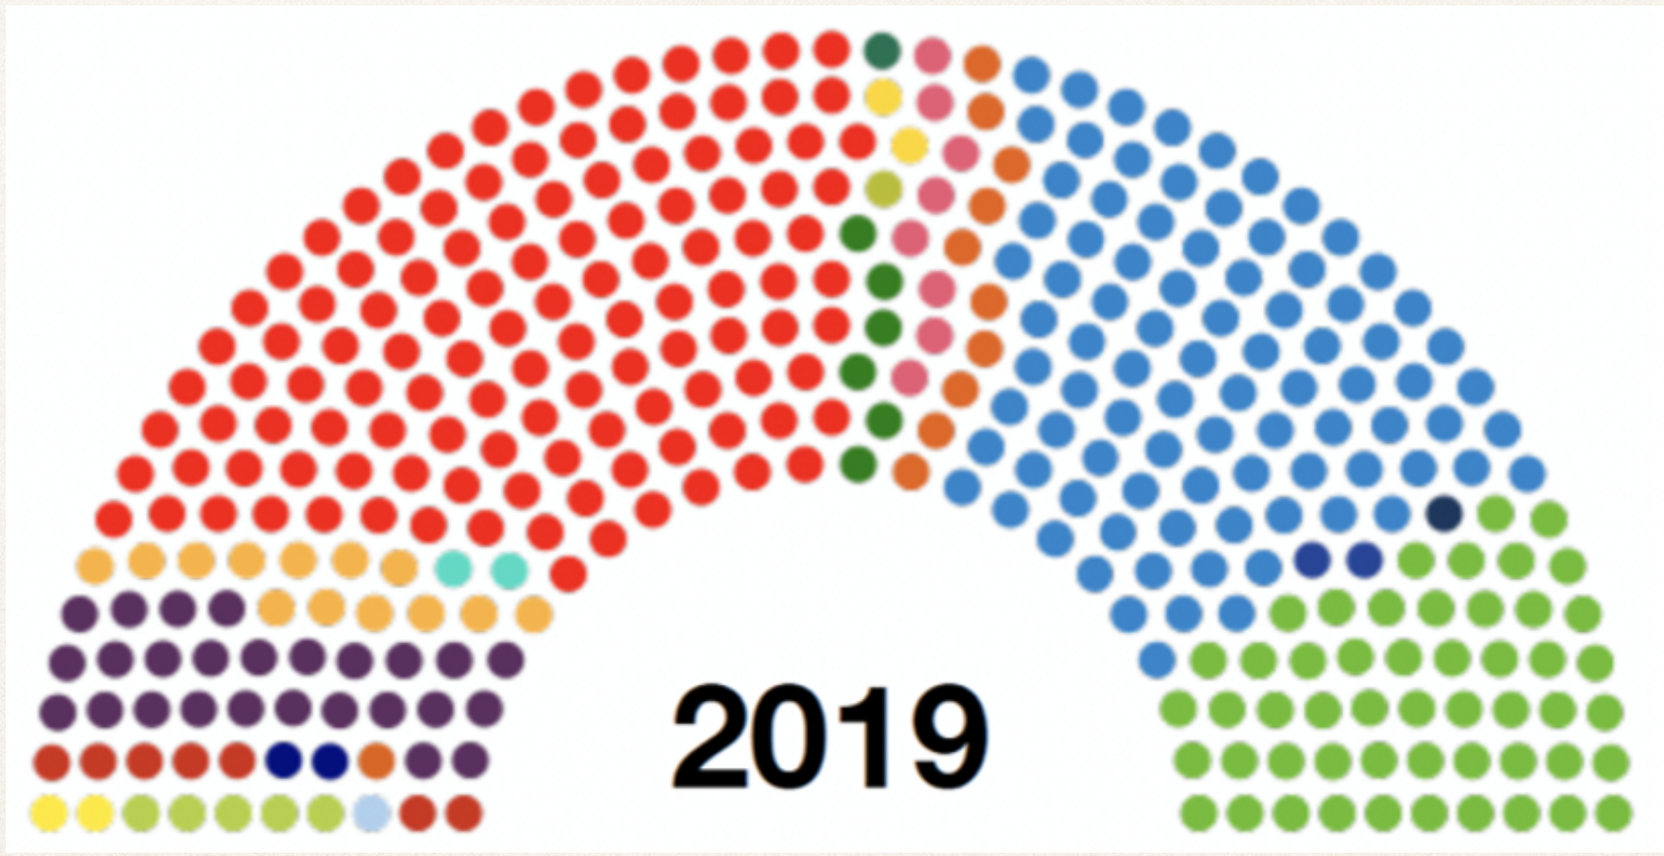

PODEMOS

Similar in **2023**, fractured parliament with extremes growing over centrist parties

VOX

We first analysed Twitter discourse from politicians (easy to harvest)

Used a bunch of NLP techniques, finding:
New parties' polarisation is higher, sentiment measure is able to track events (impeachment, procés…) and some parties are highly (anti-)correlated

NN predicts political leaning ~ 90% accuracy:
Adherence to messages is high

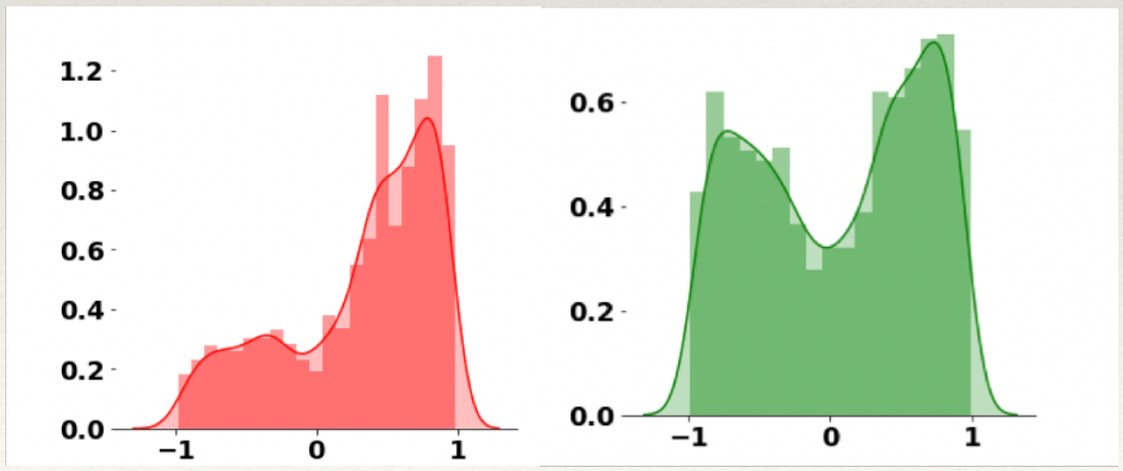

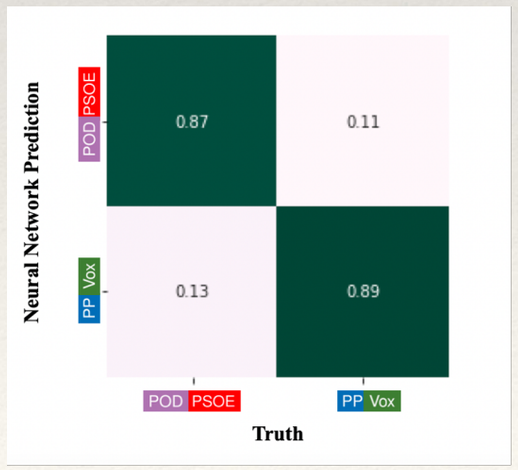

Then moved onto what they *officially* say

CORTES GENERALES

DIARIO DE SESIONES DEL CONGRESO DE LOS DIPUTADOS

PLENO Y DIPUTACIÓN PERMANENTE

Año 2023 | XIV LEGISLATURA | Núm. 239 | Pág. 1

PRESIDENCIA DE LA EXCMA. SRA. D.ª MERITXELL BATET LAMAÑA

Sesión plenaria núm. 230
(extraordinaria)

celebrada el martes 24 de enero de 2023

Página

ORDEN DEL DÍA:

Convalidación o derogación de reales decretos-leyes.

— **Real Decreto-ley 20/2022, de 27 de diciembre, de medidas de respuesta a las consecuencias económicas y sociales de la Guerra de Ucrania y de apoyo a la reconstrucción de la isla de La Palma y a otras situaciones de vulnerabilidad. (BOE núm. 311, de 28 de diciembre de 2022). (Número de expediente 130/000092)** 4

— **Real Decreto-ley 1/2023, de 10 de enero, de medidas urgentes en materia de incentivos a la contratación laboral y mejora de la protección social de las personas artistas. (BOE núm. 9, de 11 de enero de 2023). (Número de expediente 130/000093)** 23

Votaciones 42

Comparecencia del Gobierno ante el Pleno de la Cámara, de conformidad con lo dispuesto por el artículo 203 del Reglamento.

