

JAVIER DUARTE
NOVEMBER 12, 2019
UNIVERSITY OF KANSAS

DEEP LEARNING AT THE LHC

THE LARGE HADRON COLLIDER



SUISSE
FRANCE

CMS

LHCb

ATLAS

CERN Meyrin

CERN Prévessin

SPS 7 km

ALICE


LHC 27 km

THE LARGE HADRON COLLIDER



proton-proton collider @ 13 TeV center-of-mass energy

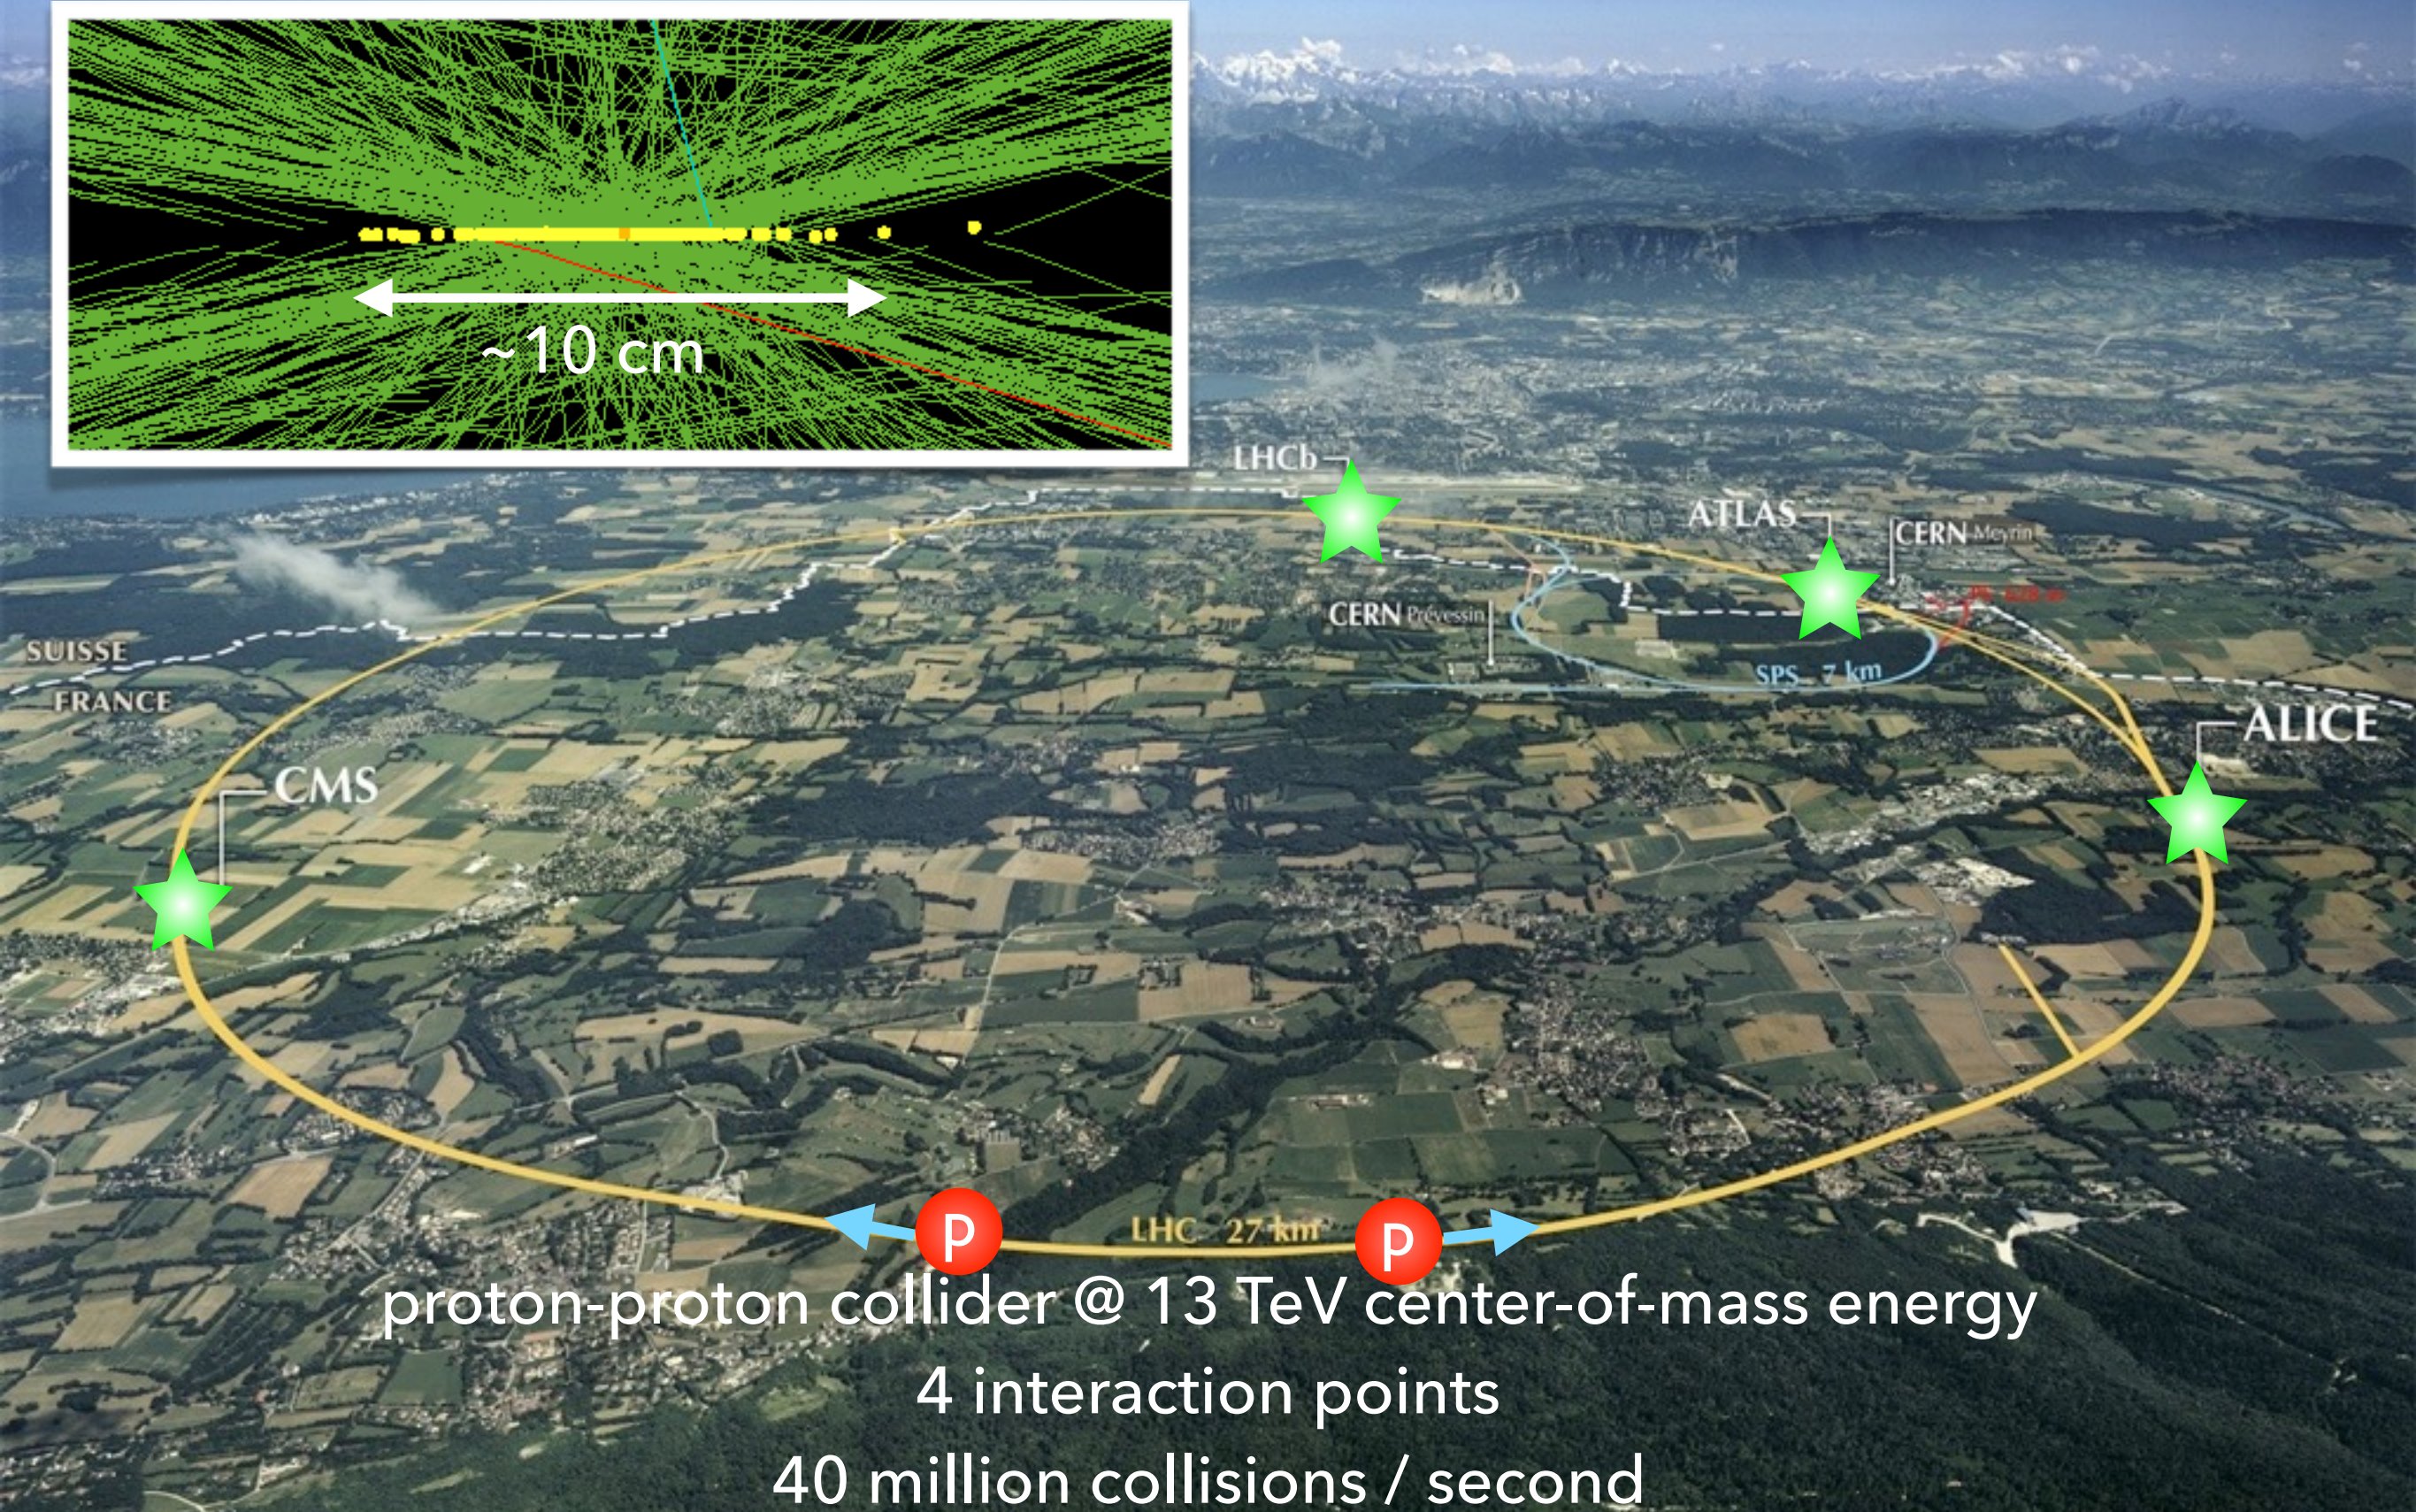

THE LARGE HADRON COLLIDER



LHC 27 km

proton-proton collider @ 13 TeV center-of-mass energy
4 interaction points

THE LARGE HADRON COLLIDER

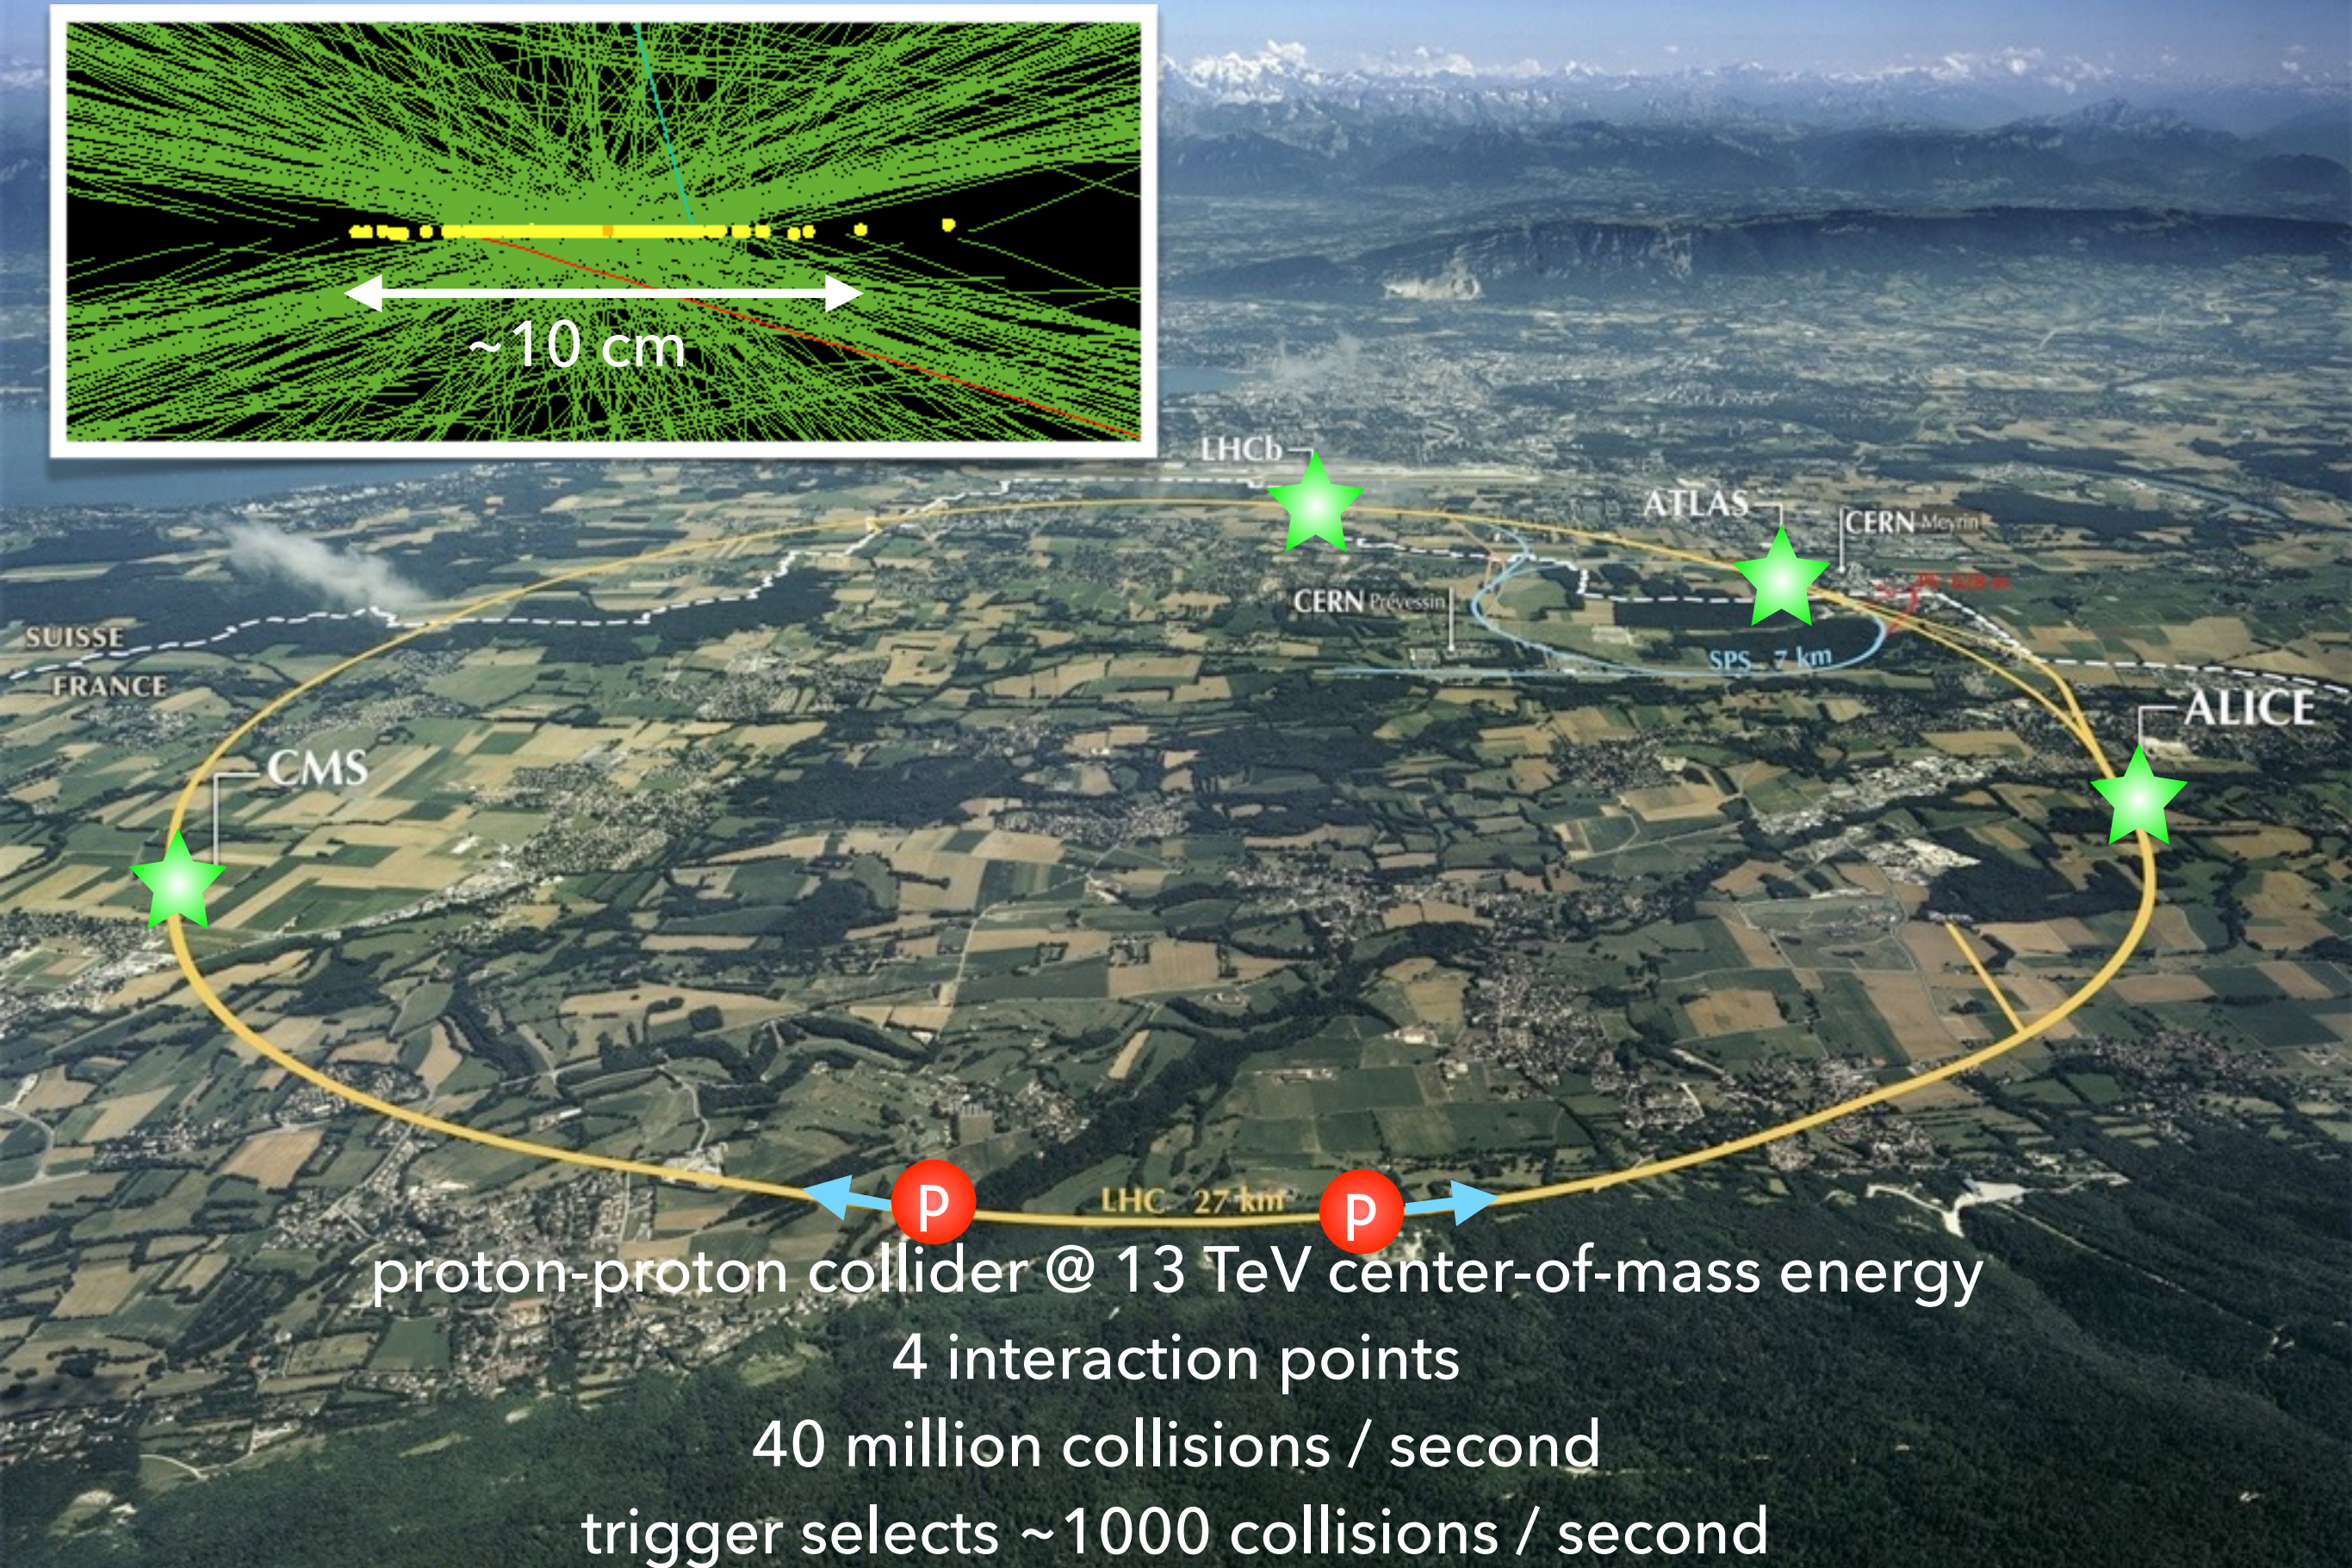



proton-proton collider @ 13 TeV center-of-mass energy

4 interaction points

40 million collisions / second

THE LARGE HADRON COLLIDER



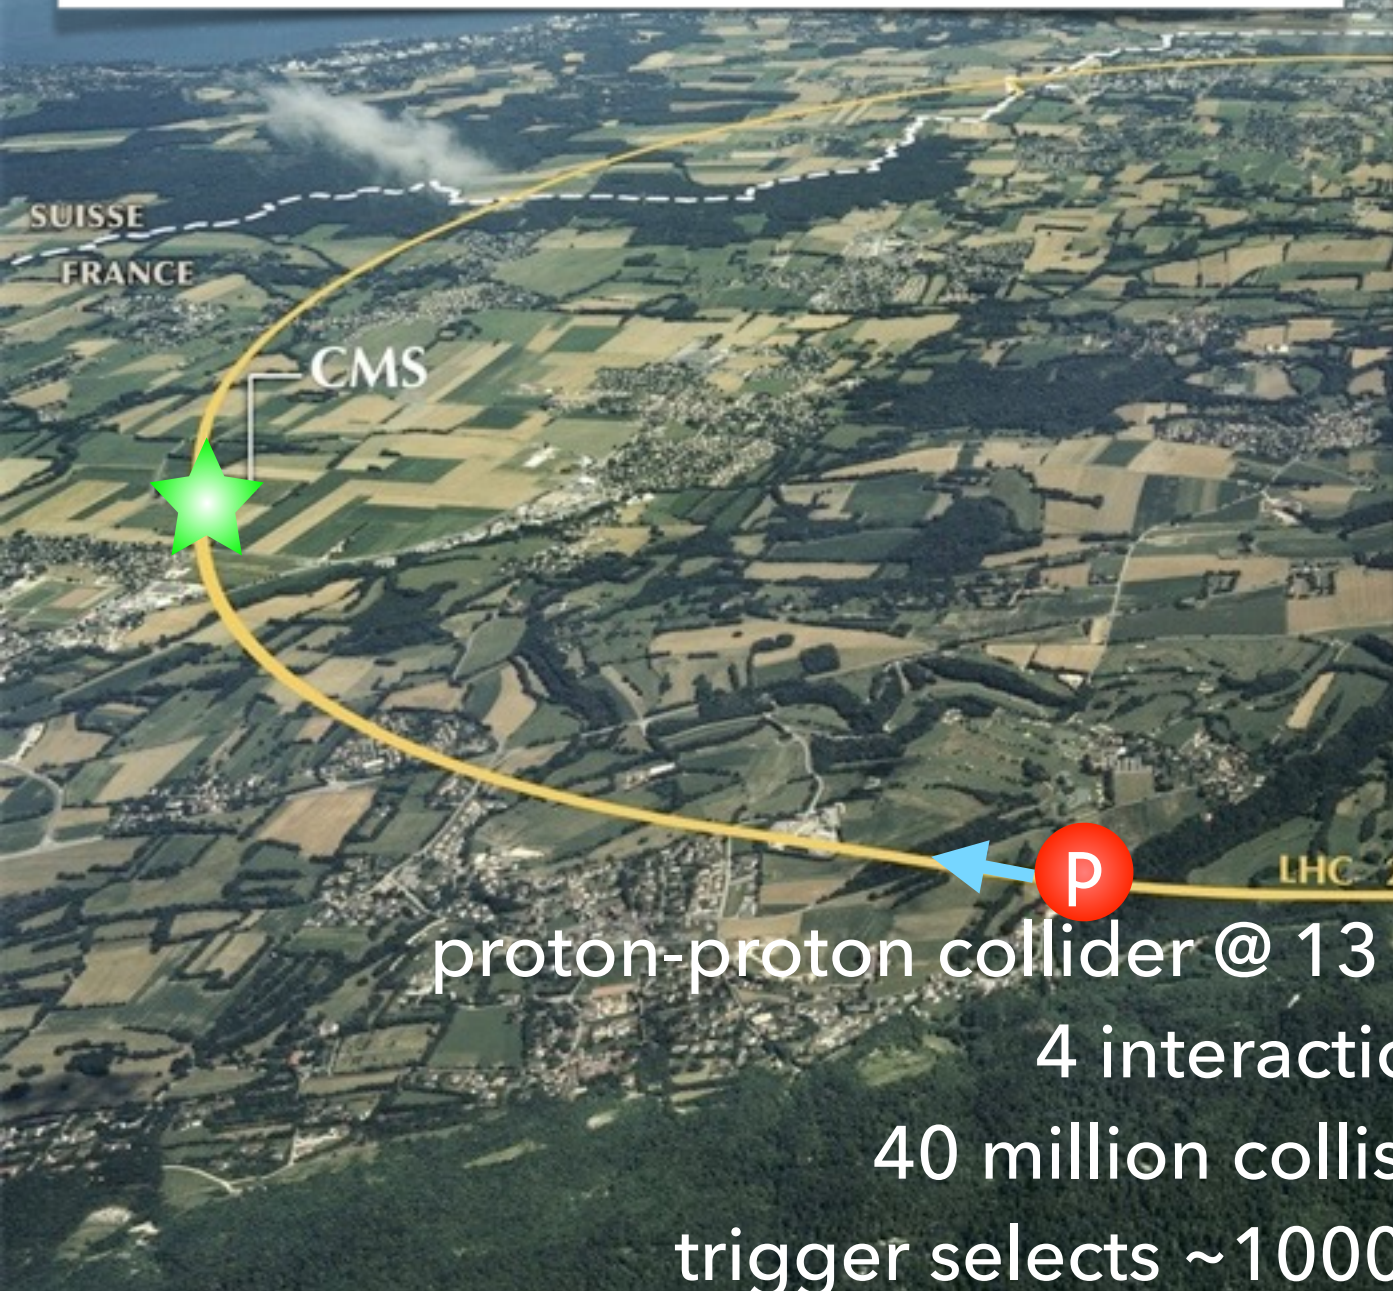


proton-proton collider @ 13 TeV center-of-mass energy

4 interaction points


40 million collisions / second

trigger selects ~1000 collisions / second

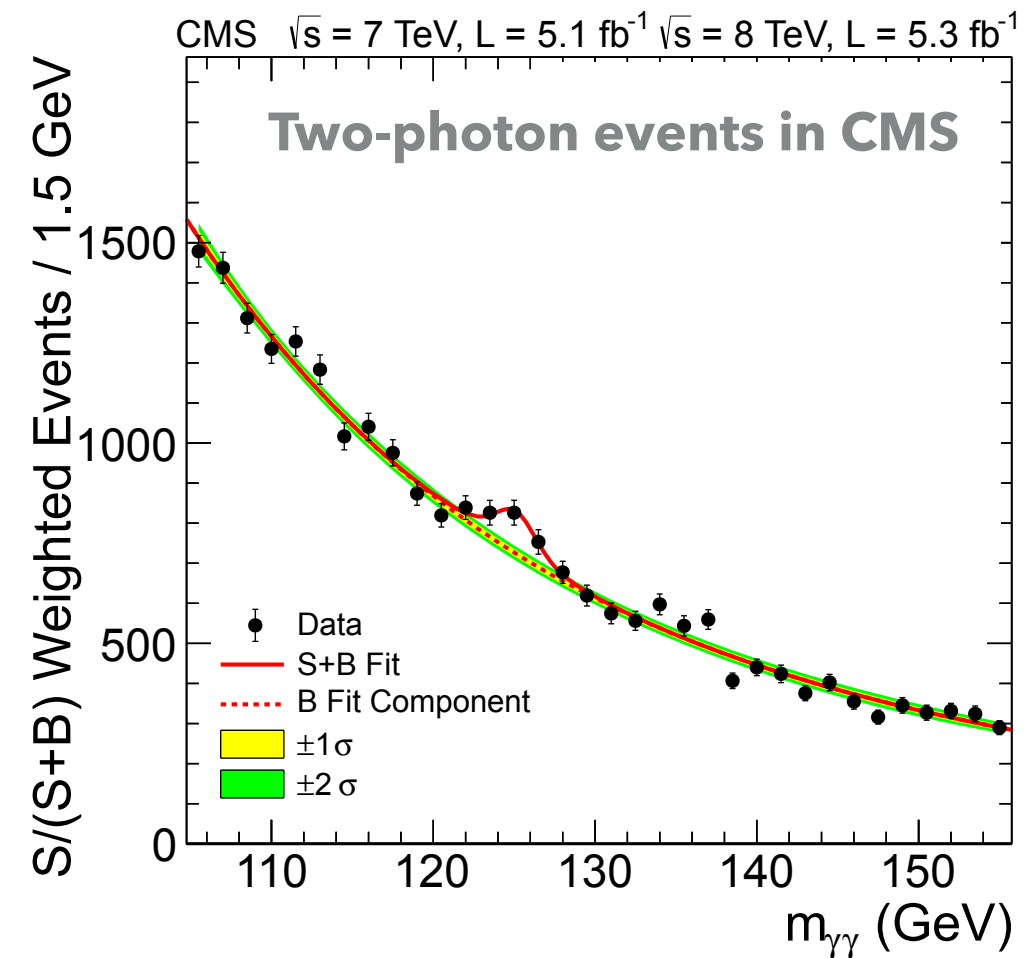
THE LARGE HADRON COLLIDER





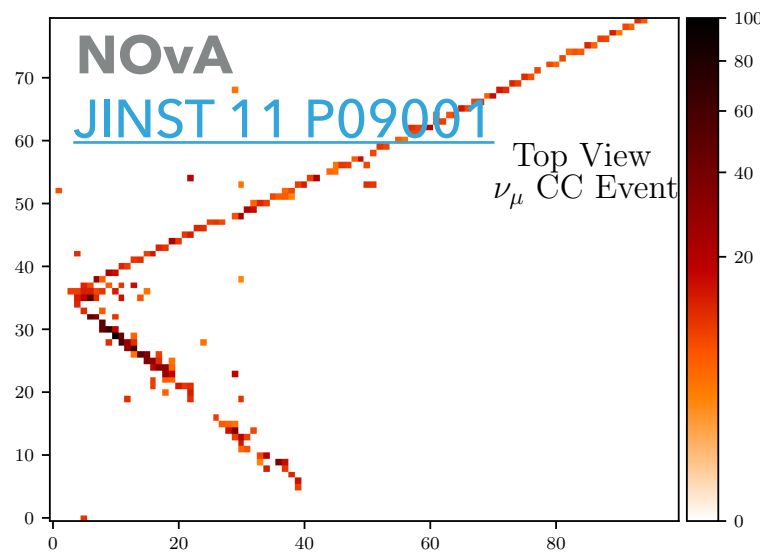

proton-proton collider @ 13 TeV center-of-mass energy
4 interaction points
40 million collisions / second
trigger selects ~1000 collisions / second



- ▶ **Machine learning** was vital to make big discoveries like the Higgs boson on July 4, 2012

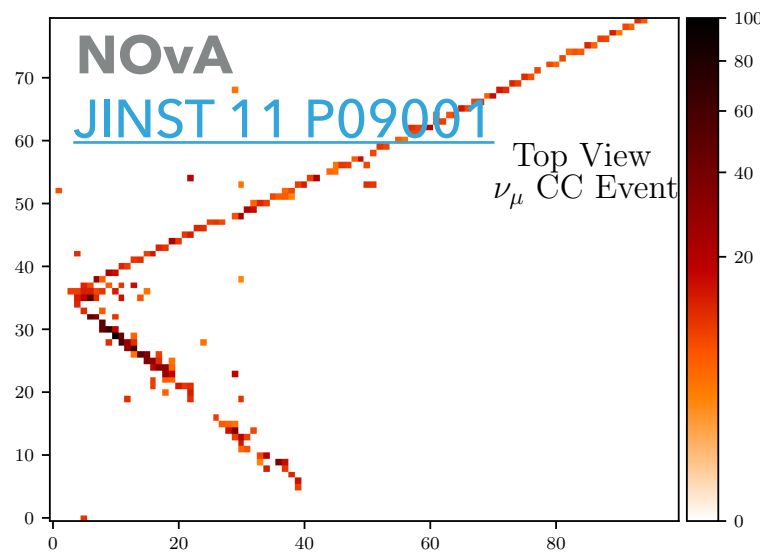



- ▶ **Machine learning** was vital to make big discoveries like the Higgs boson on July 4, 2012
- ▶ Today, ML is **enabling** new detection techniques, measurements, and searches





HEP.TrkX

- ▶ **Machine learning** was vital to make big discoveries like the Higgs boson on July 4, 2012
- ▶ Today, ML is **enabling** new detection techniques, measurements, and searches



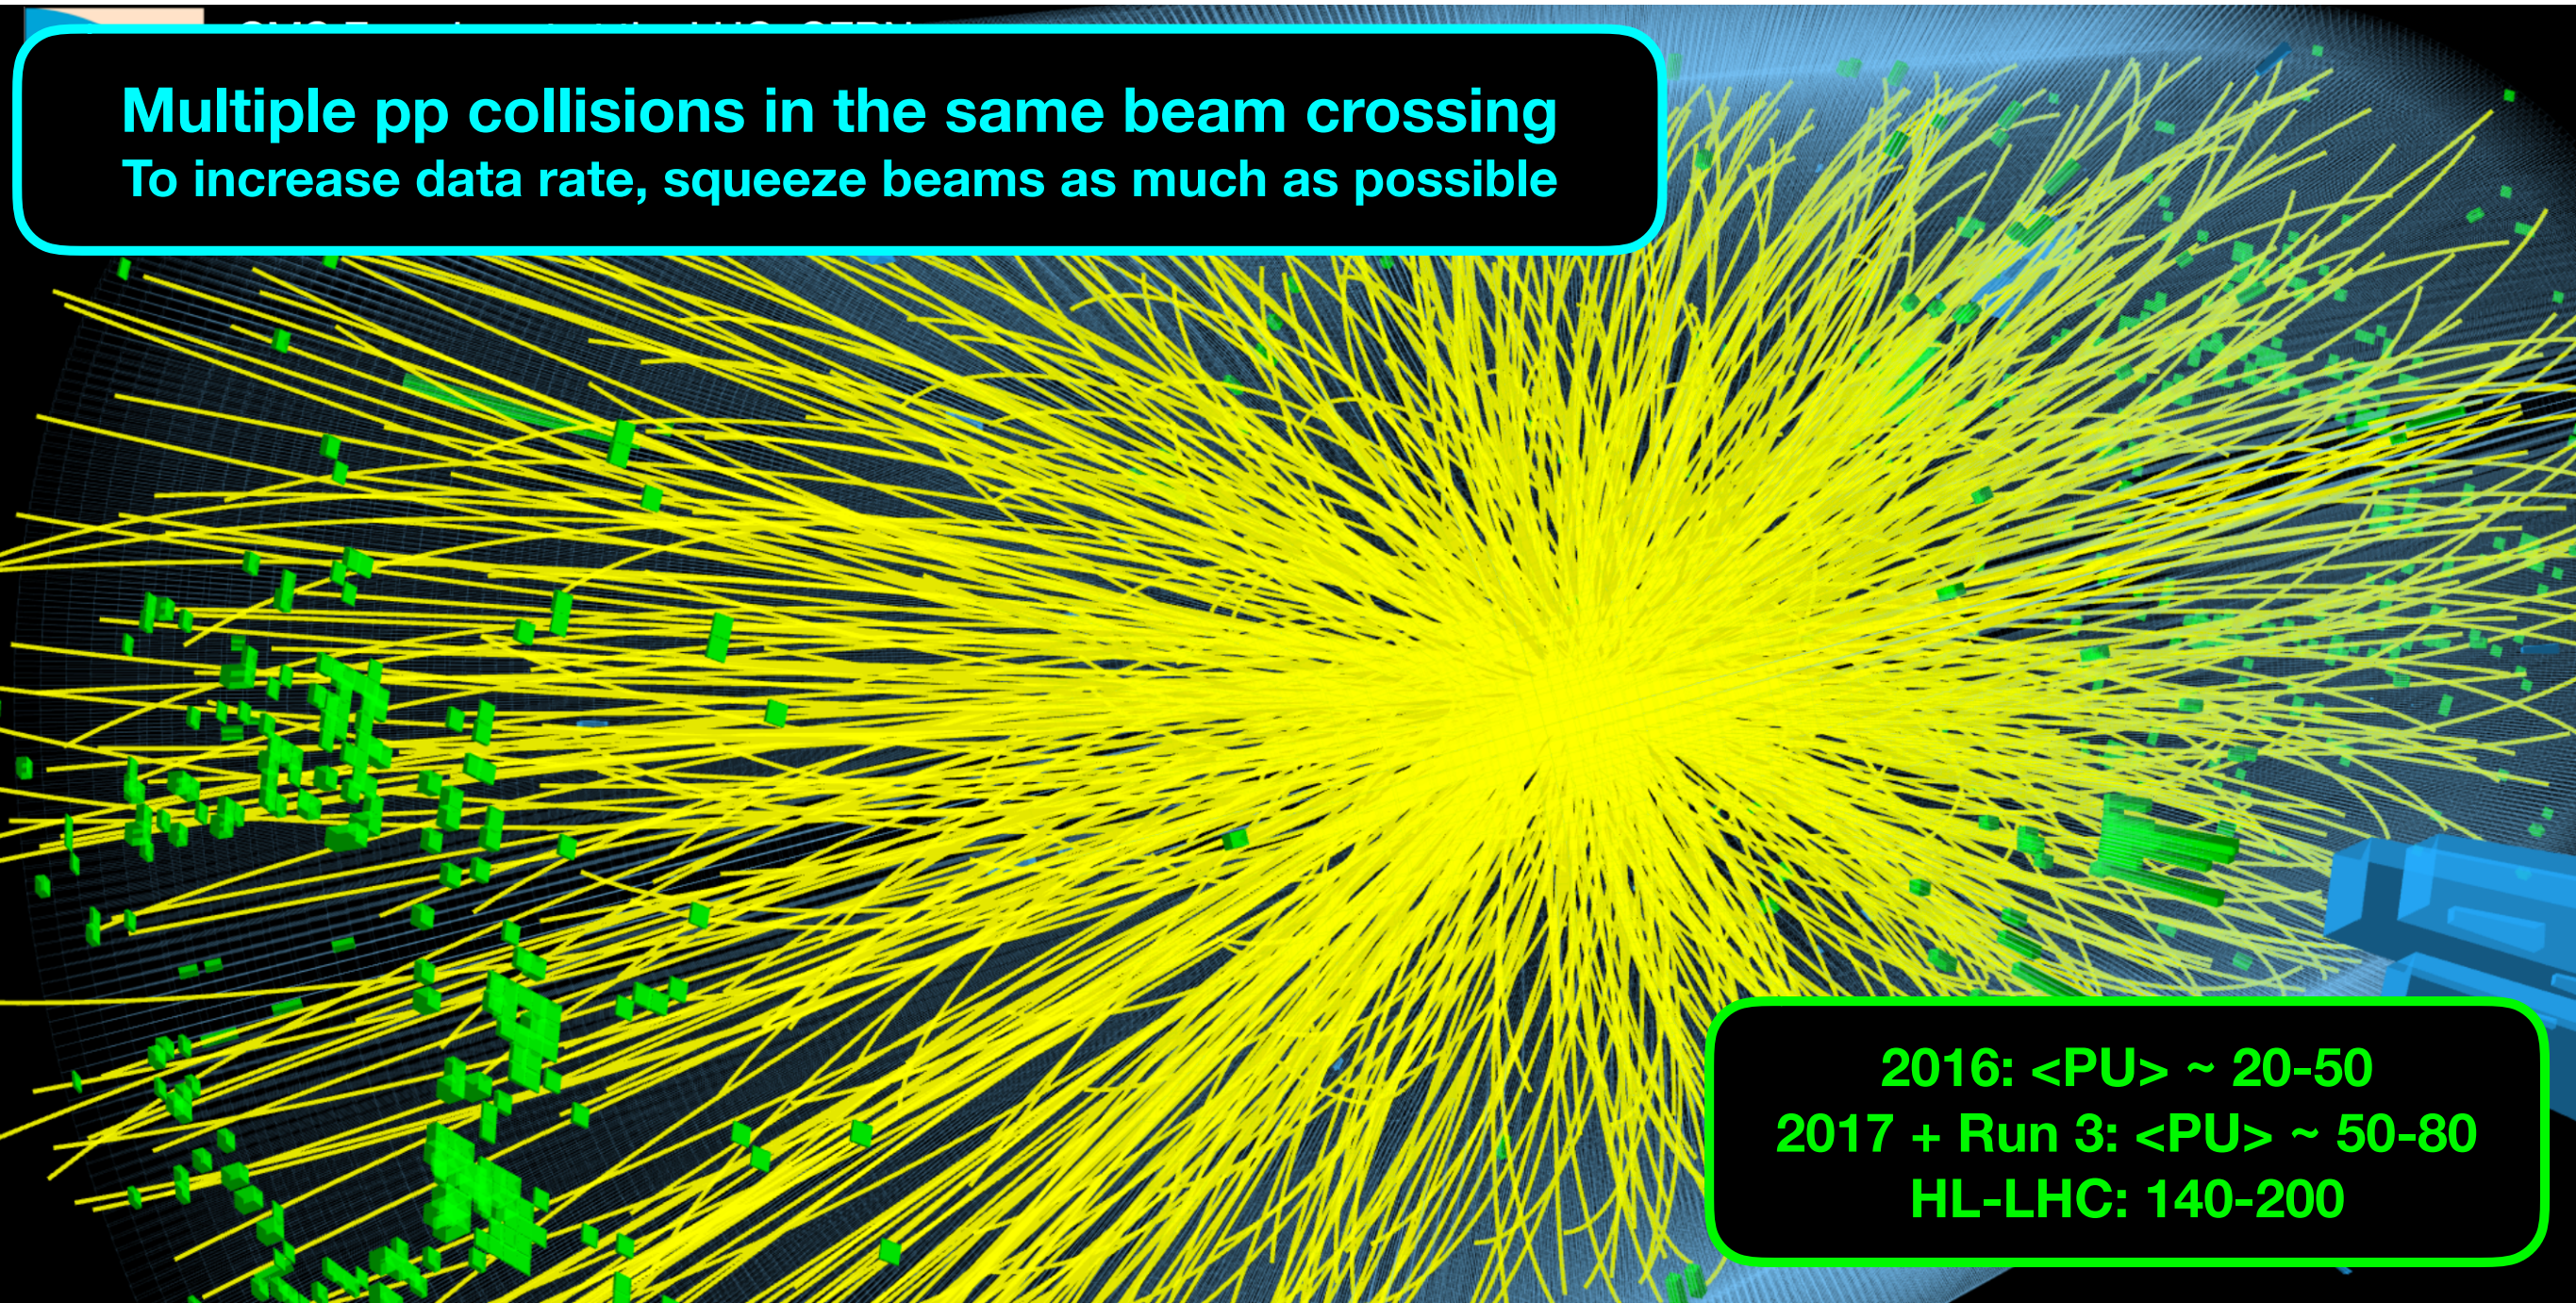
- ▶ At the same time, we must **plan** how we will overcome challenges in the **next generation of experiments**

- ▶ ML may be a way out




HEP.TrkX

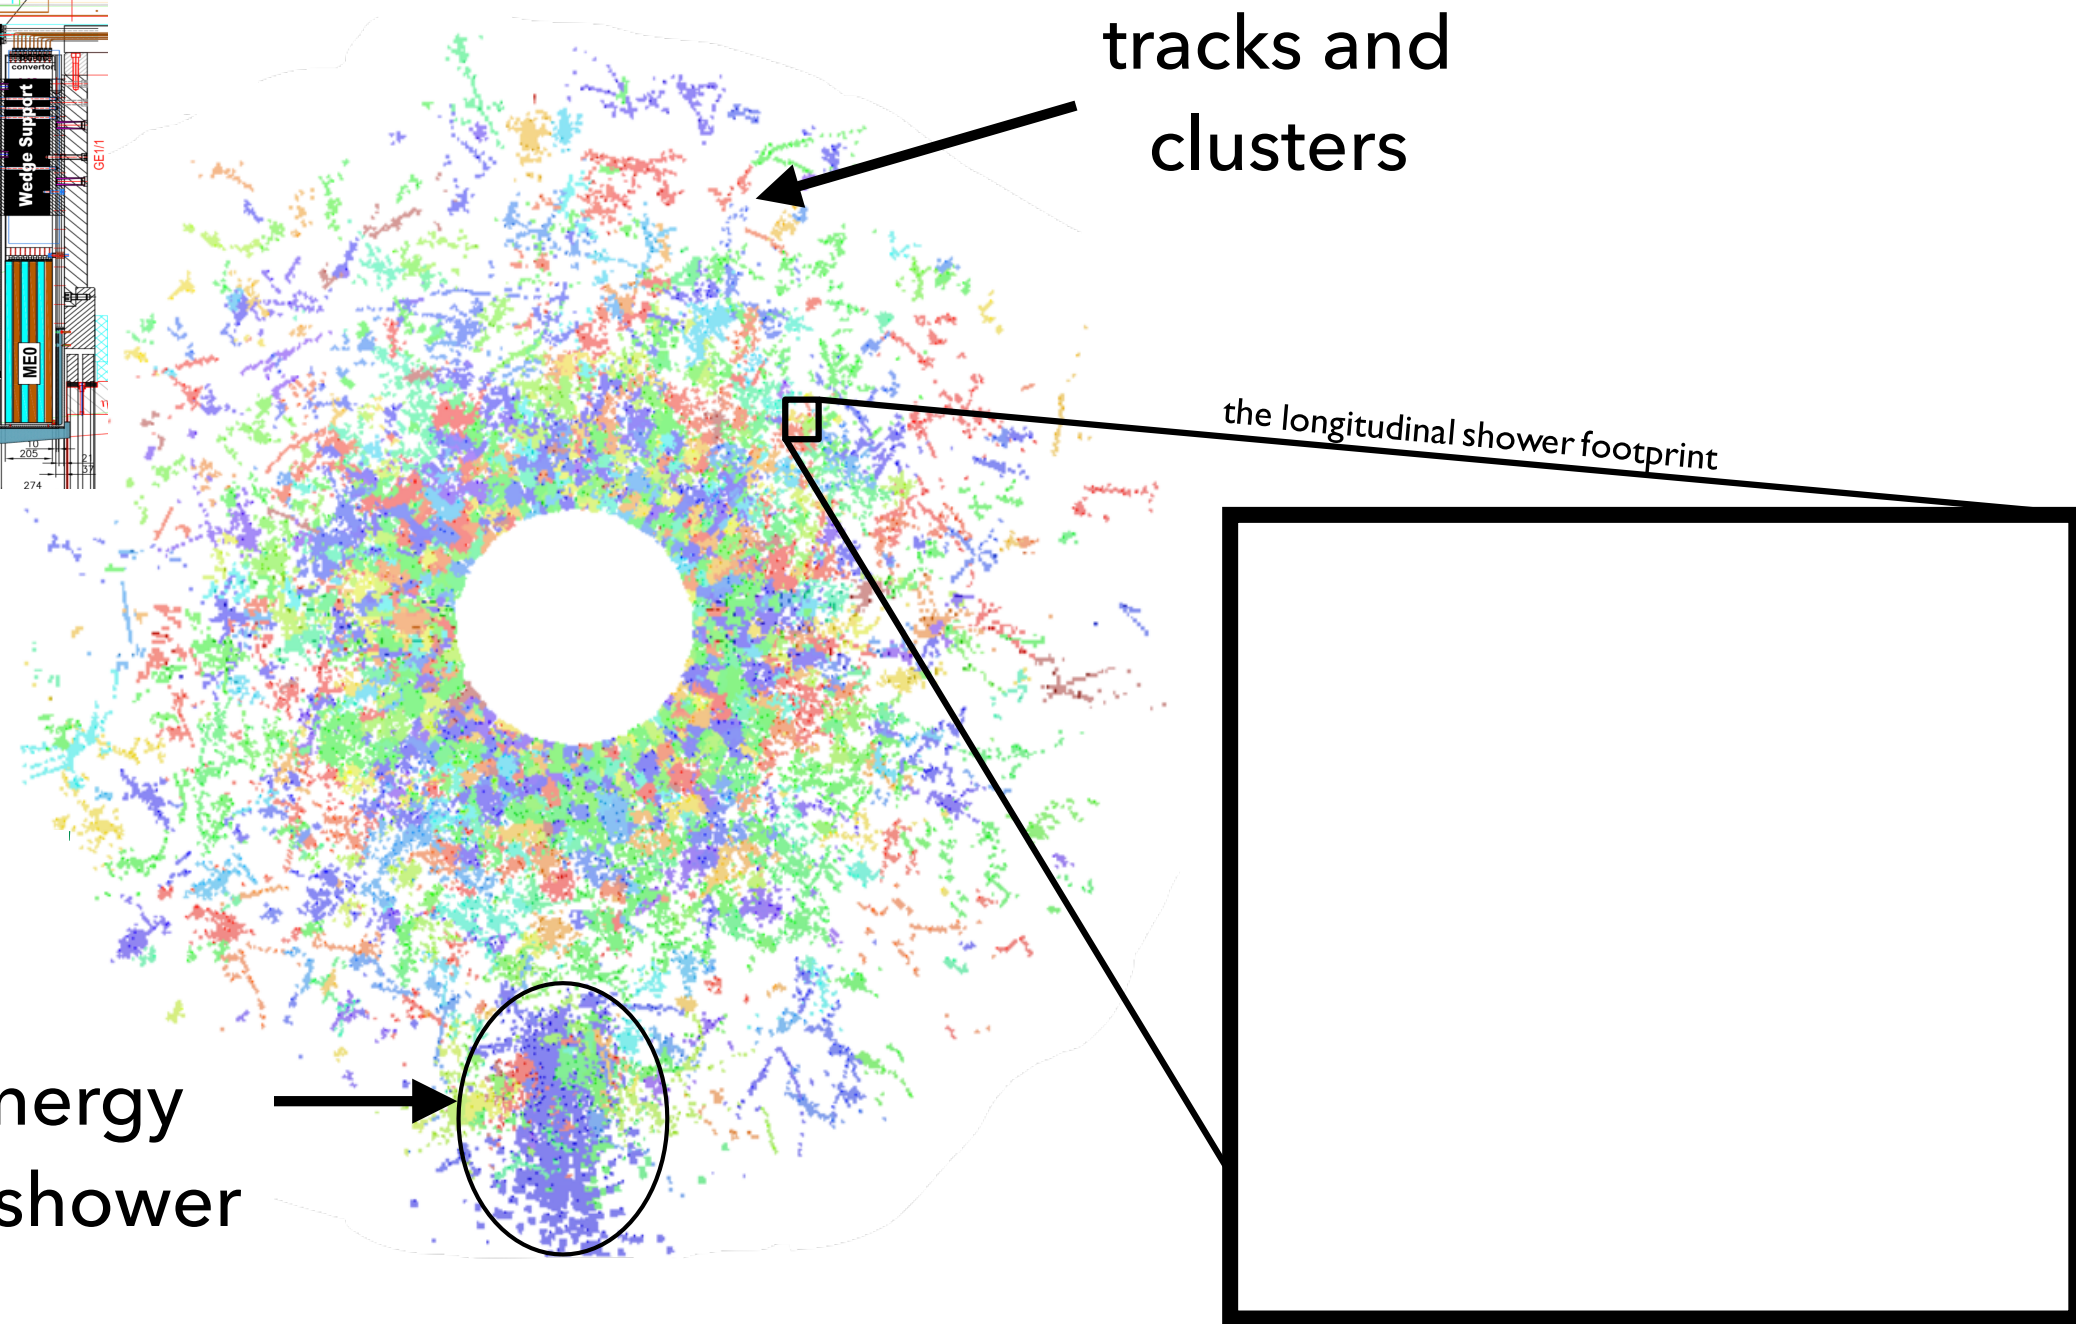
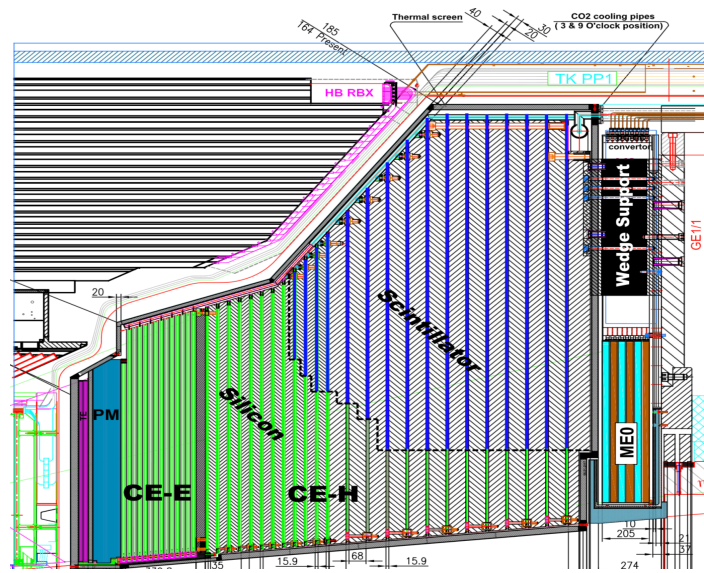
Multiple pp collisions in the same beam crossing
To increase data rate, squeeze beams as much as possible



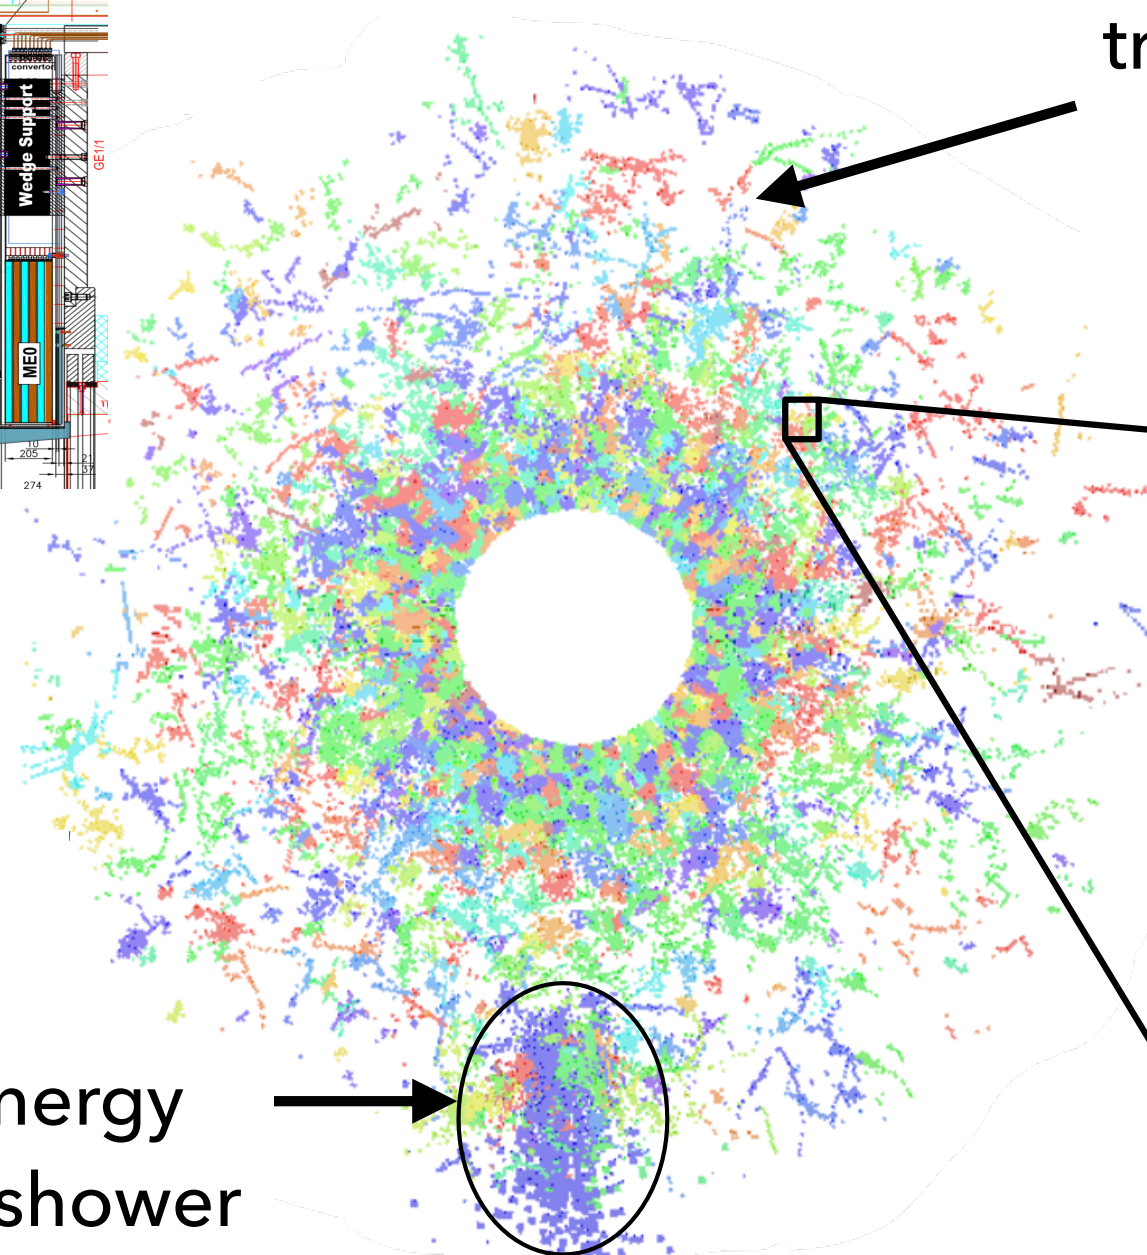
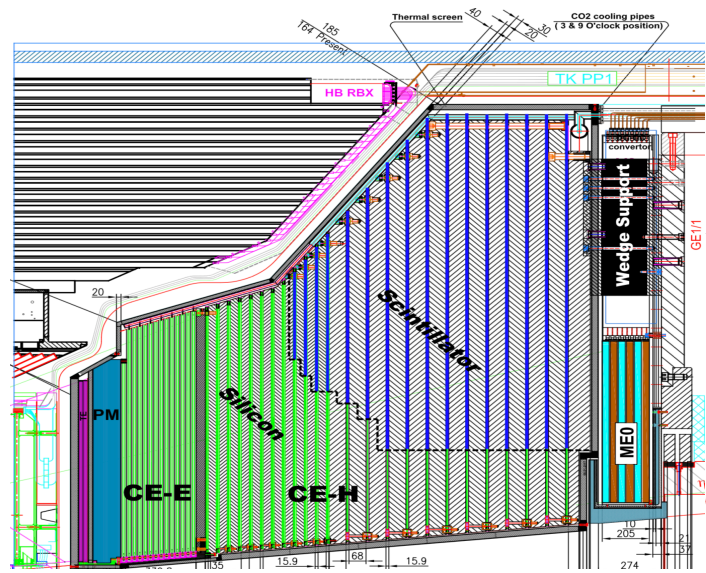
- ▶ At high luminosity, many collisions happen simultaneously (pileup)!
- ▶ Pileup makes our data more complex and noisy



- ▶ High Granularity Calorimeter will provide 3D information of a particle shower as it evolves



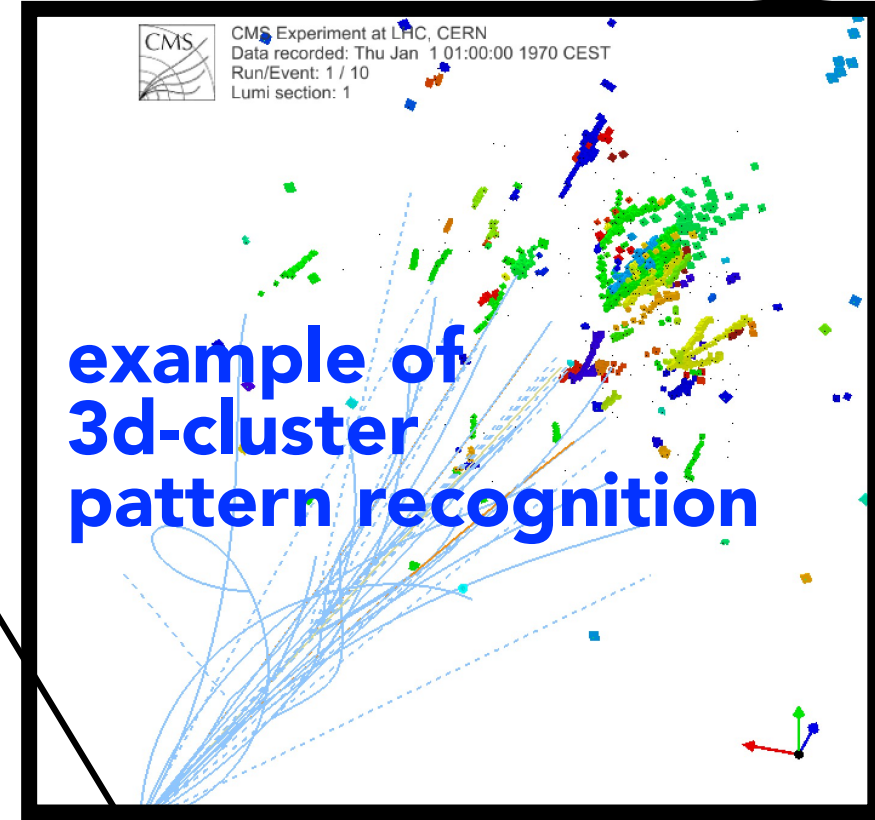
- ▶ High Granularity Calorimeter will provide 3D information of a particle shower as it evolves



tracks and clusters

the longitudinal shower footprint

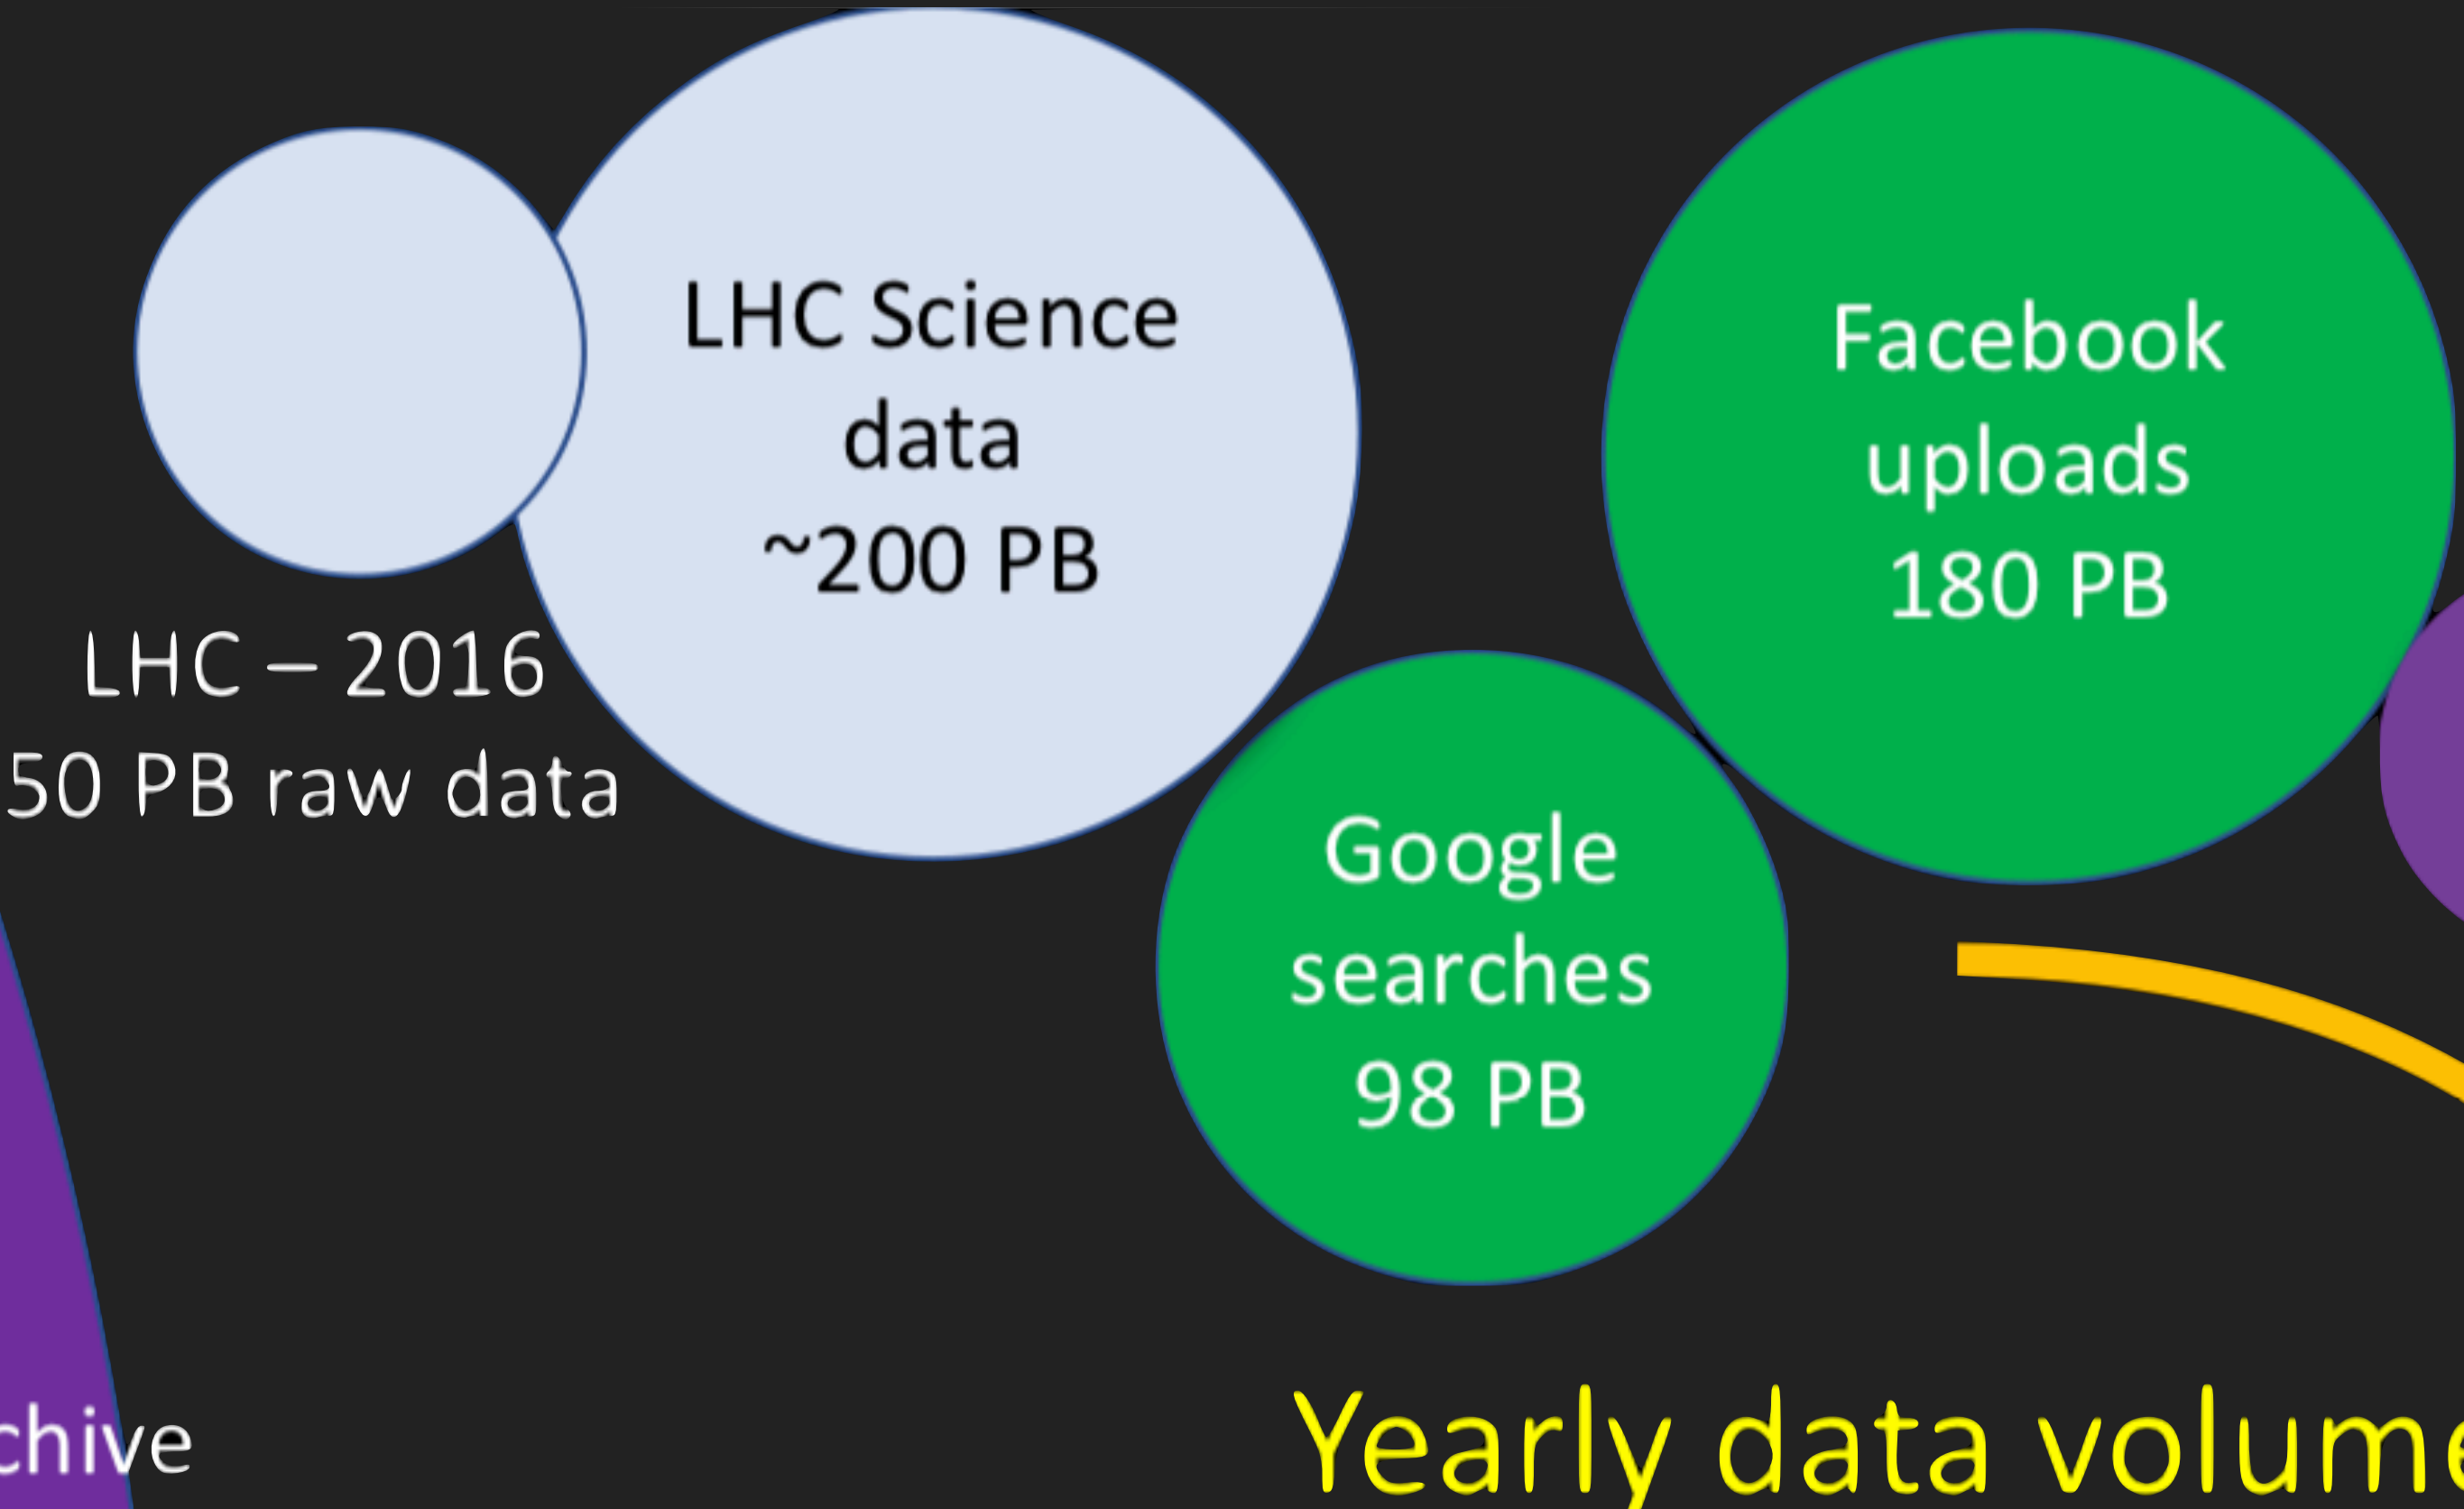
high energy particle shower



example of 3d-cluster pattern recognition

- ▶ High Granularity Calorimeter will provide 3D information of a particle shower as it evolves