— **Comparecencia, a petición propia, del presidente del Gobierno, ante el Pleno del Congreso de los Diputados, para informar del Consejo Europeo celebrado los días 20 y 21 de octubre y del Consejo Europeo celebrado el día 15 de diciembre de 2022, así como de las medidas adoptadas por el Gobierno para hacer frente a las consecuencias económicas y sociales de la guerra en Ucrania. (Número de expediente 210/000122)** ... 43

— **Comparecencia del presidente del Gobierno ante el Pleno de la Cámara, solicitada por la señora Gamarra Ruiz-Clavijo y otros 87 Diputados, para que informe de los resultados de las conclusiones del Consejo Europeo. (Número de expediente 210/000113)** 44

cve: DSCD-14-PL-239

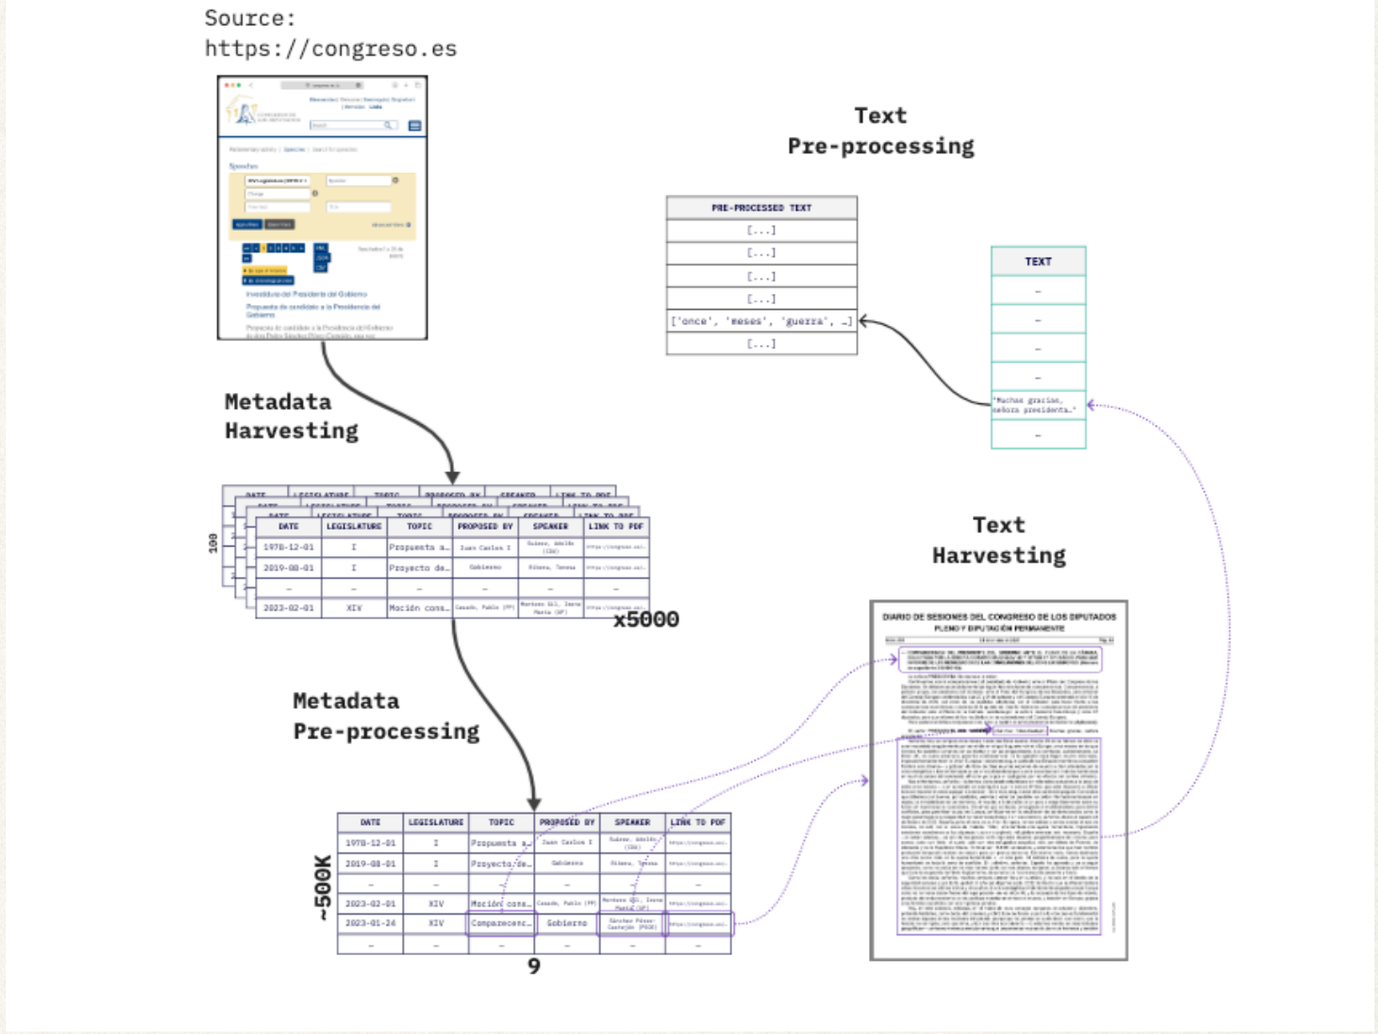

Webcrawler, clearning, preprocessing…

Result:

dataset from 1979 till 2024, available in Zenodo

Full-monty NLP incl transformers

Many results, e.g. ask to project all speeches in a 2D space directions in the internal (latent) space with more information