CHALLENGE: BIG DATA



LHC Science
data
~200 PB

Facebook
uploads
180 PB

Google
searches
98 PB

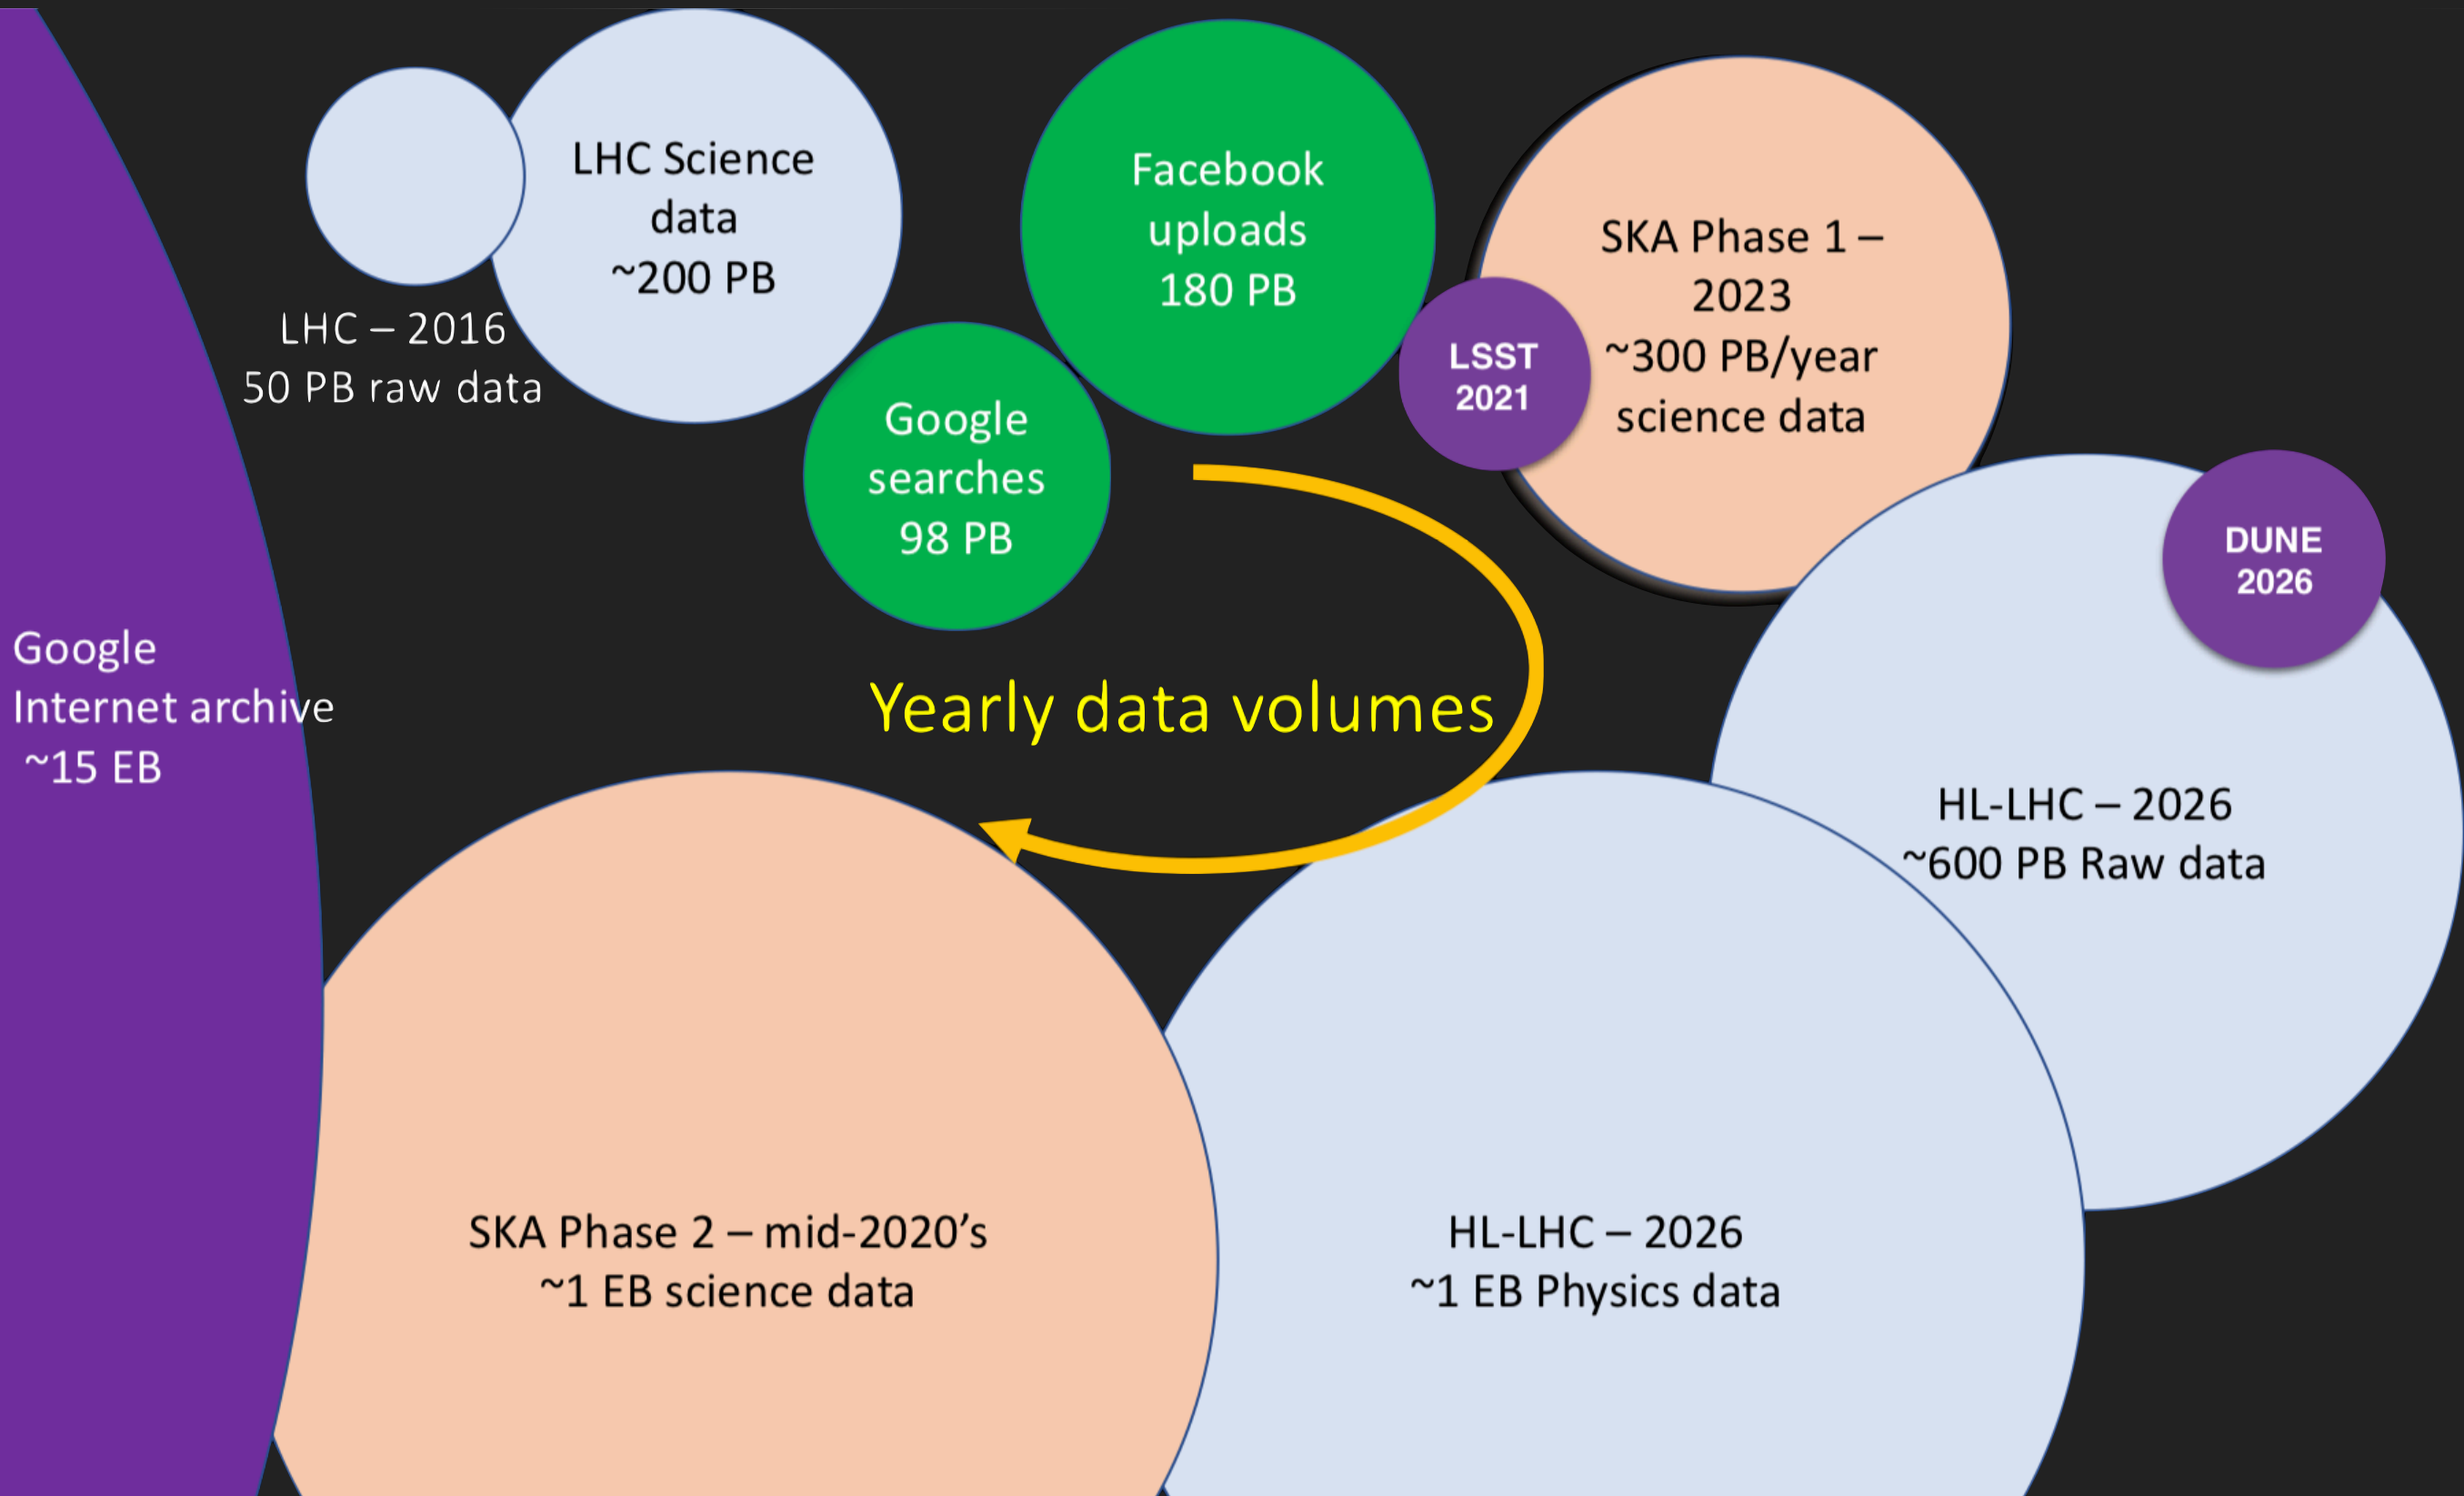
LHC - 2016
50 PB raw data


chive

Yearly data volume

CHALLENGE: BIG DATA

- ▶ HL-LHC will reach 1 exabyte of data per year



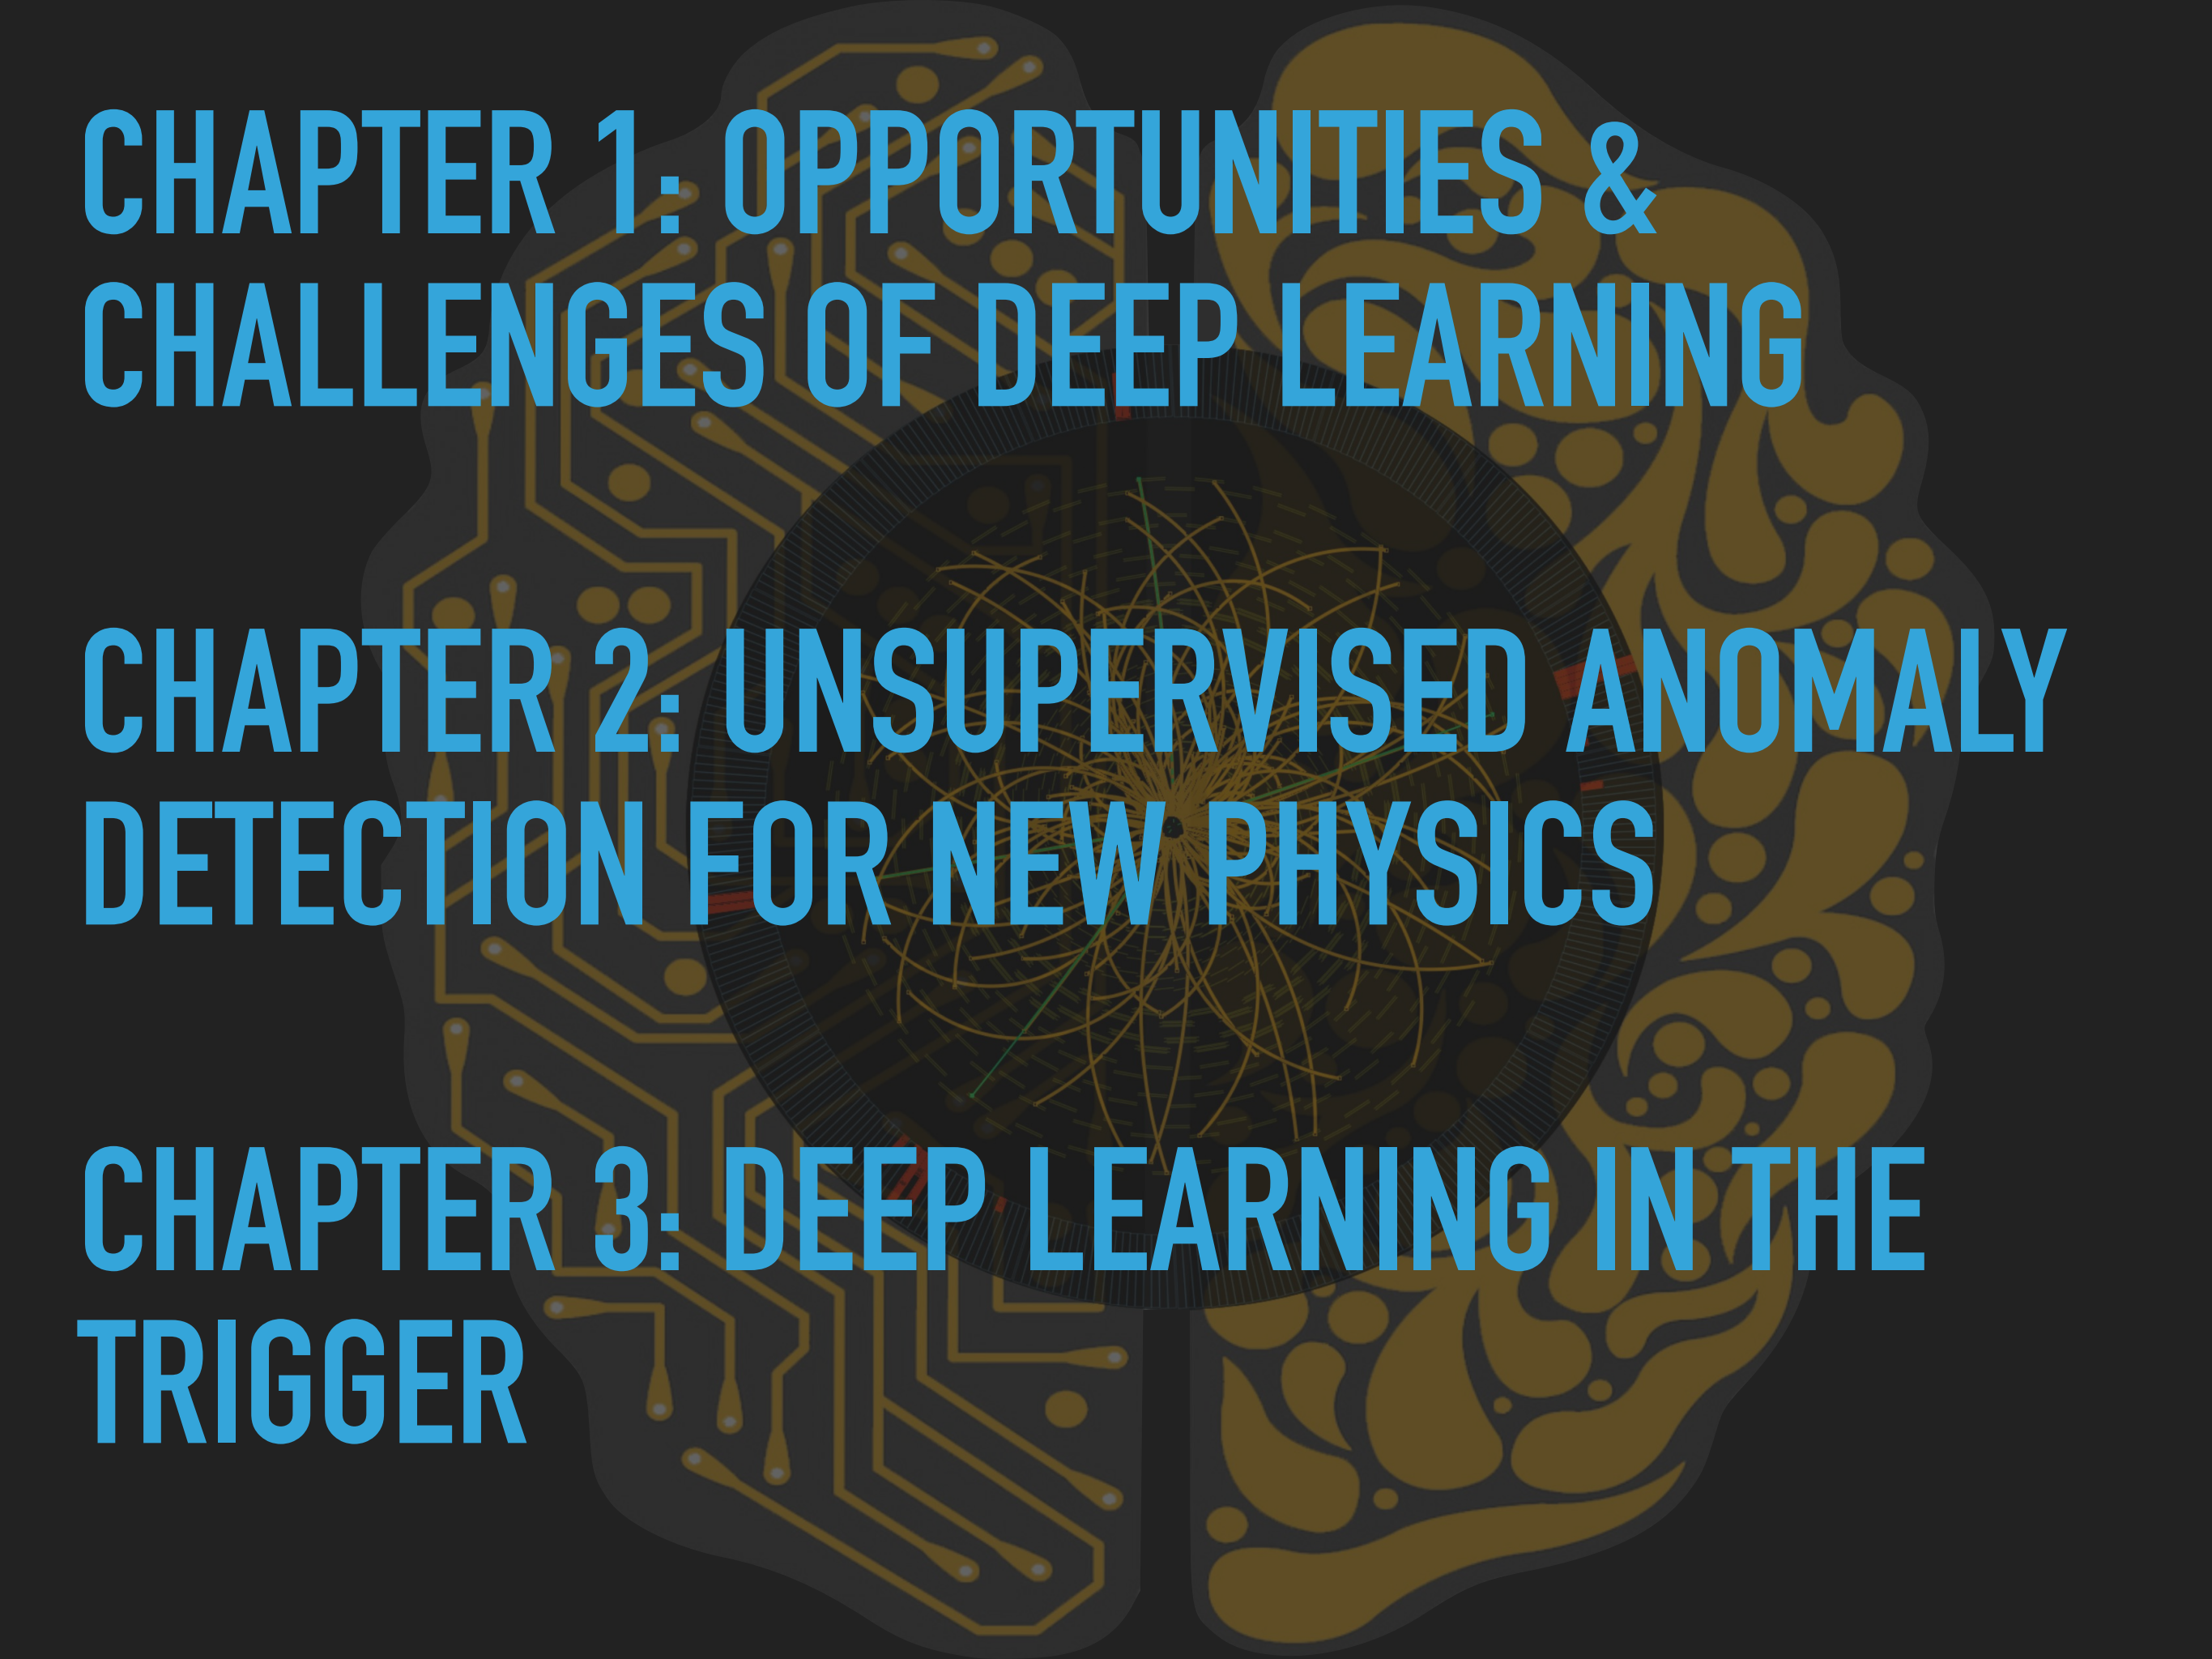


CHAPTER 1: OPPORTUNITIES & CHALLENGES OF DEEP LEARNING



CHAPTER 1: OPPORTUNITIES & CHALLENGES OF DEEP LEARNING

CHAPTER 2: UNSUPERVISED ANOMALY DETECTION FOR NEW PHYSICS



**CHAPTER 1: OPPORTUNITIES &
CHALLENGES OF DEEP LEARNING**

**CHAPTER 2: UNSUPERVISED ANOMALY
DETECTION FOR NEW PHYSICS**

**CHAPTER 3: DEEP LEARNING IN THE
TRIGGER**


CHAPTER 1: OPPORTUNITIES & CHALLENGES OF DEEP LEARNING

CHAPTER 2: UNSUPERVISED ANOMALY DETECTION FOR NEW PHYSICS

CHAPTER 3: DEEP LEARNING IN THE TRIGGER



THE LHC'S FAVORITE WAY TO MAKE HIGGS BOSONS

signal
 $gg \rightarrow H \rightarrow bb$




AN OVERWHELMING BACKGROUND

signal
 $gg \rightarrow H \rightarrow bb$



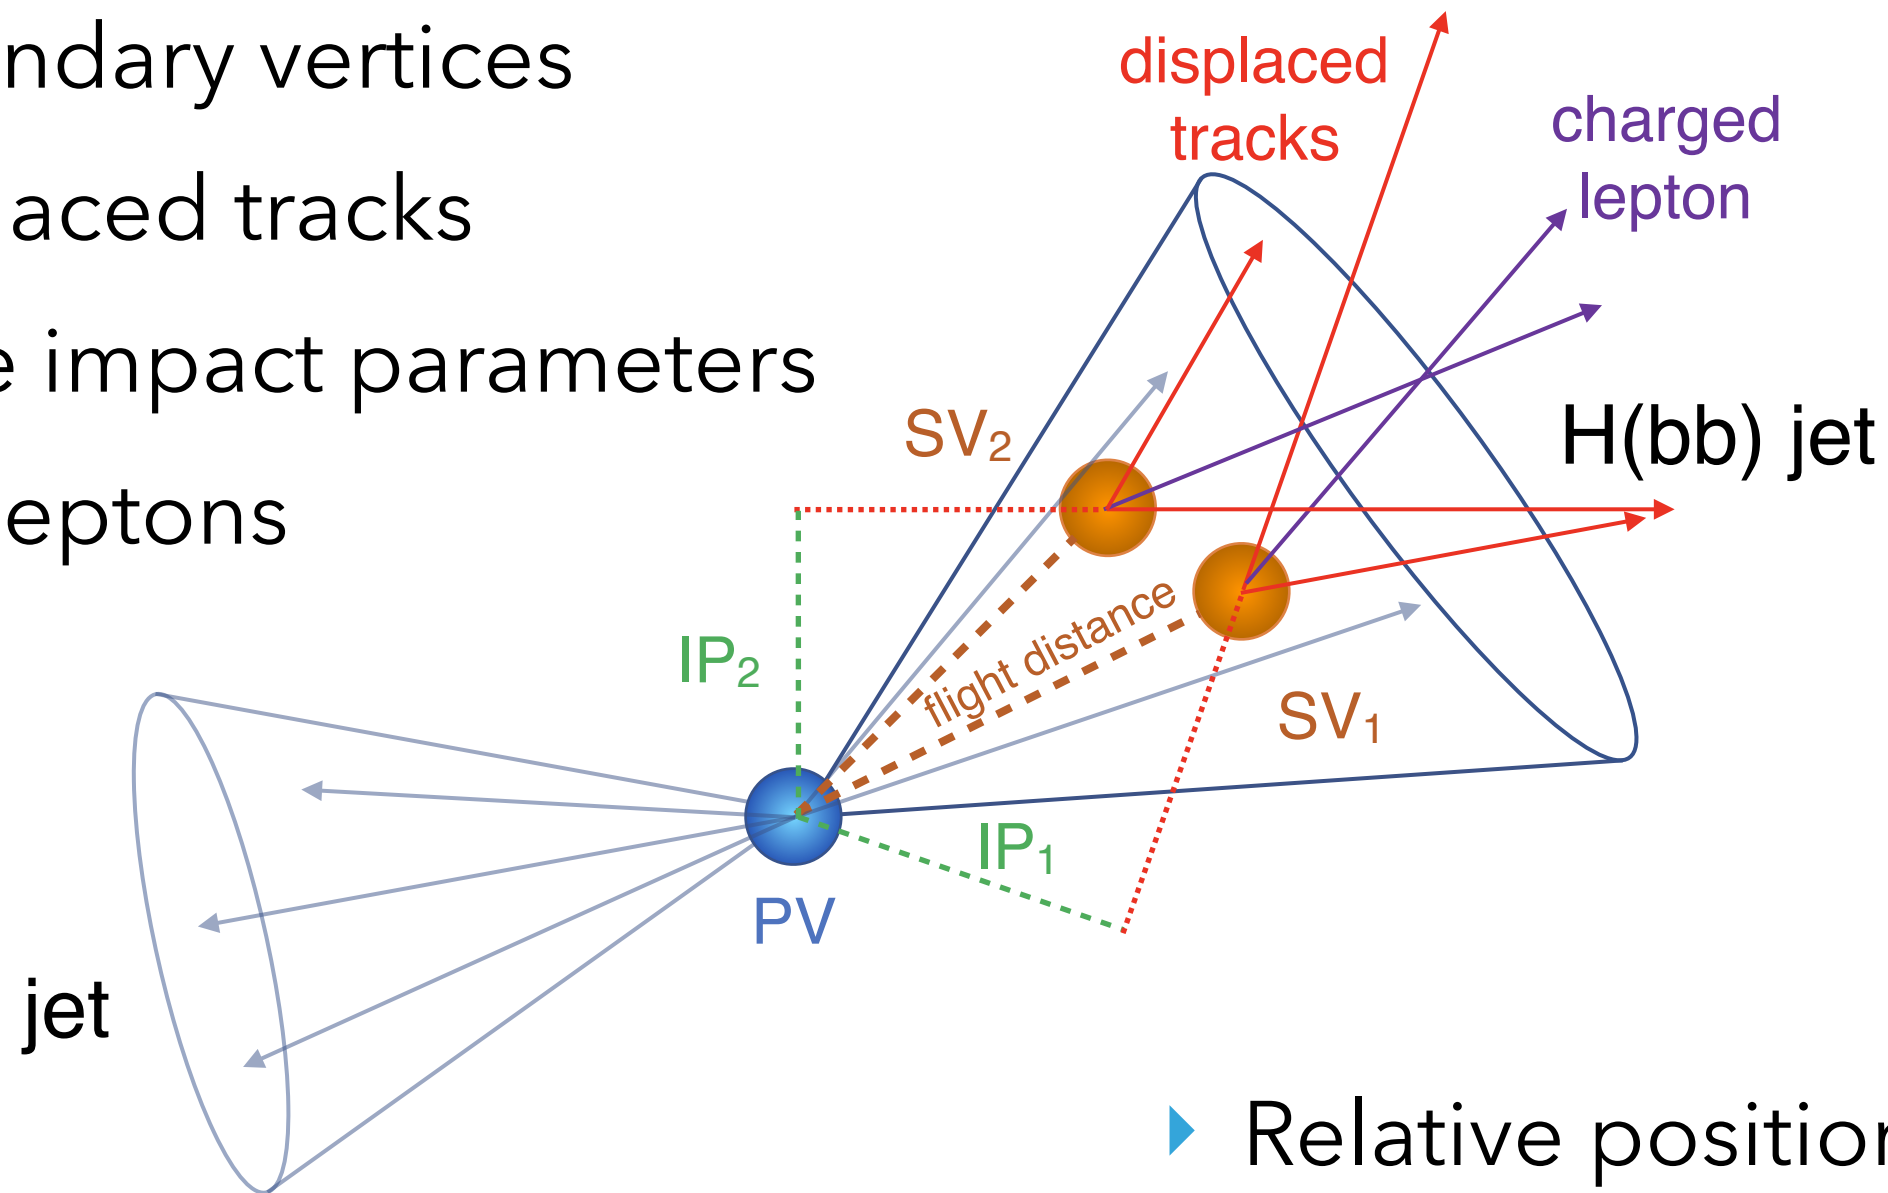
background
 $gg \rightarrow bb$

AN OVERWHELMING BACKGROUND





b hadrons have long lifetimes:
travel $O(\text{mm})$ before decay!

- ▶ Handles:
 - ▶ secondary vertices
 - ▶ displaced tracks
 - ▶ large impact parameters
 - ▶ soft leptons



- ▶ Relative positions of SVs

- ▶ Classic fully connected architecture
- ▶ Each **input** multiplied by a **weight**
- ▶ **Weighted** values are summed, **bias** is added
- ▶ Nonlinear **activation function** is applied
- ▶ Trained by varying the **parameters** to minimize a loss function (quantifies how many mistakes the network makes)



A sufficiently "wide" neural network can approximate any function!

- ▶ **Step 0:** Define the problem (choice of loss function)

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal)

p = output of our NN (probability of signal)

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal)

if $p \sim y$, $L \sim 0$ (correct!)

p = output of our NN (probability of signal)

if $p \sim 1-y$, $L \sim \infty$ (incorrect!)

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal) if $p \sim y$, $L \sim 0$ (correct!)
 $p =$ output of our NN (probability of signal) if $p \sim 1-y$, $L \sim \infty$ (incorrect!)

- ▶ Step 1: Acquire lots of labeled data and split into training and testing sets

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal) if $p \sim y$, $L \sim 0$ (correct!)
 $p =$ output of our NN (probability of signal) if $p \sim 1-y$, $L \sim \infty$ (incorrect!)

- ▶ Step 1: Acquire lots of labeled data and split into training and testing sets
- ▶ Step 2: Select input features

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal) if $p \sim y$, $L \sim 0$ (correct!)
 $p =$ output of our NN (probability of signal) if $p \sim 1-y$, $L \sim \infty$ (incorrect!)

- ▶ Step 1: Acquire lots of labeled data and split into training and testing sets
- ▶ Step 2: Select input features
- ▶ Step 3: Explore/train different neural network architectures

- ▶ **Step 0:** Define the problem (choice of loss function)

$$L = -y \log(p) + (1-y) \log(1-p)$$

$y = 0$ (background) or 1 (signal) if $p \sim y$, $L \sim 0$ (correct!)
 $p =$ output of our NN (probability of signal) if $p \sim 1-y$, $L \sim \infty$ (incorrect!)

- ▶ Step 1: Acquire lots of labeled data and split into training and testing sets
- ▶ Step 2: Select input features
- ▶ Step 3: Explore/train different neural network architectures
- ▶ **Step 4:** Evaluate performance

Sample with jet, track and secondary vertex properties for Hbb tagging ML studies HiggsToBBNTuple_HiggsToBB_QCD_RunII_13TeV_MC

Duarte, Javier;

Dataset Derived Datascience CMS CERN-LHC

Parent Dataset: /BulkGravTohhTohbbhb_narrow_M-600_13TeV-madgraph/RunIISummer16MiniAODv2-PUMoriond17_80X_mcRun2_asymptotic_2016_TracheIV_v6_ext1-v1/MINIAODSIM

Description

The dataset consists of particle jets extracted from simulated proton-proton collision events at a center-of-mass energy of 13 TeV generated with Pythia 8. It has been produced for developing machine-learning algorithms to differentiate jets originating from a Higgs boson decaying to a bottom quark-antiquark pair (Hbb) from quark or gluon jets originating from quantum chromodynamic (QCD) multijet production.

The reconstructed jets are clustered using the anti-kT algorithm with R=0.8 from particle flow (PF) candidates (AK8 jets). The standard L1+L2+L3+residual jet energy corrections are applied to the jets and pileup contamination is mitigated using the charged hadron subtraction (CHS) algorithm. Features of the AK8 jets with transverse momentum $p_T > 200$ GeV and pseudorapidity $|\eta| < 2.4$ are provided. Selected features of inclusive (both charged and neutral) PF candidates with $p_T > 0.95$ GeV associated to the AK8 jet are provided. Additional features of charged PF candidates (formed primarily by a charged particle track) with $p_T > 0.95$ GeV associated to the AK8 jet are also provided. Finally, additional features of reconstructed secondary vertices (SVs) associated to the AK8 jet (within $\Delta R < 0.8$) are also provided.

▶ Derived datasets (ROOT & HDF5):

<http://opendata-dev.web.cern.ch/record/12102>

▶ 182 files, 245 GB, 18 million total entries (jets)


- ▶ event features, e.g. MET, ρ (average density)
- ▶ jet features, e.g. mass, p_T , N-subjettiness variables
- ▶ particle candidate features, e.g. p_T , η , ϕ (for up to 100 particles)
- ▶ charged particle / track features, e.g. impact parameter (for up to 60 tracks)
- ▶ secondary vertex features, e.g. flight distance (for up to 5 vertices)

<https://github.com/cernopendata-datascience/HiggsToBBMachineLearning>

- ▶ Train fully connected neural network with high level features in ~30 lines of code

Layer (type)	Output Shape	Param #
input (InputLayer)	(None, 27)	0
bn_1 (BatchNormalization)	(None, 27)	108
dense_1 (Dense)	(None, 64)	1792
dense_2 (Dense)	(None, 32)	2080
dense_3 (Dense)	(None, 32)	1056
output (Dense)	(None, 2)	66

Total params: 5,102
Trainable params: 5,048
Non-trainable params: 54



track
inputs

secondary
vertex
inputs

expert
inputs

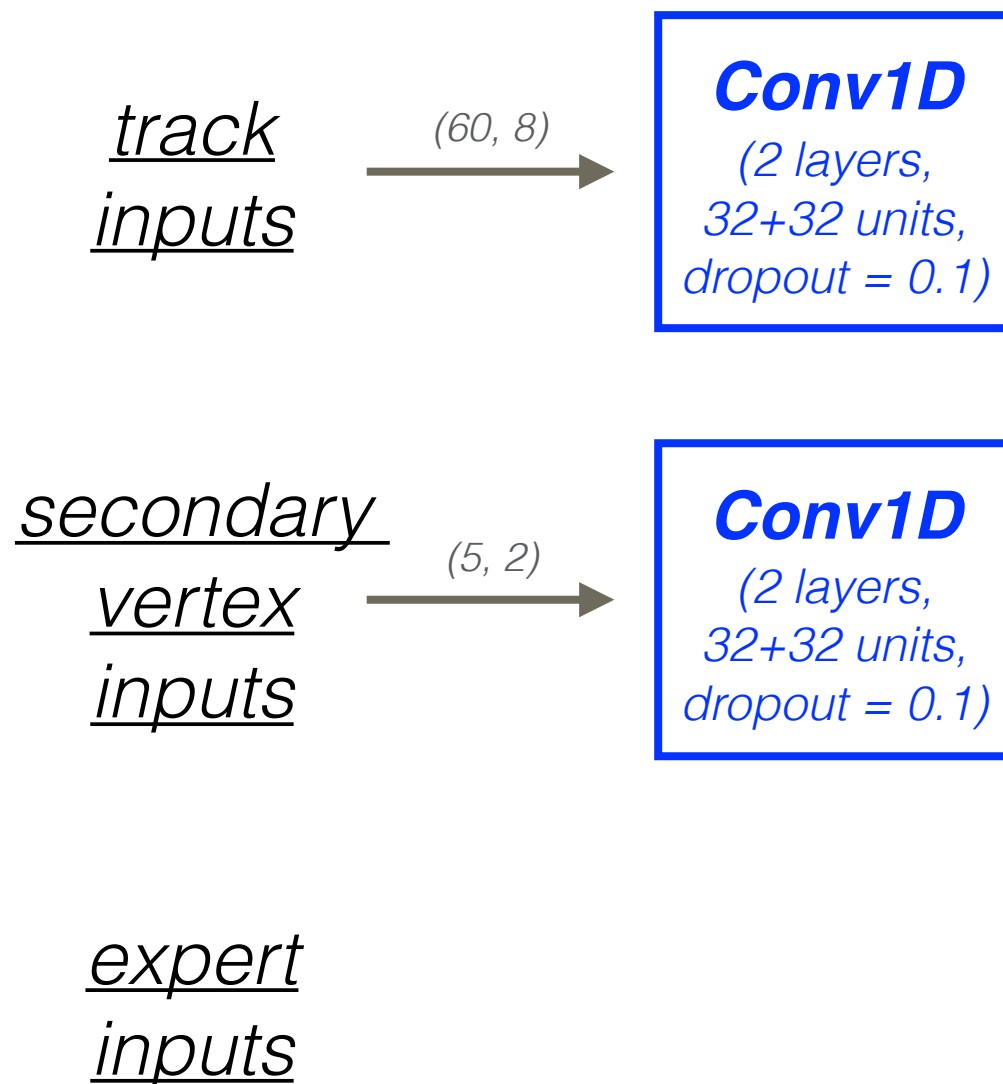
- ▶ Process low-level track and SV inputs as ***ordered lists***

track
inputs $\xrightarrow{(60, 8)}$

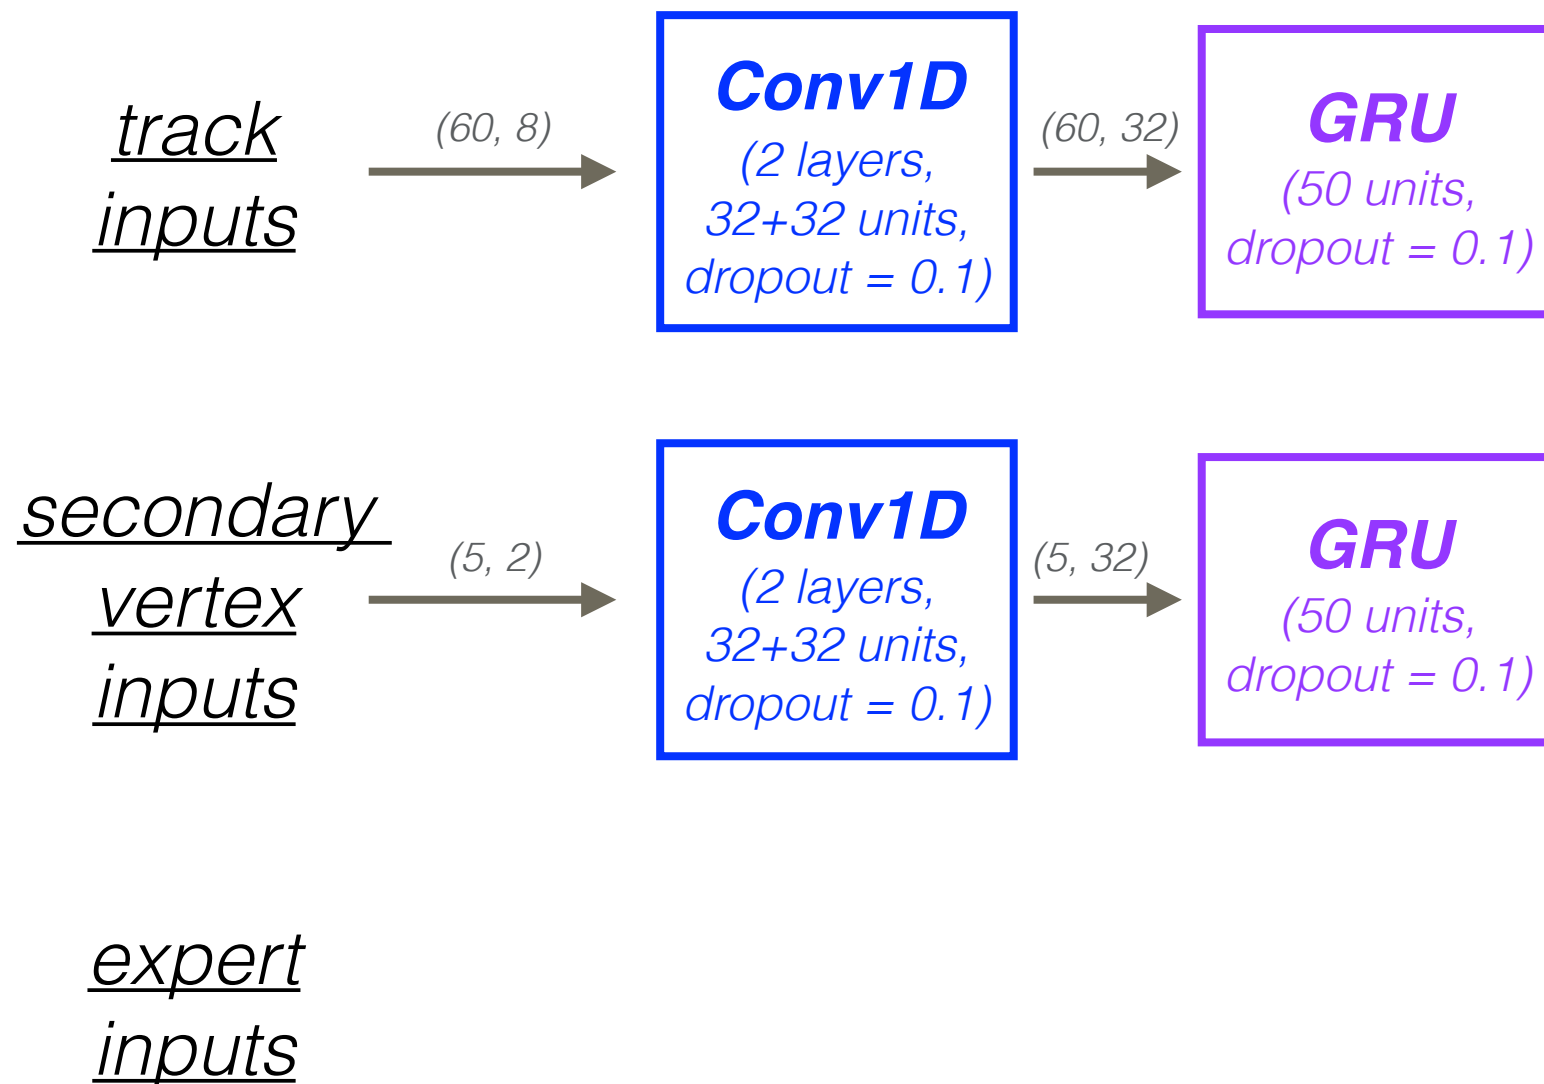
secondary
vertex
inputs $\xrightarrow{(5, 2)}$

expert
inputs

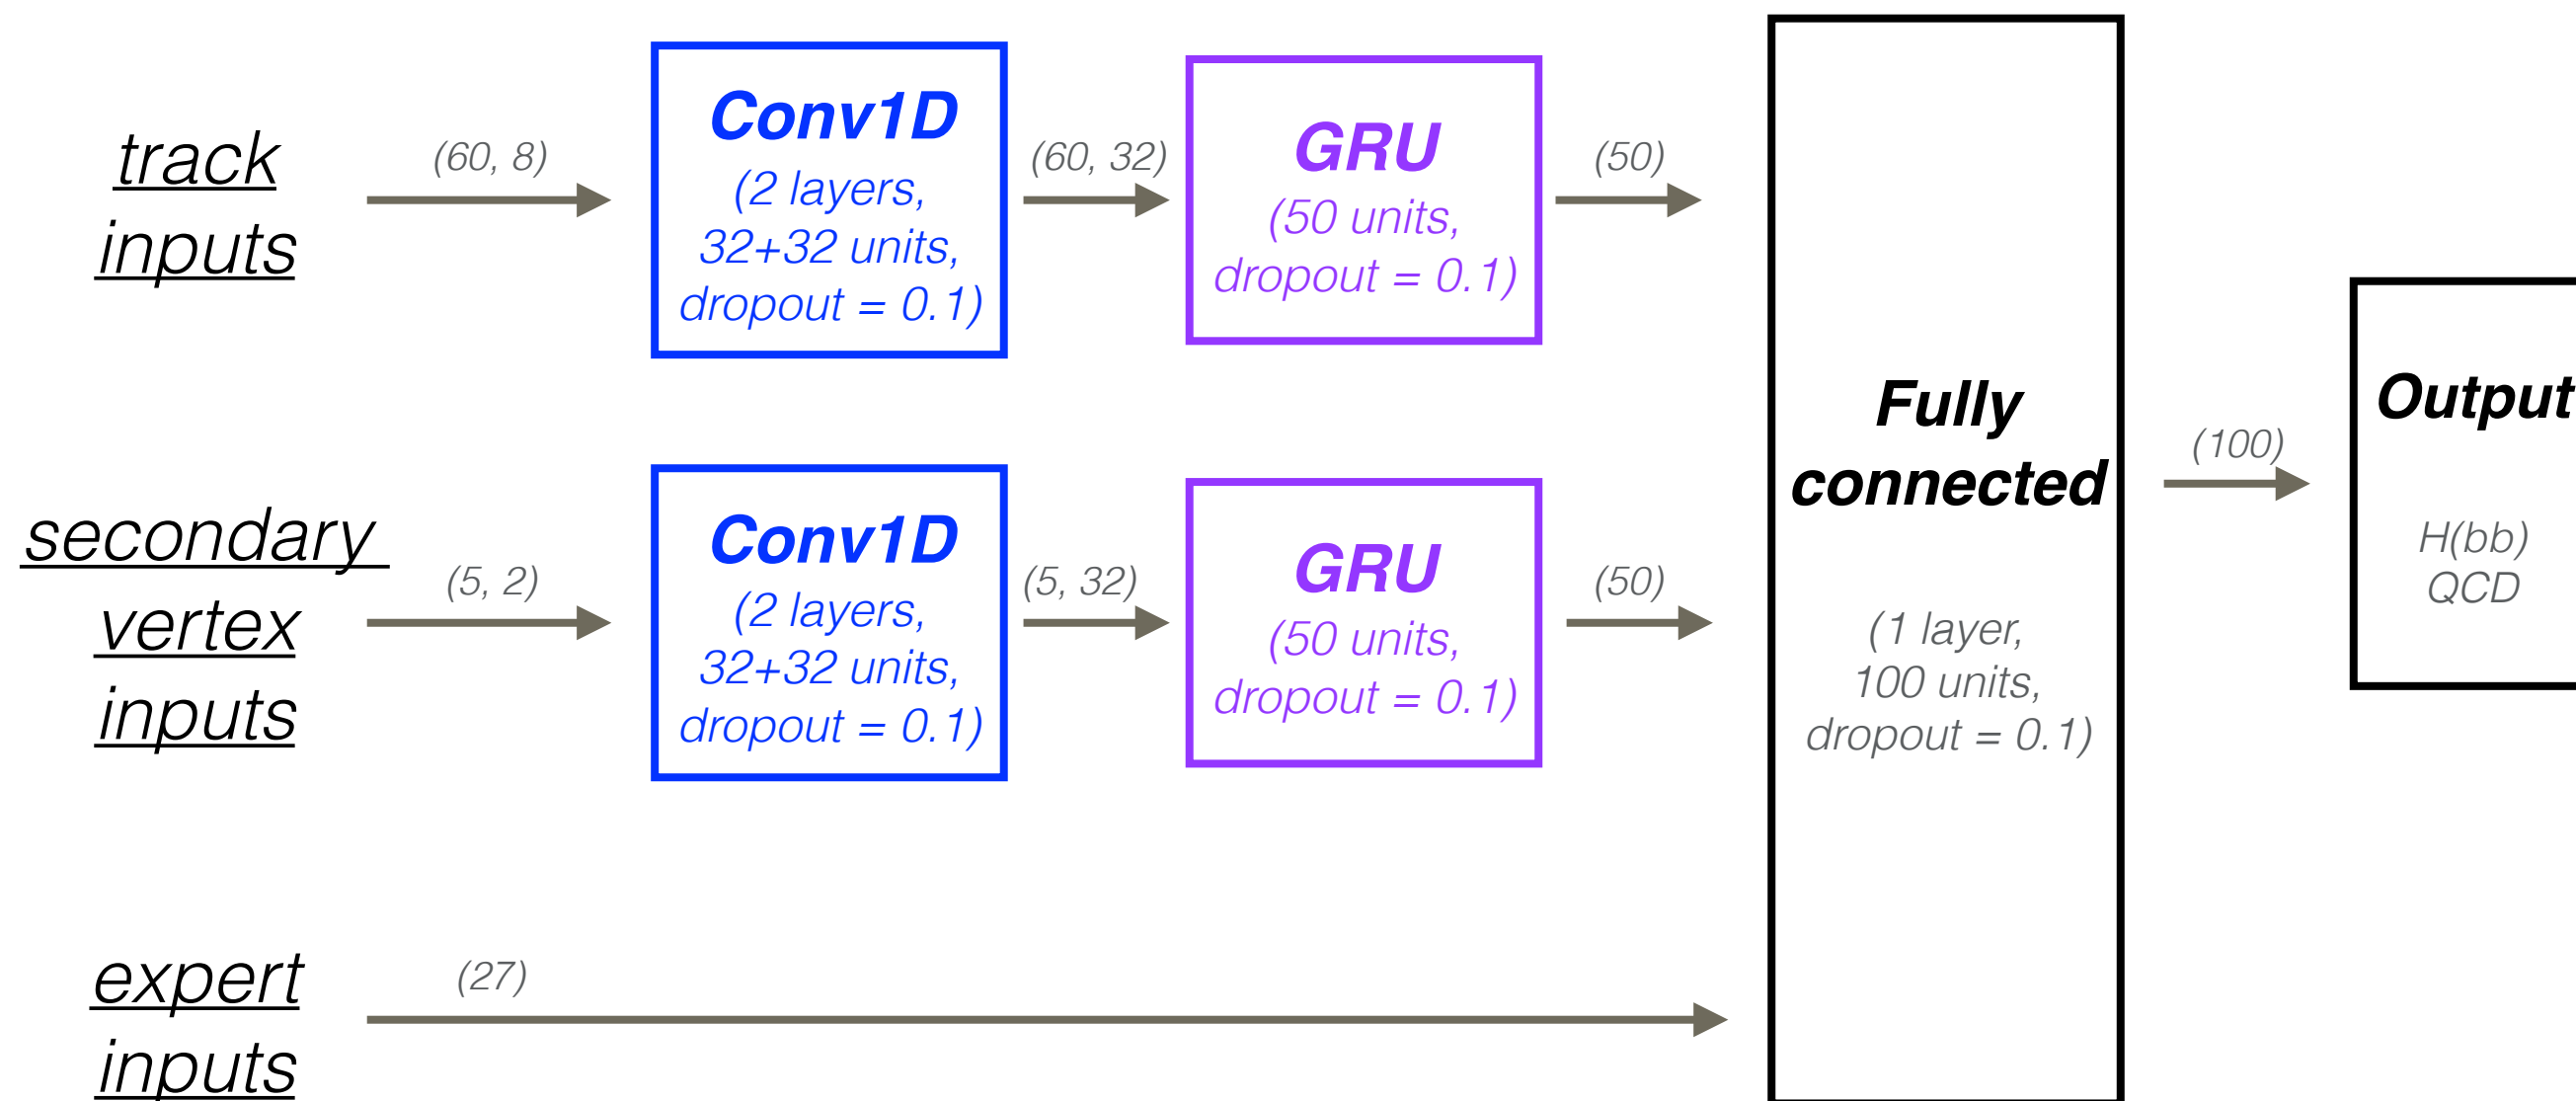
- ▶ Process low-level track and SV inputs as **ordered lists**
 - ▶ **Convolutional** NN layers: share parameters across inputs, ...




- ▶ Process low-level track and SV inputs as **ordered lists**
 - ▶ **Convolutional** NN layers: share parameters across inputs, ...
 - ▶ **Recurrent** NN layers: performs dimensional reduction, ...




- ▶ Process low-level track and SV inputs as **ordered lists**
 - ▶ **Convolutional** NN layers: share parameters across inputs, ...
 - ▶ **Recurrent** NN layers: performs dimensional reduction, ...
- ▶ Combine in final layer with expert inputs







- ▶ Success for deep learning!




► Success for deep learning!



- ▶ An unintended consequence: network “learns” the jet mass




MITIGATING THE MASS SCULPTING



MITIGATING THE MASS SCULPTING

- ▶ How can we quantify the mass sculpting?




MITIGATING THE MASS SCULPTING

▶ How can we quantify the mass sculpting?

▶ Kullback-Liebler divergence

$$D_{\text{KL}} = h(x) \log \left(\frac{h(x)}{q(x)} \right)$$




MITIGATING THE MASS SCULPTING

▶ How can we quantify the mass sculpting?

▶ Kullback-Liebler divergence

$$D_{\text{KL}} = h(x) \log \left(\frac{h(x)}{q(x)} \right)$$

▶ How can we mitigate the mass sculpting?



MITIGATING THE MASS SCULPTING

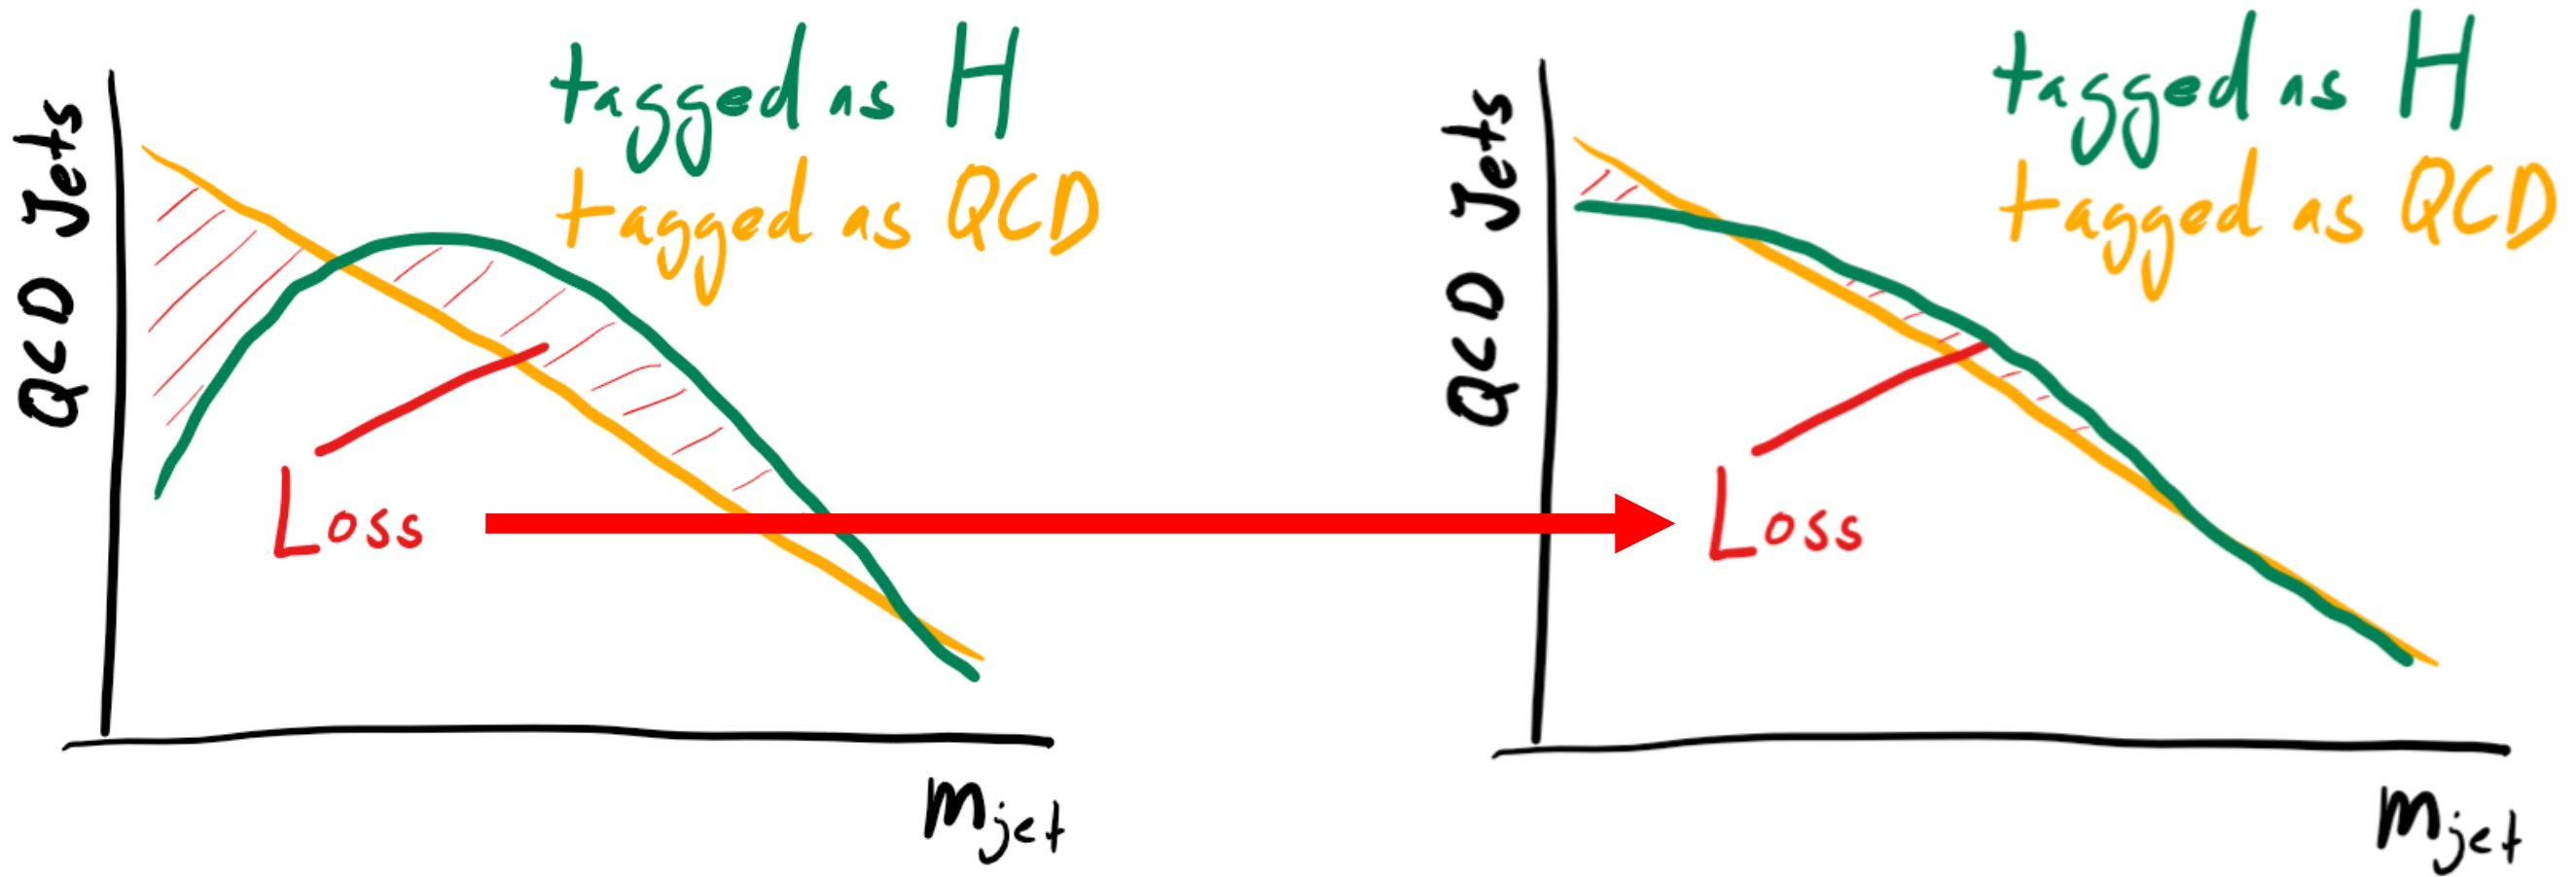
▶ How can we quantify the mass sculpting?


▶ Kullback-Liebler divergence $D_{\text{KL}} = h(x) \log \left(\frac{h(x)}{q(x)} \right)$

▶ How can we mitigate the mass sculpting?


▶ Add it to the loss function as a "penalty"

$$L = L_{\text{disc}} + \lambda D_{\text{KL}}$$






- ▶ Penalty term mitigates the mass sculpting



- ▶ Penalty term mitigates the mass sculpting
- ▶ Small trade-off with performance



CAN WE DO EVEN BETTER?



CAN WE DO EVEN BETTER?


- ▶ Ordered lists of particles not the most natural representation of a jet
 - ▶ What if we consider each jet as a **graph** of interconnected particles?




CAN WE DO EVEN BETTER?


- ▶ Ordered lists of particles not the most natural representation of a jet
 - ▶ What if we consider each jet as a **graph** of interconnected particles?
- ▶ **Geometric deep learning** (a.k.a graph neural networks, interaction networks, message-passing neural networks) is the extension of deep learning to deal with data structured as a graph or on a manifold
 - ▶ See Interaction Networks for Learning about Objects, Relations, and Physics [[arXiv:1612.0222](#)], Neural Message Passing for Quantum Chemistry [[arXiv:1704.01212](#)], Dynamic Graph CNN for Learning on Point Clouds [[arXiv:1801.07829](#)], Fast Graph Representation Learning with PyTorch Geometric [[arXiv:1903.0242](#)]



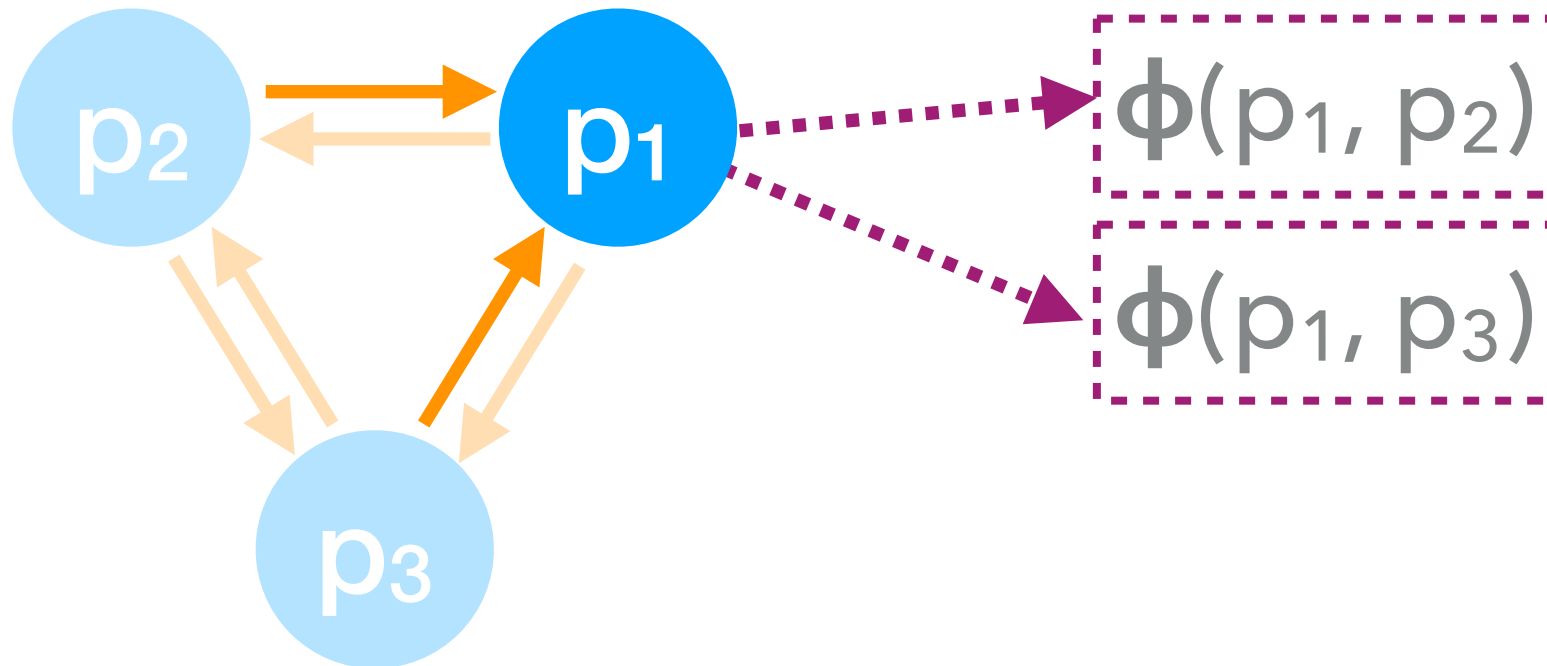





- ▶ A **graph** of objects and their connections is defined



- ▶ A **graph** of objects and their connections is defined




- ▶ A **graph** of objects and their connections is defined
- ▶ **NN** is evaluated on **pairs of connected objects** to produce a message



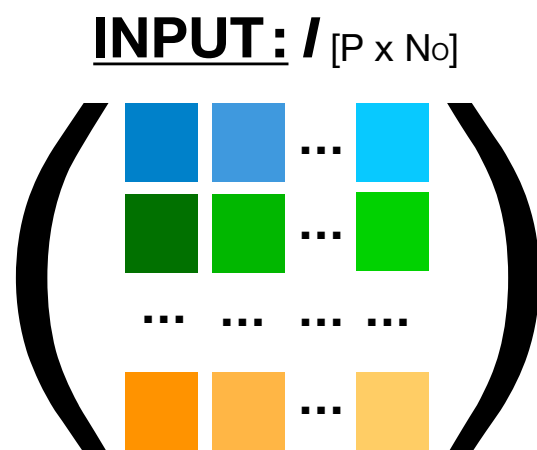
- ▶ A **graph** of objects and their connections is defined
- ▶ **NN** is evaluated on **pairs of connected objects** to produce a message
- ▶ **Messages** are communicated from nearest neighbors (and *summed**) to **update** each object's **hidden state**

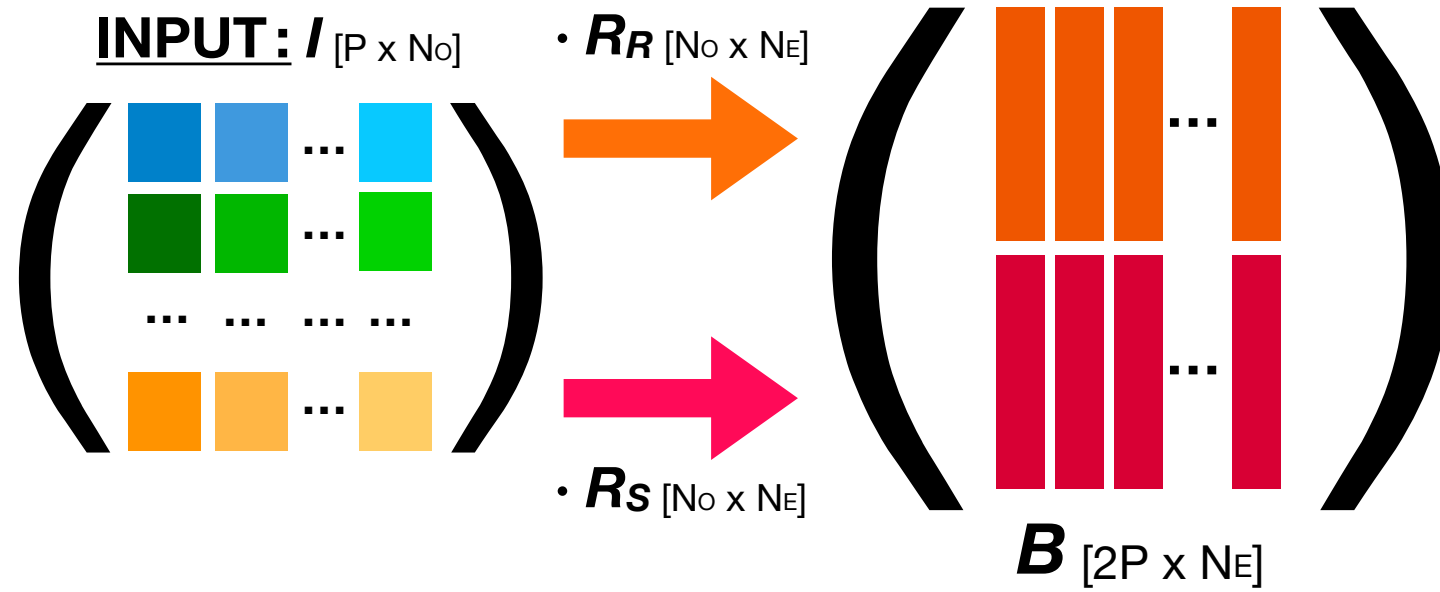
*sum preserves permutation invariance

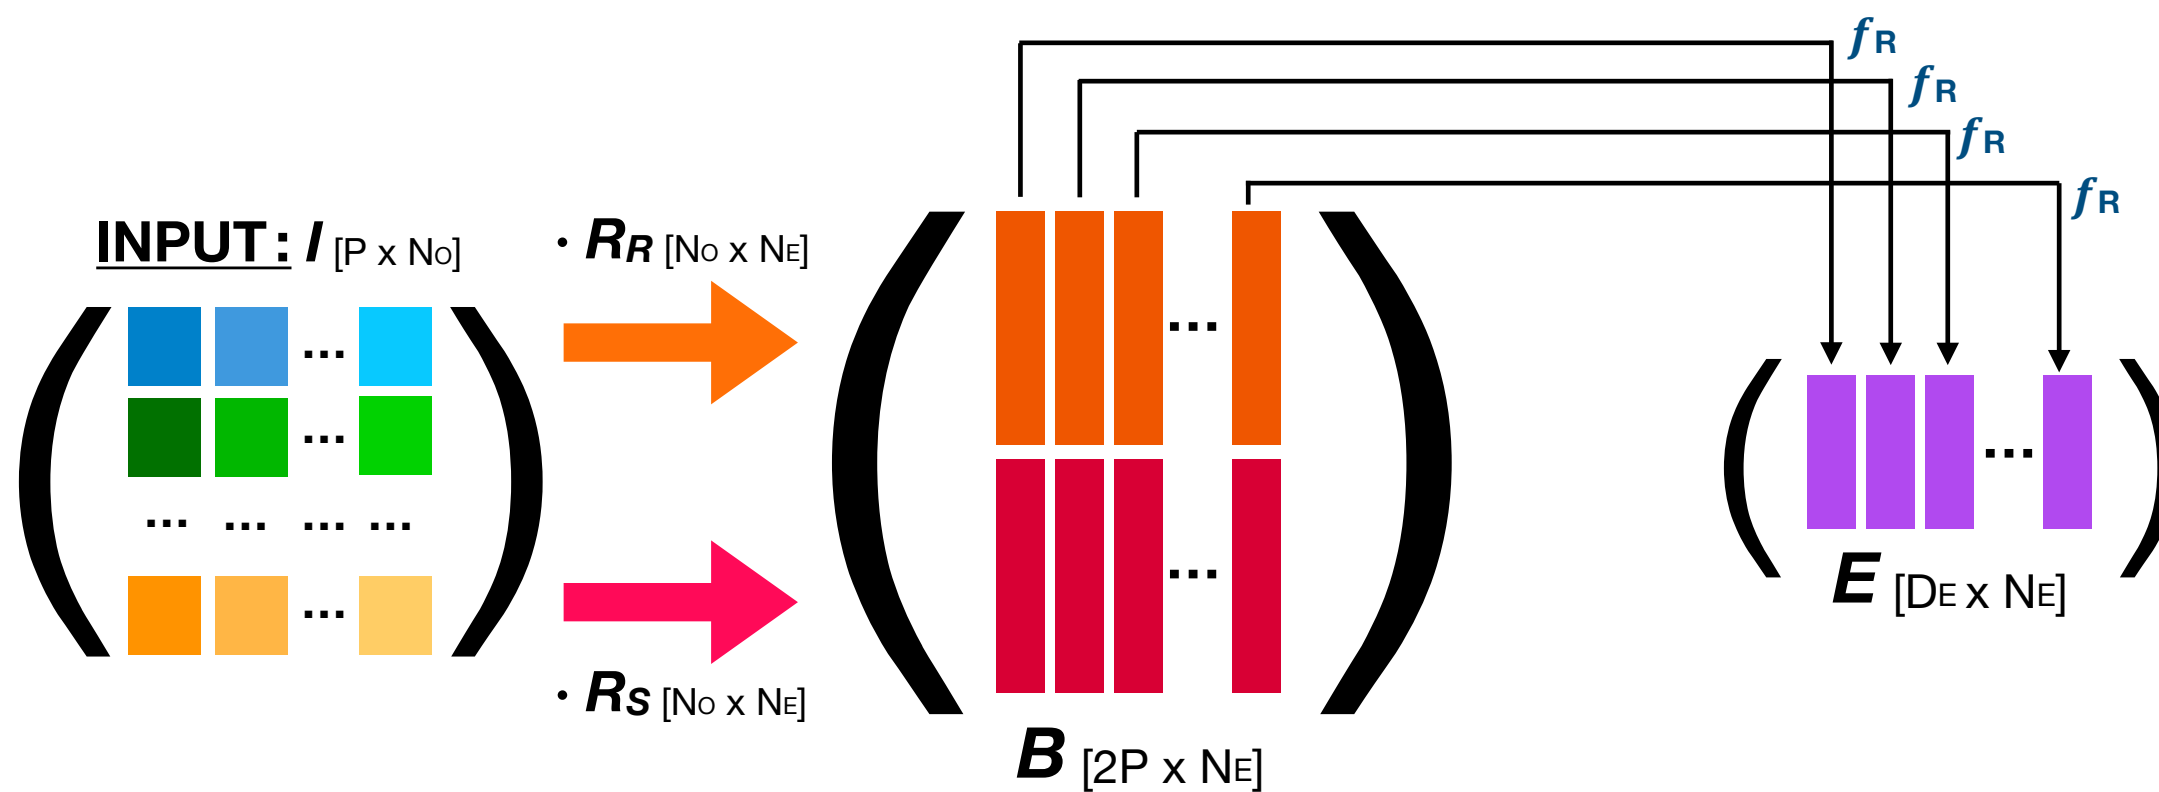


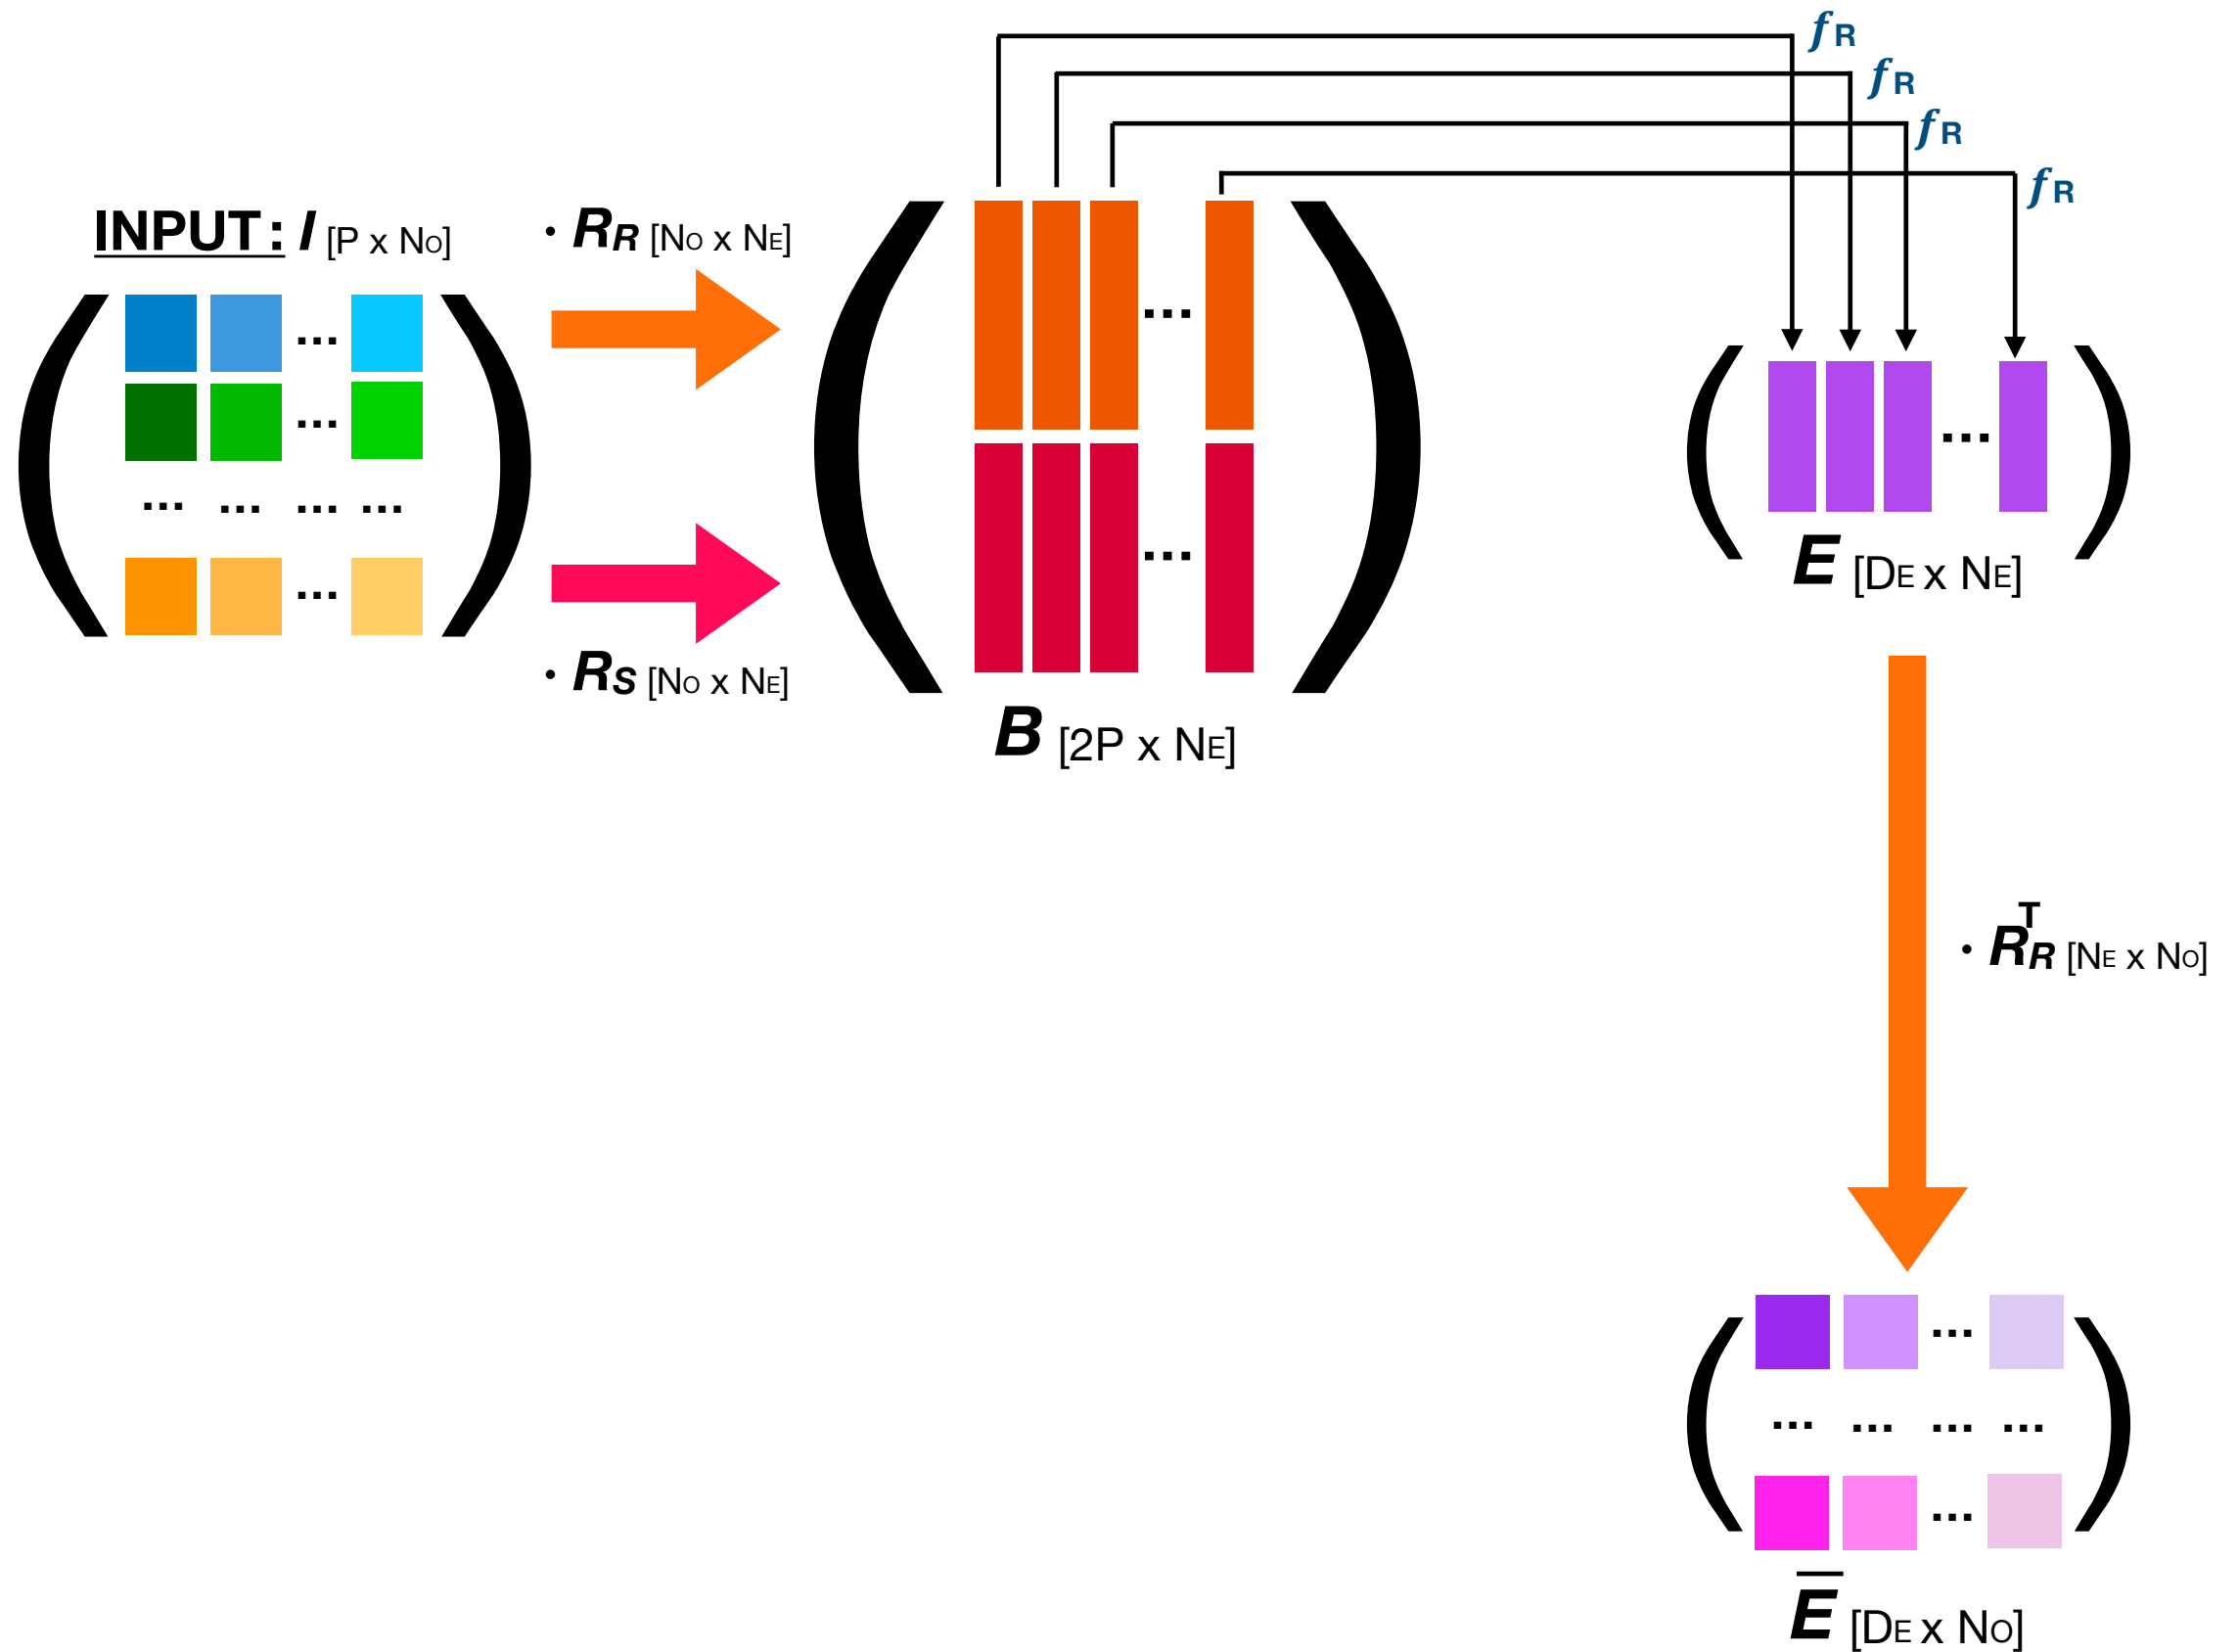
- ▶ A **graph** of objects and their connections is defined
- ▶ **NN** is evaluated on **pairs of connected objects** to produce a message
- ▶ **Messages** are communicated from nearest neighbors (and *summed**) to **update** each object's **hidden state**
- ▶ A **single output** is computed based on the *summed** hidden states of all objects in the graph

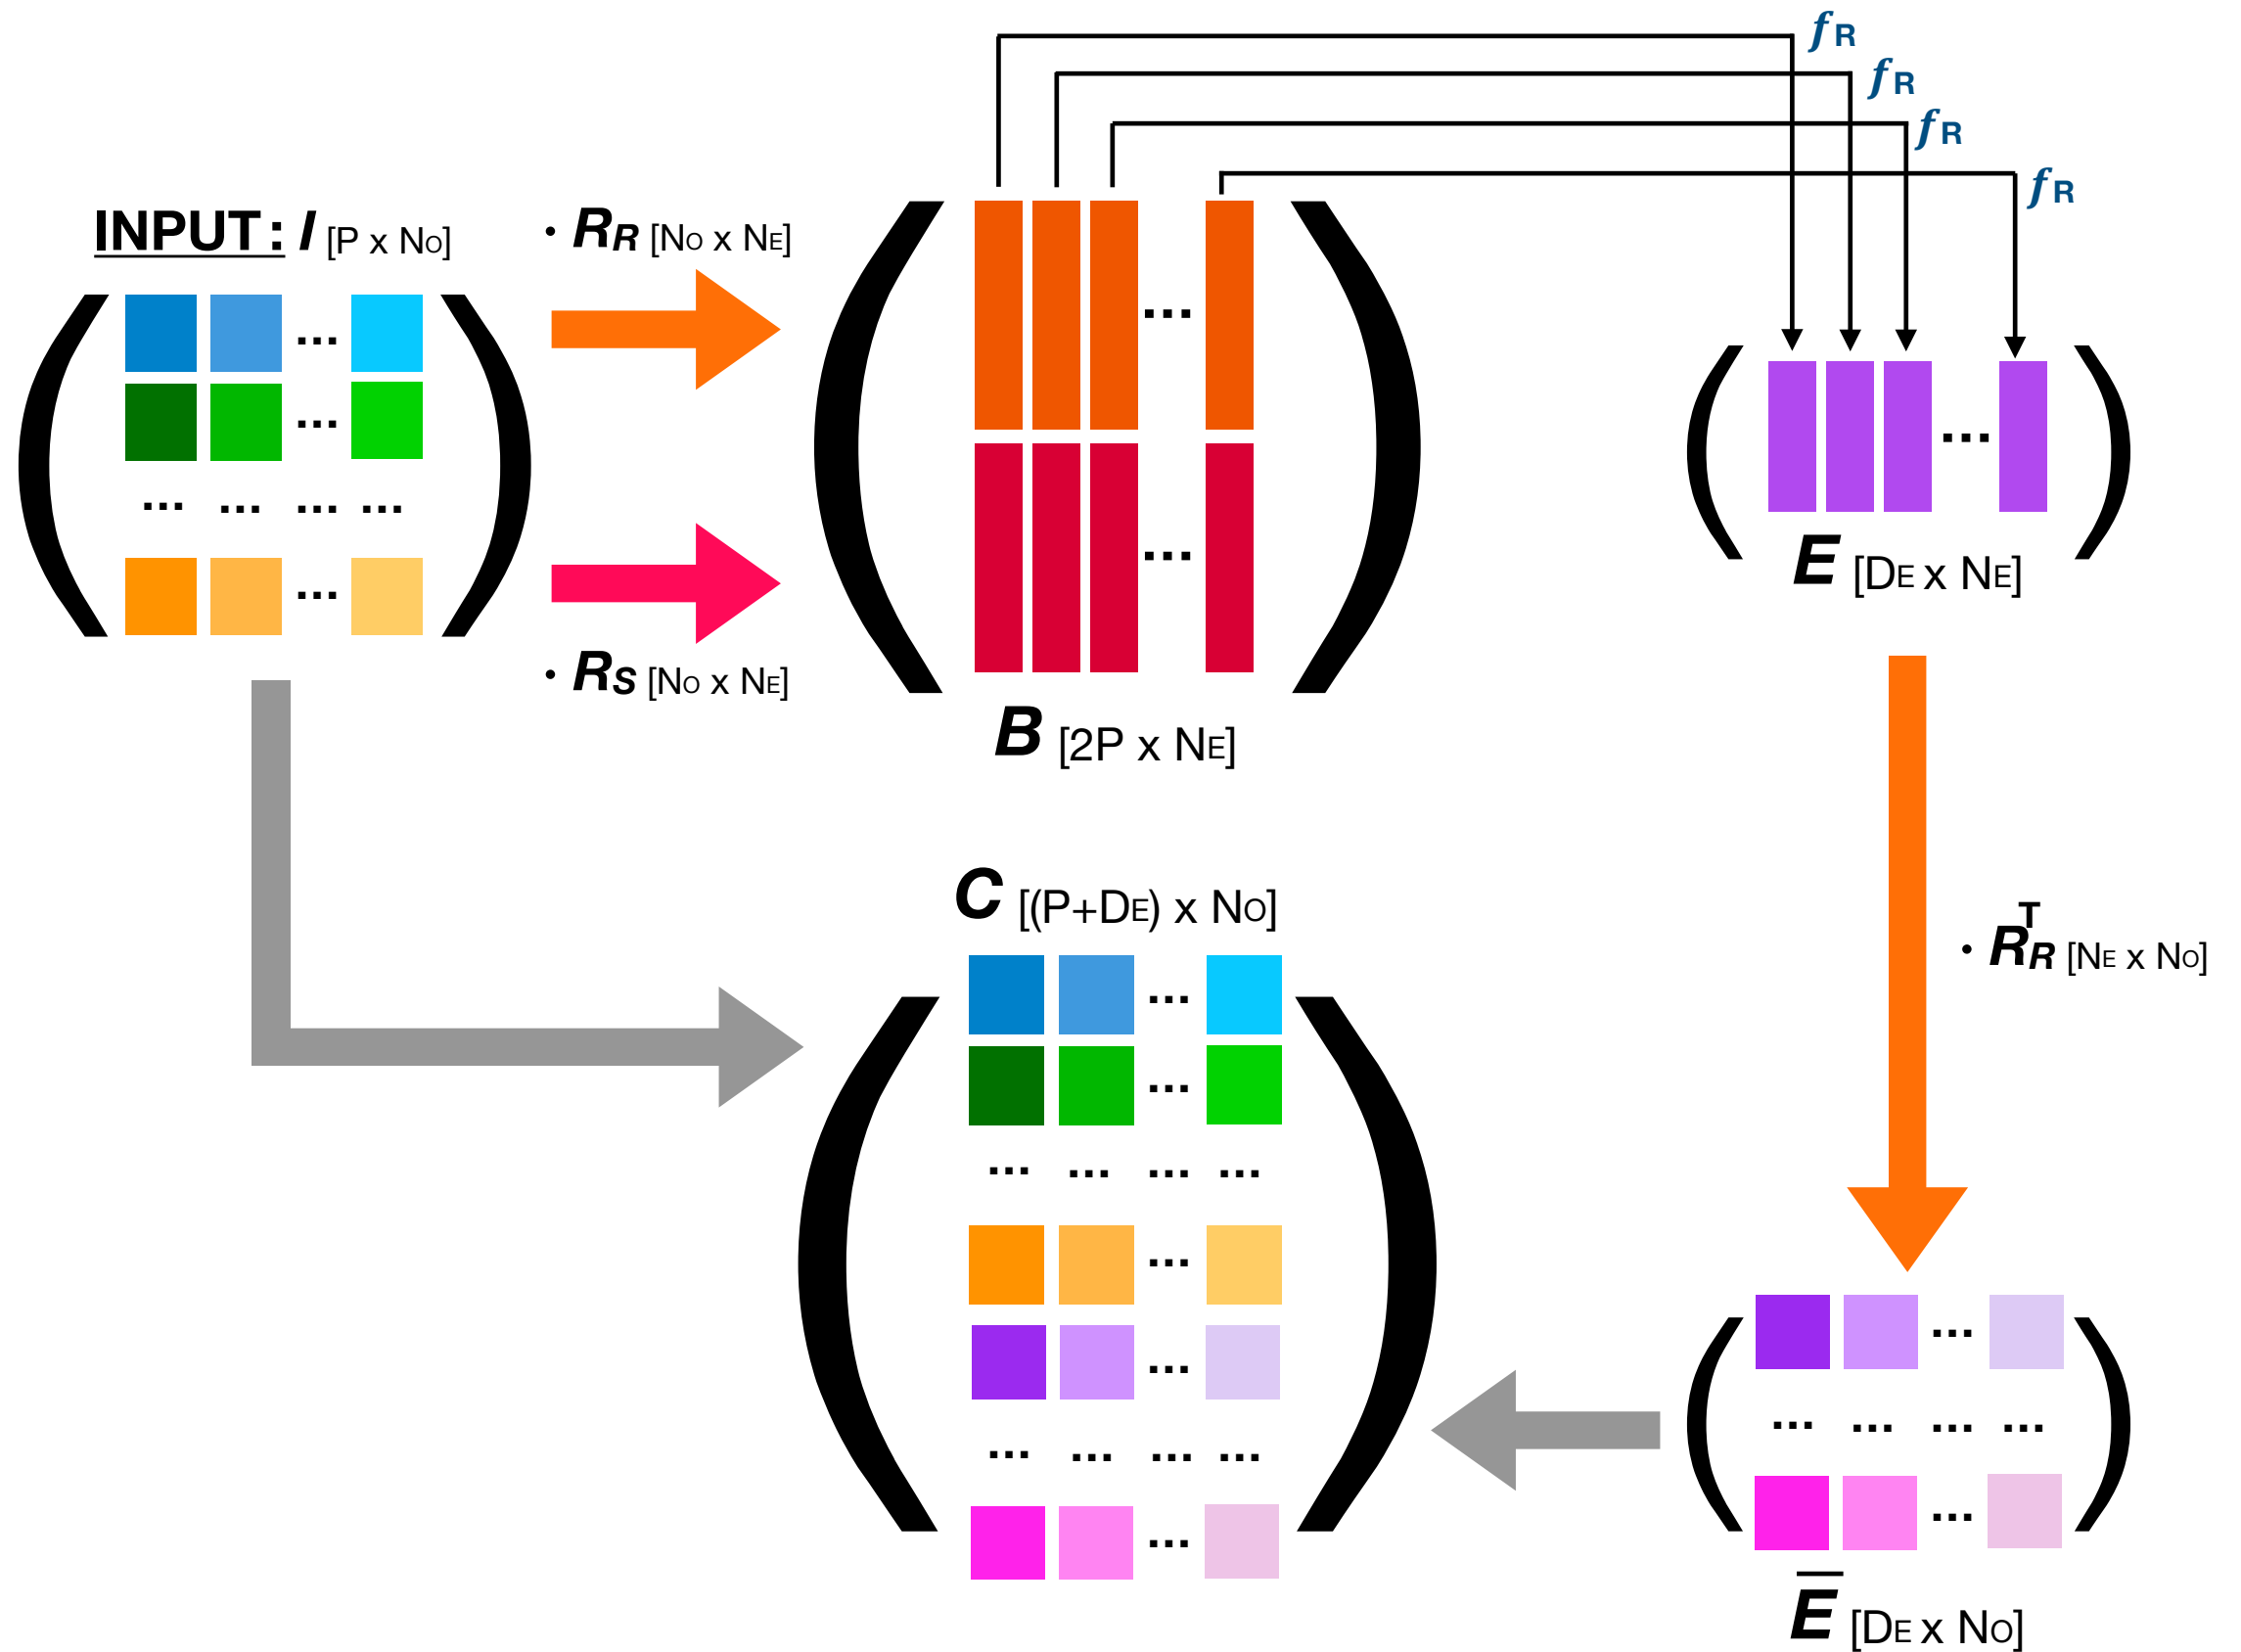
*sum preserves permutation invariance

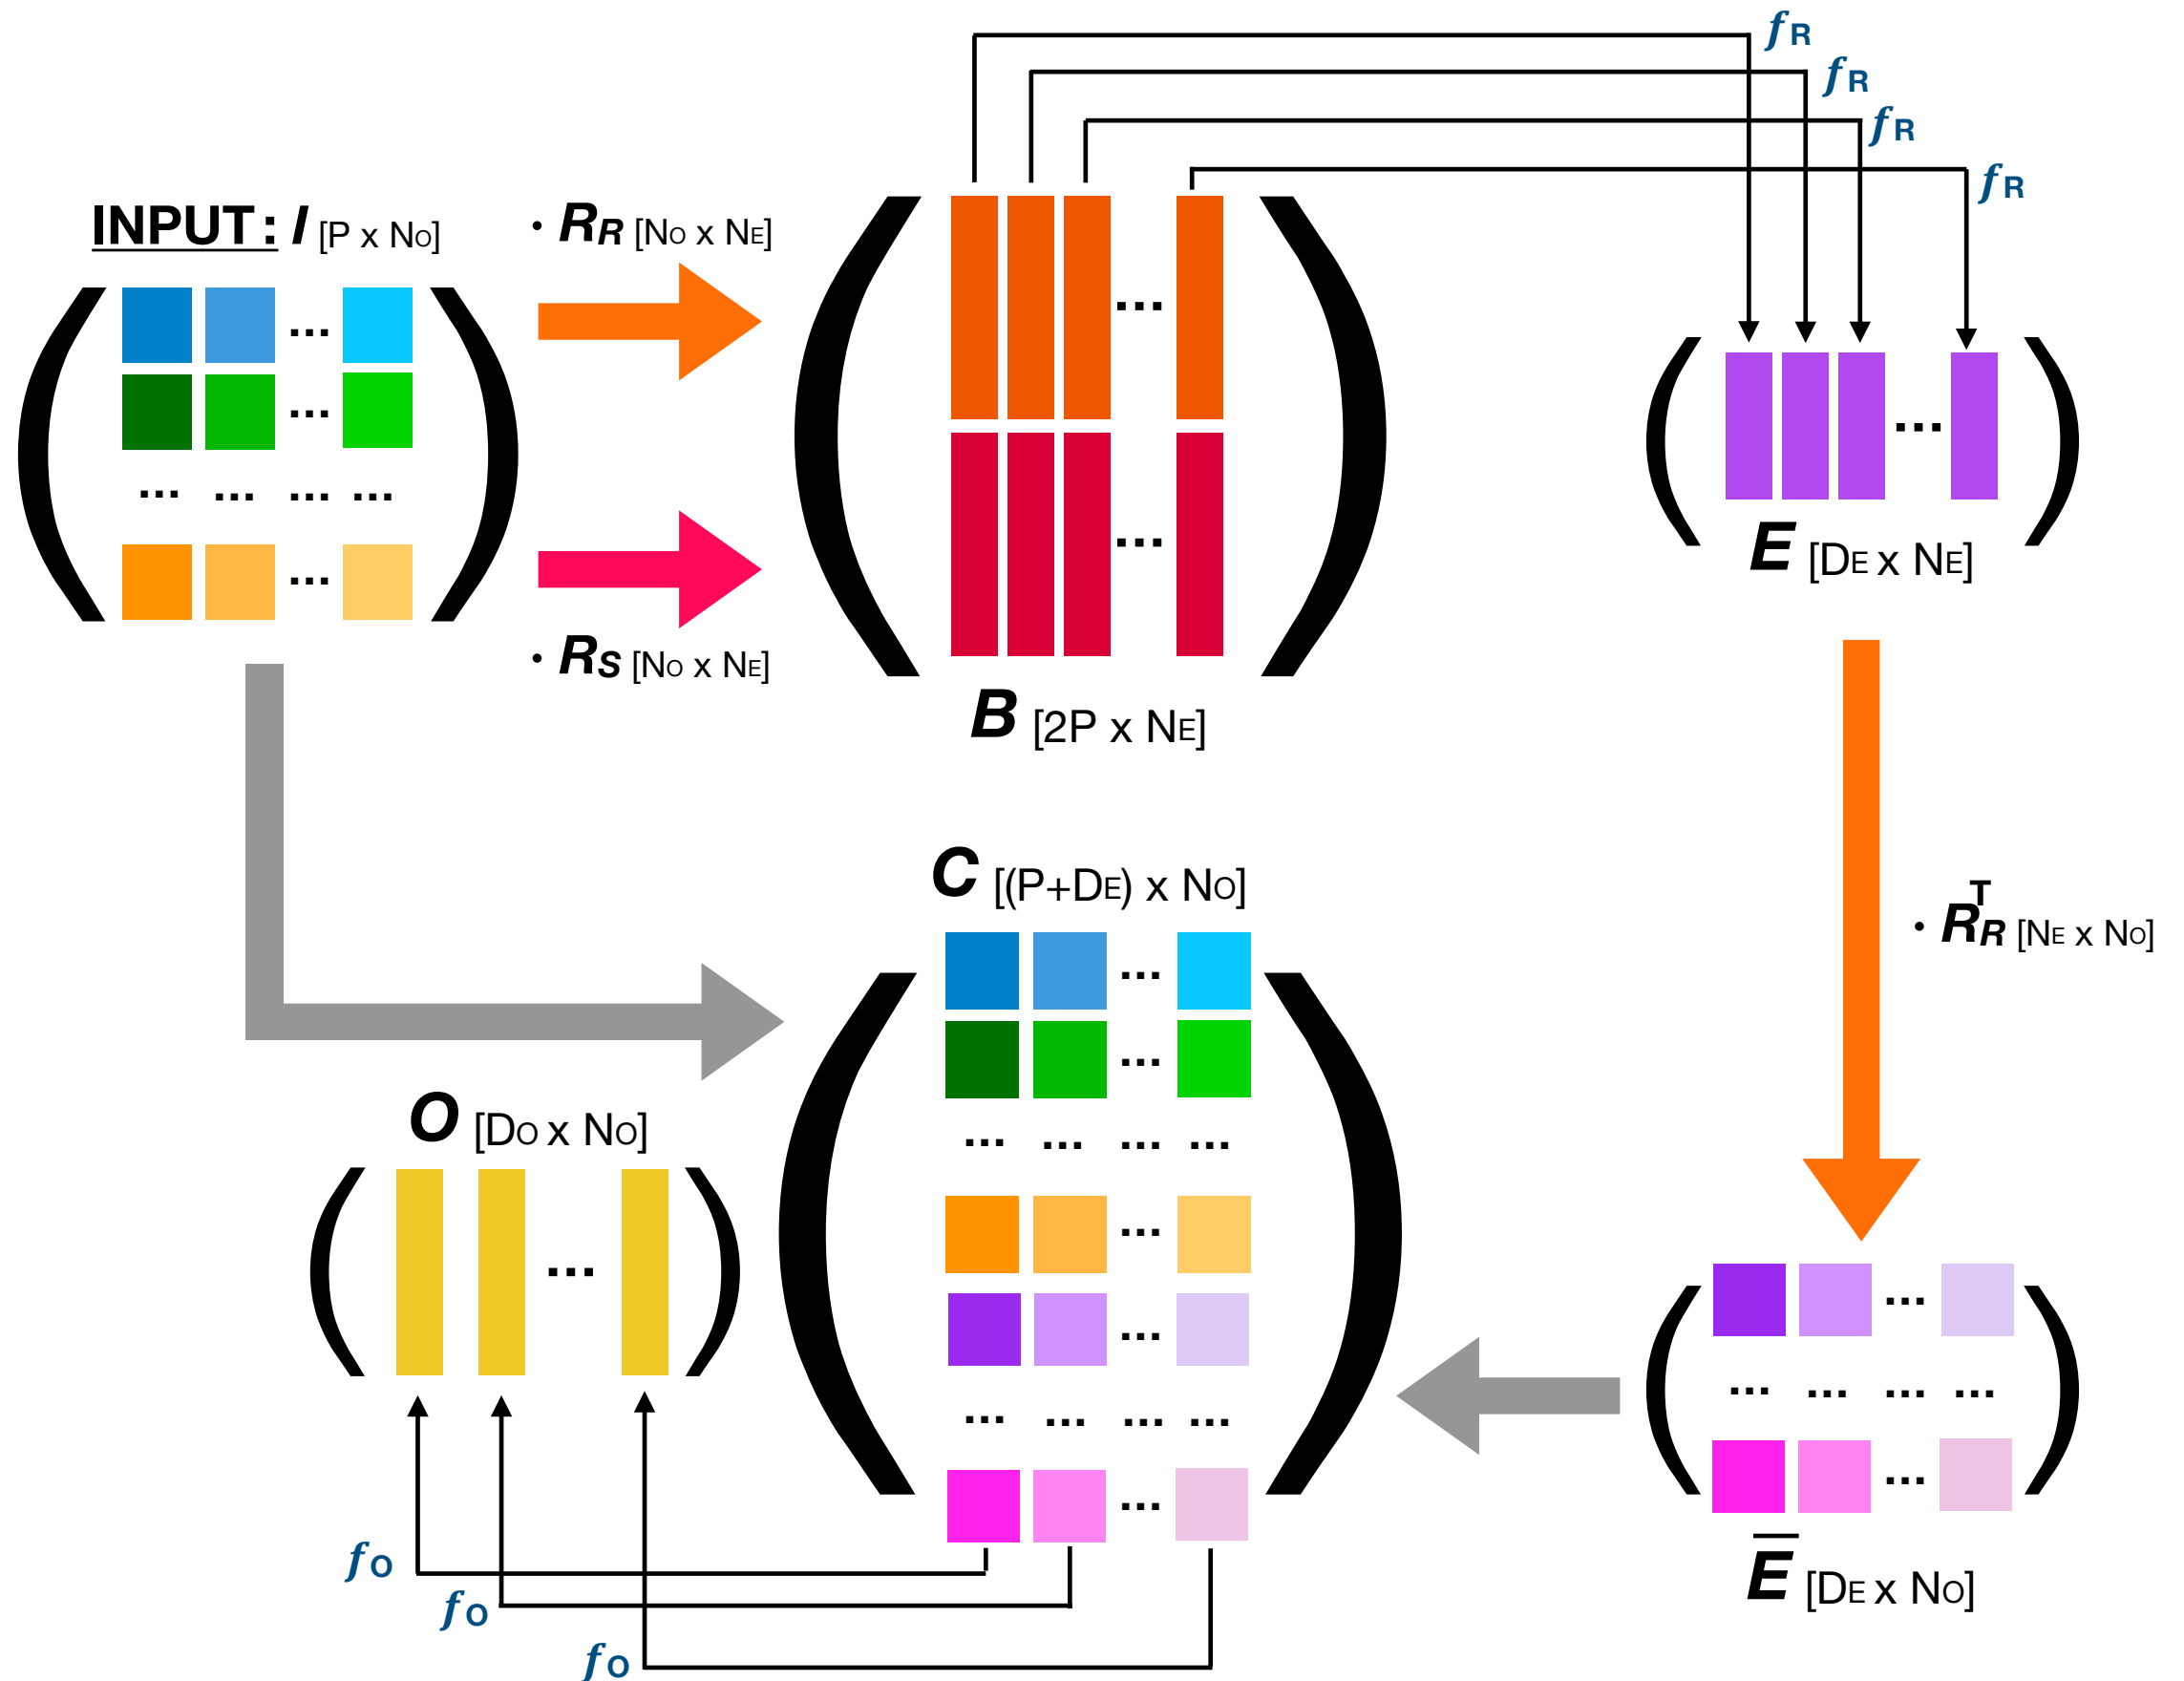


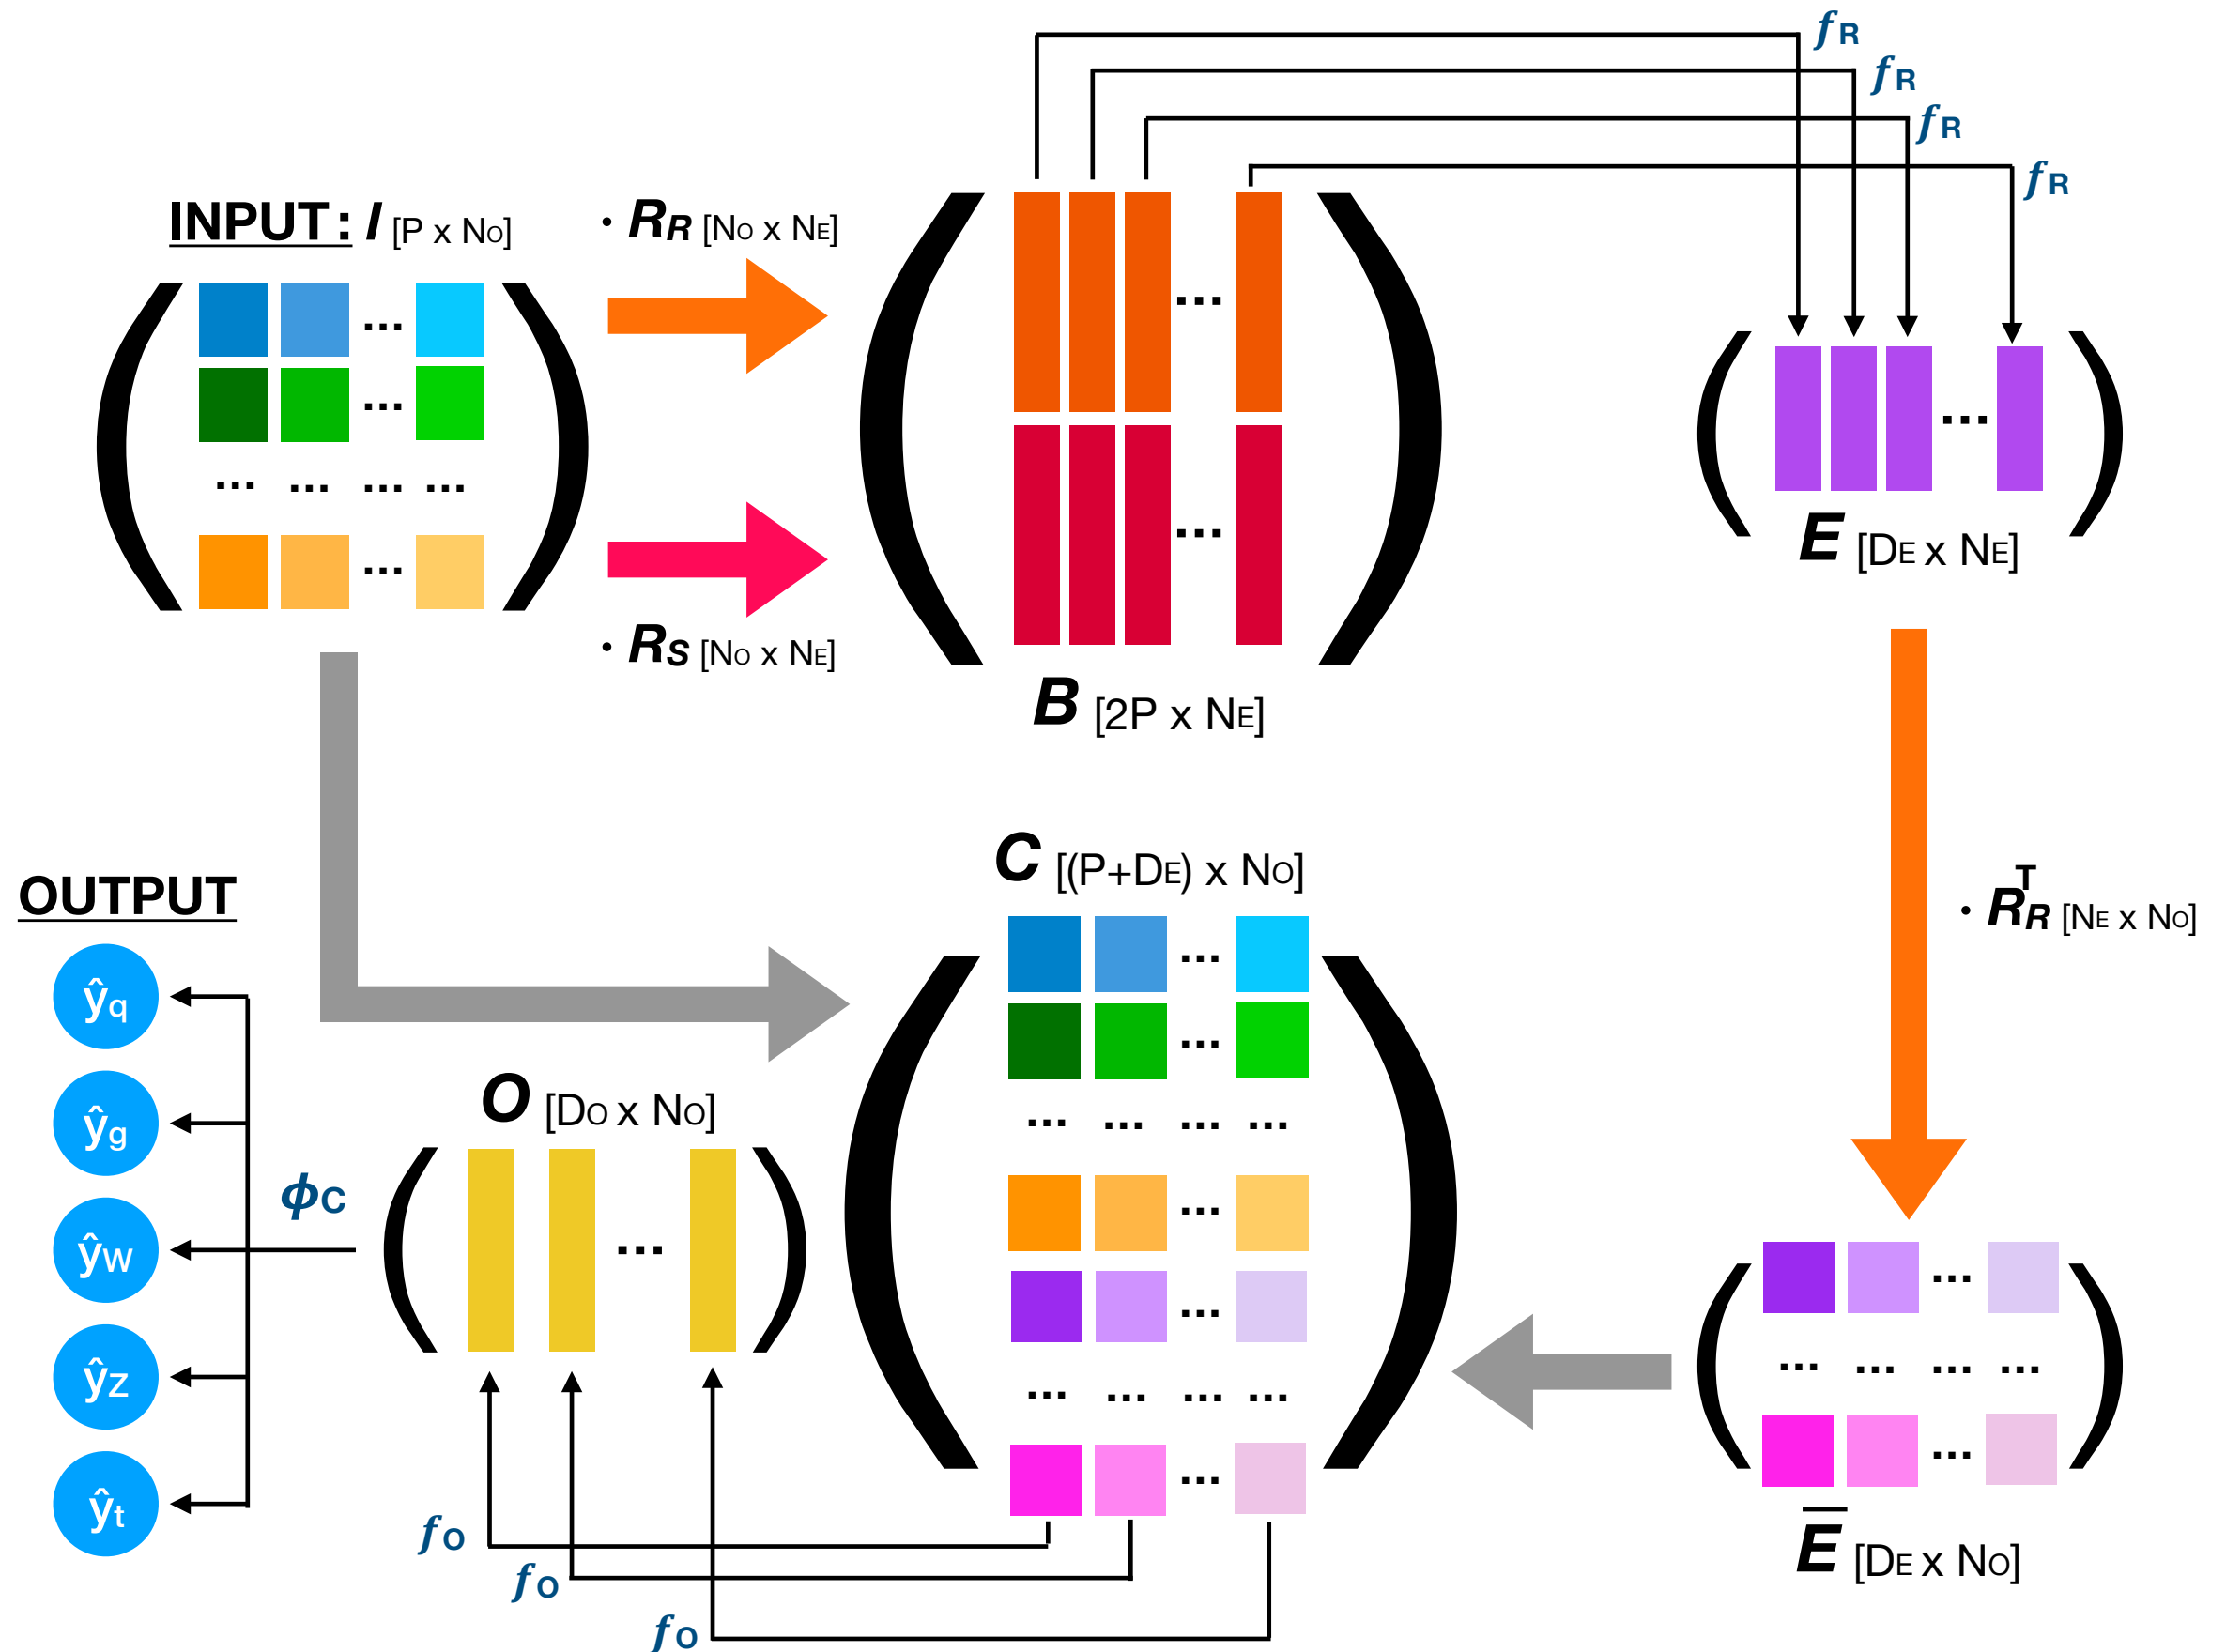
















- ▶ **Performance gain**
- ▶ GNNs have many other applications in HEP
 - ▶ tracking [[arXiv:1810.06111](https://arxiv.org/abs/1810.06111)]
 - ▶ clustering [[arXiv:1902.07987](https://arxiv.org/abs/1902.07987)]
 - ▶ **detector linking (i.e. particle flow)**
 - ▶ exotic particle tagging
 - ▶ anomaly detection
 - ▶ detector simulation

input per event:

set of elements $e = \{ECAL, HCAL, TRK, \dots\}$



output per event:

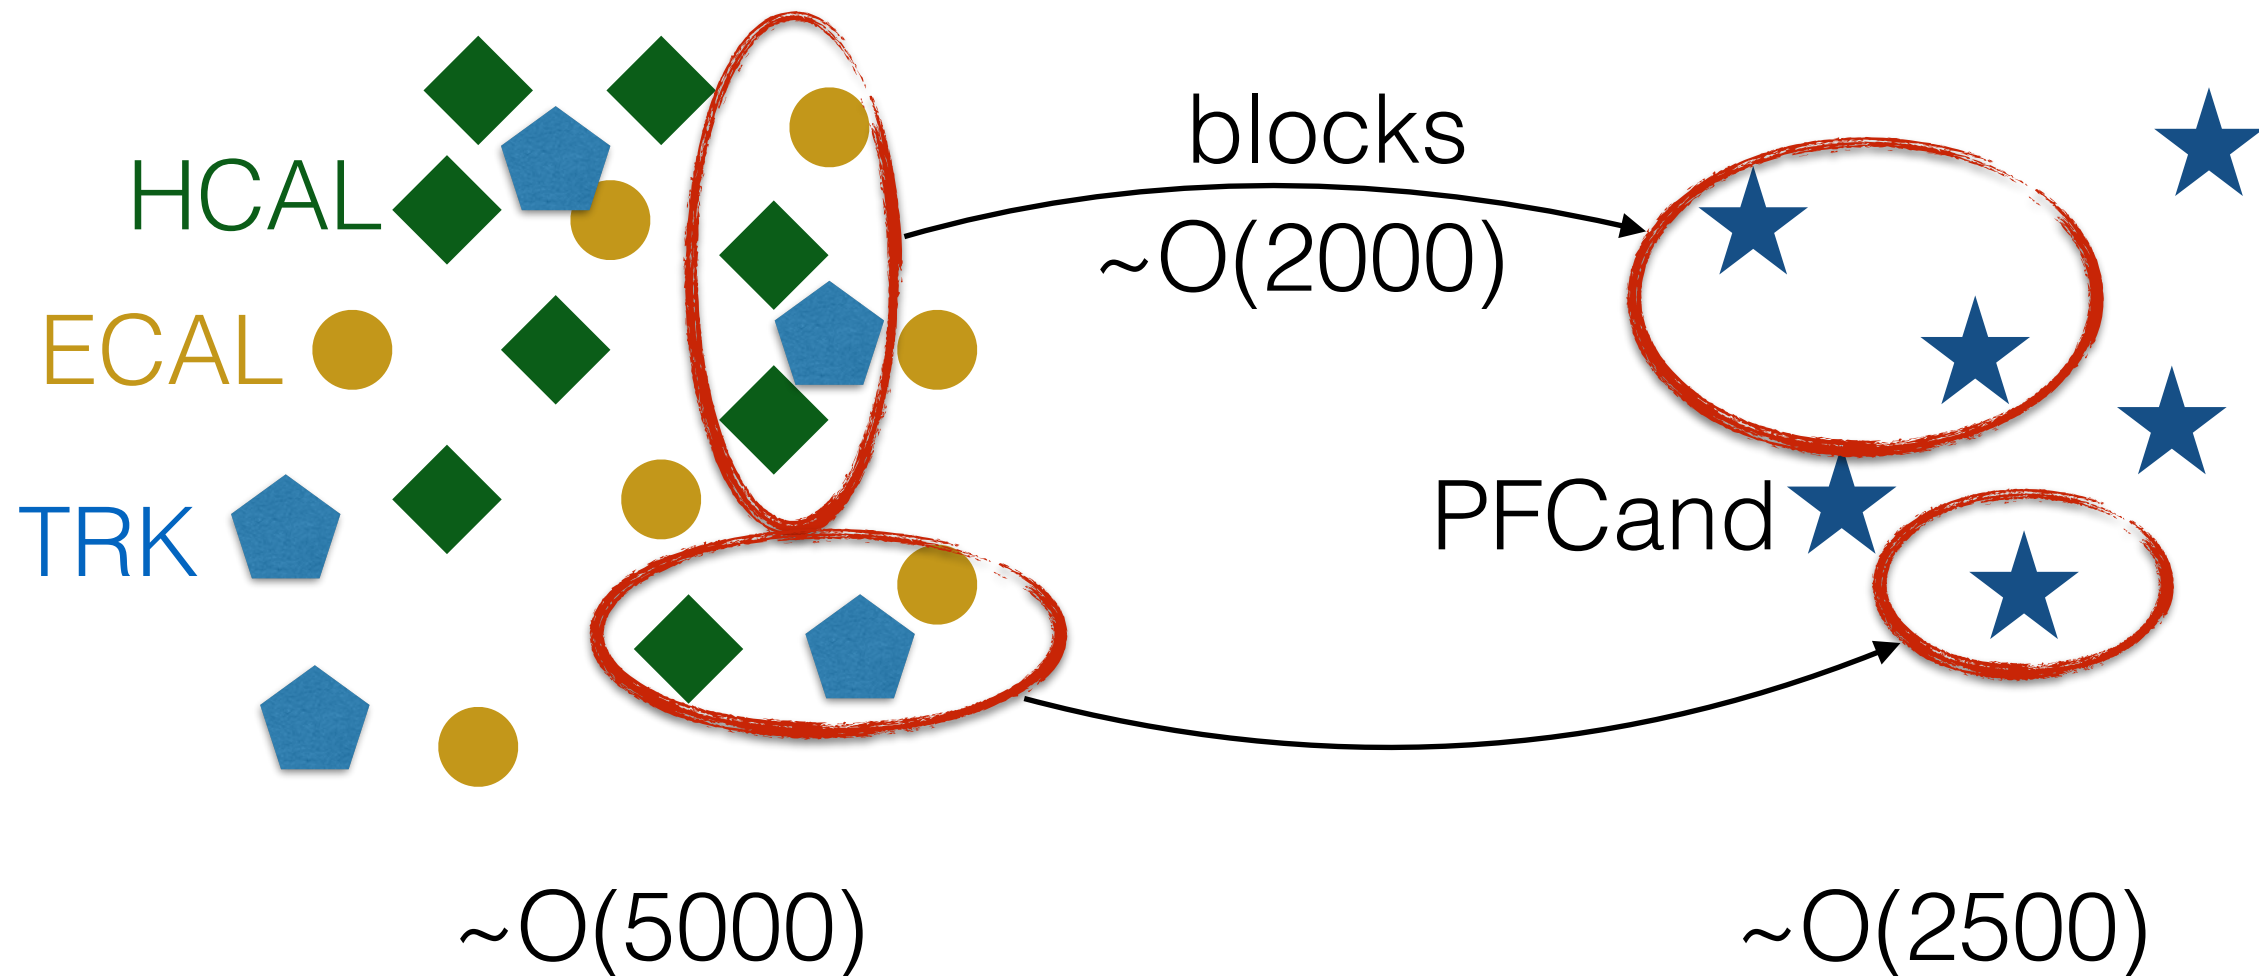
set of candidates $\mathbf{c} = \{\text{PFCand}, \dots\}$



$\sim O(2500)$


output per event:

set of candidates $\mathbf{c} = \{\text{PFCand}, \dots\}$




block: candidate associated to elements
A few elements \rightarrow a few candidates

Input Graph



Input Graph

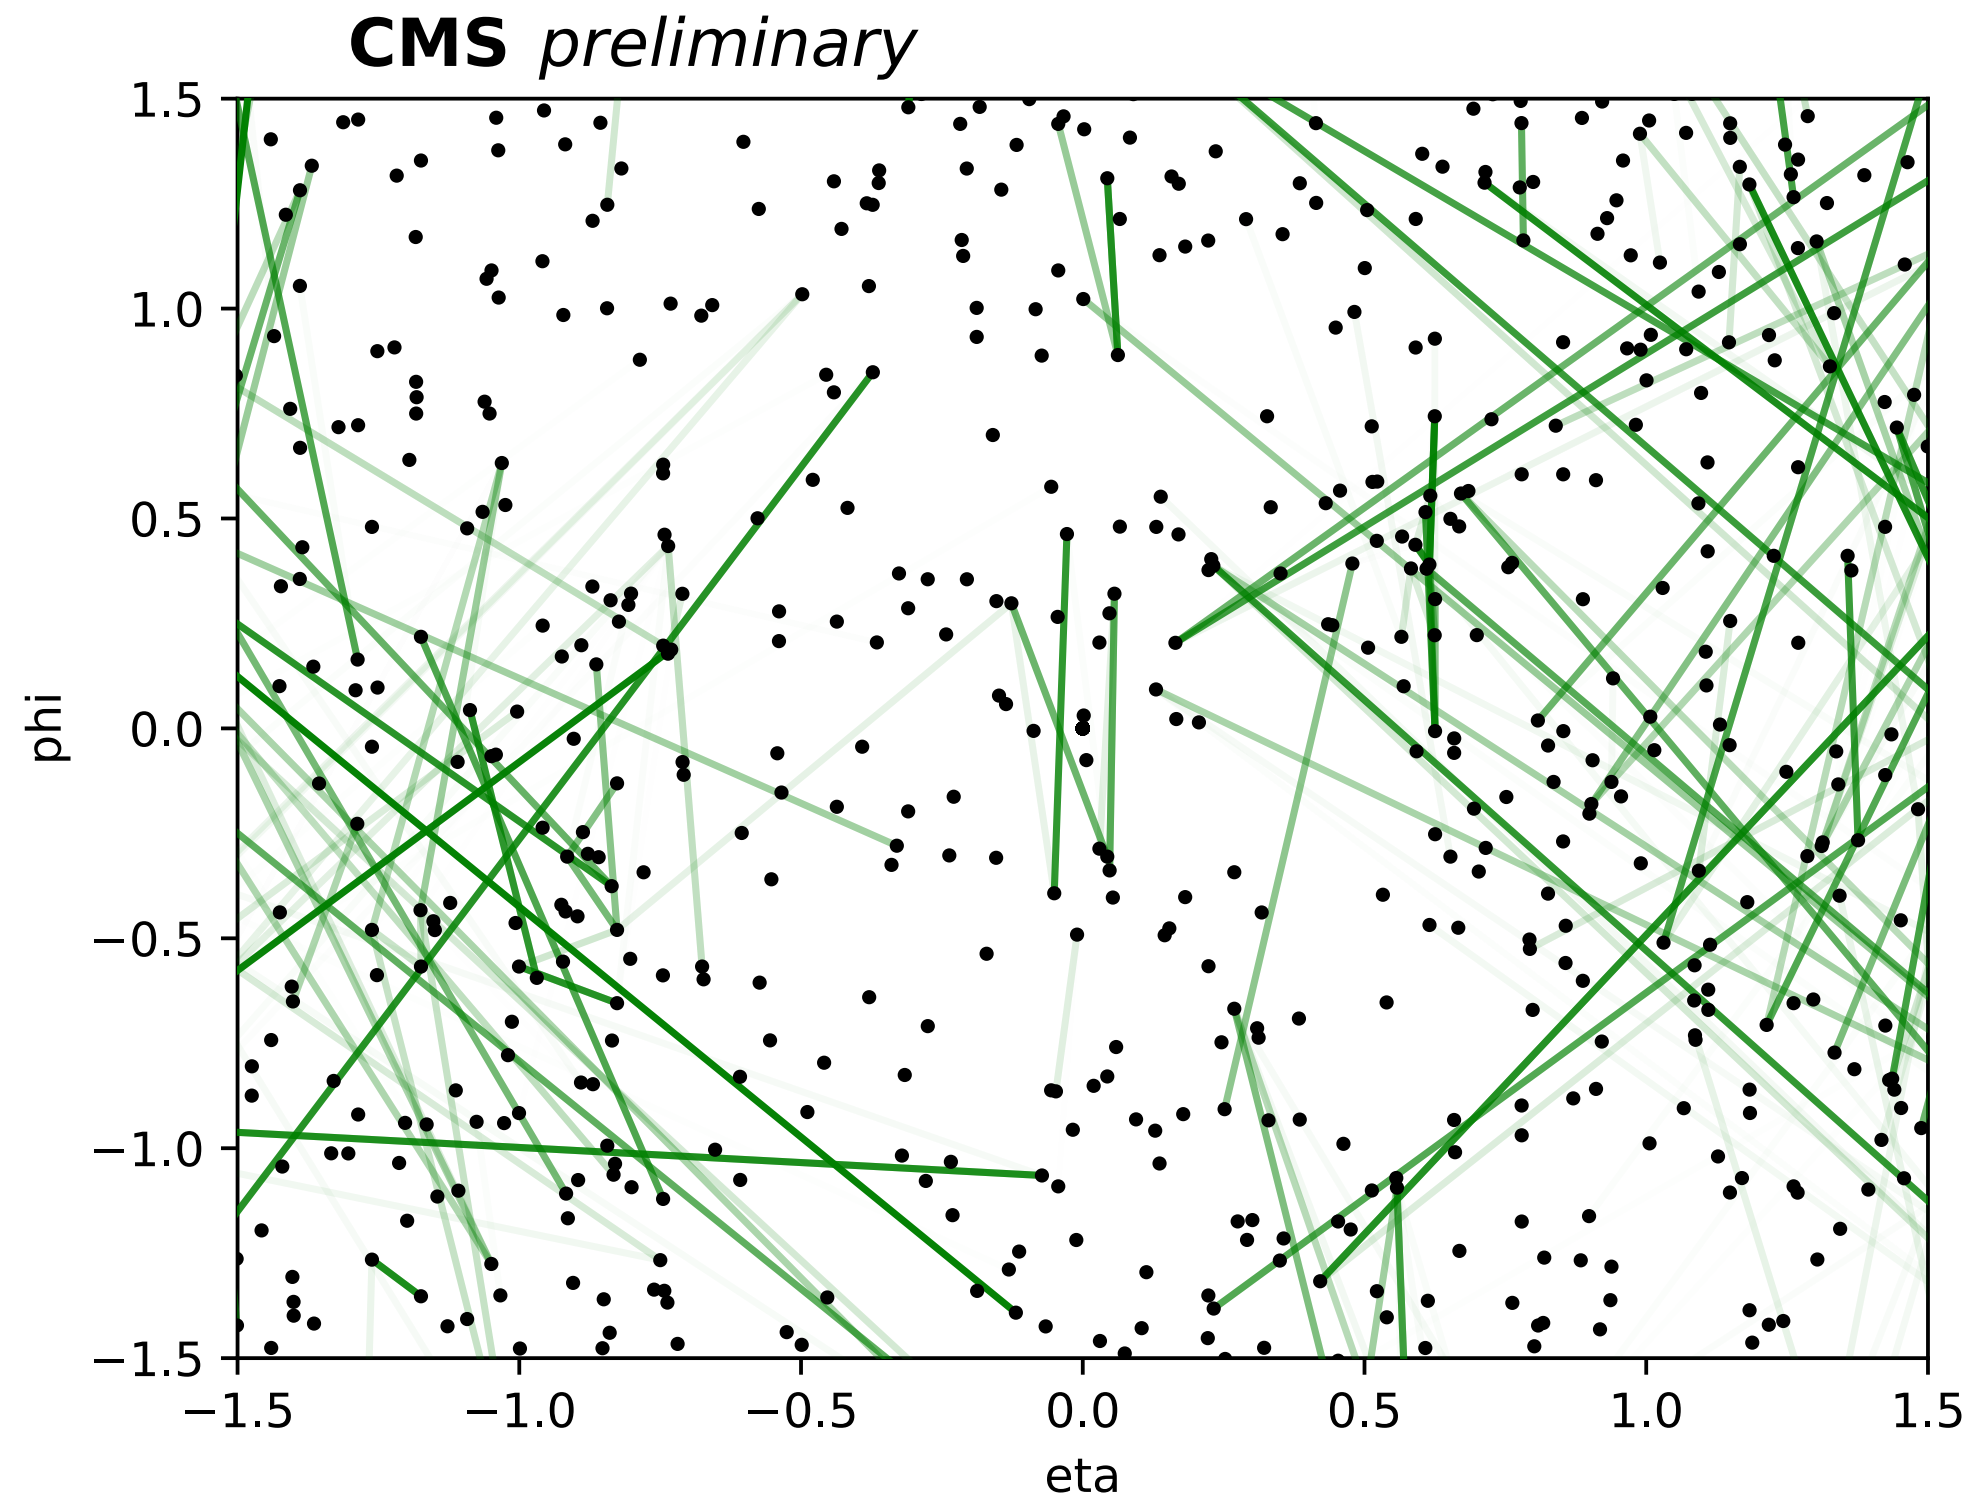
Truth Graph



Input Graph

Truth Graph

Output Graph







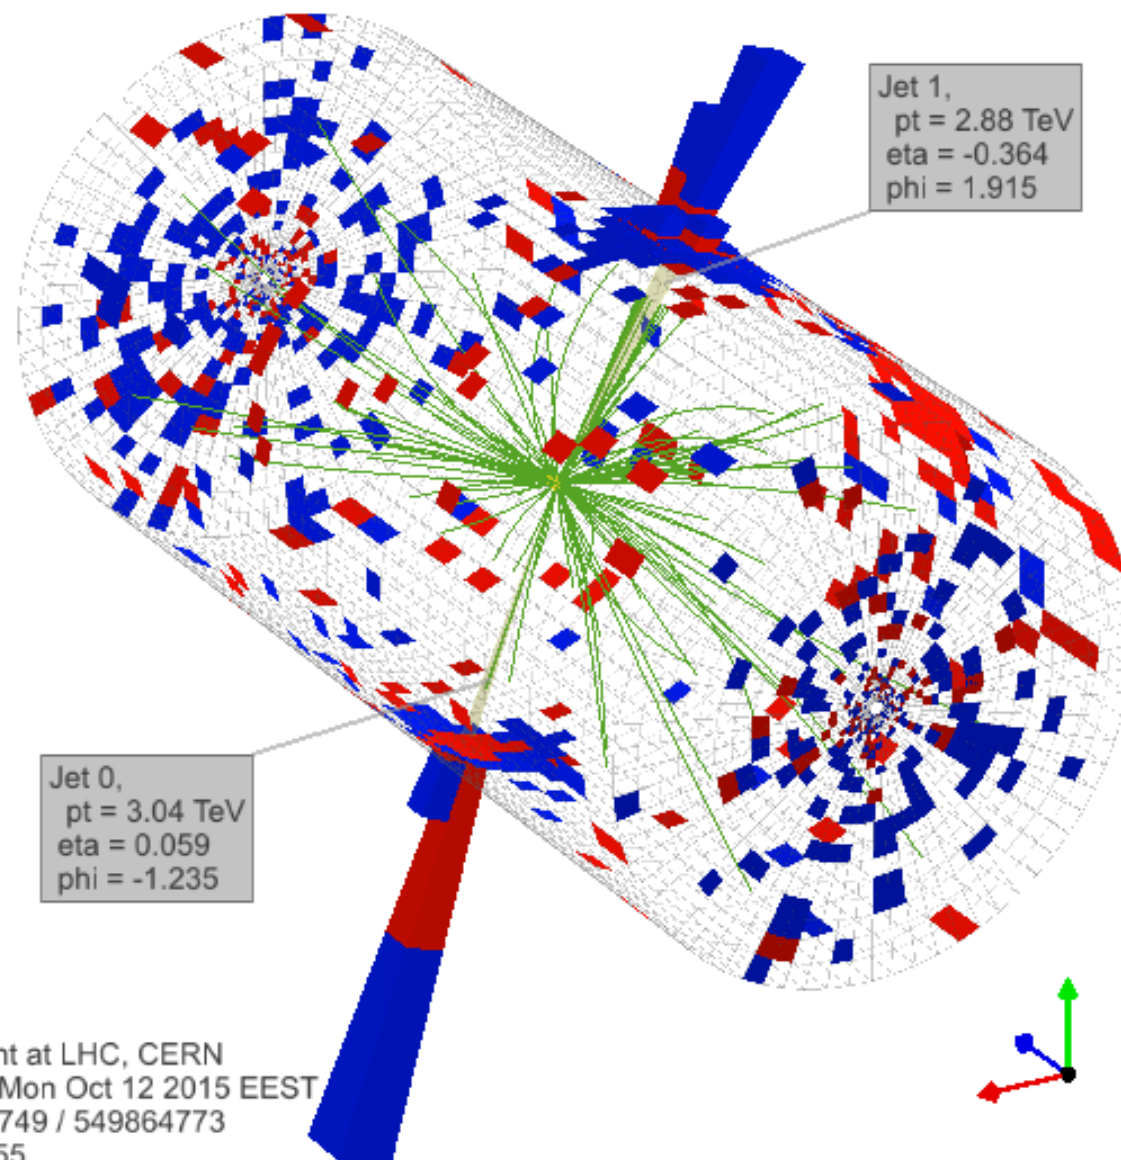
CHAPTER 1: OPPORTUNITIES &
CHALLENGES OF GEOMETRIC DEEP
LEARNING

**CHAPTER 2: UNSUPERVISED ANOMALY
DETECTION FOR NEW PHYSICS**

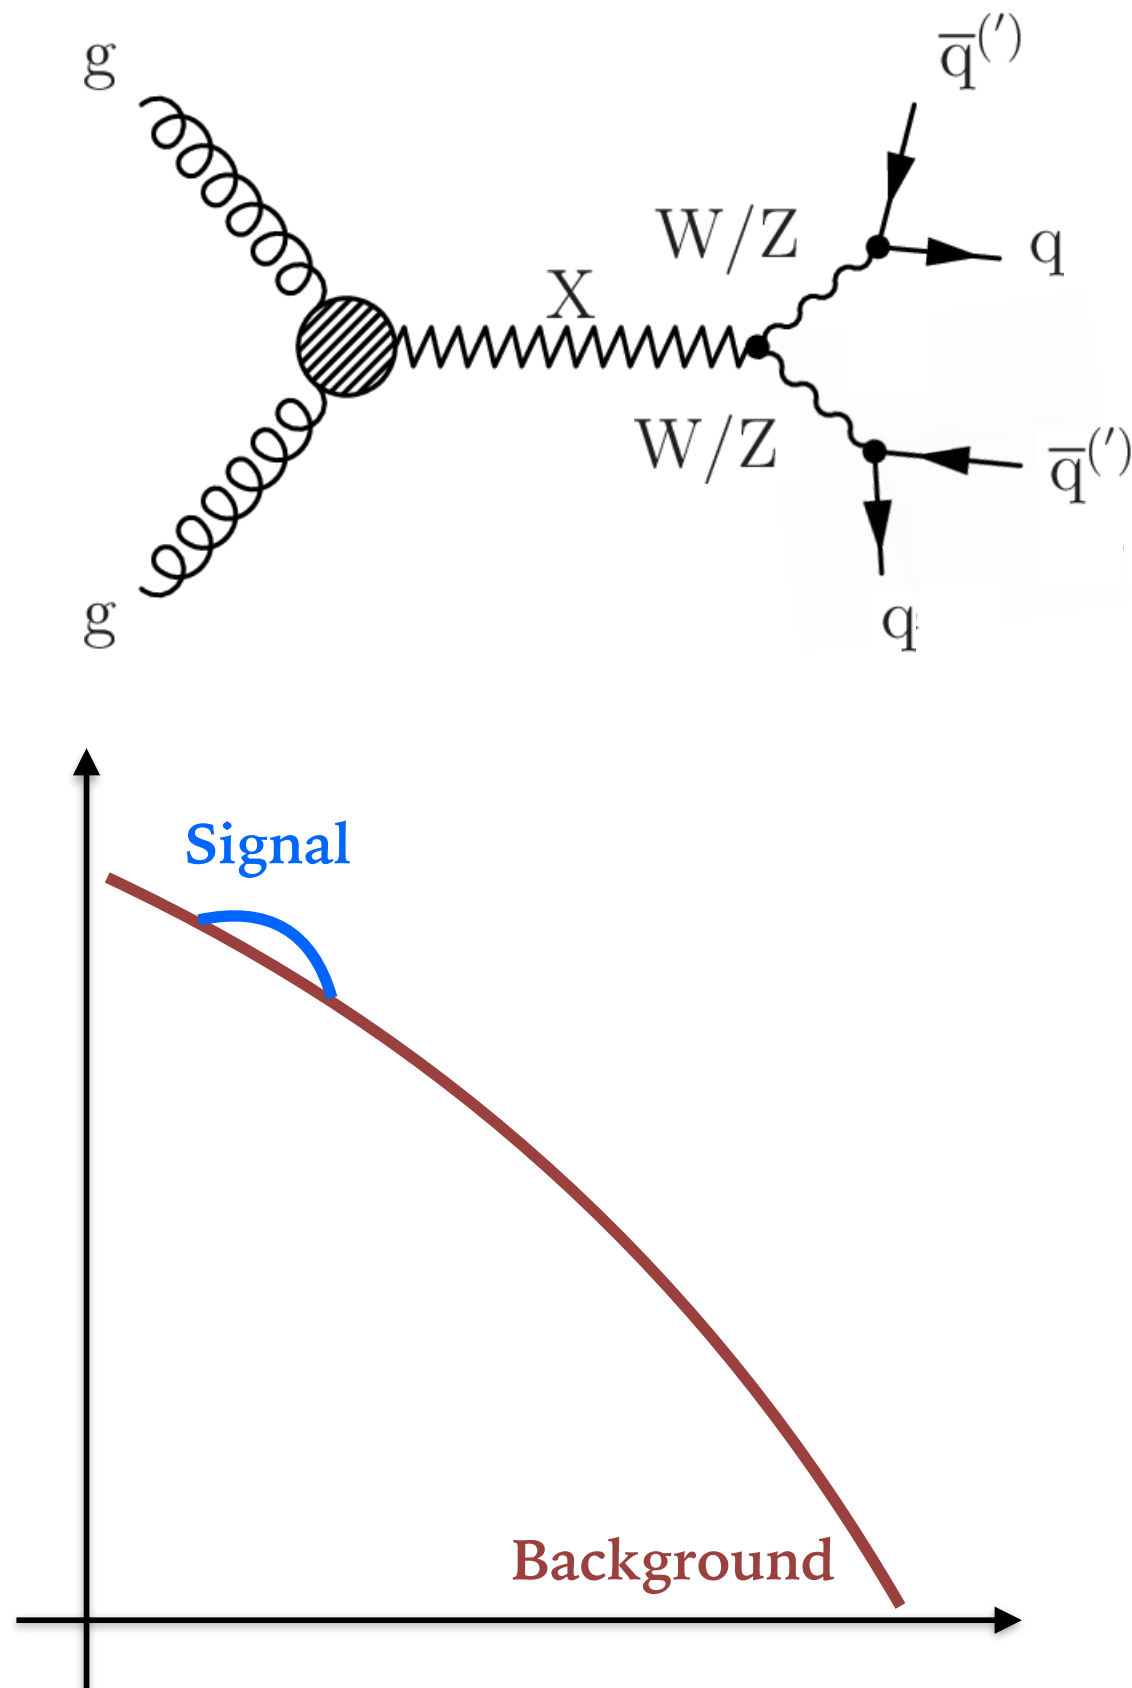
CHAPTER 3: DEEP LEARNING IN THE
TRIGGER



CMS Experiment at LHC, CERN
Data recorded: Mon Oct 12 03:26:27 2015 EEST
Run/Event: 258749 / 549864773
Lumi section: 355
Dijet Mass: 6.14 TeV


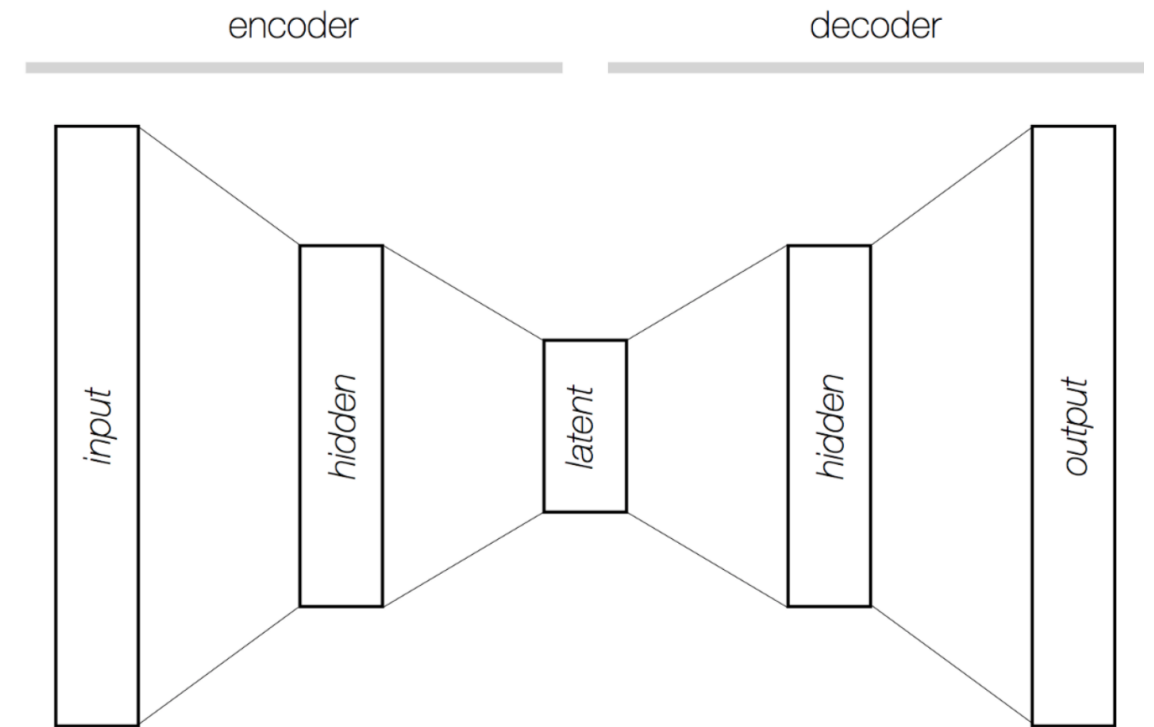


CMS Experiment at LHC, CERN
Data recorded: Mon Oct 12 2015 EEST
Run/Event: 258749 / 549864773
Lumi section: 355
Dijet Mass: 6.14 TeV




- ▶ Look for new heavy particle decaying to two (wide) jets
 - ▶ Compute invariant mass of two high- p_T (wide) jets
 - ▶ Look for a bump (indicating a new resonance) over a smoothly falling background
- ▶ Problems
 - ▶ Very large background
 - ▶ Many signal models; should we create an ML algorithm to identify each one?
- ▶ Can we use unsupervised methods for model-agnostic new physics searches?

- ▶ Map an input onto itself passing through a latent representation
- ▶ Unsupervised algorithm, used for data compression, generation, clustering, etc.
- ▶ Anomaly: any event whose output is "far" from the input





See: [O. Cerri Lepton Photon 2019](#), [K. Wozniak CHEP 2019](#)



Training:


Fit the VAE params to minimize the input-output distance






Evaluate:


One-sided hypothesis test on the input-output distance






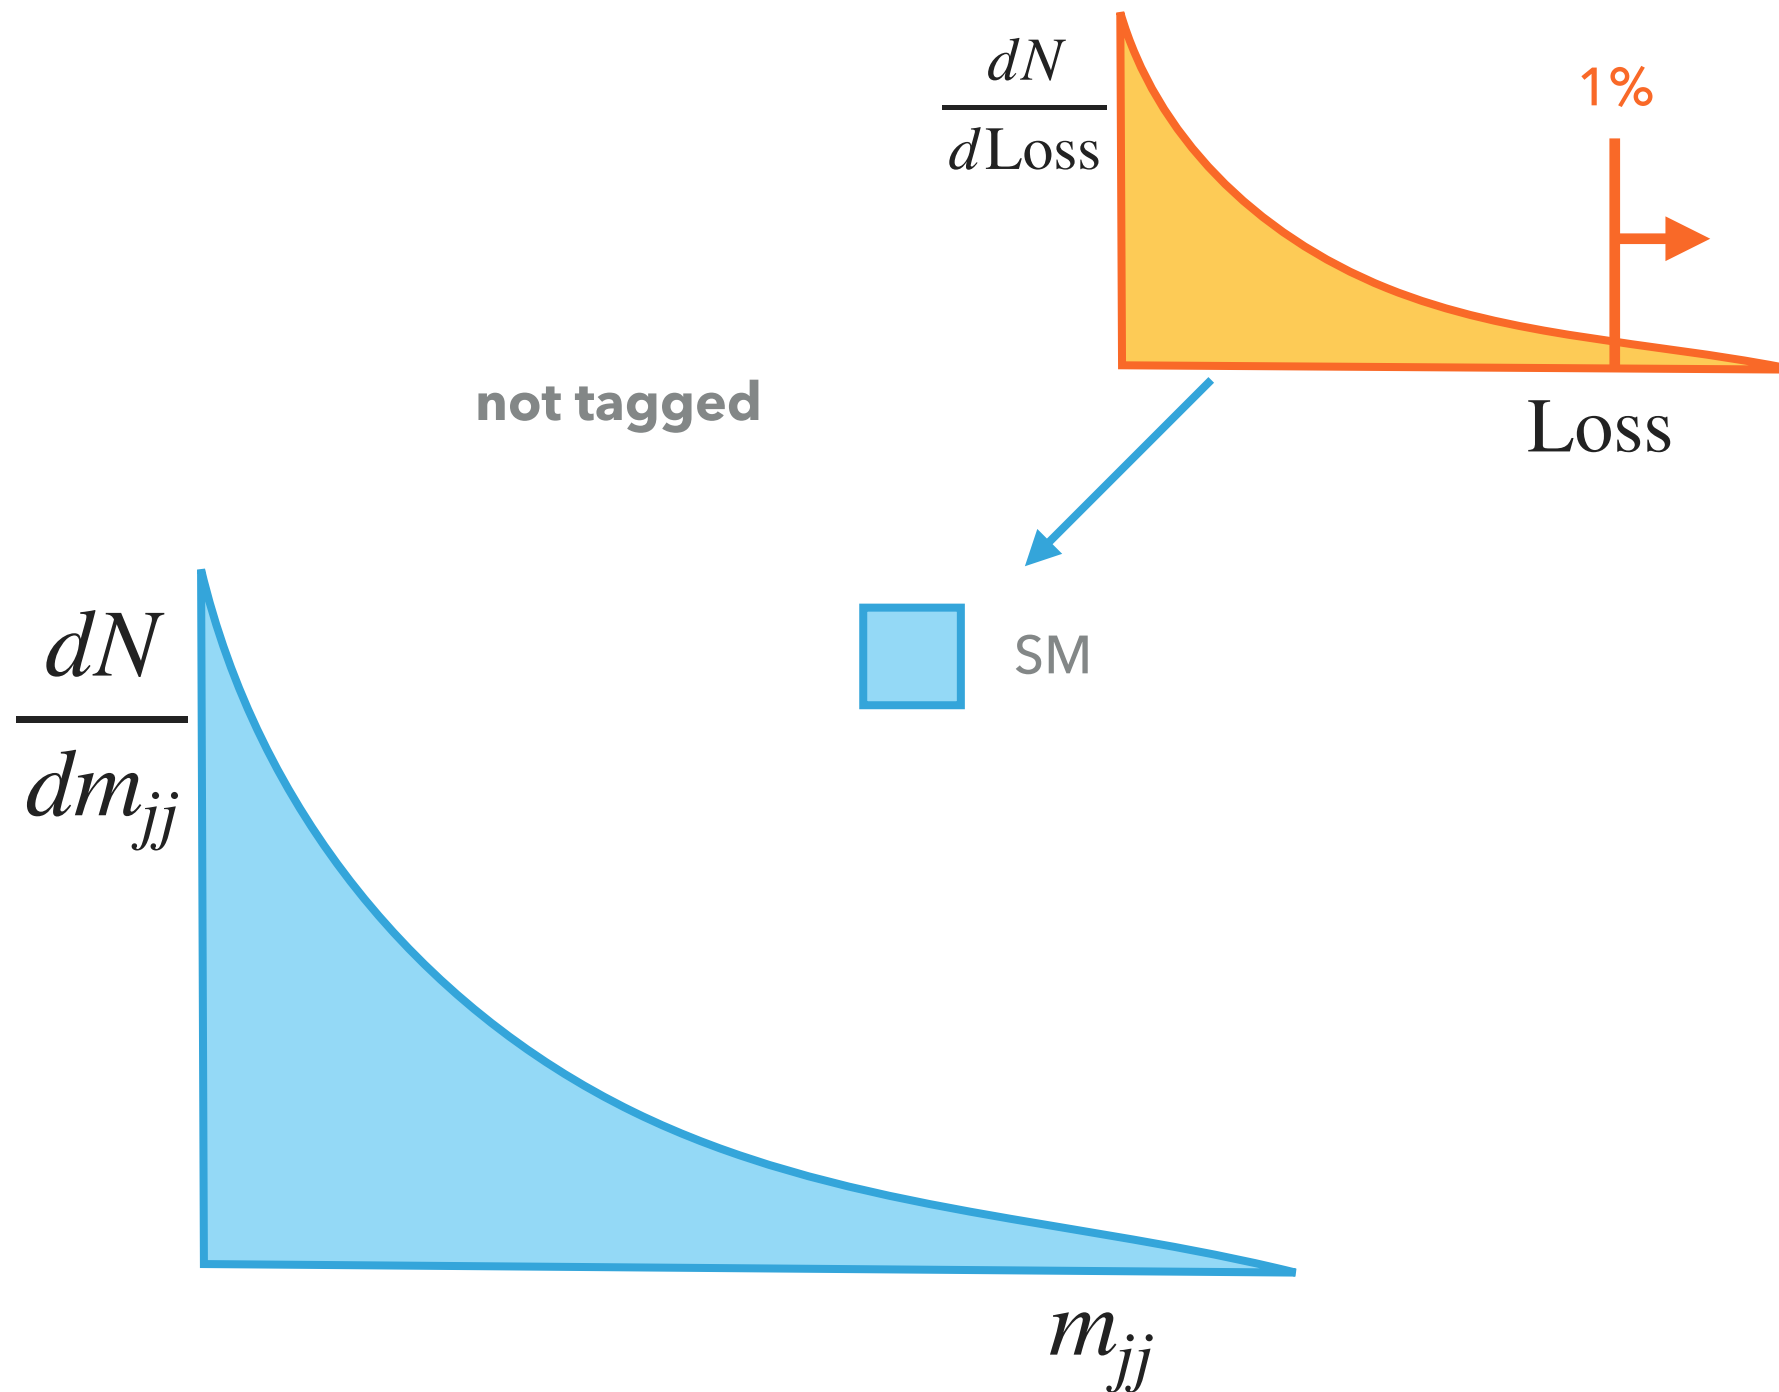
Evaluate:

One-sided hypothesis test on the input-output distance

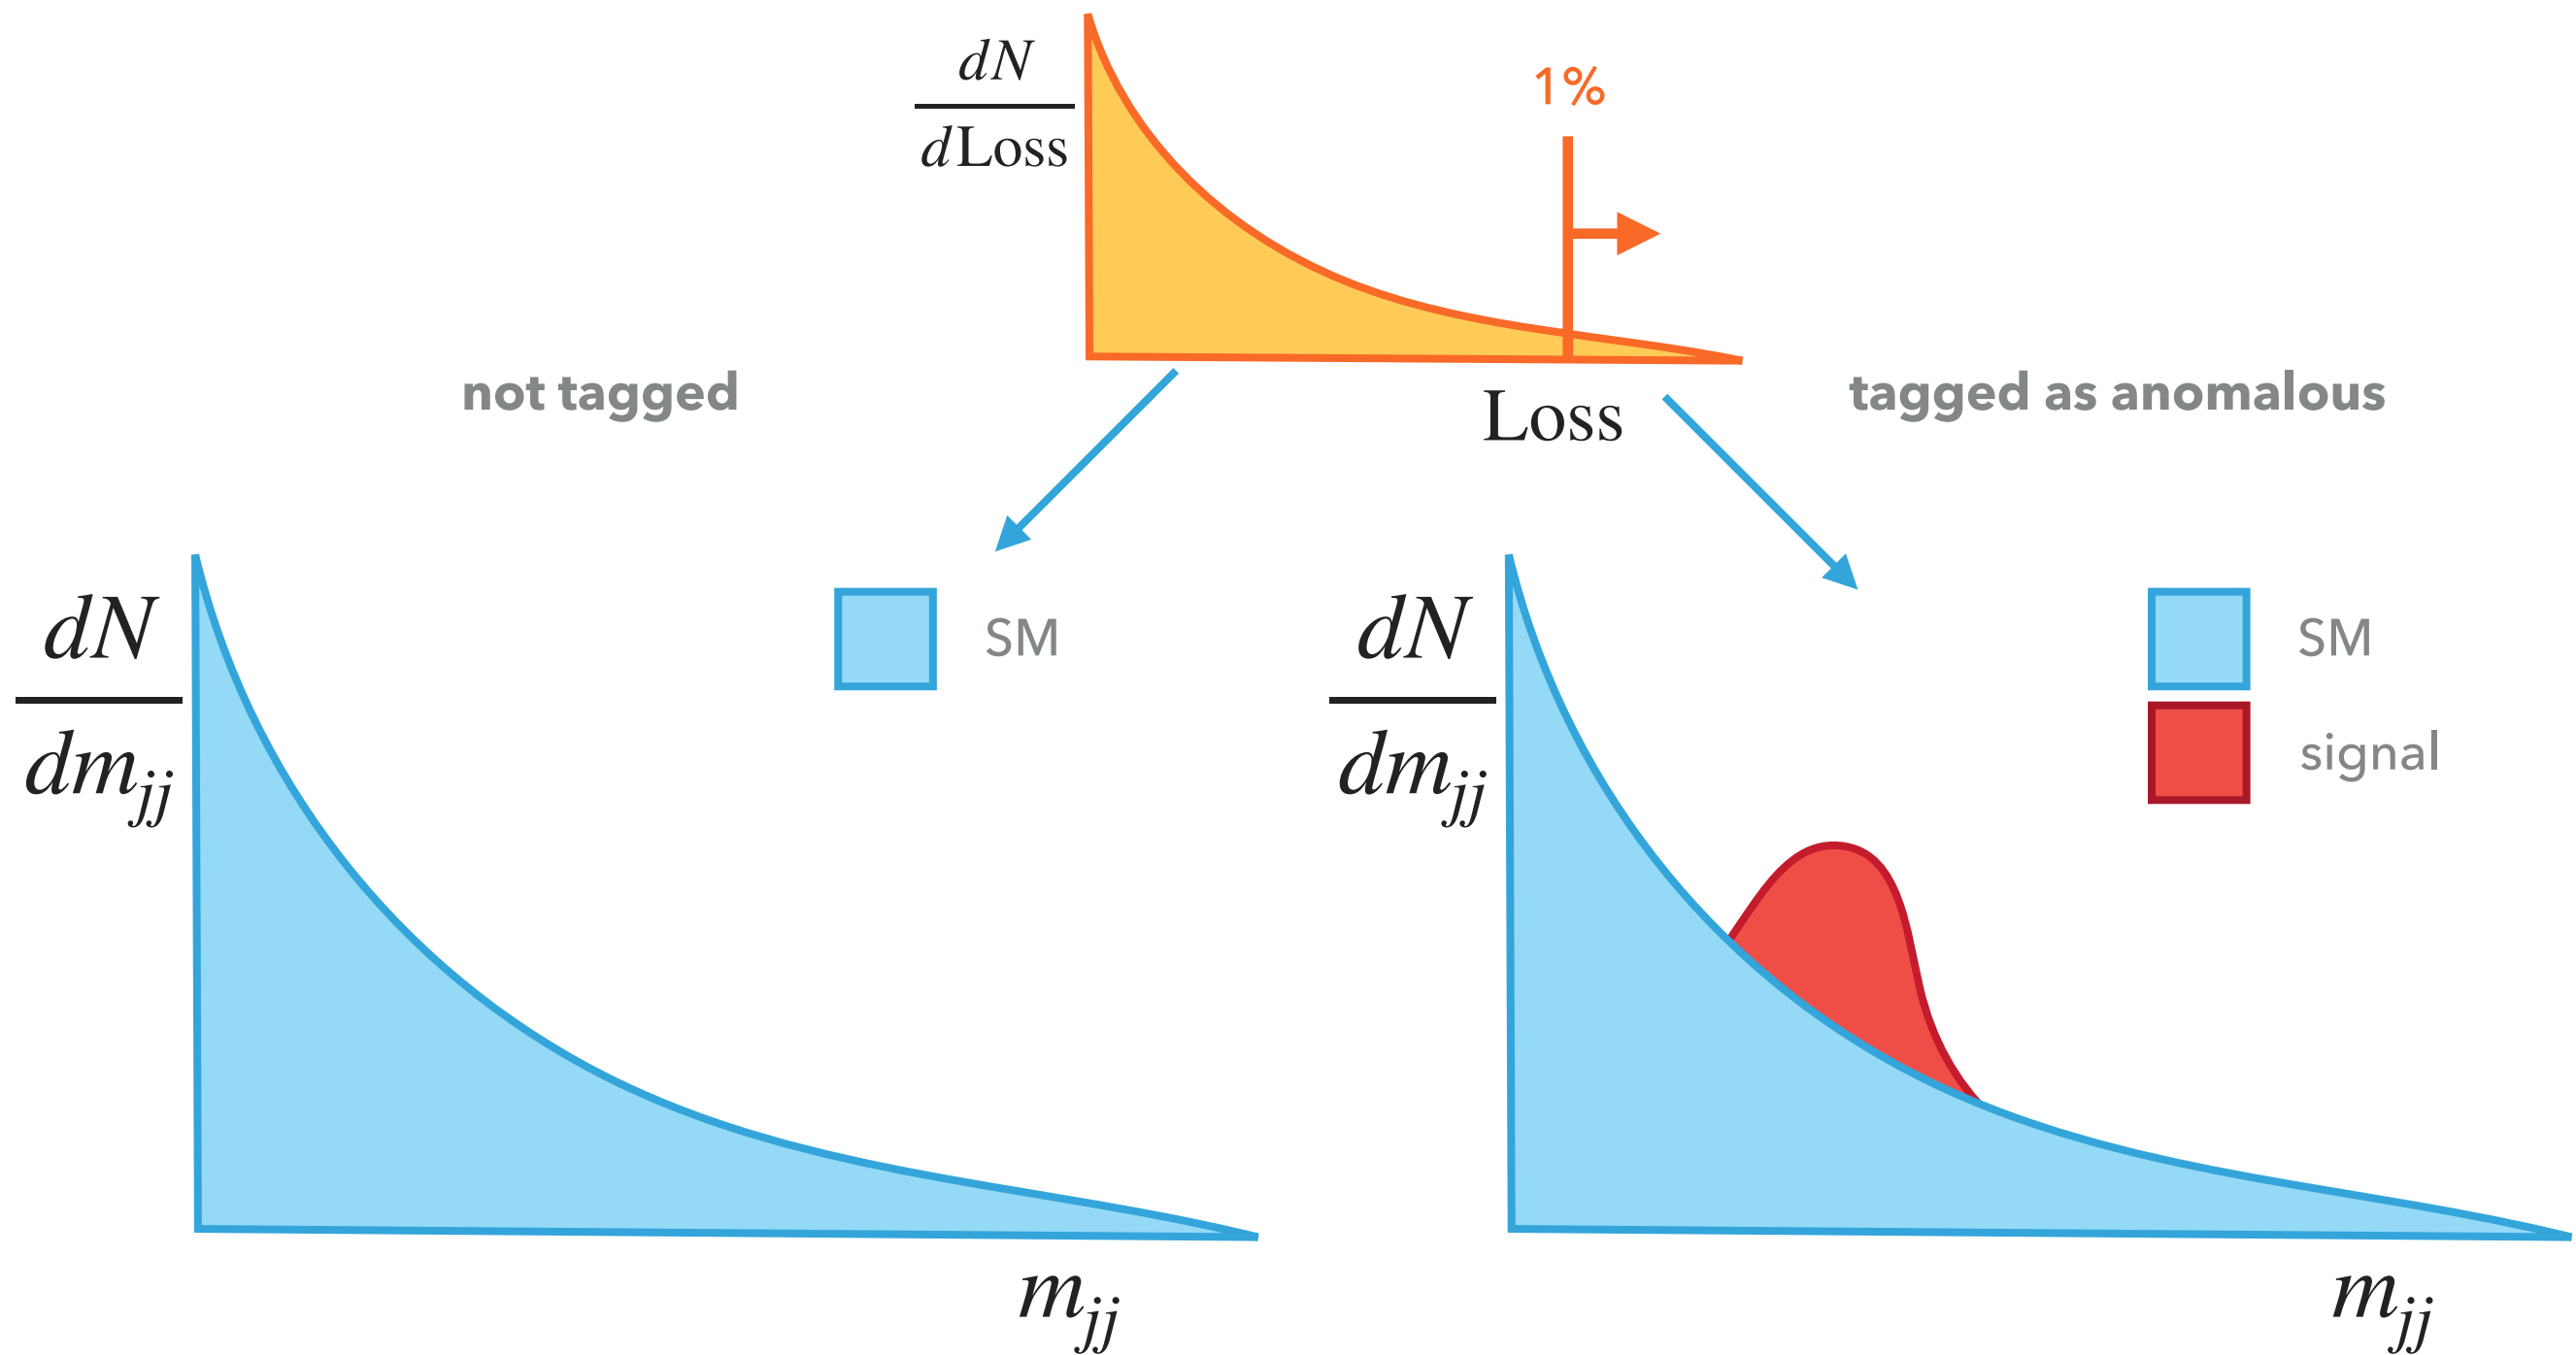


- ▶ Dataset:
 - ▶ QCD dijet simulation (Pythia + Delphes)
- ▶ Input:
 - ▶ anti- k_T $R=0.8$ jets
 - ▶ transformed to binned, p_T -weighted **jet images**
- ▶ Training in **control region**:
 $1.4 < |\Delta\eta| < 2.4$
- ▶ Application in **signal region**: $|\Delta\eta| < 1.4$






- ▶ Cut on the loss to keep 1% of background events
- ▶ Use untagged events to constrain background shape in tagged region






- ▶ Cut on the loss to keep 1% of background events
- ▶ Use untagged events to constrain background shape in tagged region



- ▶ Note background m_{jj} distribution is not preserved after applying selection on loss (***sculpting!***)
- ▶ But we can apply a m_{jj} dependent threshold on the loss to preserve shape of background



- ▶ Note background m_{jj} distribution is not preserved after applying selection on loss (***sculpting!***)
- ▶ But we can apply a m_{jj} dependent threshold on the loss to preserve shape of background



- ▶ Comparison between standard dijet search and VAE-assisted search
- ▶ Sensitivity boosted from 3σ to 4σ
- ▶ Can we apply this technique in the "trigger" algorithm in hardware?