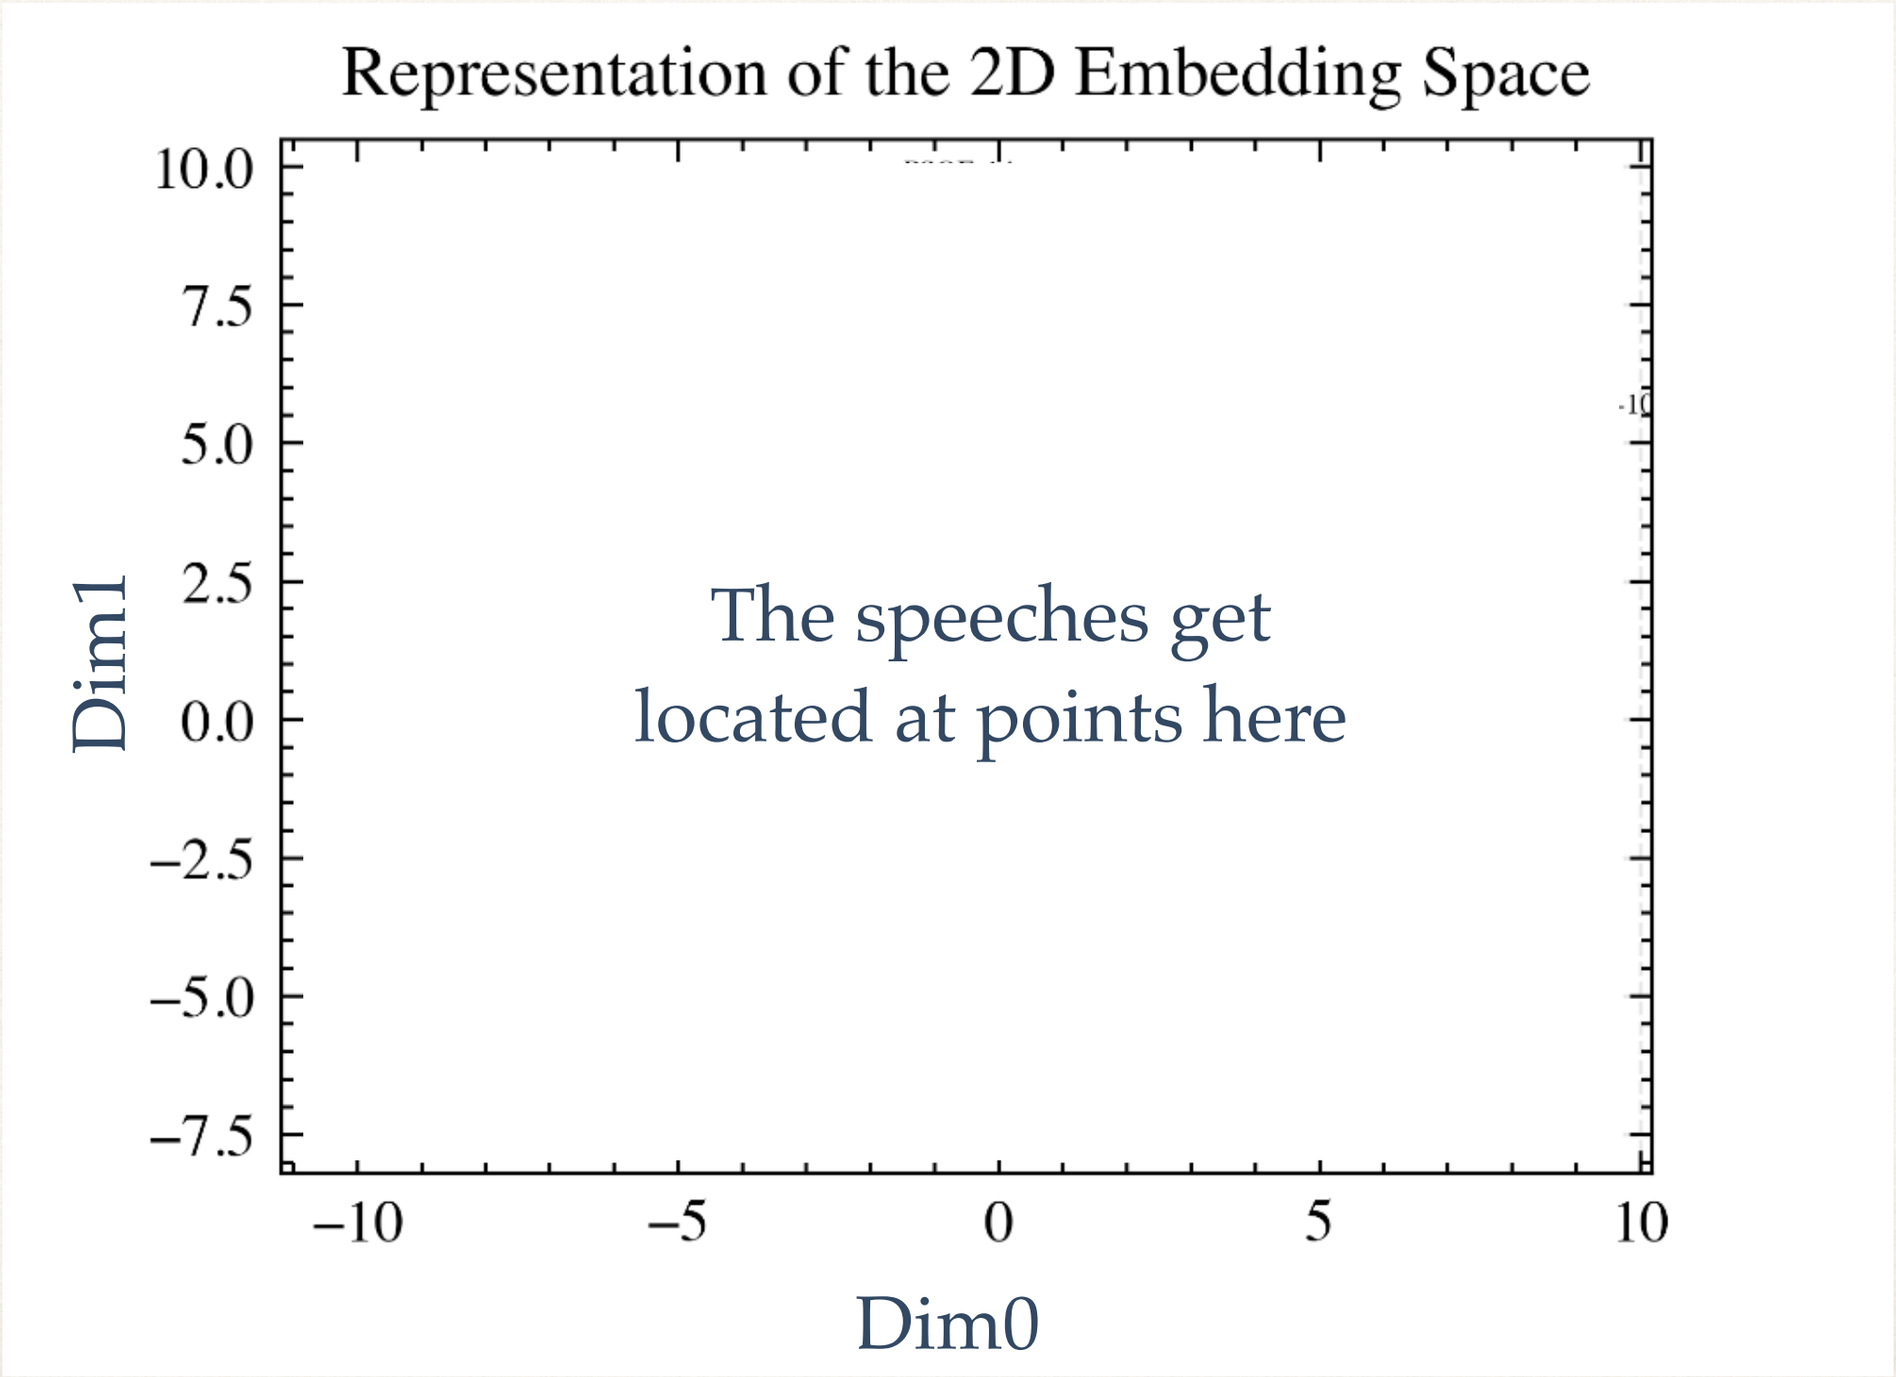

Representation in 2D: clustering

We then use the labels we had kept invisible to the AI to ID political party and term

CLUSTERING occurs, the AI has *understood* party's ideology

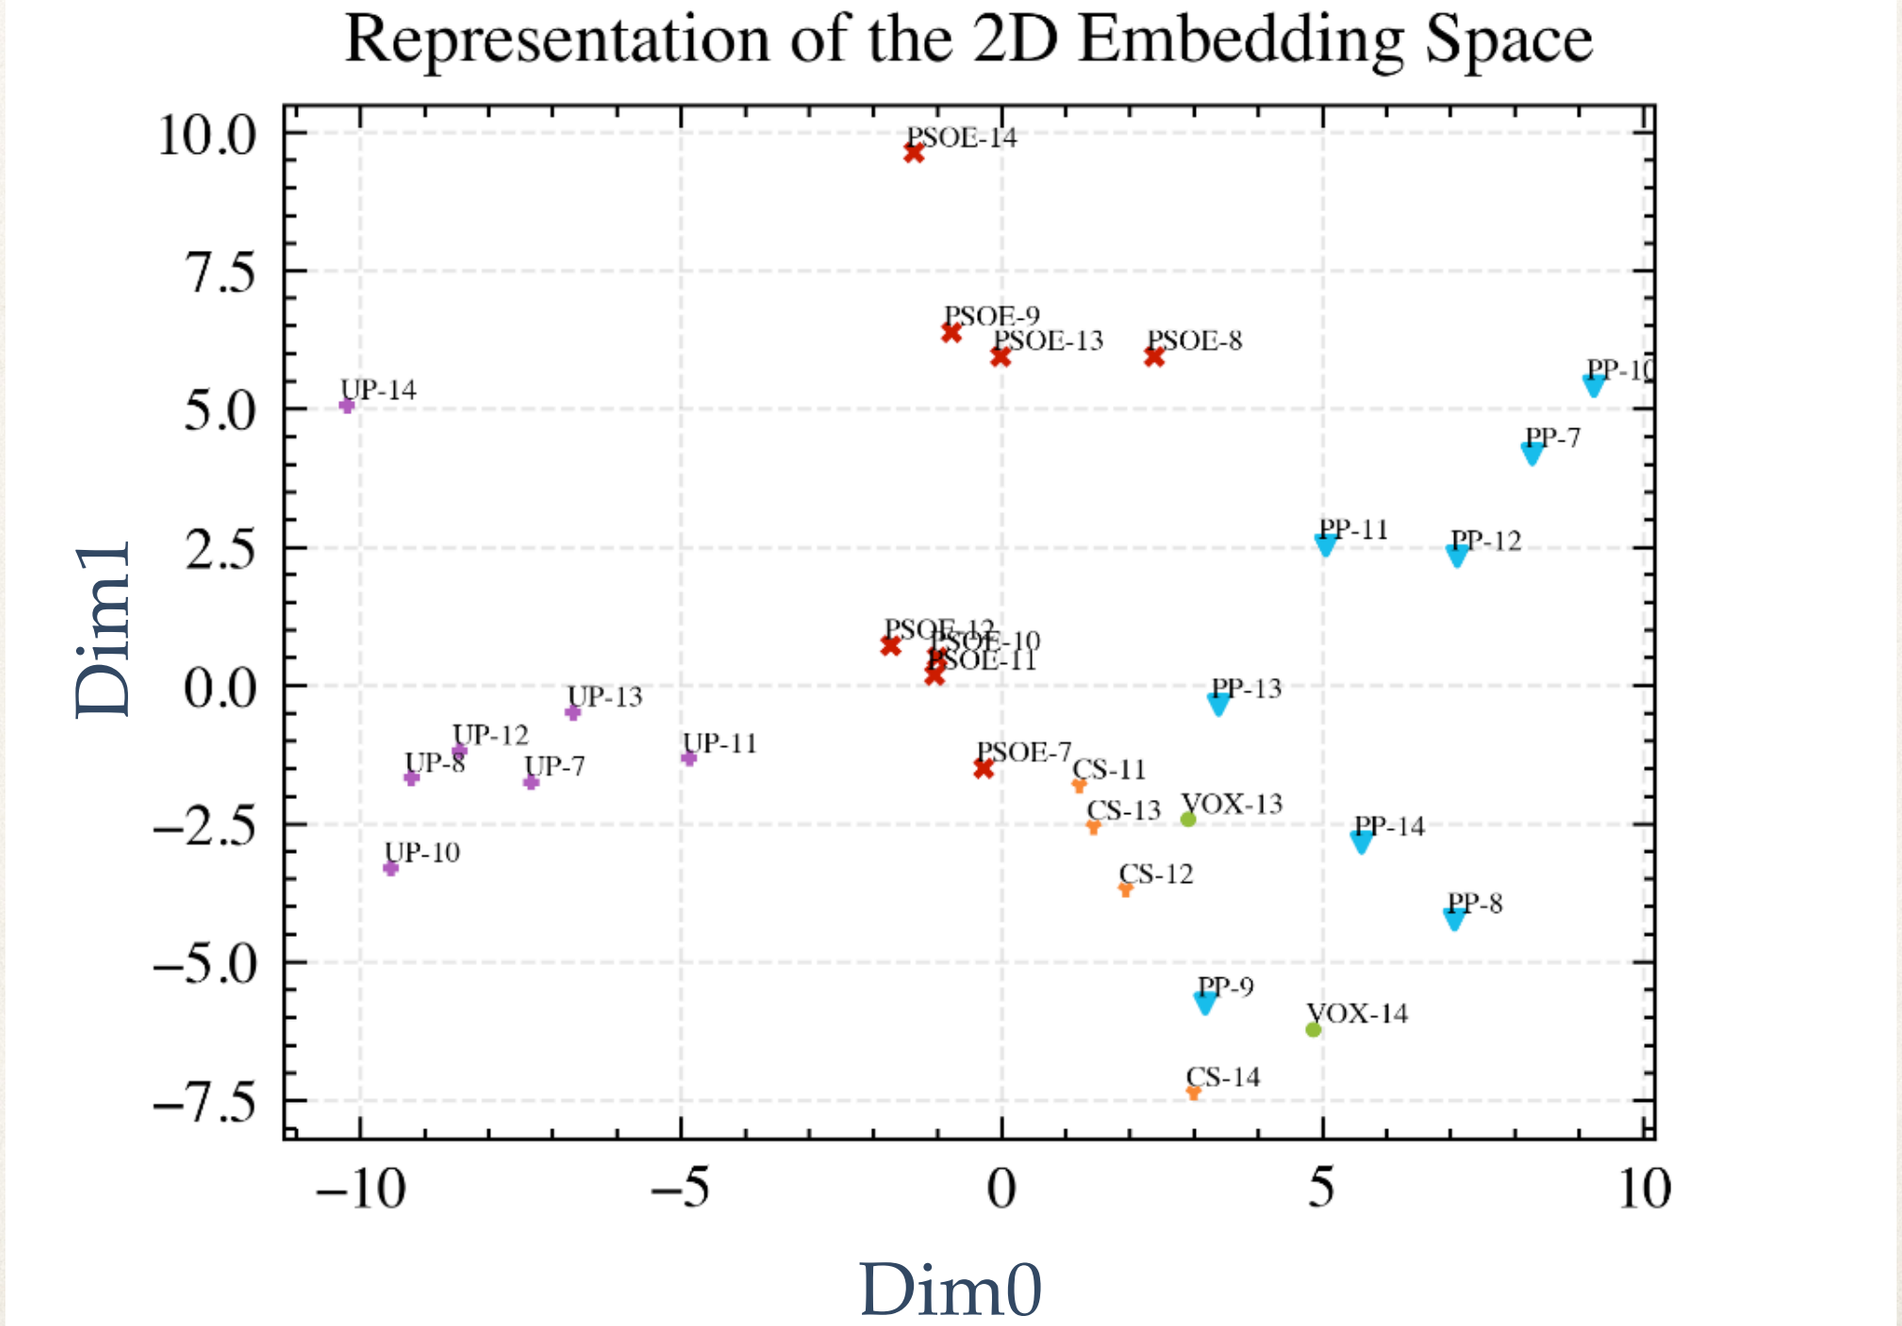

Representation in 2D: meaning

We can then ask if Dim0 and Dim1 have any meaning,
the AI has *understood* divisions of left/right and government/opposition