The background features a stylized brain graphic. The left side of the brain is filled with yellow circuitry patterns, while the right side is filled with yellow organic, cellular-like shapes. In the center of the brain, there is a circular area containing a complex network of yellow and green lines, resembling a neural network or a data visualization. The overall color palette is dark grey, yellow, and light blue.


CHAPTER 1: OPPORTUNITIES &
CHALLENGES OF GEOMETRIC DEEP
LEARNING

CHAPTER 2: UNSUPERVISED ANOMALY
DETECTION FOR NEW PHYSICS


**CHAPTER 3: DEEP LEARNING IN THE
TRIGGER**

UPGRADED LEVEL-1 TRIGGER







- ▶ Level-1 Trigger:
40 MHz → 750 kHz
- ▶ Reconstruct and filter
2% of events in $\sim 12 \mu\text{s}$




- ▶ Level-1 Trigger:
40 MHz → 750 kHz
- ▶ Reconstruct and filter
2% of events in $\sim 12 \mu\text{s}$
- ▶ Latency necessitates all
FPGA design






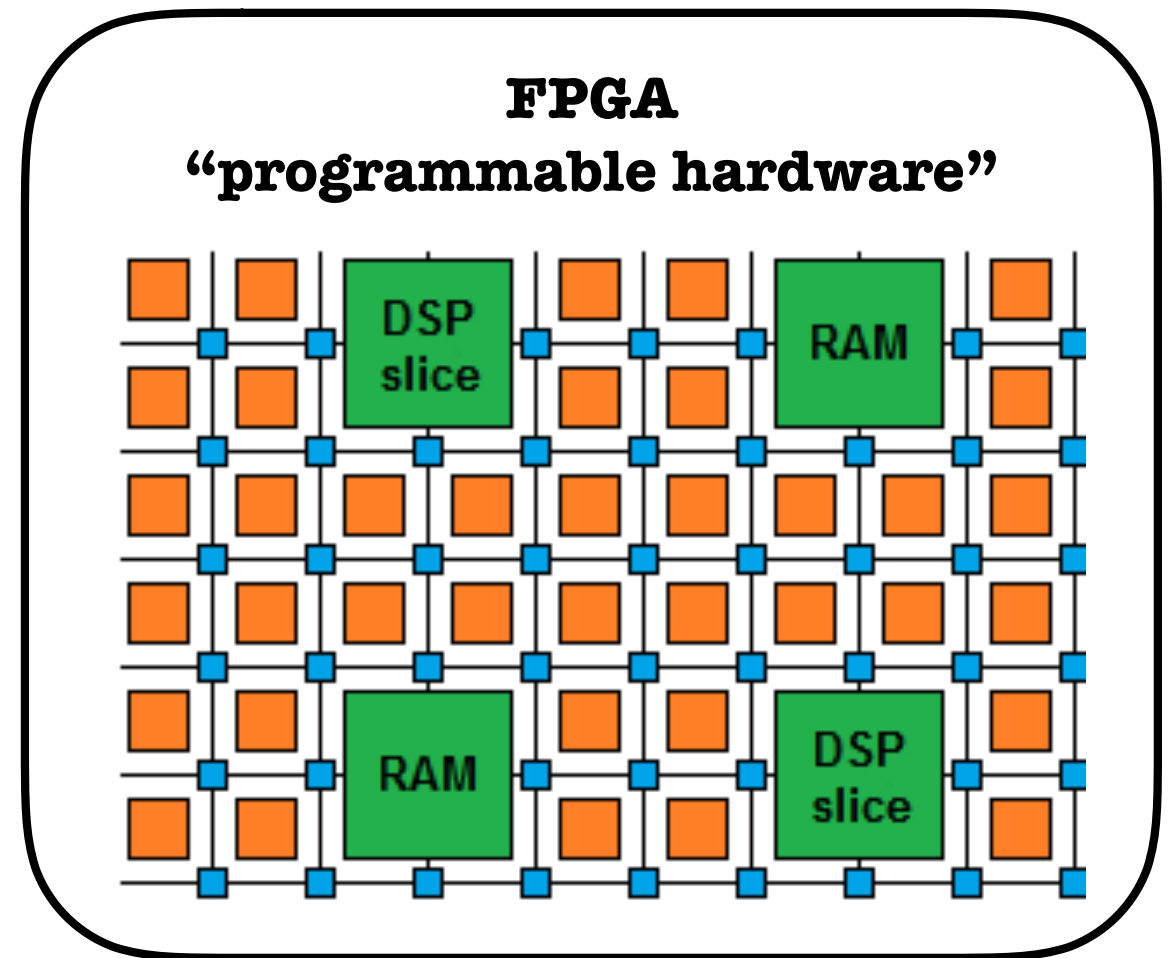
- ▶ Level-1 Trigger:
40 MHz → 750 kHz
- ▶ Reconstruct and filter
2% of events in $\sim 12 \mu\text{s}$
- ▶ Latency necessitates all
FPGA design



- ▶ **Pros:** reprogrammable,
high throughput,
massively parallel, & low
power




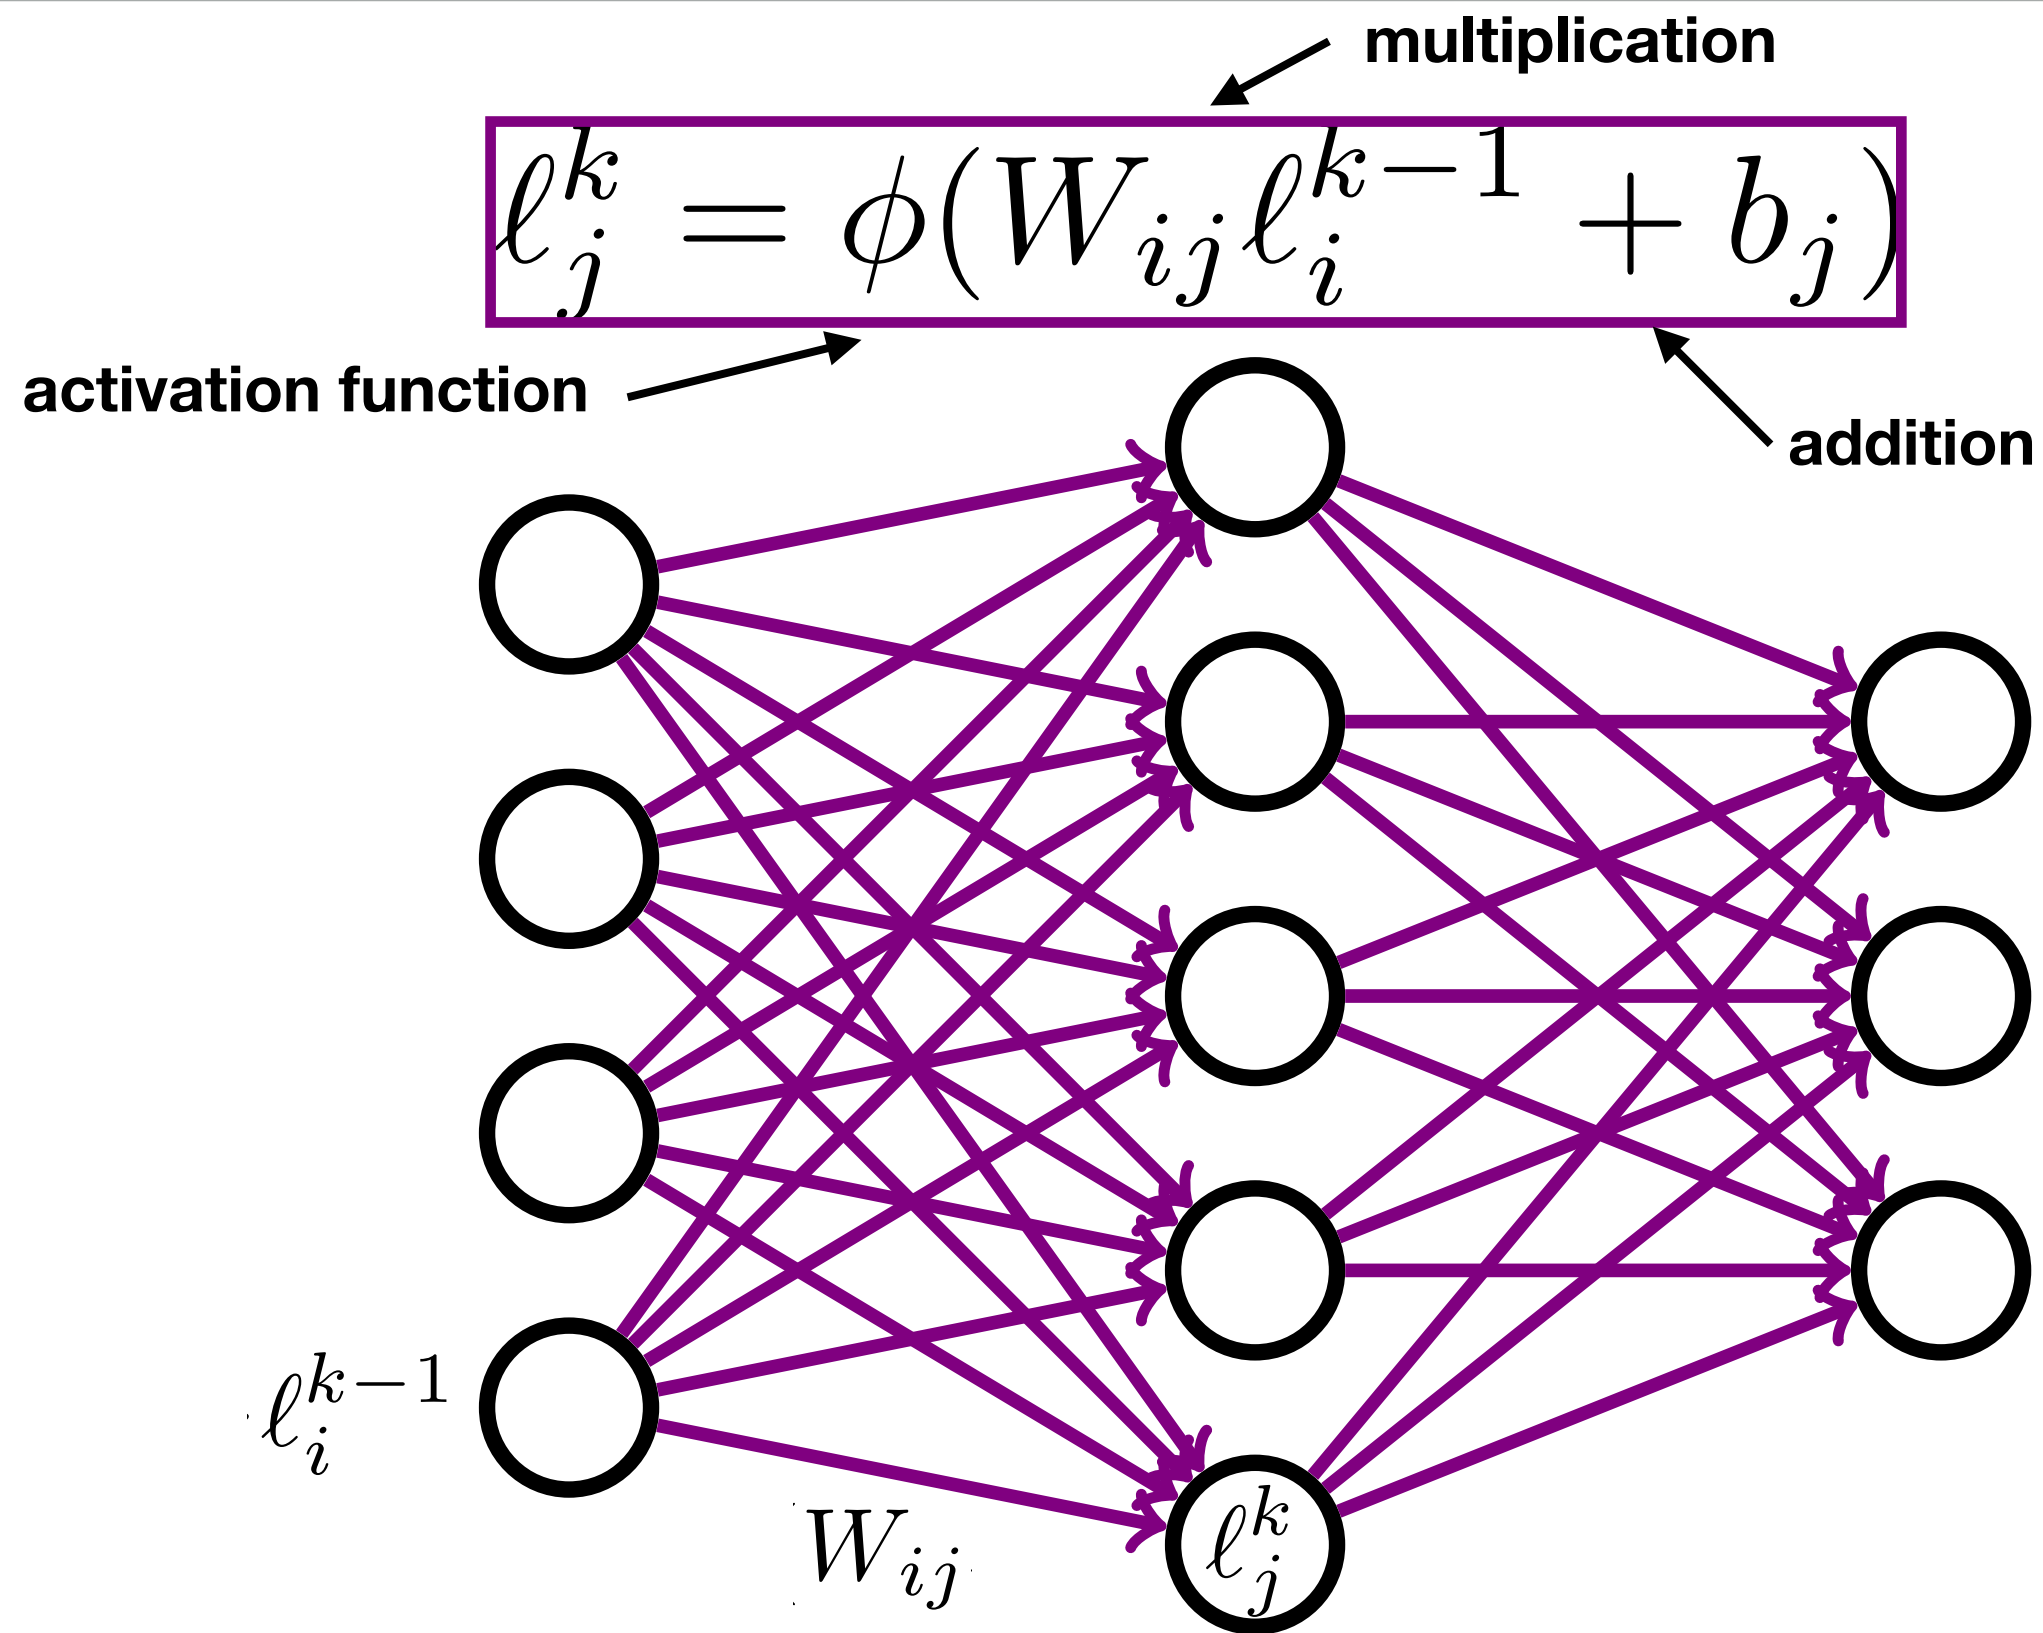
- ▶ Level-1 Trigger:
40 MHz → 750 kHz
- ▶ Reconstruct and filter
2% of events in $\sim 12 \mu\text{s}$
- ▶ Latency necessitates all
FPGA design



- ▶ **Pros:** reprogrammable,
high throughput,
massively parallel, & low
power
- ▶ **Con:** requires domain
knowledge to program

$$l_j^k = \phi(W_{ij}l_i^{k-1} + b_j)$$






$$l_j^k = \phi(W_{ij}l_i^{k-1} + b_j)$$


multiplication

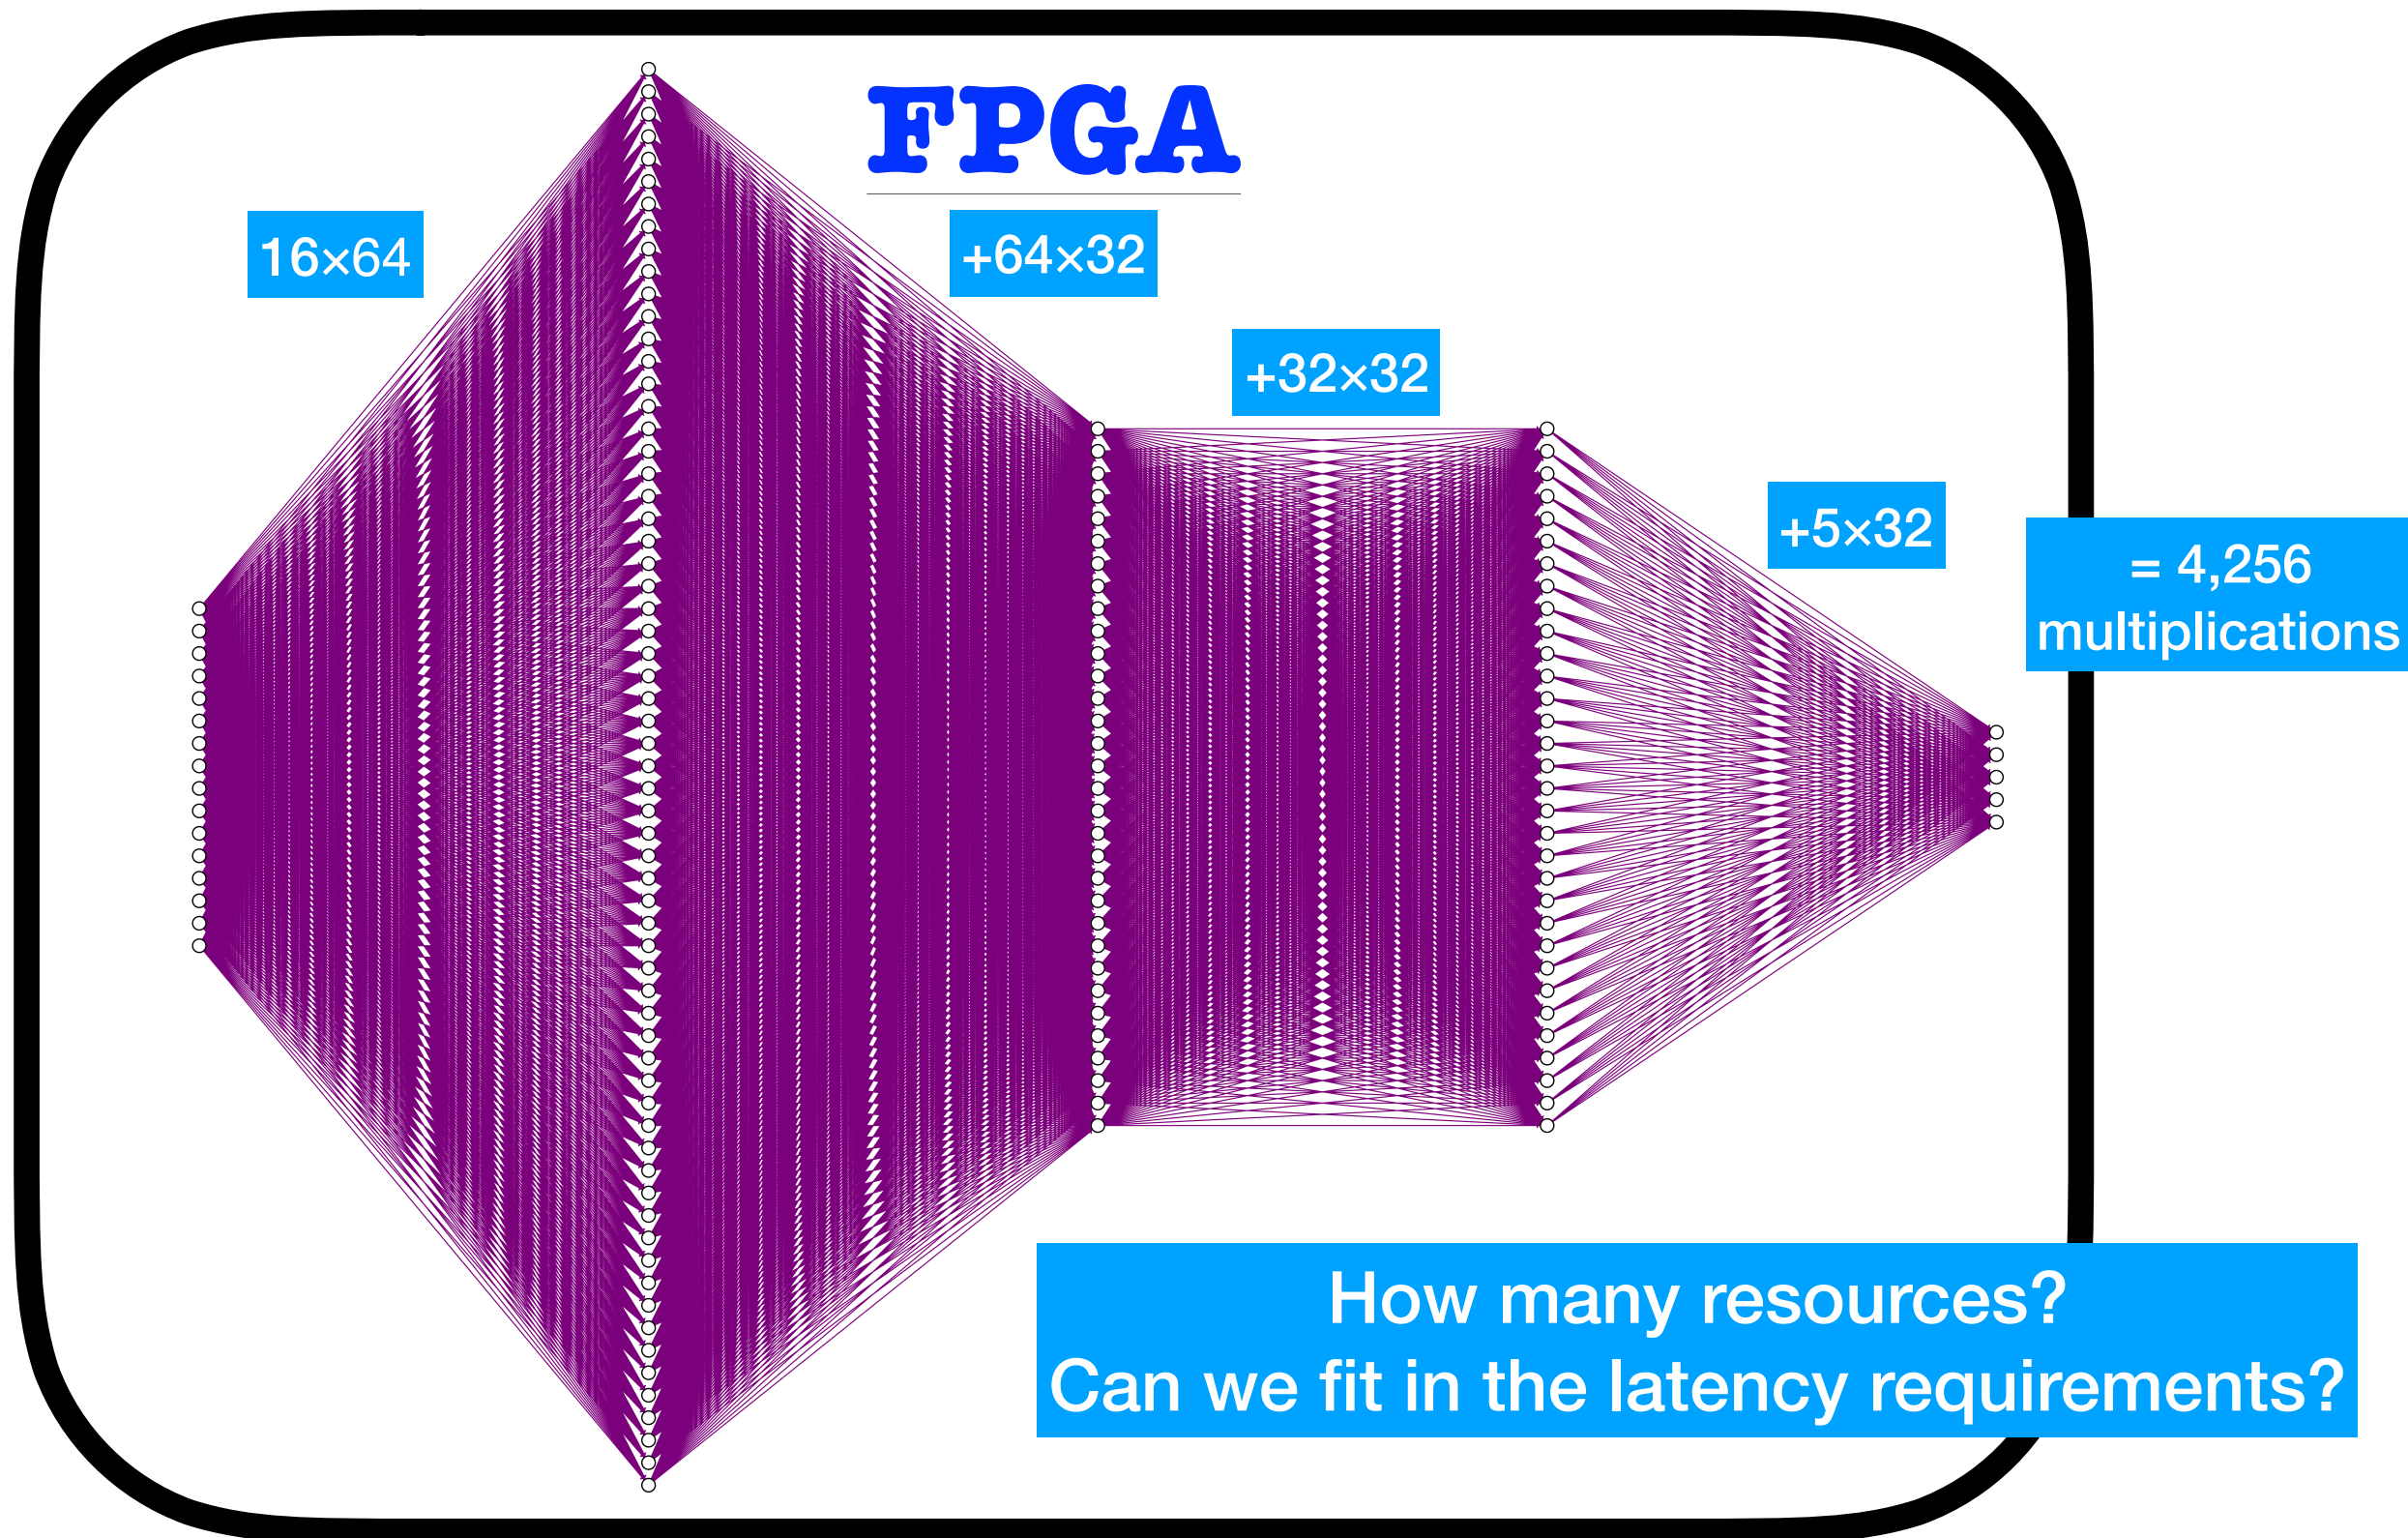
activation function

addition




Maps nicely onto FPGA resources: high IO, DSPs, LUTs, etc.






1. Compression

- ▶ Maintain high performance while removing redundant synapses and neurons



1. Compression

- ▶ Maintain high performance while removing redundant synapses and neurons




2. Quantization

- ▶ Reduce precision from 32-bit floating point to 20-bit, 8-bit, ...

1. Compression

- ▶ Maintain high performance while removing redundant synapses and neurons




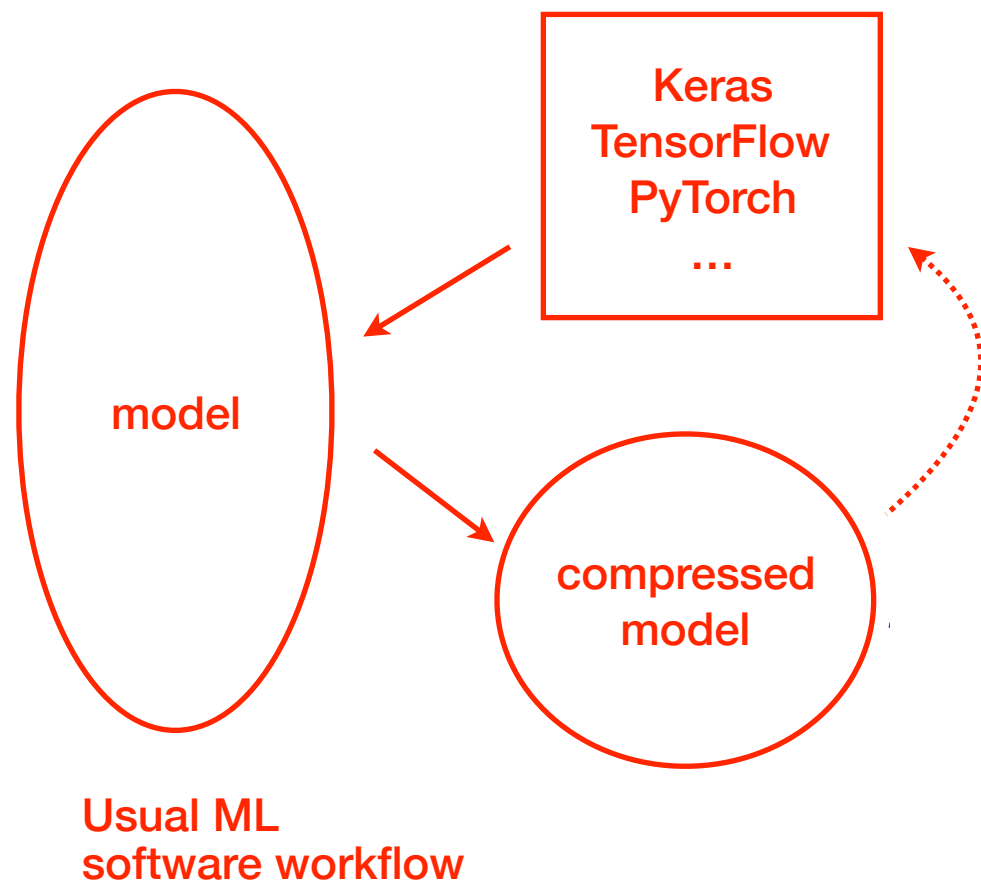
2. Quantization

- ▶ Reduce precision from 32-bit floating point to 20-bit, 8-bit, ...

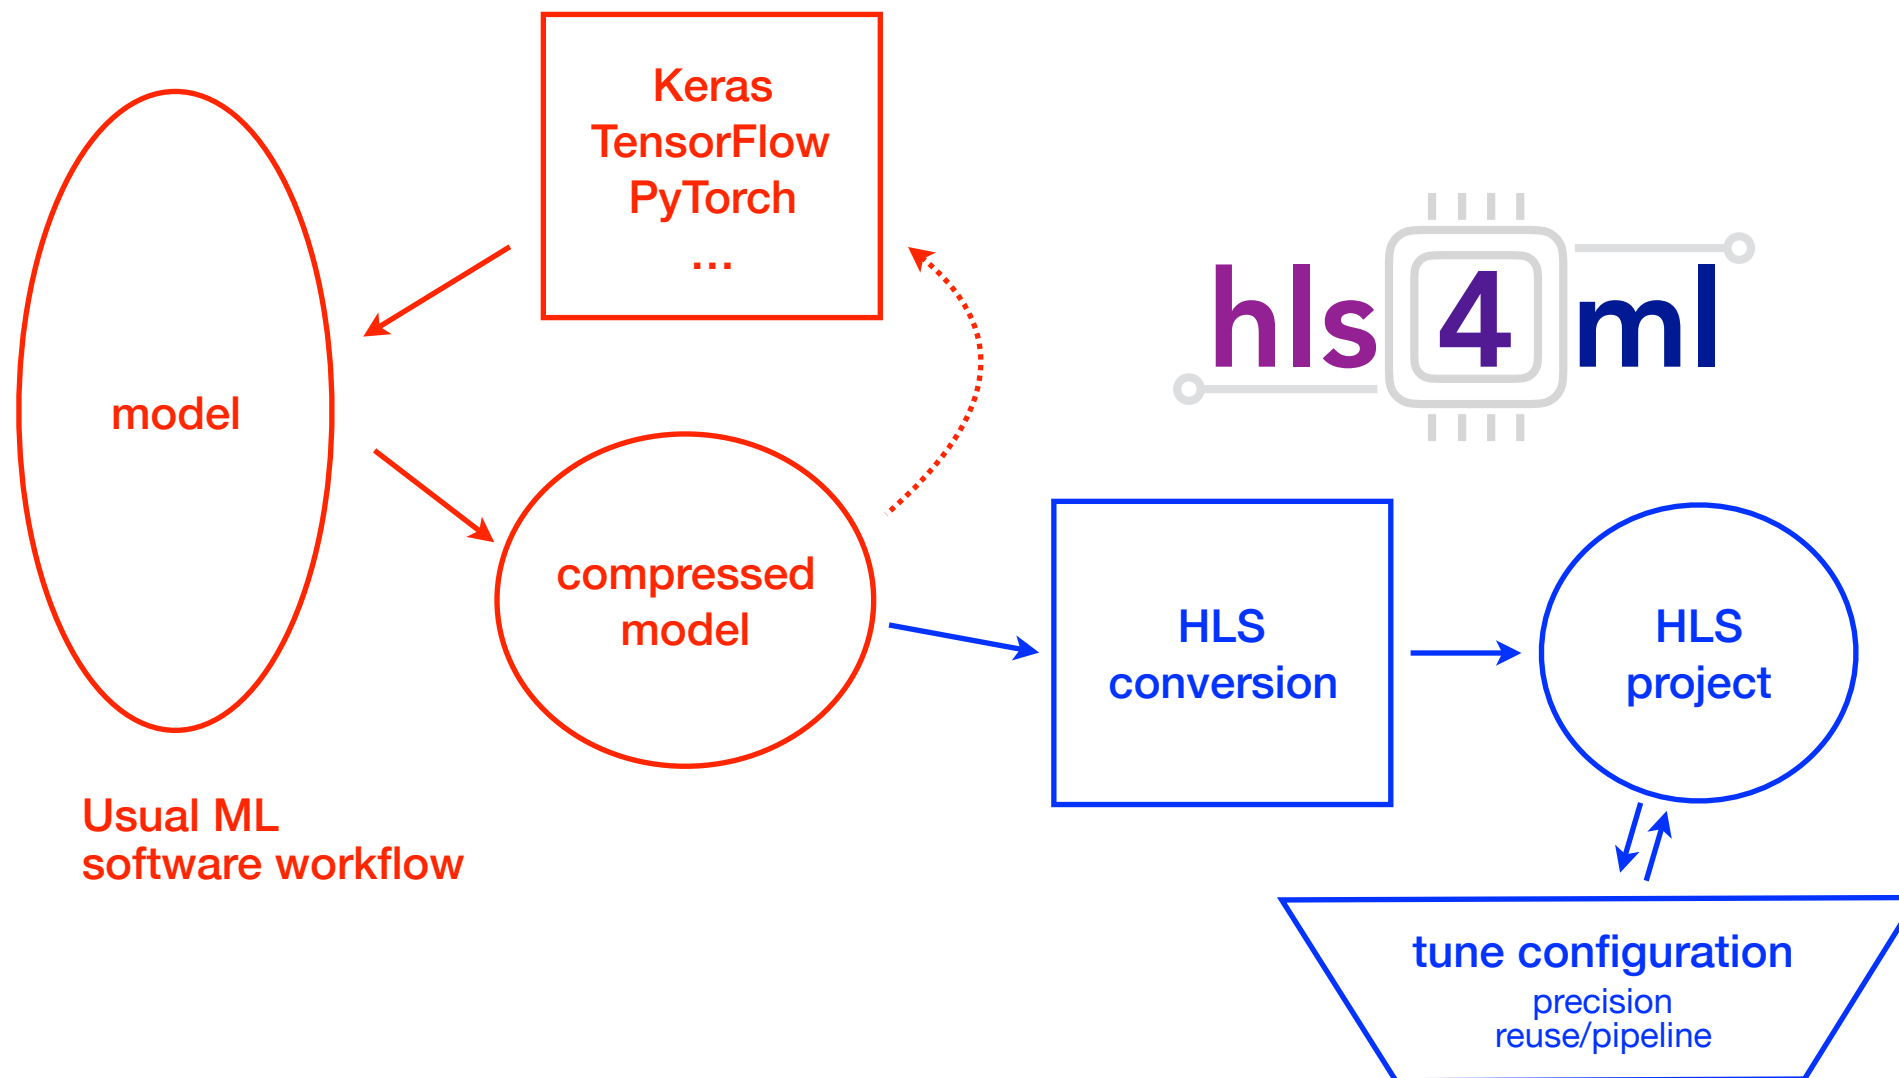
3. Parallelization/Reuse

- ▶ Balance parallelization (how fast) with resources needed (how costly)

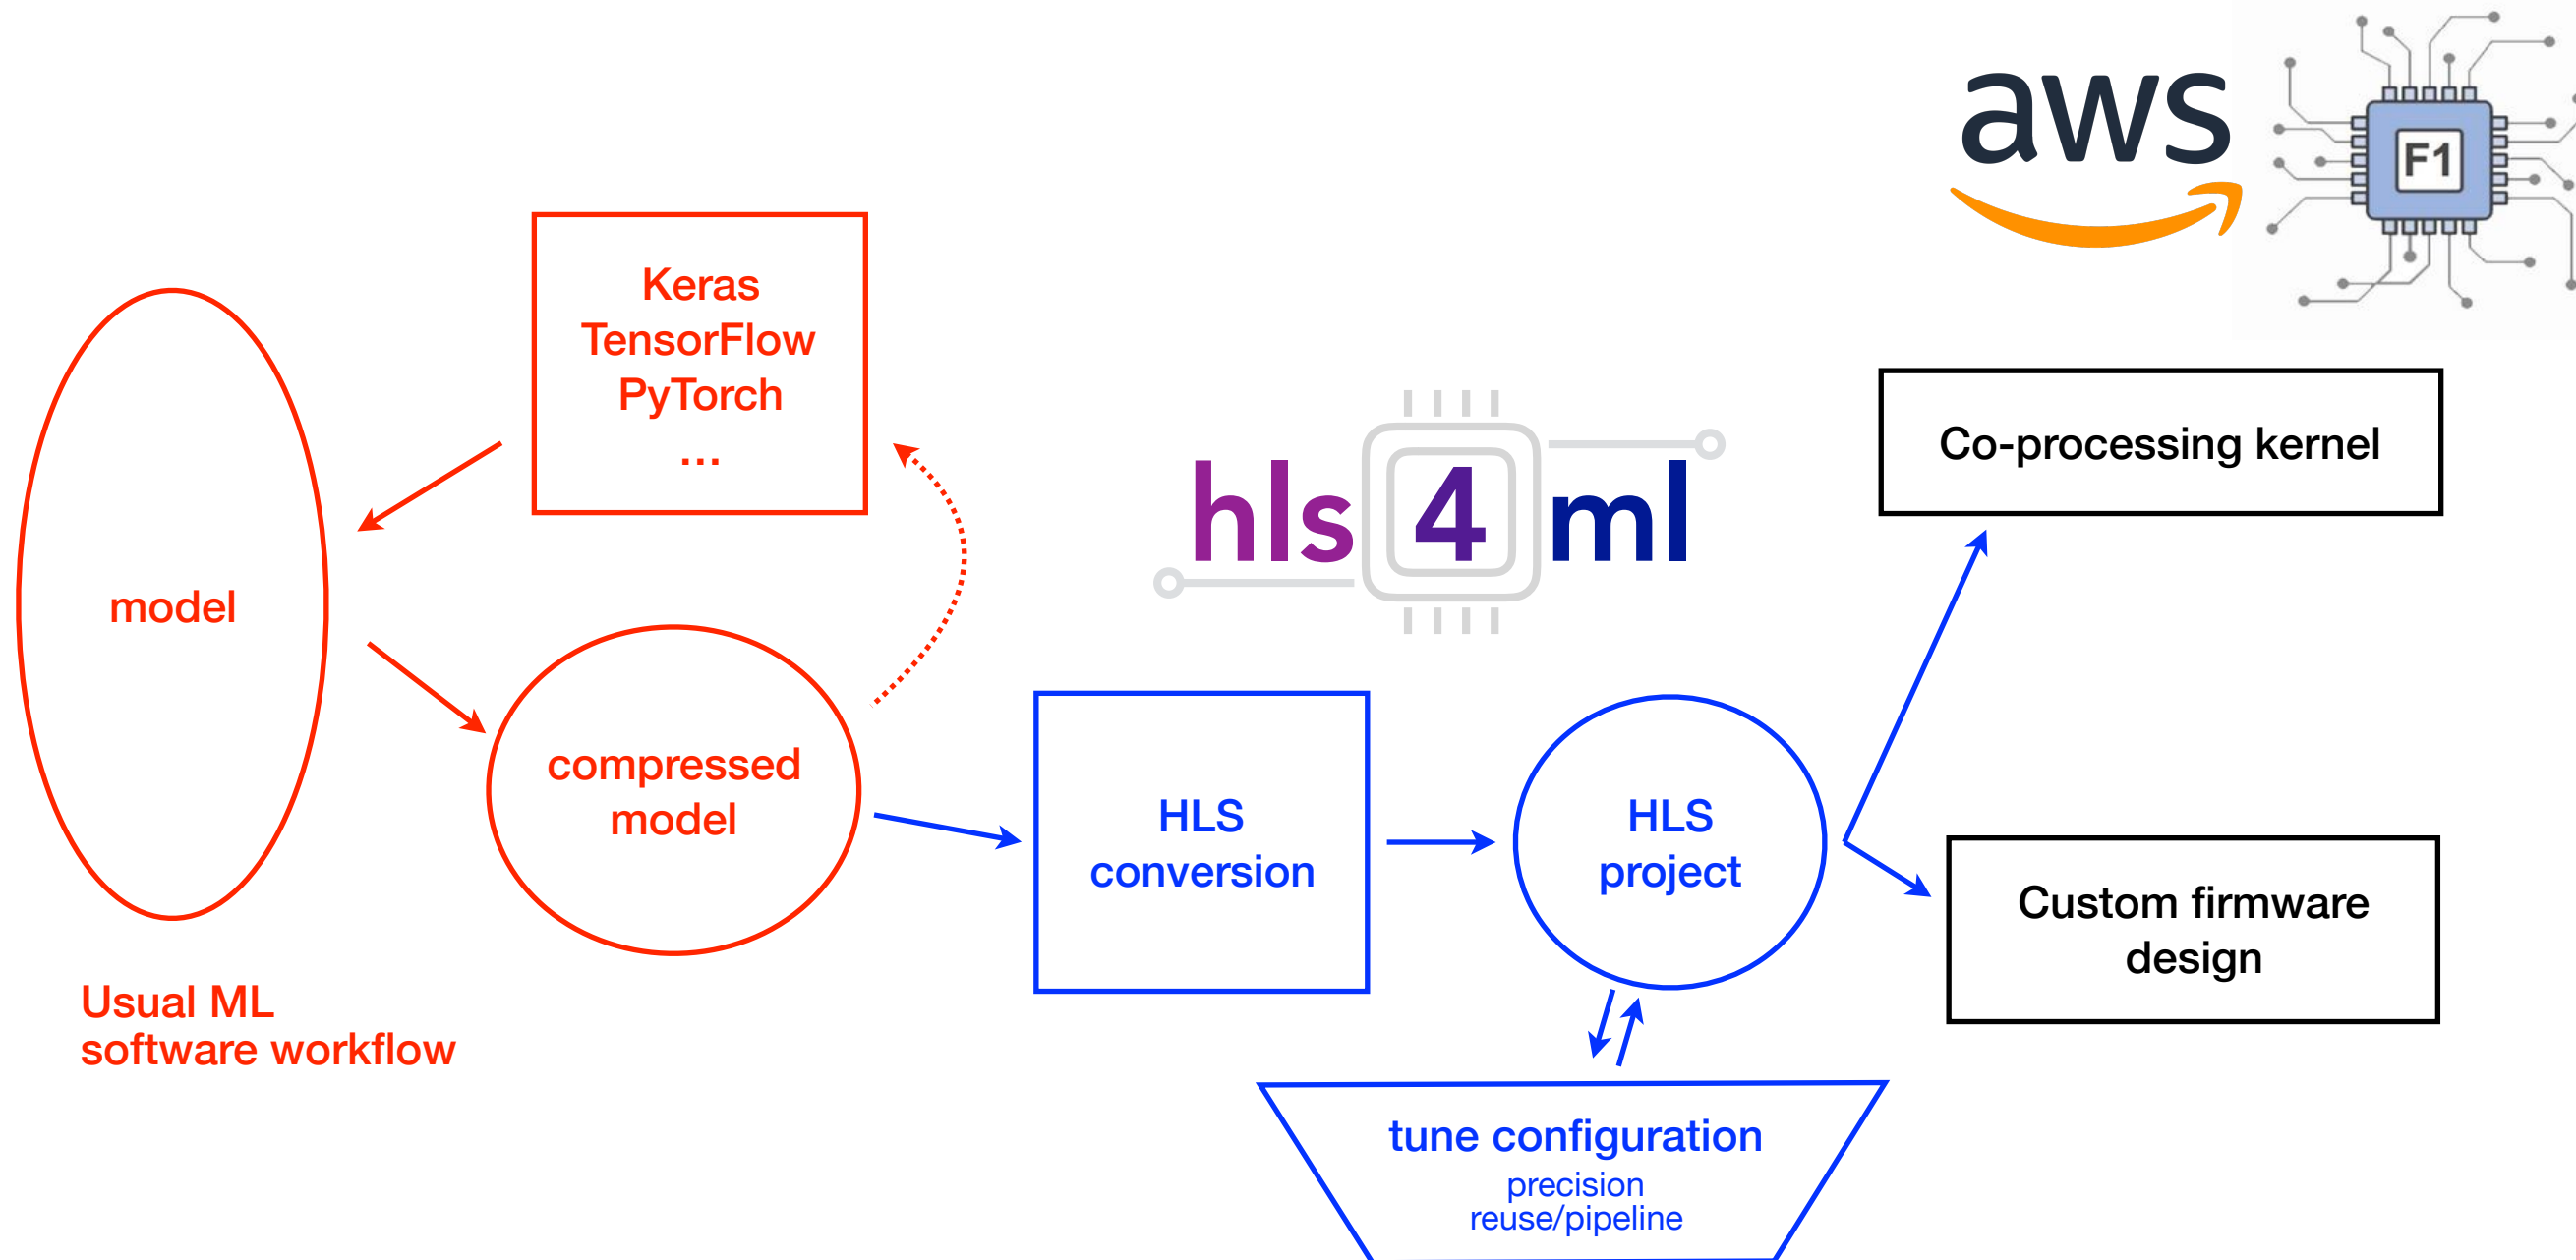
- ▶ [hls4ml](#) for physicists or ML experts to translate **ML algorithms** into **FPGA firmware**