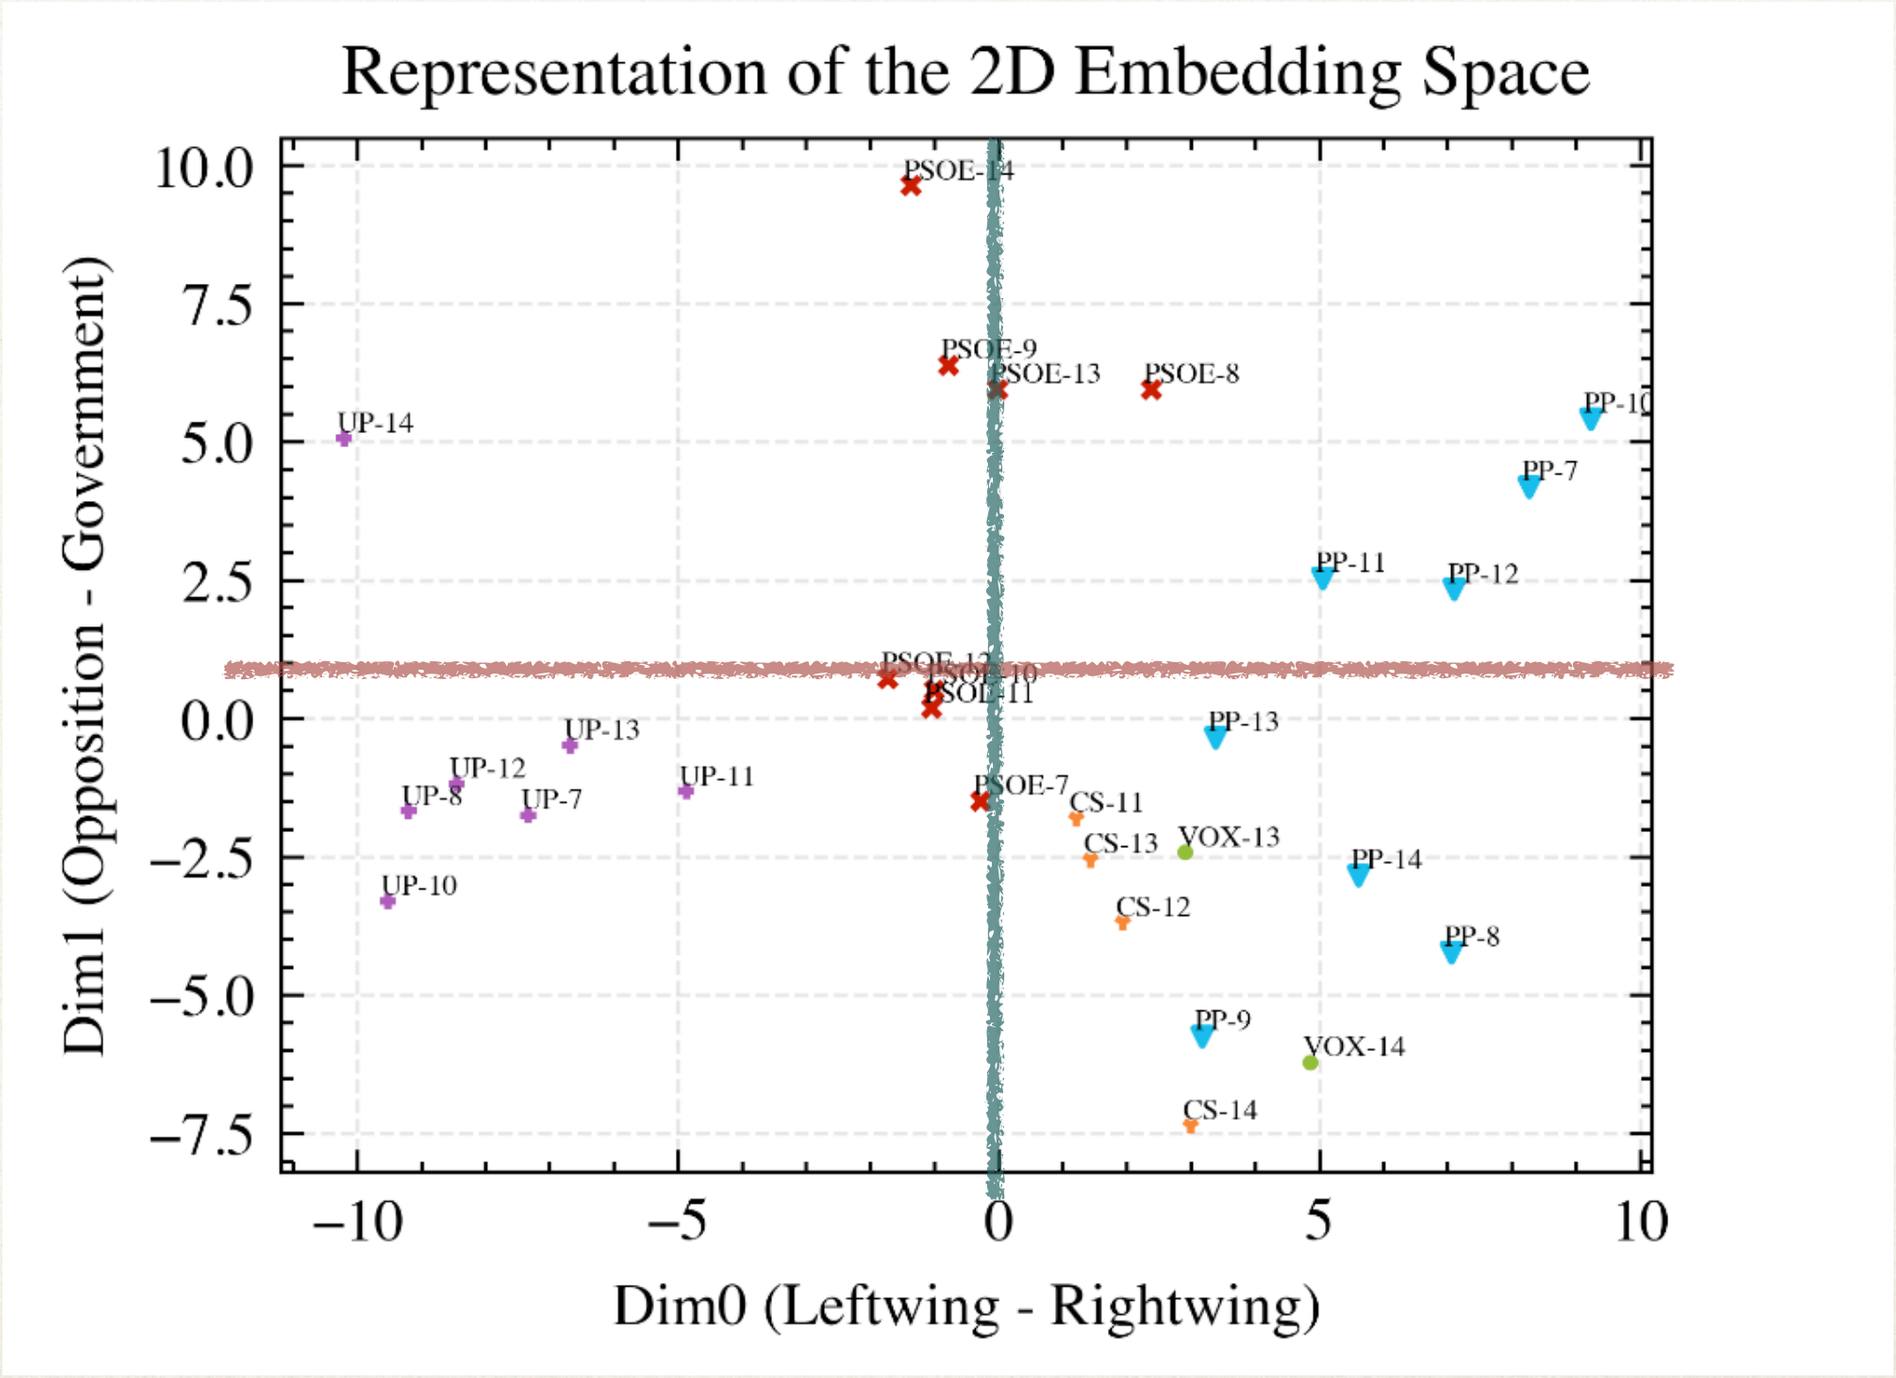

Map of concepts

We can go further, and assign in the 2D space points
concepts that represent them
and trace them back to the party's positions

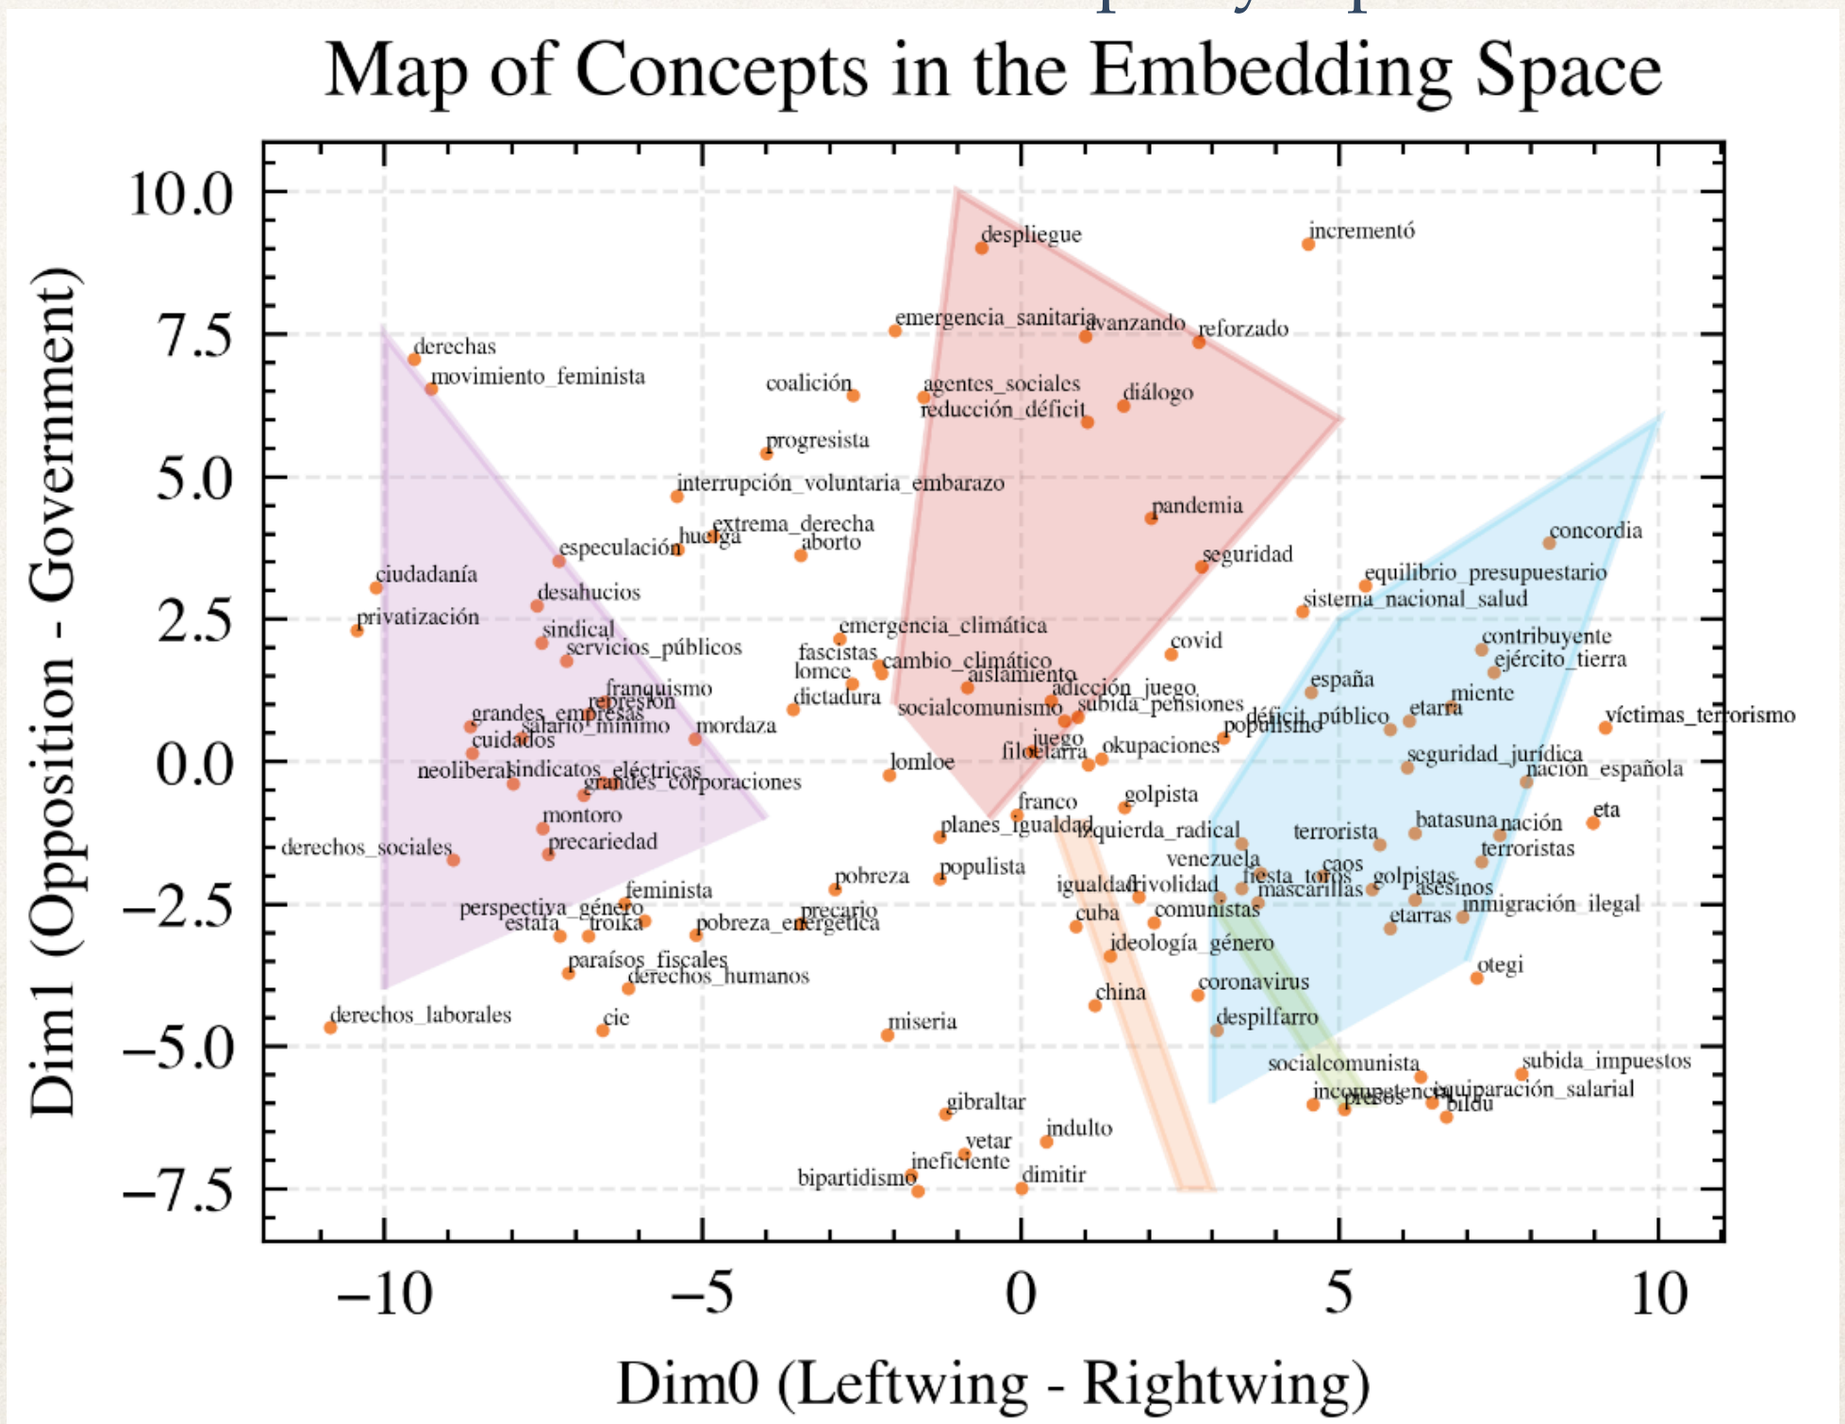

Map of concepts

We can go further, and assign in the 2D space points concepts that represent them and trace them back to the party's positions