- ▶ [hls4ml](#) for physicists or ML experts to translate **ML algorithms** into **FPGA firmware**



- ▶ [hls4ml](#) for physicists or ML experts to translate **ML algorithms** into **FPGA firmware**



Translation



```
hls4ml convert -c keras-config.yml
```

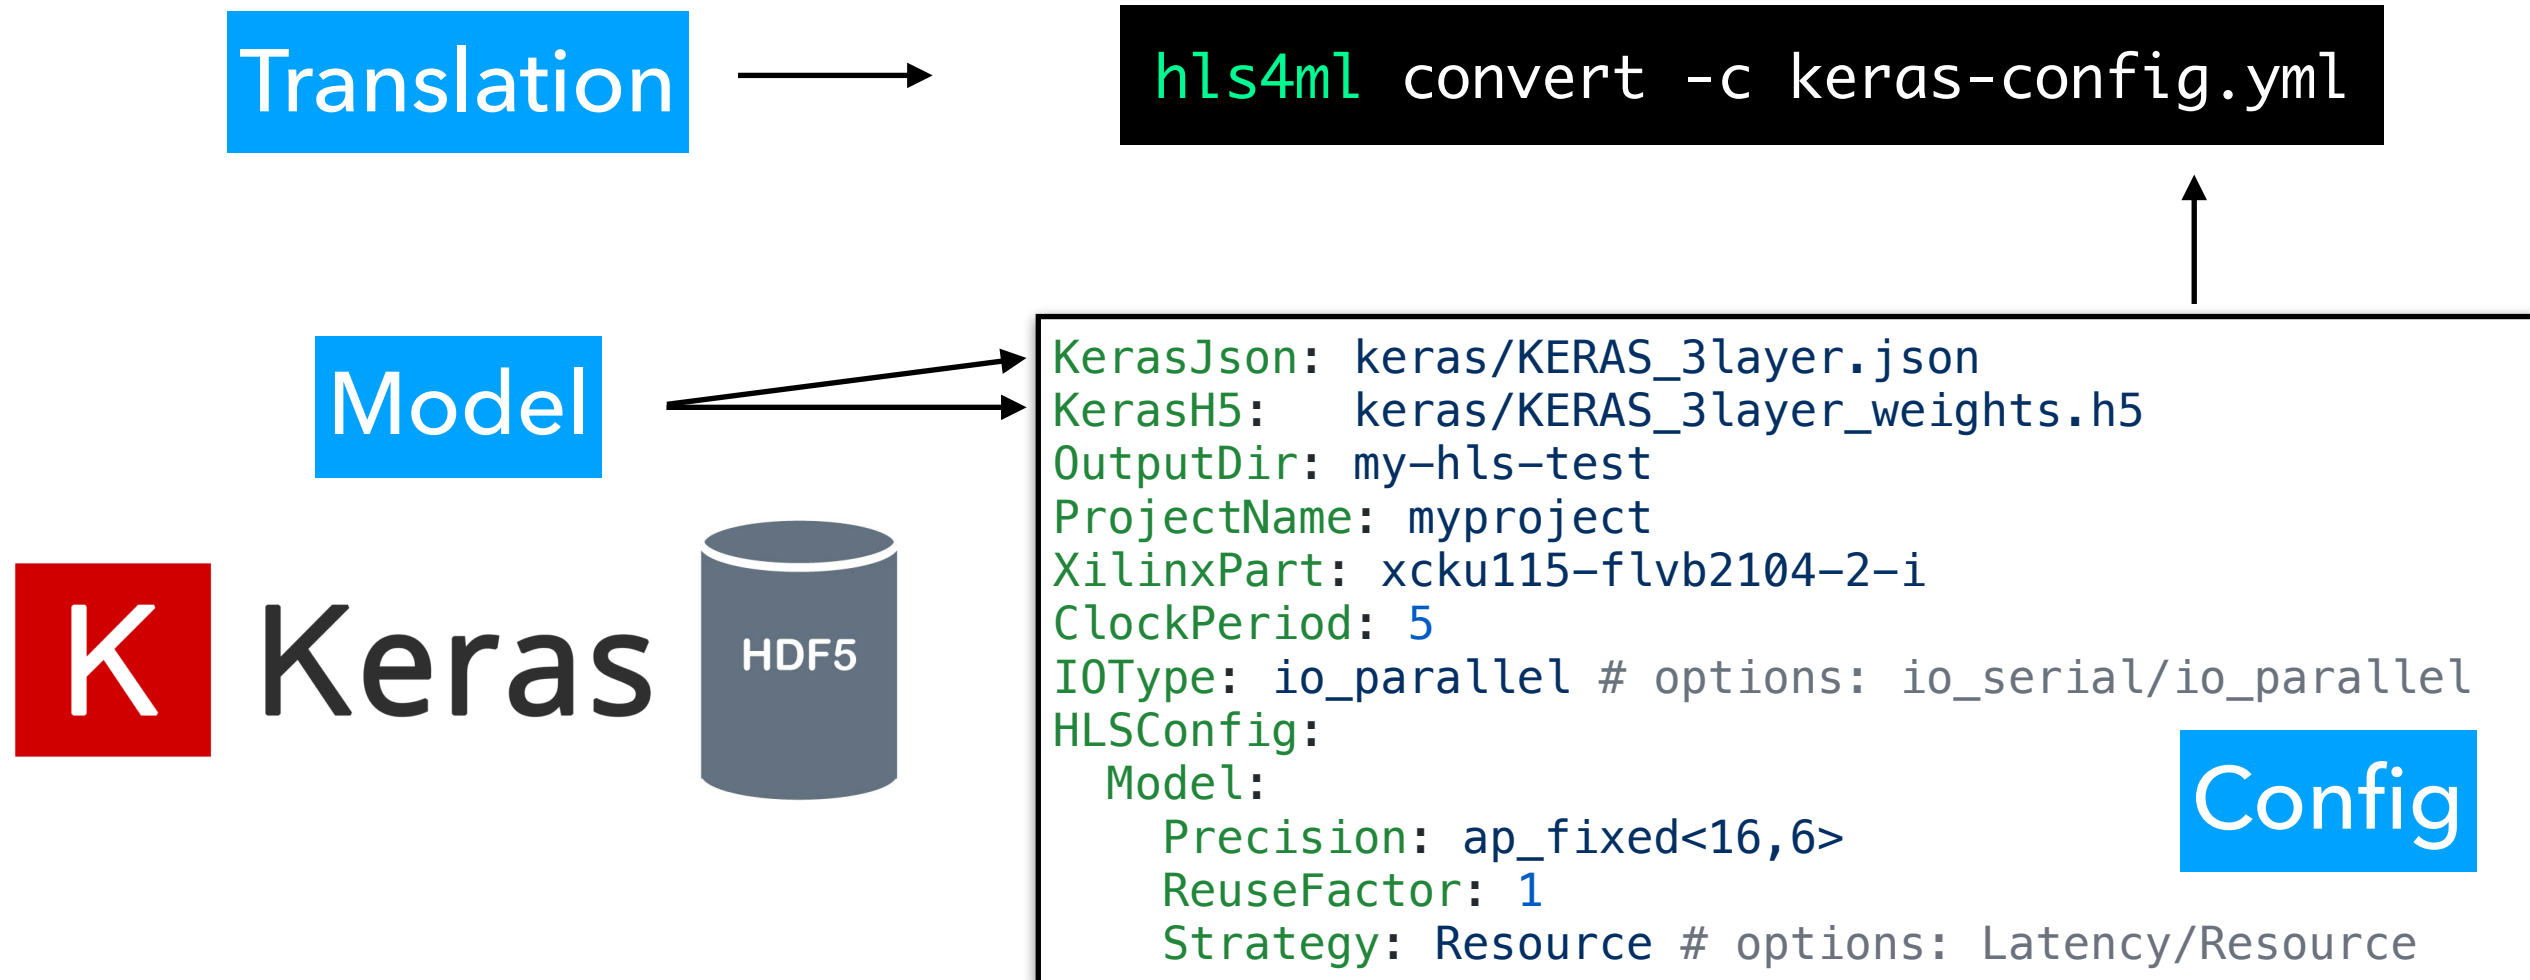
Translation



```
hls4ml convert -c keras-config.yml
```



```
KerasJson: keras/KERAS_3layer.json  
KerasH5:   keras/KERAS_3layer_weights.h5  
OutputDir: my-hls-test  
ProjectName: myproject  
XilinxPart: xcku115-flvb2104-2-i  
ClockPeriod: 5  
IOType: io_parallel # options: io_serial/io_parallel  
HLSConfig:  
  Model:  
    Precision: ap_fixed<16,6>  
    ReuseFactor: 1  
    Strategy: Resource # options: Latency/Resource
```

Translation



```
hls4ml convert -c keras-config.yml
```

Model



Keras



HDF5

```
KerasJson: keras/KERAS_3layer.json
KerasH5:   keras/KERAS_3layer_weights.h5
OutputDir: my-hls-test
ProjectName: myproject
XilinxPart: xcku115-flvb2104-2-i
ClockPeriod: 5
IOType: io_parallel # options: io_serial/io_parallel
HLSConfig:
  Model:
    Precision: ap_fixed<16,6>
    ReuseFactor: 1
    Strategy: Resource # options: Latency/Resource
```

Config

- ▶ **IOType:** parallel or serial
- ▶ **ReuseFactor:** how much to parallelize
- ▶ **Precision:** inputs, weights, biases
- ▶ **Strategy:**
 - ▶ Resource for large NN
 - ▶ Latency for small NN (fully pipelined)

Translation



```
hls4ml convert -c keras-config.yml
```

Model



Keras



```
KerasJson: keras/KERAS_3layer.json
KerasH5:   keras/KERAS_3layer_weights.h5
OutputDir: my-hls-test
ProjectName: myproject
XilinxPart: xcku115-flvb2104-2-i
ClockPeriod: 5
IOType: io_parallel # options: io_serial/io_parallel
HLSConfig:
  Model:
    Precision: ap_fixed<16,6>
    ReuseFactor: 1
    Strategy: Resource # options: Latency/Resource
```


Config

- ▶ **IOType:** parallel or serial
- ▶ **ReuseFactor:** how much to parallelize
- ▶ **Precision:** inputs, weights, biases
- ▶ **Strategy:**
 - ▶ Resource for large NN
 - ▶ Latency for small NN (fully pipelined)

Build HLS project





```
hls4ml build -p my-hls-test -a
```



- ▶ Train with **L₁ regularization** (down-weights unimportant synapses)



$$L_{\lambda}(\mathbf{w}) = L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1 \quad \|\mathbf{w}\|_1 = \sum_i |w_i|$$



- ▶ Train with **L₁ regularization** (down-weights unimportant synapses)

$$L_{\lambda}(\mathbf{w}) = L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1 \quad \|\mathbf{w}\|_1 = \sum_i |w_i|$$

- ▶ Remove **smallest** weights





- ▶ Train with **L₁ regularization** (down-weights unimportant synapses)

$$L_{\lambda}(\mathbf{w}) = L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$

$$\|\mathbf{w}\|_1 = \sum_i |w_i|$$

- ▶ Remove **smallest** weights
- ▶ Iterate

70% REDUCTION OF
WEIGHTS WITH NO
LOSS IN PERF.




- ▶ Train with **L₁ regularization** (down-weights unimportant synapses)


$$L_{\lambda}(\mathbf{w}) = L(\mathbf{w}) + \lambda \|\mathbf{w}\|_1$$

$$\|\mathbf{w}\|_1 = \sum_i |w_i|$$


- ▶ Remove **smallest** weights
- ▶ Iterate

70% REDUCTION OF
WEIGHTS WITH NO
LOSS IN PERF.






ap_fixed<width,integer>
0101.1011101010
integer fractional
width



- ▶ Big reduction in DSPs (multipliers) with compression




- ▶ Big reduction in DSPs (multipliers) with compression
- ▶ Easily fits on 1 FPGA **after compression**


- ▶ Inference of ML algorithms possible in **O(100 ns)** on **1 FPGA** with **hls4ml!**

-
- ▶ Inference of ML algorithms possible in **O(100 ns)** on **1 FPGA** with **hls4ml!**
 - ▶ Applications across CMS, ATLAS, DUNE, and accelerator controls


- ▶ Inference of ML algorithms possible in **$O(100\text{ ns})$** on **1 FPGA** with **hls4ml!**
 - ▶ Applications across CMS, ATLAS, DUNE, and accelerator controls
 - ▶ E.g. muon p_T determination in the CMS endcap with a DNN:
runs in 160 ns on an FPGA and **reduces the fake muon rate** by **up to 80%**




- ▶ Inference of ML algorithms possible in **$O(100\text{ ns})$** on **1 FPGA** with **hls4ml!**
 - ▶ Applications across CMS, ATLAS, DUNE, and accelerator controls
 - ▶ E.g. muon p_T determination in the CMS endcap with a DNN:
runs in 160 ns on an FPGA and **reduces the fake muon rate by up to 80%**

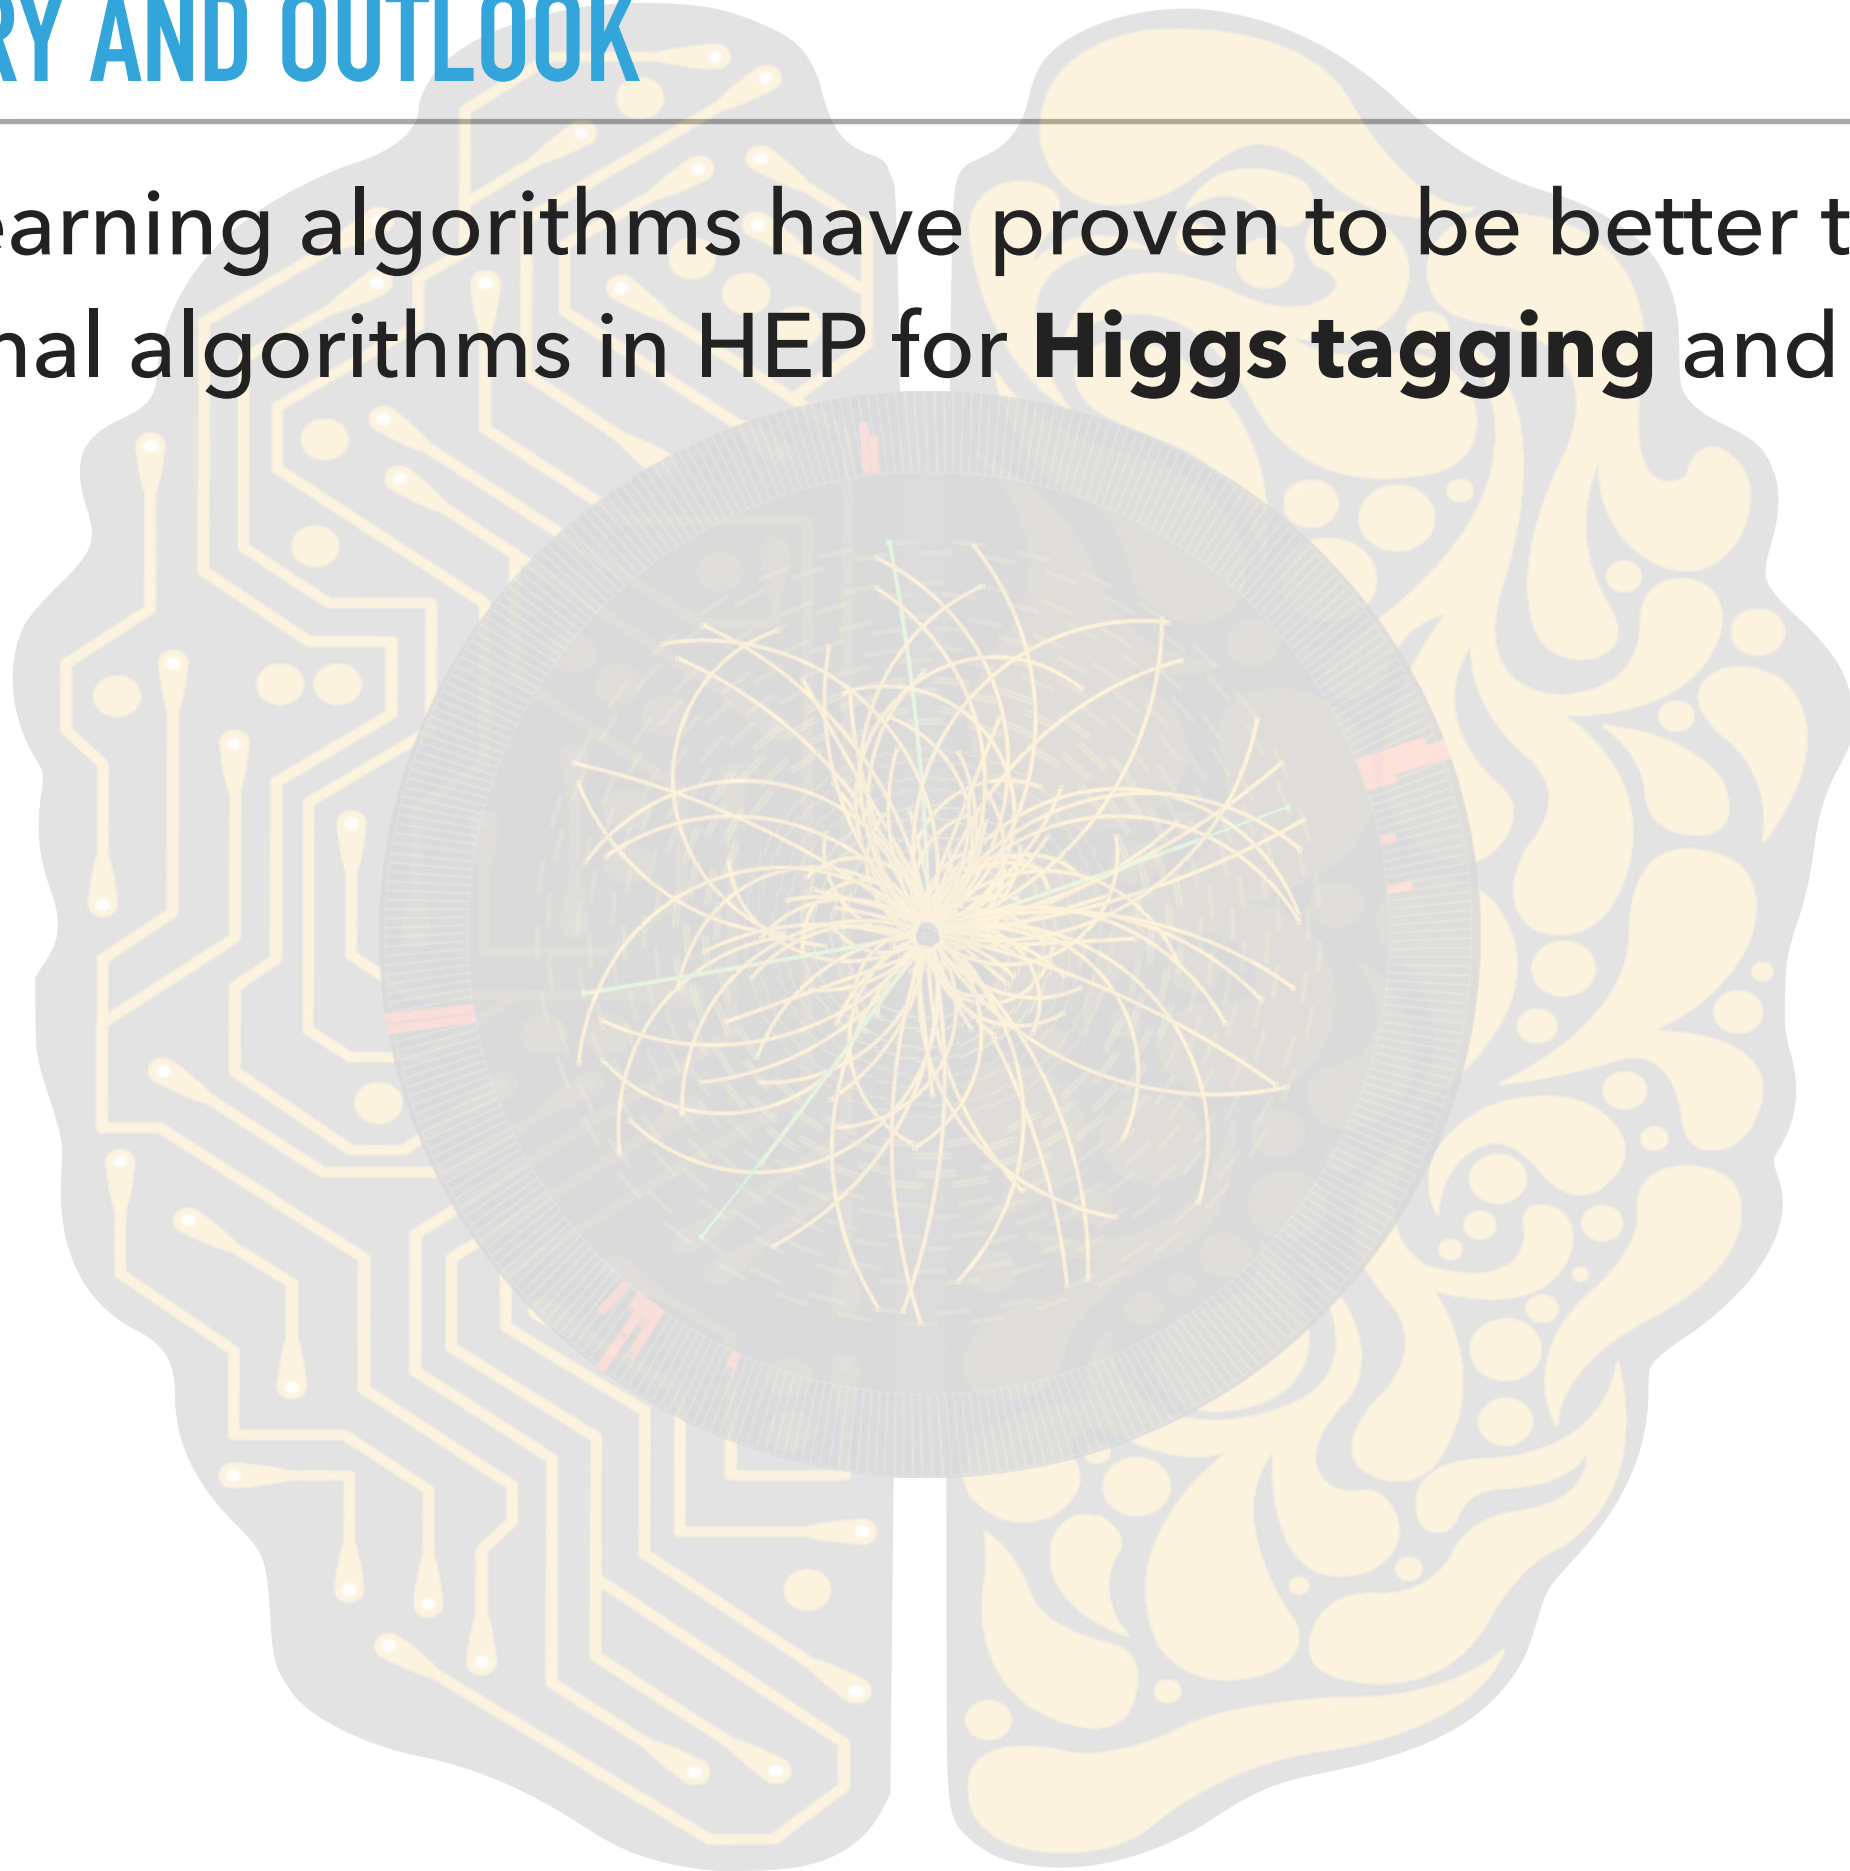


- ▶ Currently supported:
 - ▶ Small and large dense NNs
 - ▶ Binary and ternary NNs
 - ▶ Small 1D/2D CNNs
- ▶ Planned support
 - ▶ Big 1D/2D CNNs
 - ▶ Graph NNs
 - ▶ Other HLS/RTL backends

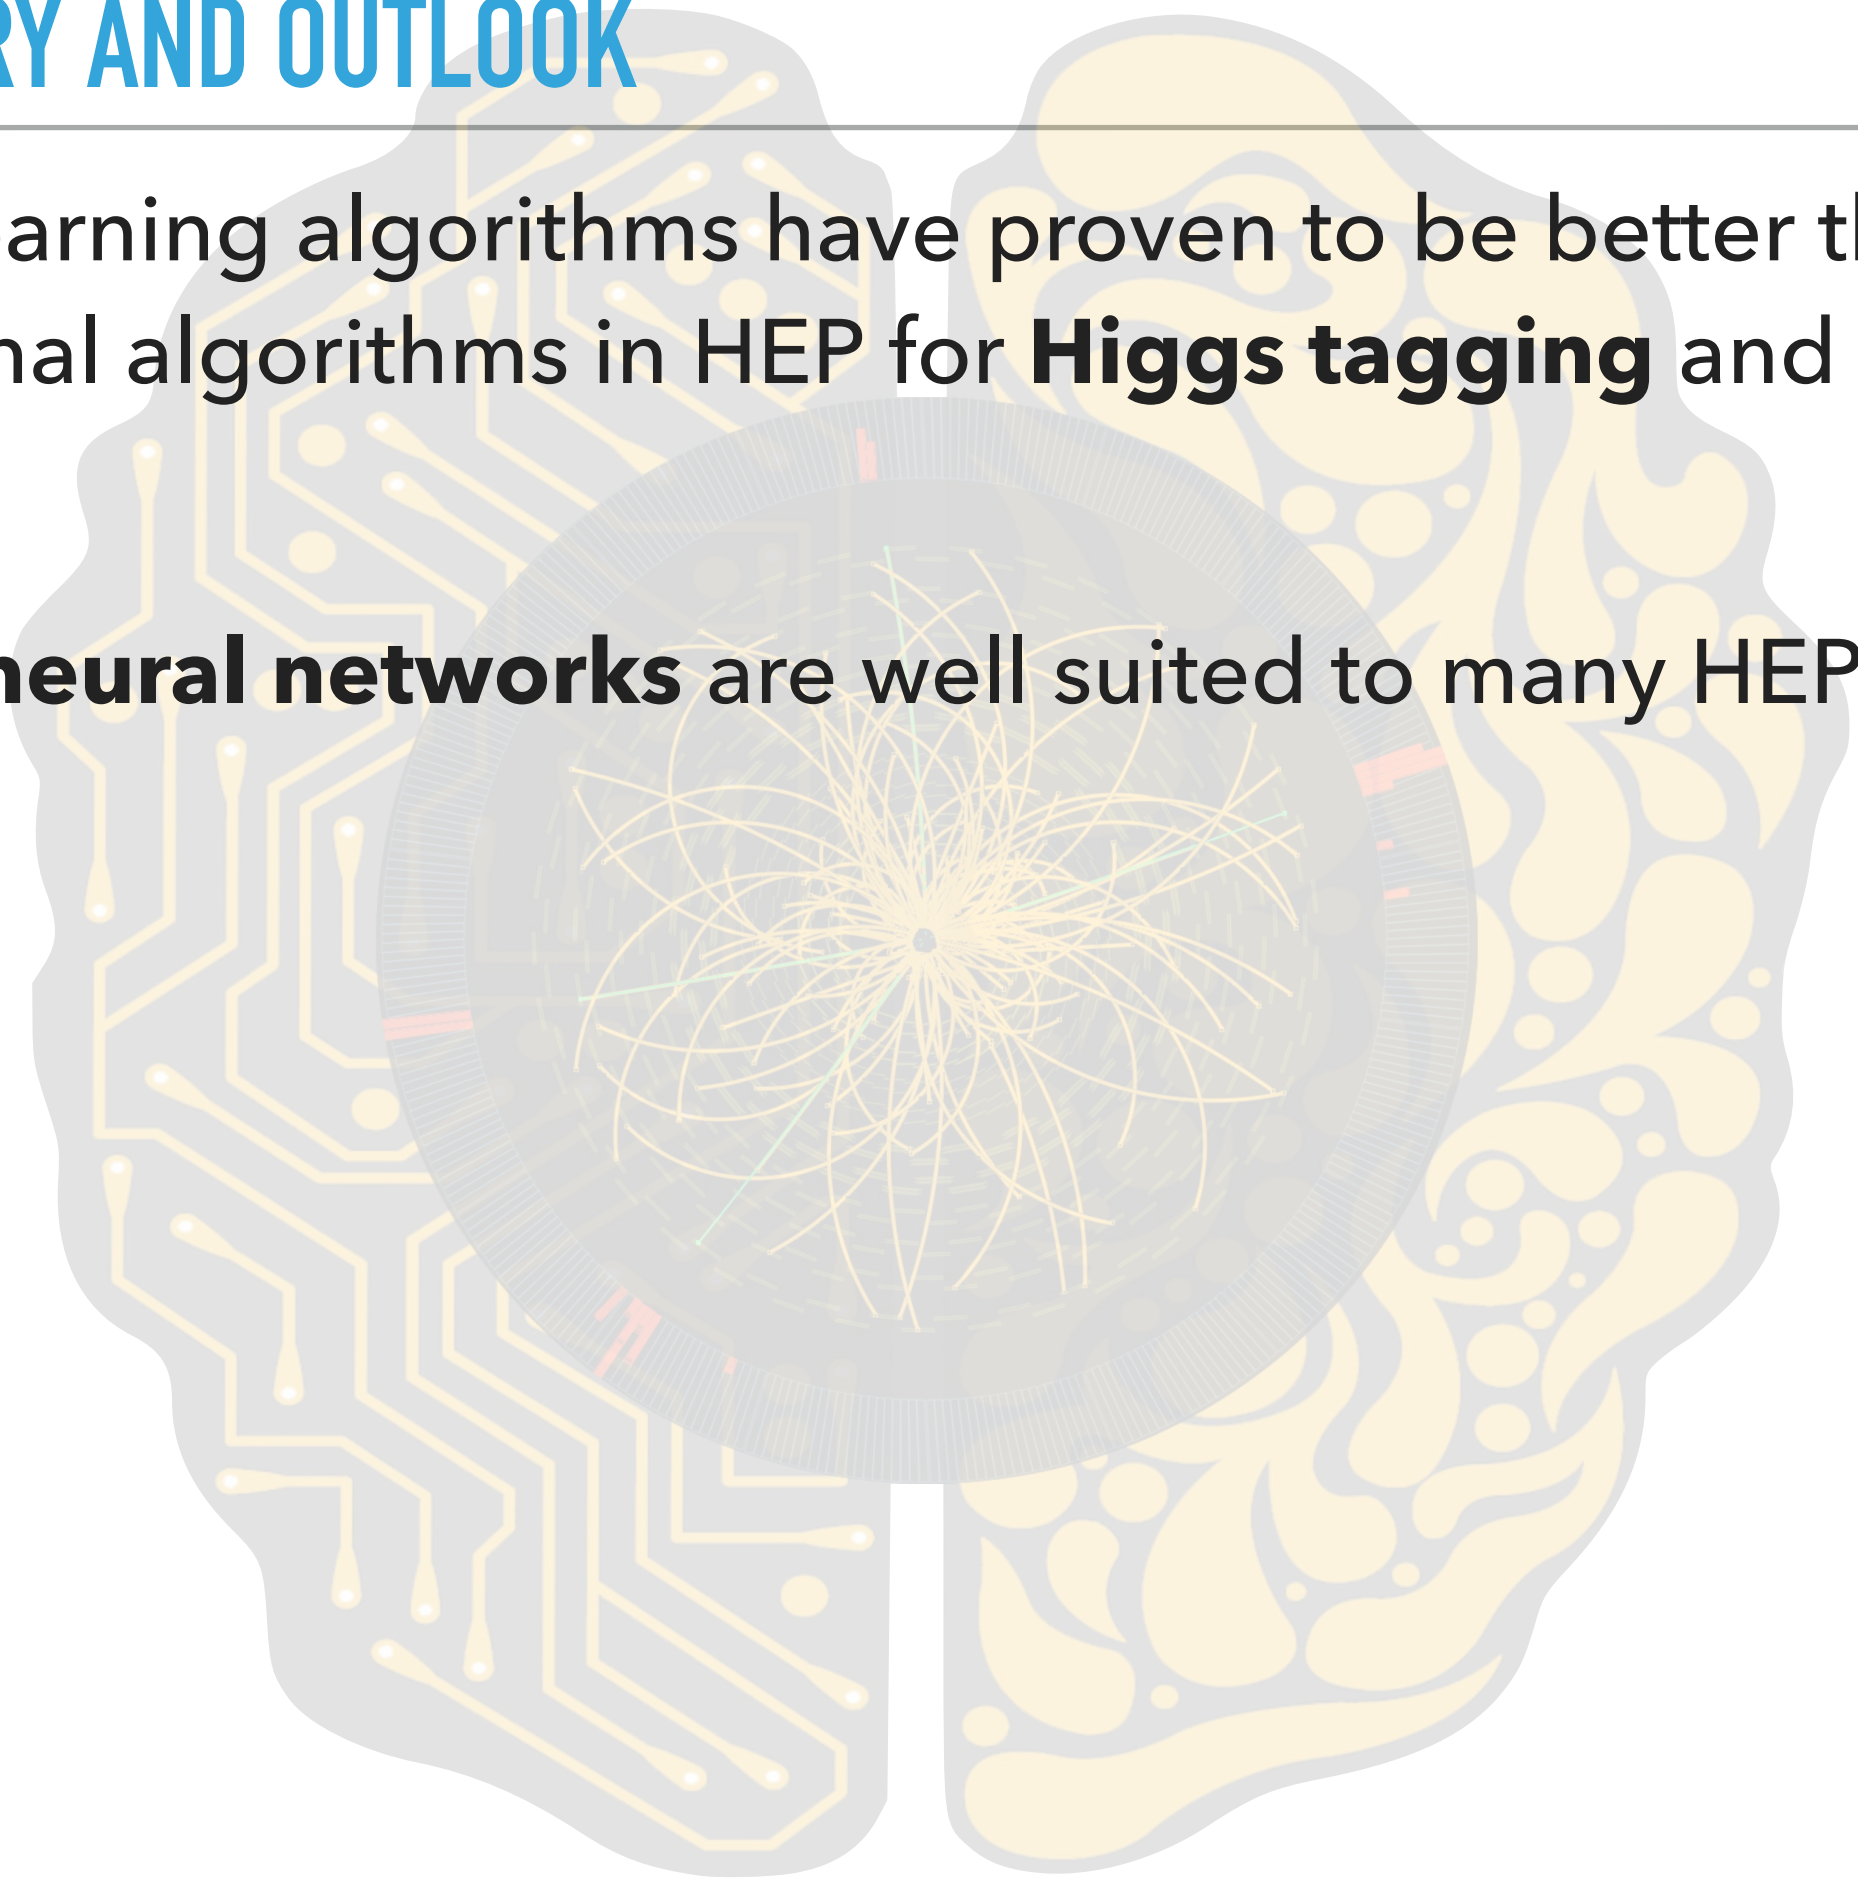


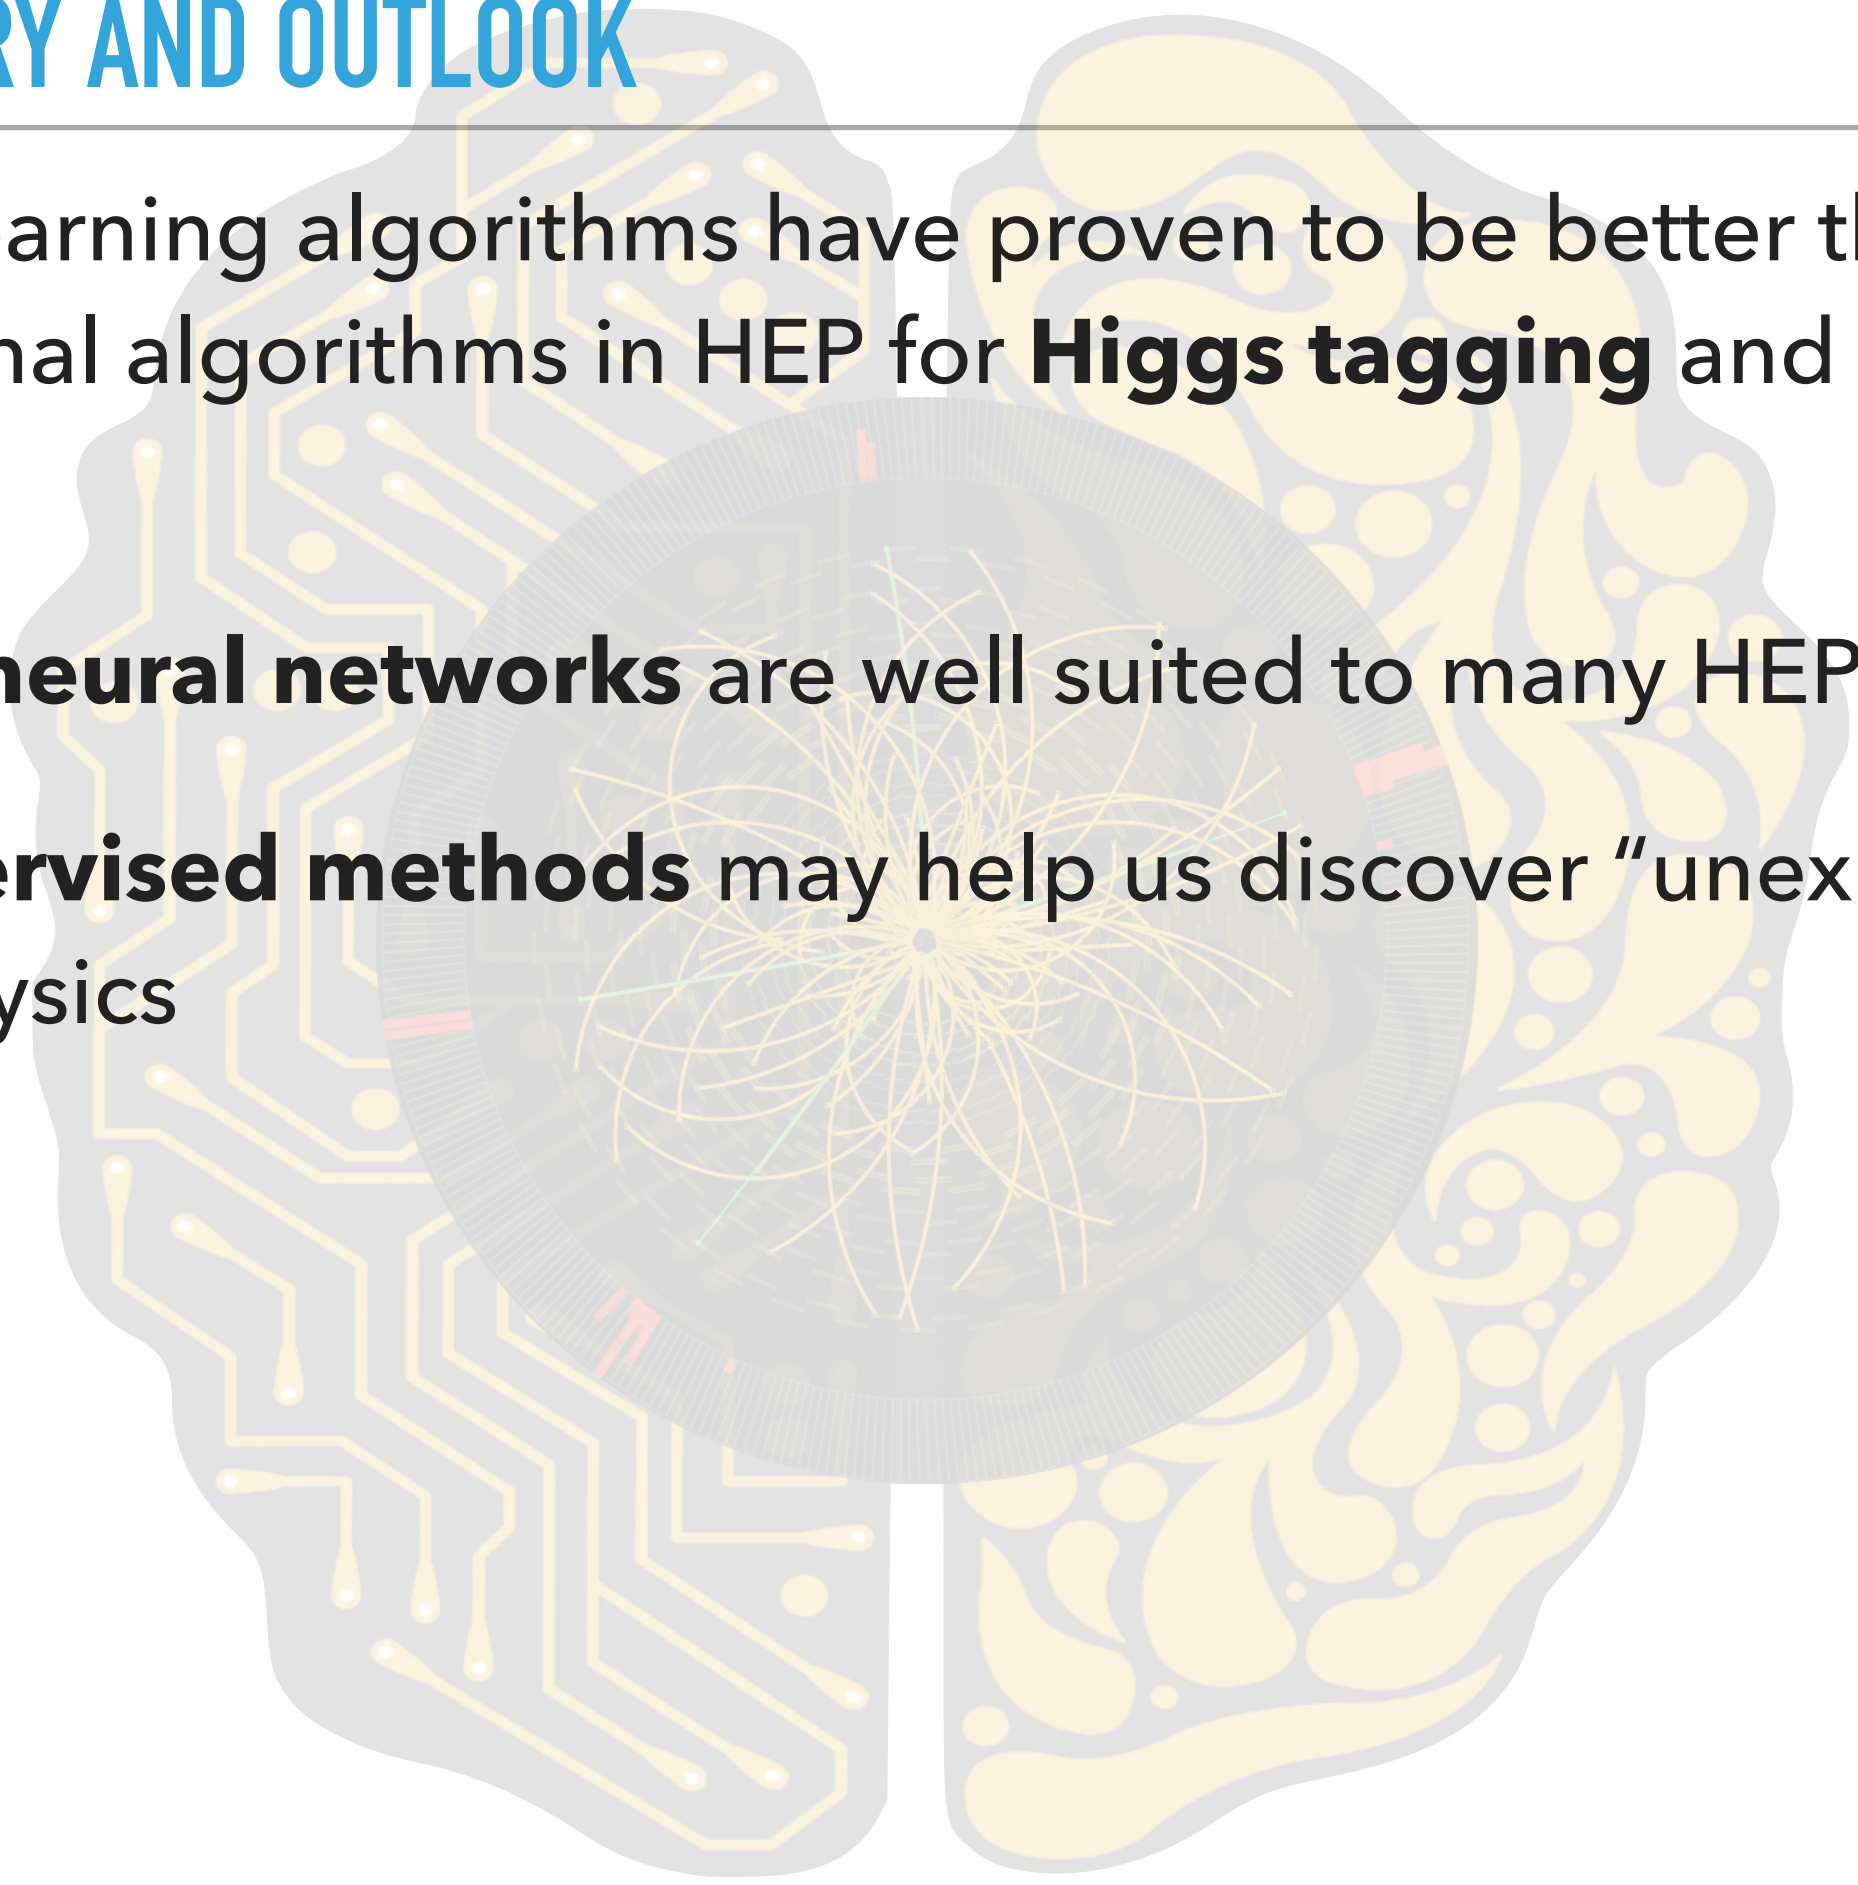


- ▶ Deep learning algorithms have proven to be better than traditional algorithms in HEP for **Higgs tagging** and much more



- ▶ Deep learning algorithms have proven to be better than traditional algorithms in HEP for **Higgs tagging** and much more
- ▶ **Graph neural networks** are well suited to many HEP tasks



- ▶ Deep learning algorithms have proven to be better than traditional algorithms in HEP for **Higgs tagging** and much more
 - ▶ **Graph neural networks** are well suited to many HEP tasks
 - ▶ **Unsupervised methods** may help us discover “unexpected” new physics
- 
- A stylized graphic of a human brain in shades of grey and yellow. The left hemisphere is overlaid with a yellow circuit board pattern. The right hemisphere is overlaid with a yellow particle detector pattern, showing a central vertex with many tracks radiating outwards, similar to a bubble chamber or silicon detector readout.