Map of Concepts in the Embedding Space

PC1 (Ideological axis)		PC2 (Rhetorical style axis)	
Left (Negative)	**Right (Positive)**	**Bottom (Negative)**	**Top (Positive)**
Equality	Security	Populist Language	Institutional Language
Social Justice	National Unity	Emotive Expressions	Procedural Discourse
Feminism	Constitution	Political Attacks	Legislative Vocabulary
Labor Rights	Territorial Integrity	Crisis Framing	Administrative Language
Environmentalism	Traditional Values	Polarizing Terms	Formal Register

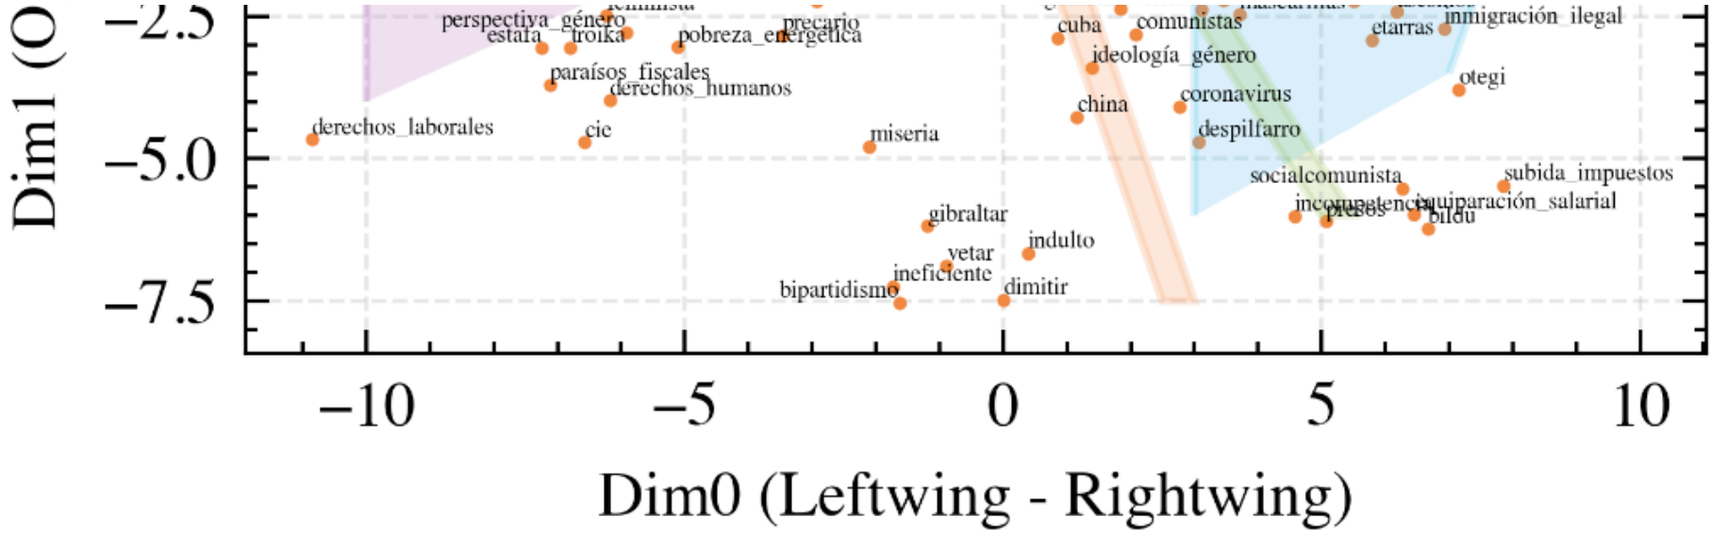

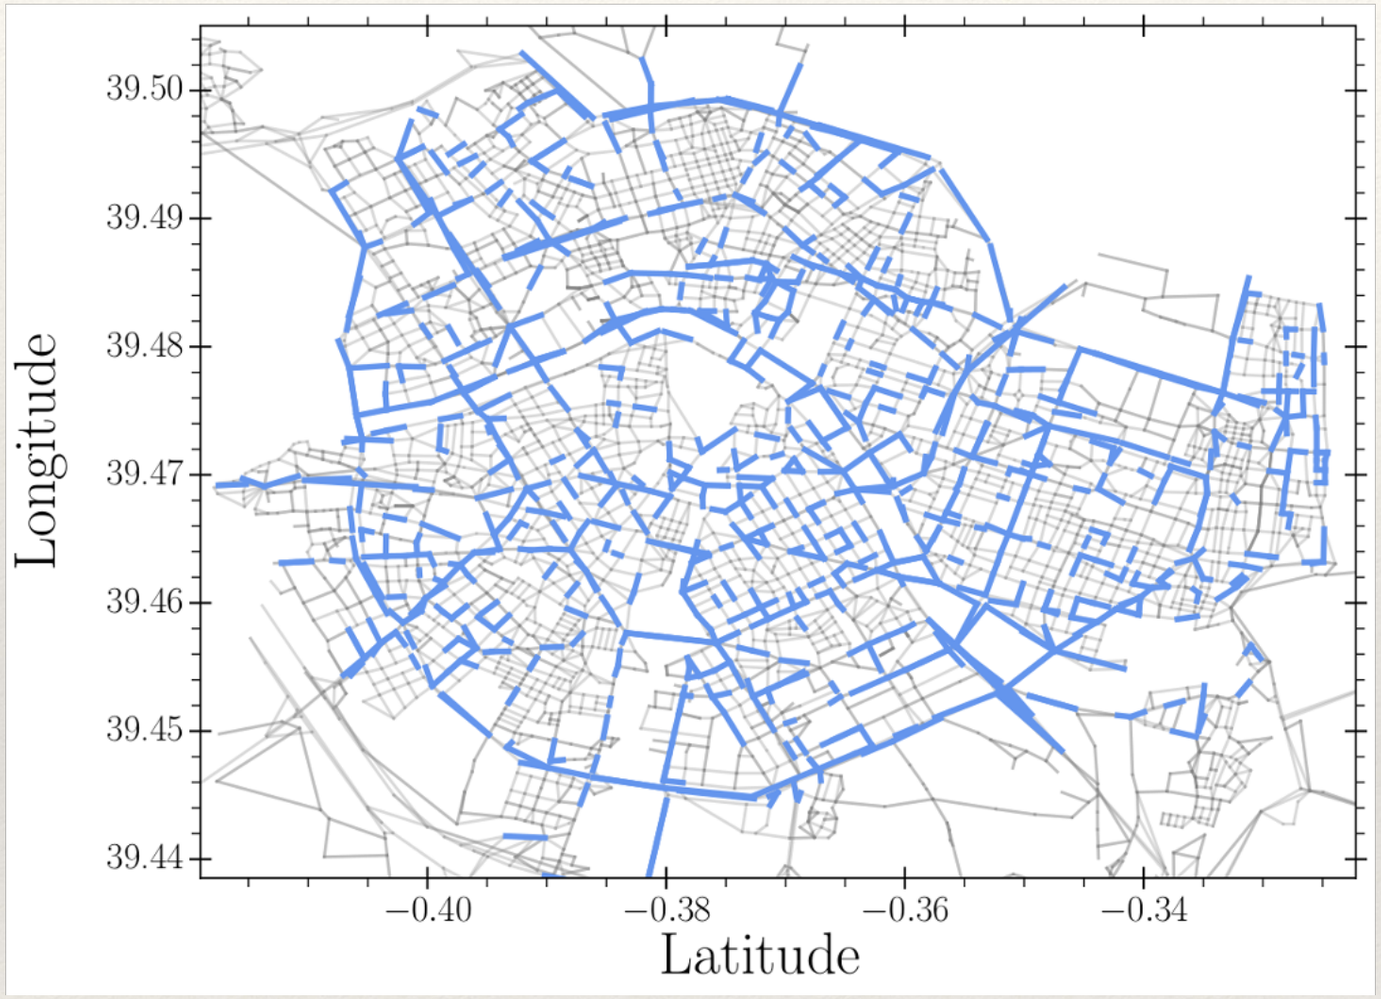

Traffic and pollution

Based on
Applied AI Letters (2025)
Neural Computing and Applications (2025)

Valencia

Nice city

One of the most sensorized in the world

About 9000 road segments collecting data on passing vehicles every 5 mins

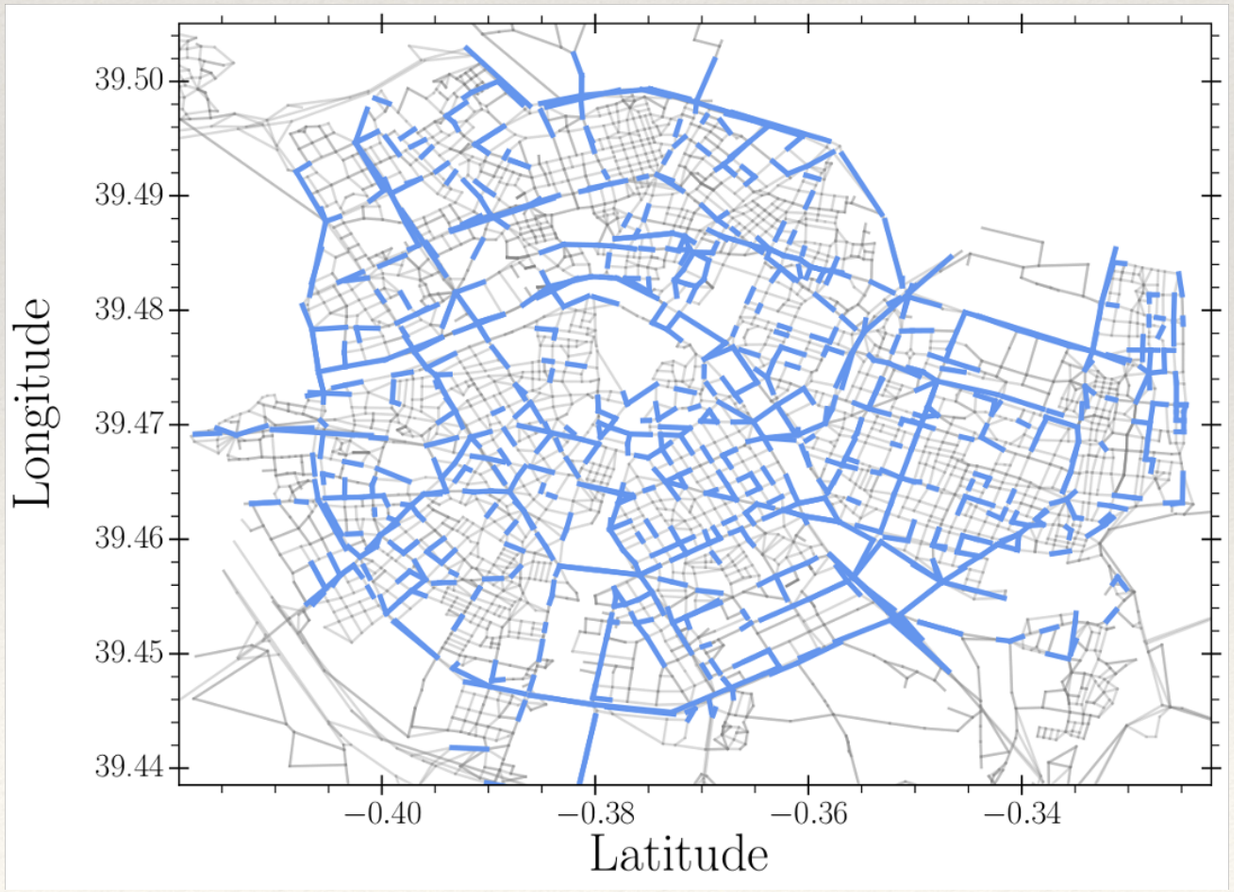

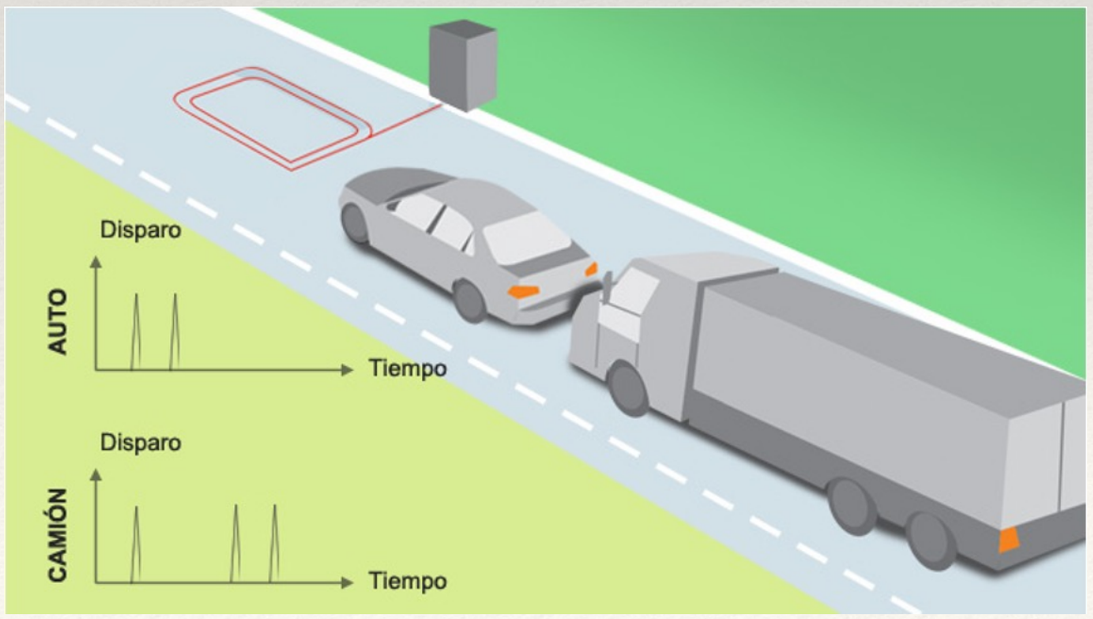

Common project across institutions

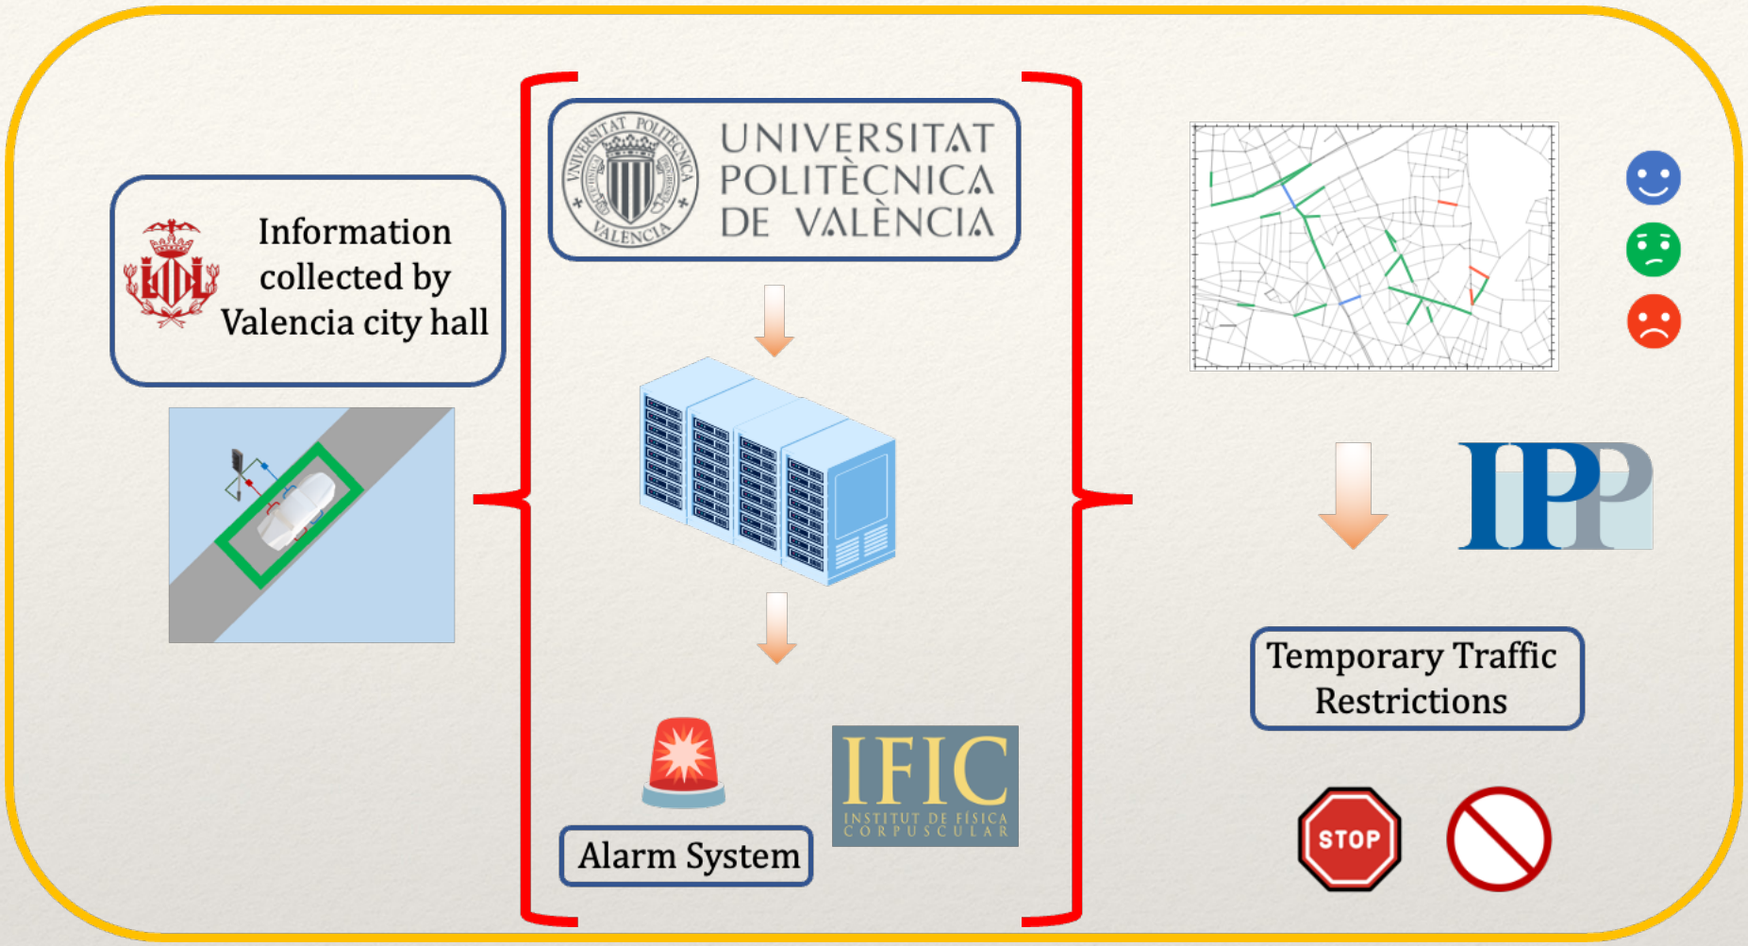

City govt. **Data+ML** **Sociologists**

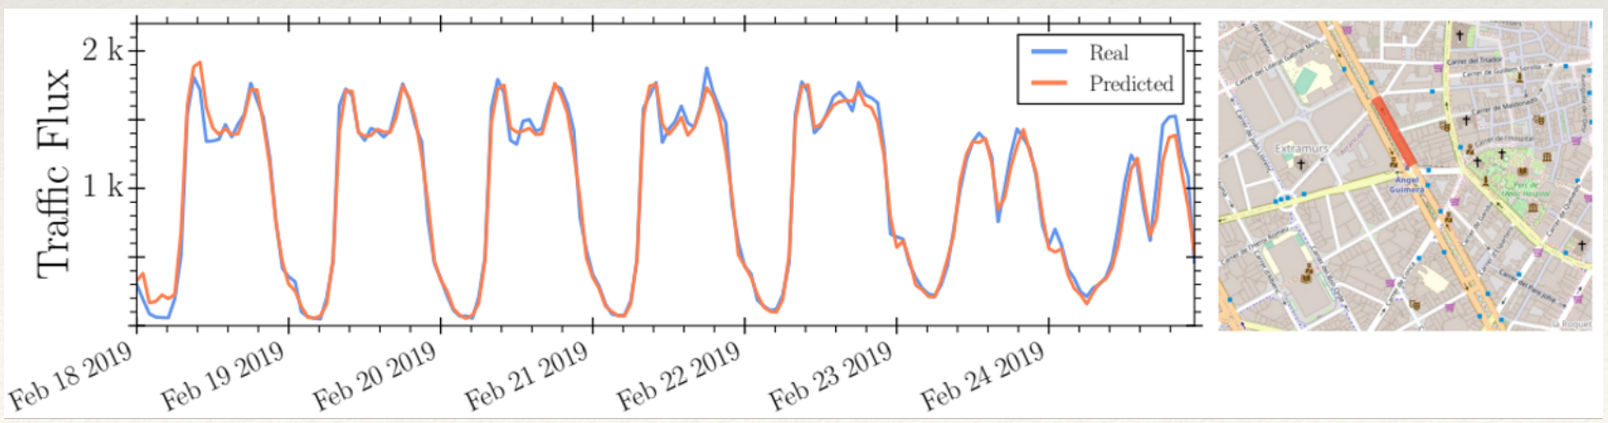

We used arrays of LSTMs (despite coverage, GraphNN or more sophisticated techniques did not perform as well)

Wrapping up…

We are just starting to explore the applications of ML in Fundamental Physics

They go beyond a mere iteration of our traditional statistical methods:
unsupervised methods, generative AI, reinforcement learning…

A very efficient blackbox is not good enough for us,
we try to *communicate* with the AI,
to find ways to understand its inner workings

We learned that an AI can identify and use higher level concepts,
and this learning can be found in subtle features of the blackbox

We have seen that through AI methods, there is interesting
cross-pollination between our area and others (music->LHC)

I believe that moving forward,
more thought needs to be placed on XAI methods
to make the AI inner workings closer to human intuition

Thanks for listening!

Questions?

Example 1 PDE-Nets

Navier Stokes, Maxwell, Schrodinger…

Laws of Nature are PDEs

But there are other areas in physics where equations are not known and even when we know them, they include assumptions

Pose an inverse problem: given an observed temporal distribution, can a NN learn the non-linear response?

$$u_t = F(x, u, \nabla u, \nabla^2 u, \ldots), \quad x \in \Omega \subset \mathbb{R}^2, \quad t \in [0, T].$$

$$u_t(t, x, y) = F(x, y, u, u_x, u_y, u_{xx}, u_{xy}, u_{yy}, \ldots), \quad (x, y) \in \Omega \subset \mathbb{R}^2, t \in [0, T].$$

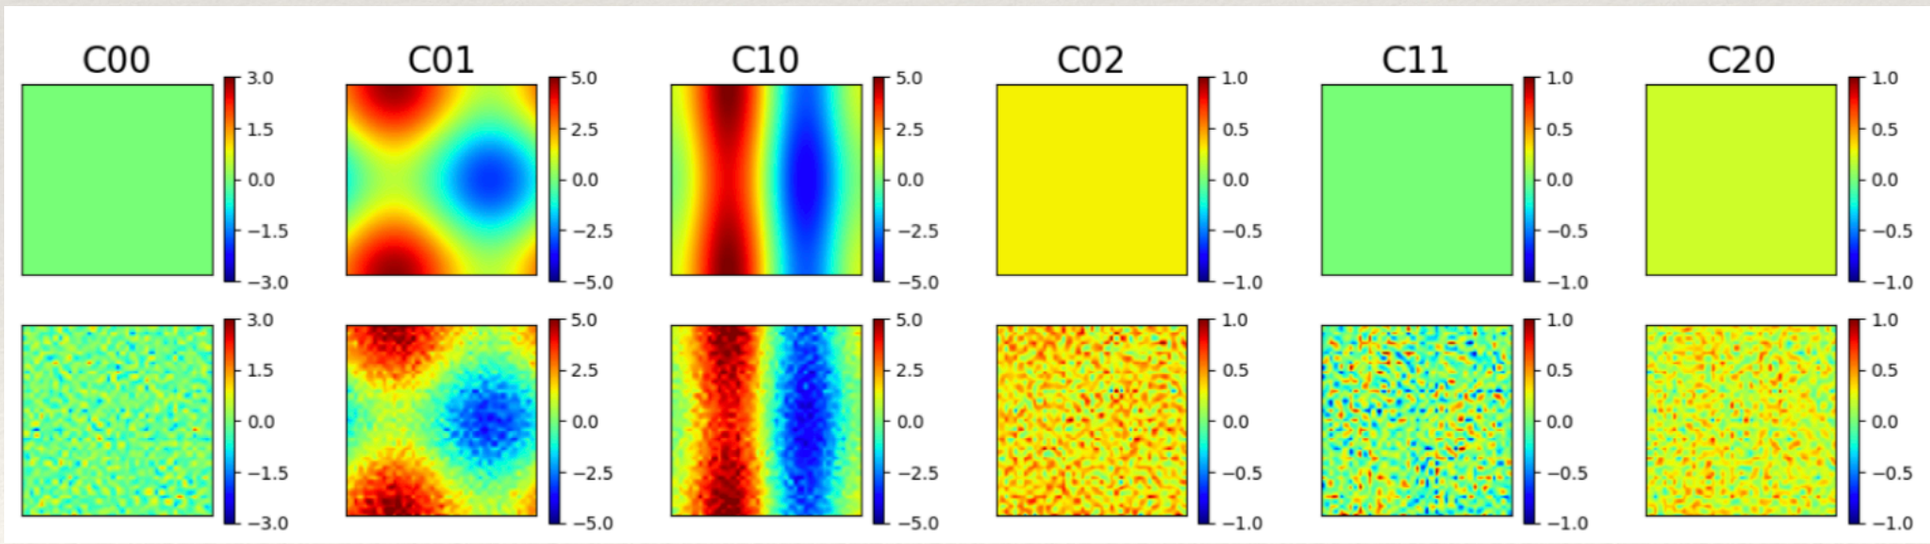

e.g PDE-Net, Long&Dong, 1812.04426