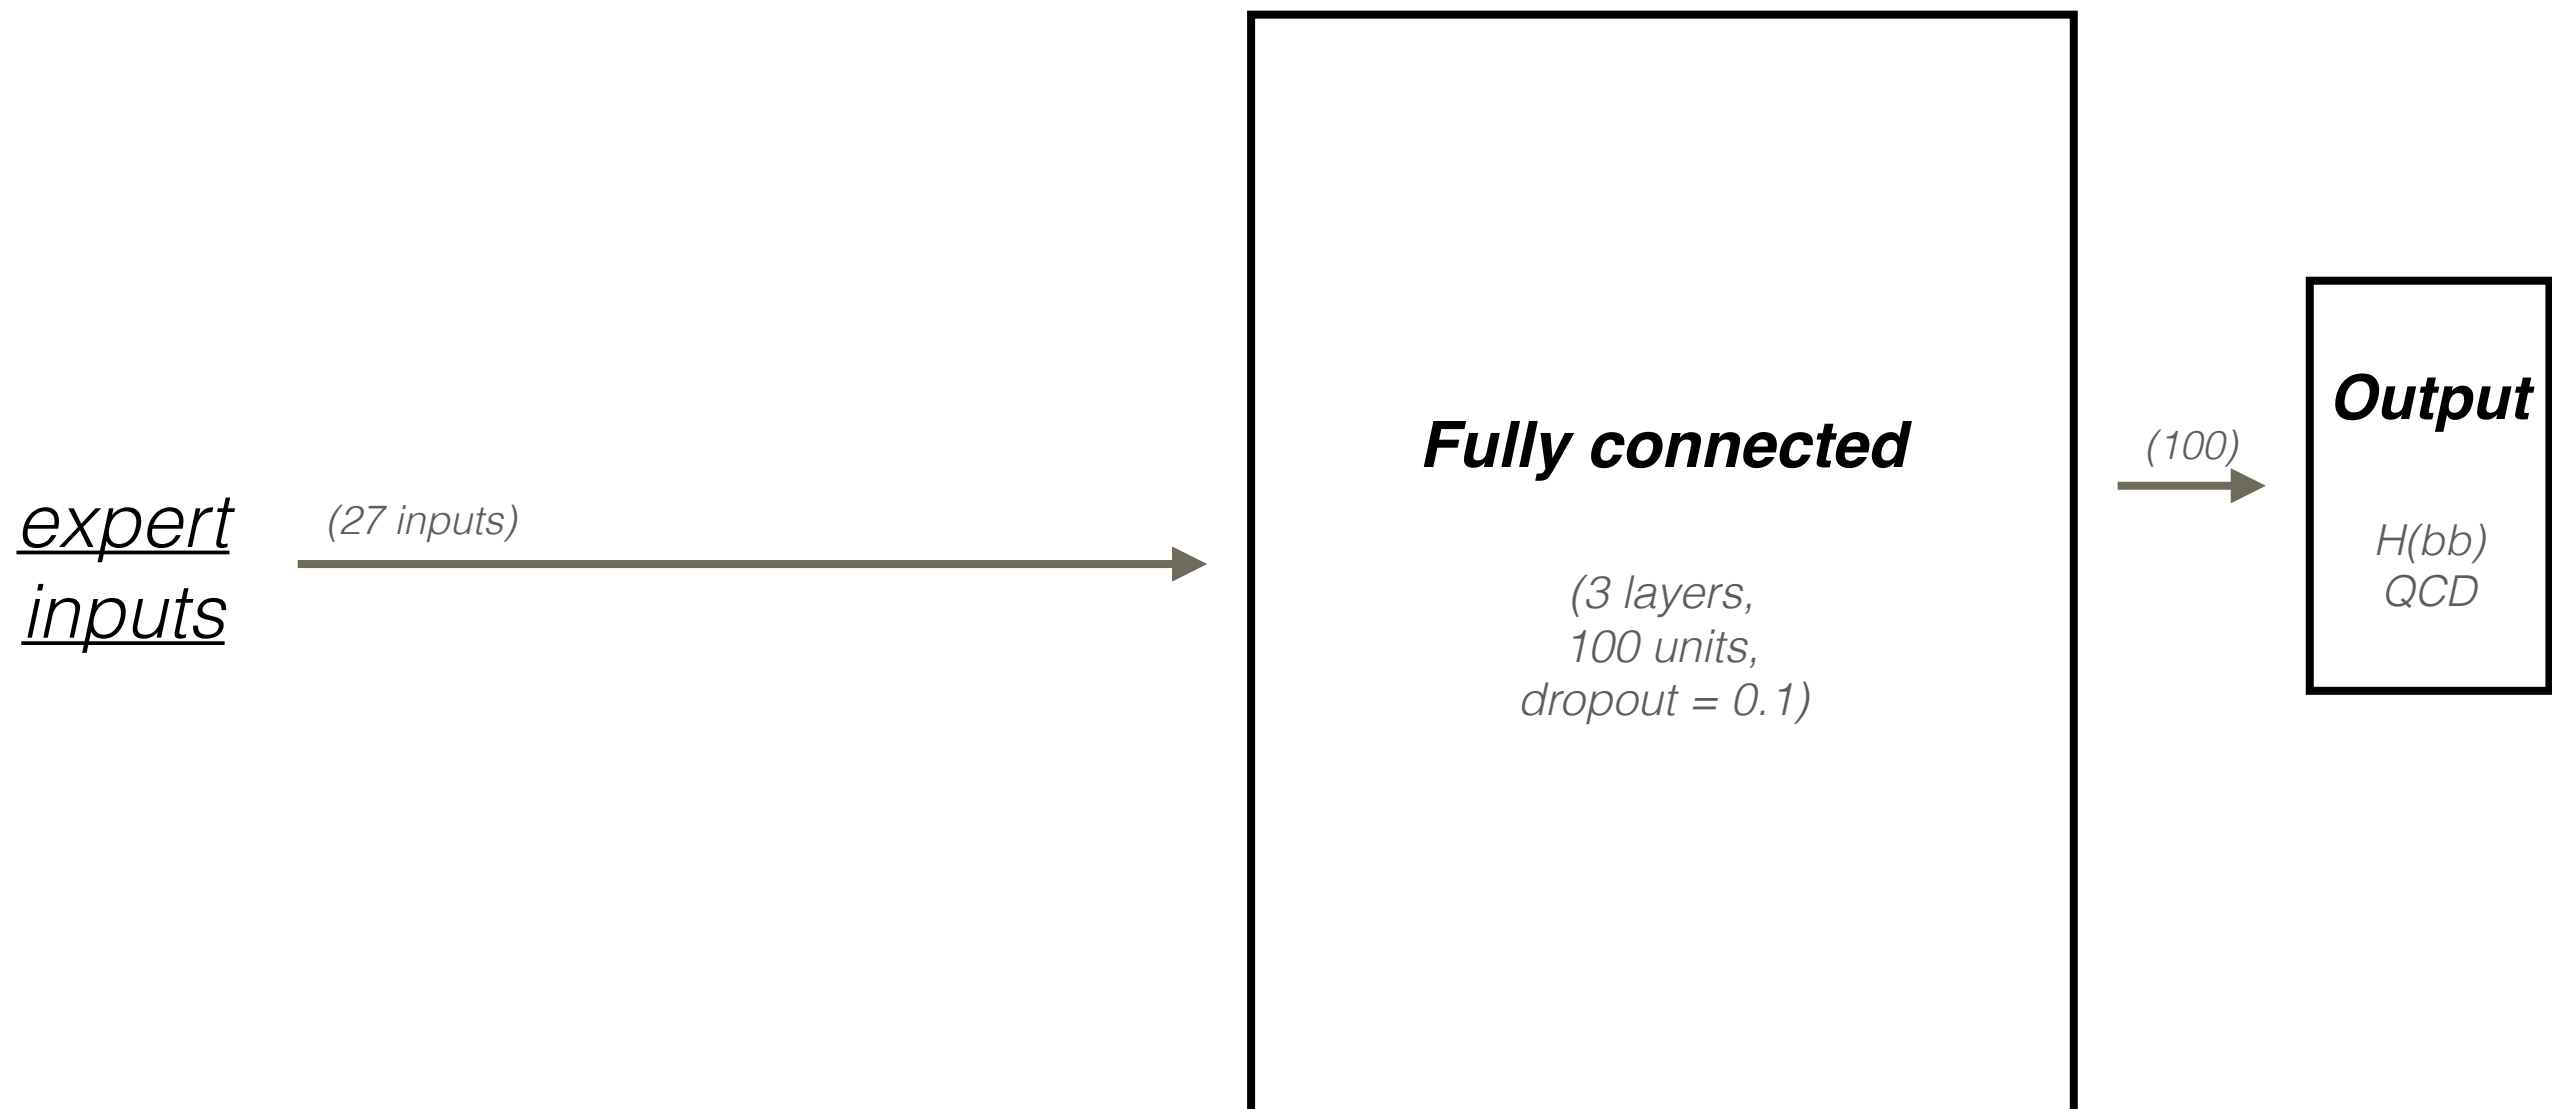
- ▶ Deep learning algorithms have proven to be better than traditional algorithms in HEP for **Higgs tagging** and much more
- ▶ **Graph neural networks** are well suited to many HEP tasks
- ▶ **Unsupervised methods** may help us discover “unexpected” new physics
- ▶ With FPGAs, ML methods can be implemented **quickly** and **efficiently**

JAVIER DUARTE
NOVEMBER 12, 2019
UNIVERSITY OF KANSAS




BACKUP

- ▶ First, we can change the architecture from a BDT to a neural network



Same inputs, simple neural network → same performance



expert
inputs

(27 in.)

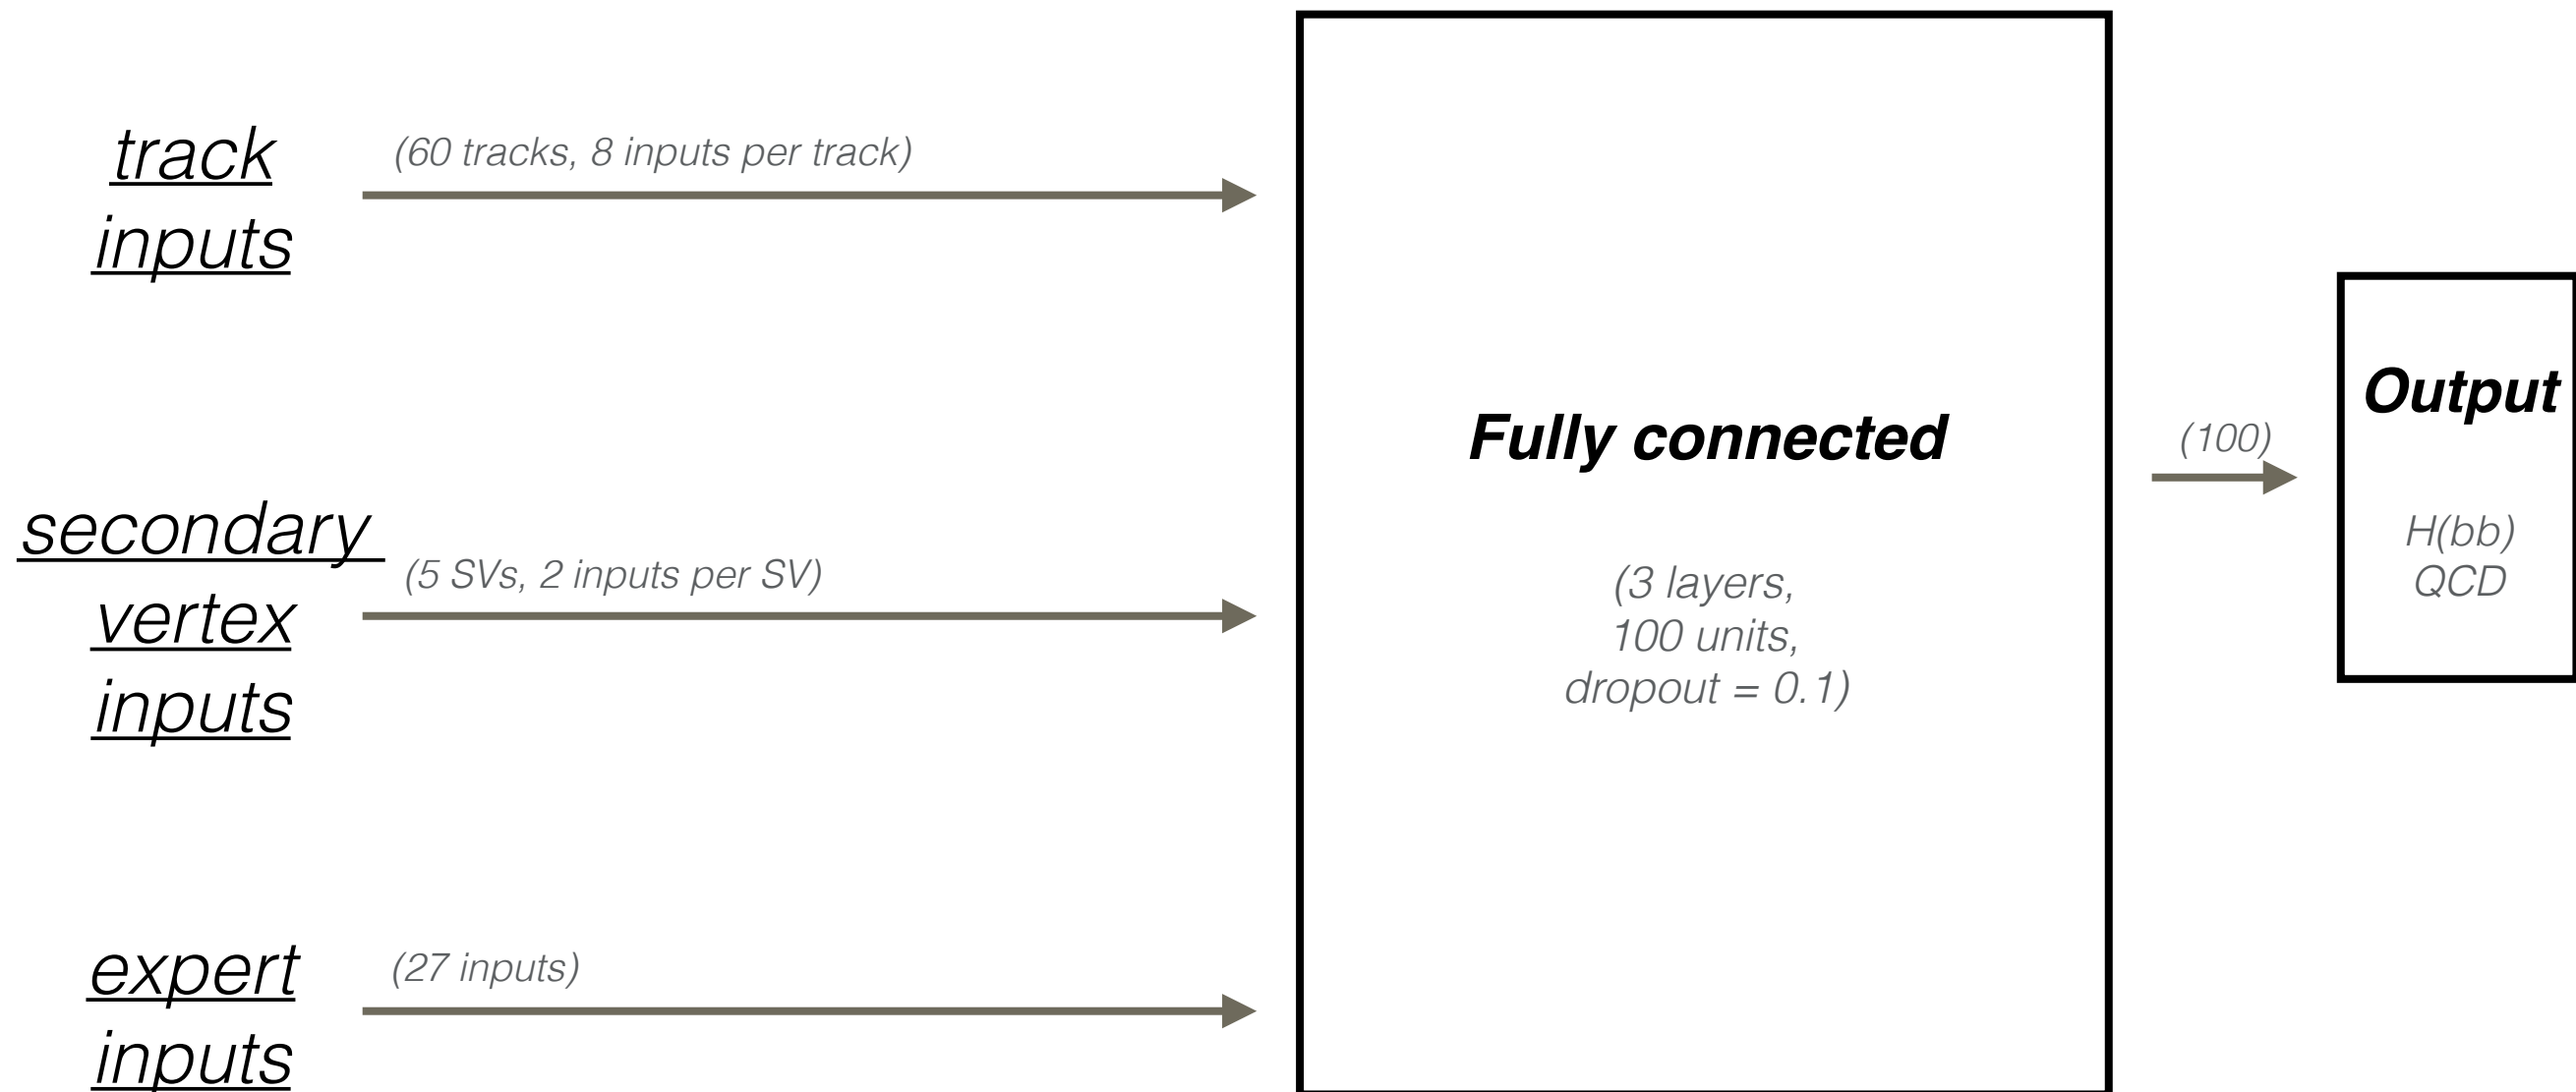
DT to a

(100)


Output

$H(bb)$
QCD

- ▶ We can switch to a neural network and add more low-level inputs based on track information and secondary vertex information: up to 517 input variables!



No big gain in performance...



We can simplify
level inputs
vertex info

track (60 in)
inputs

secondary (5 in)
vertex
inputs

expert (27 in)
inputs

more low-
secondary

(100)

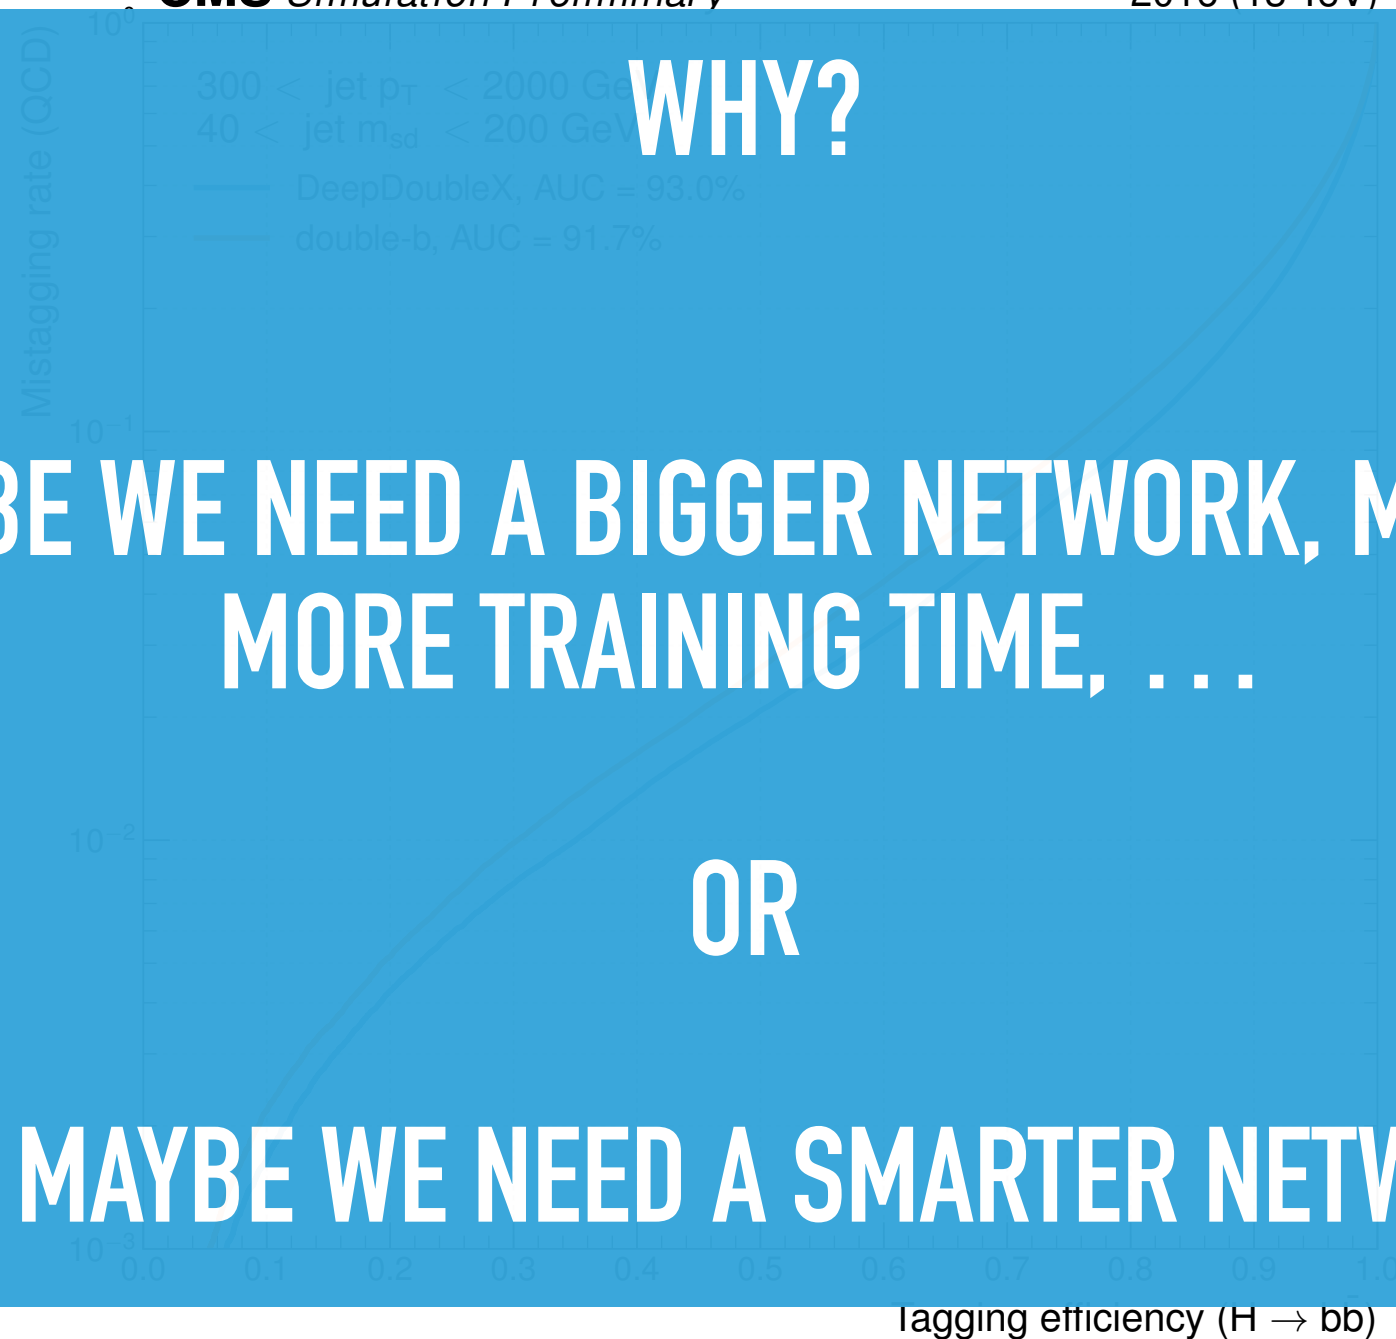
Output

H(bb)
QCD

No big gain in performance...

CMS Simulation Preliminary

2016 (13 TeV)



WHY?

(1) MAYBE WE NEED A BIGGER NETWORK, MORE DATA, MORE TRAINING TIME, ...

OR


(2) MAYBE WE NEED A SMARTER NETWORK

tagging efficiency (H → bb)

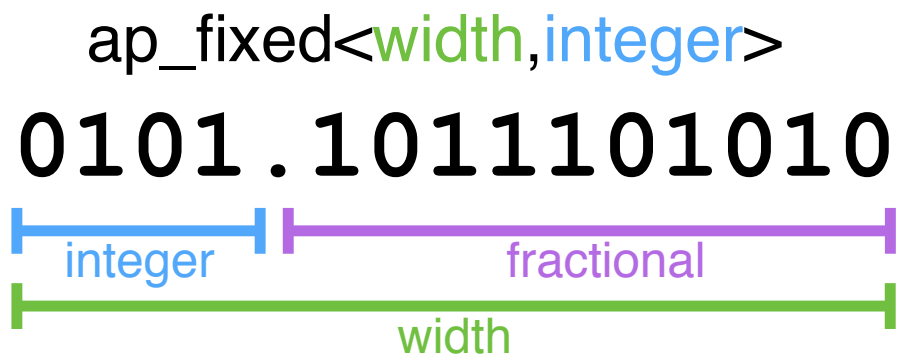
expert
inputs

(27 in)


- ▶ **Increasing** reuse factor, **decreases** resources



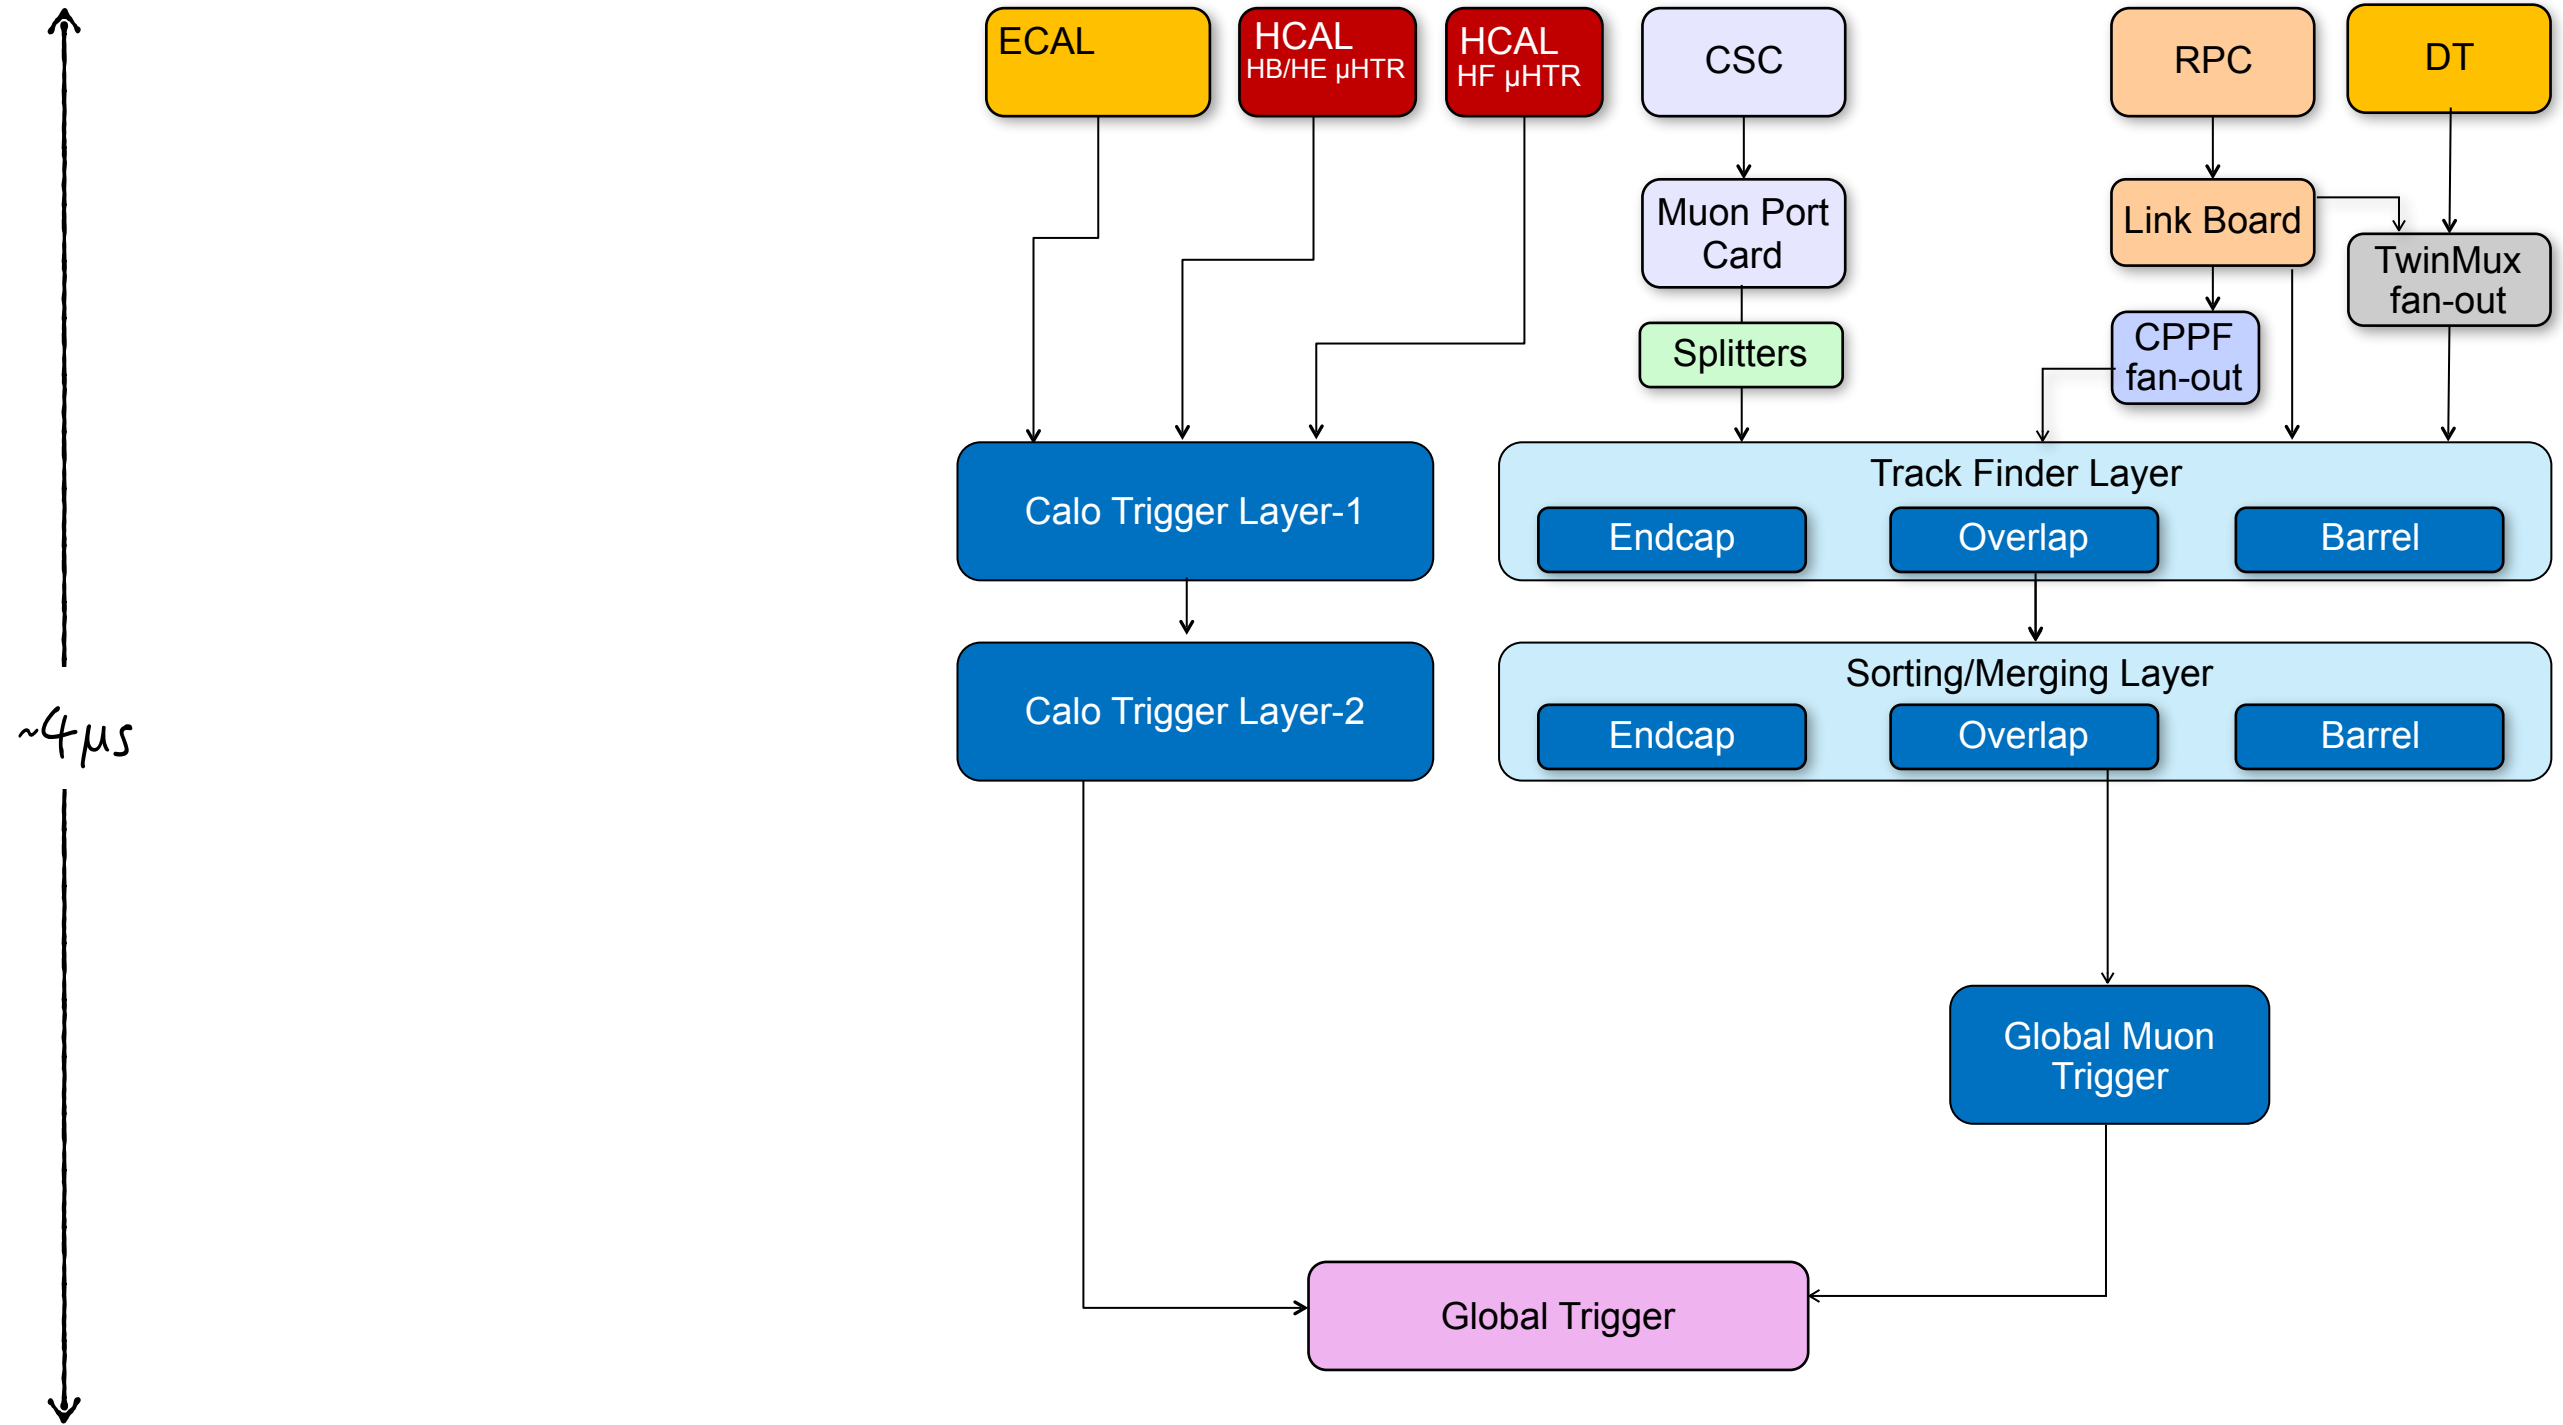
- ▶ Scan the bit width until you reach optimal performance



Full performance
with 16 bits



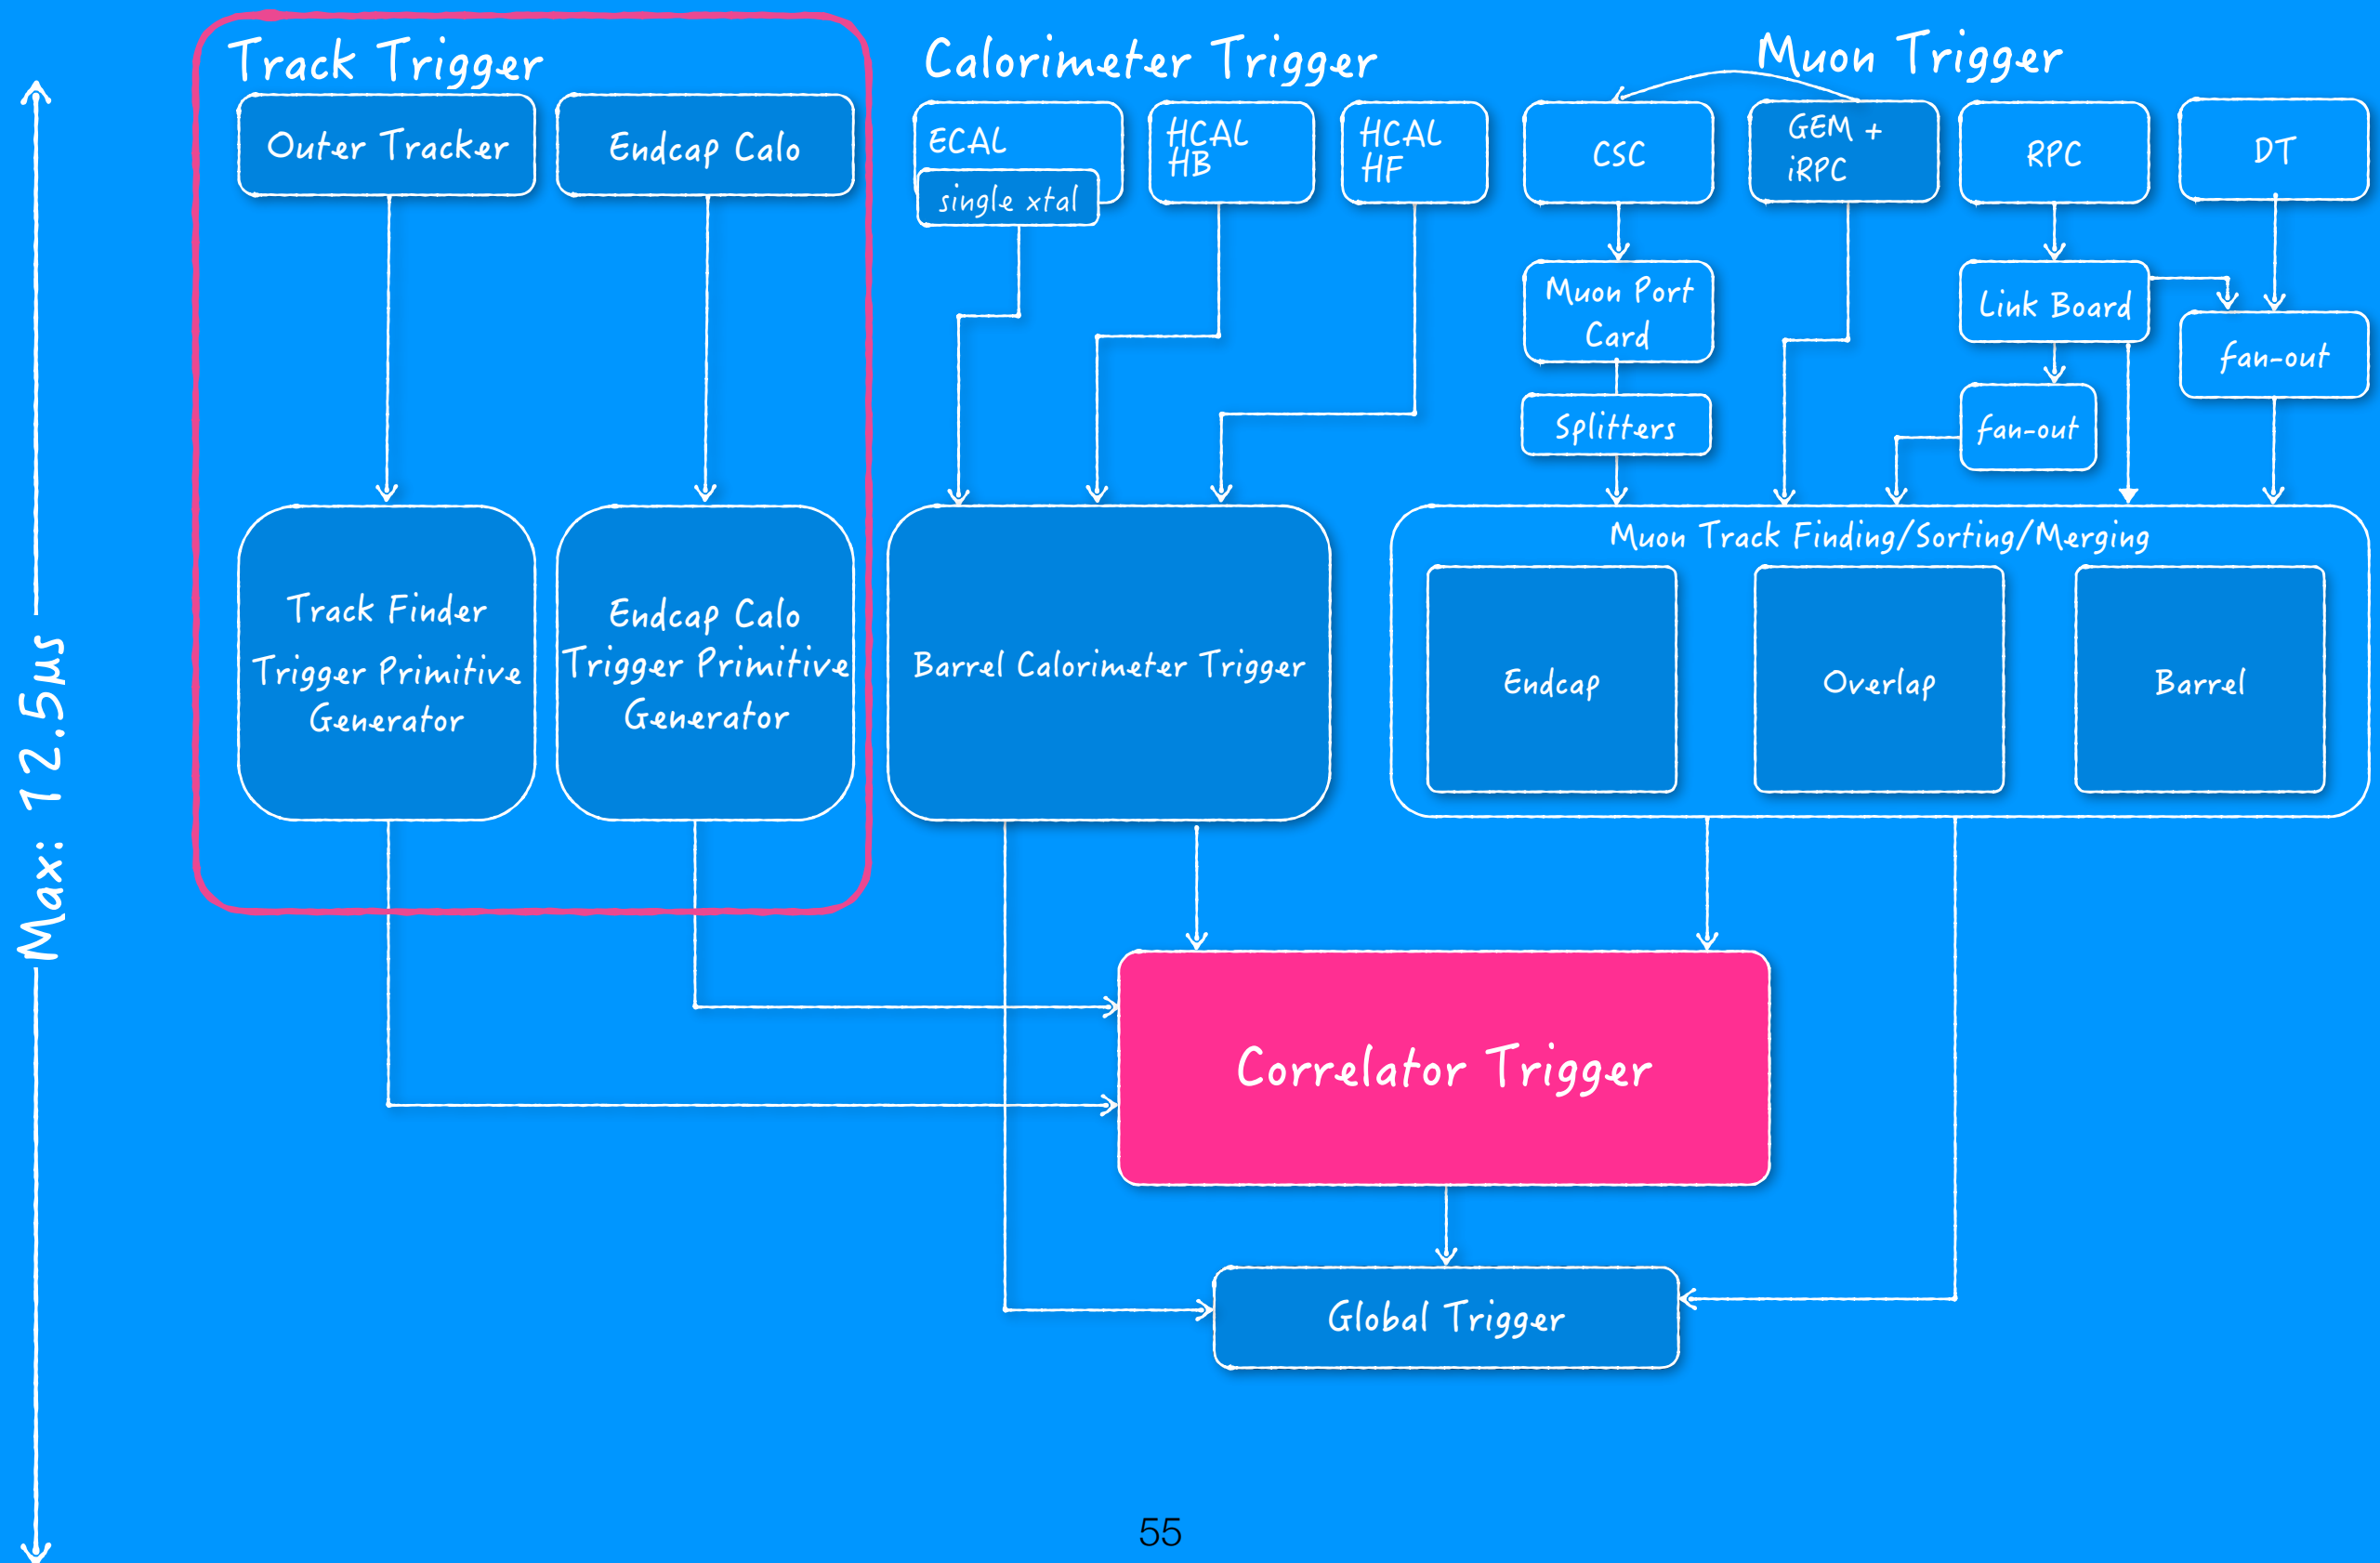
UPGRADING THE LEVEL-1 TRIGGER (BEFORE)



UPGRADING THE LEVEL-1 TRIGGER (AFTER)





More and better information available in the Level-1 trigger!

What can we do with it?




Observables


- m_{mMDT}
- $N_2^{\beta=1,2}$
- $M_2^{\beta=1,2}$
- $C_1^{\beta=0,1,2}$
- $C_2^{\beta=1,2}$
- $D_2^{\beta=1,2}$
- $D_2^{(\alpha,\beta)=(1,1),(1,2)}$
- $\sum z \log z$
- Multiplicity




- ▶ 16 expert observables provide separation between top, W/Z, and quark/gluon



Train
with L_1




Prune




2nd iteration

Retrain
with L_1



Prune




■
■
■

■
■
■


■
■
■

7th iteration

Retrain
with L_1
















Prune




A mostly complete chart of Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org

-  Backfed Input Cell
-  Input Cell
-  Noisy Input Cell
-  Hidden Cell
-  Probabilistic Hidden Cell
-  Spiking Hidden Cell
-  Output Cell
-  Match Input Output Cell
-  Recurrent Cell
-  Memory Cell
-  Different Memory Cell
-  Kernel
-  Convolution or Pool


Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Deep Feed Forward (DFF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)




Auto Encoder (AE)




Variational AE (VAE)




Denoising AE (DAE)




Sparse AE (SAE)




Markov Chain (MC)




Hopfield Network (HN)




Boltzmann Machine (BM)




Restricted BM (RBM)




Deep Belief Network (DBN)




Deep Convolutional Network (DCN)



Deconvolutional Network (DN)
















Deep Convolutional Inverse Graphics Network (DCIGN)




▶ You have a task to accomplish, which can be represented as a smooth function from your inputs to the answer you want

A mostly complete chart of Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org

-  Backfed Input Cell
-  Input Cell
-  Noisy Input Cell
-  Hidden Cell
-  Probabilistic Hidden Cell
-  Spiking Hidden Cell
-  Output Cell
-  Match Input Output Cell
-  Recurrent Cell
-  Memory Cell
-  Different Memory Cell
-  Kernel
-  Convolution or Pool


Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Deep Feed Forward (DFF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)




Auto Encoder (AE)




Variational AE (VAE)




Denoising AE (DAE)




Sparse AE (SAE)




Markov Chain (MC)




Hopfield Network (HN)




Boltzmann Machine (BM)




Restricted BM (RBM)




Deep Belief Network (DBN)




Deep Convolutional Network (DCN)



Deconvolutional Network (DN)



Deep Convolutional Inverse Graphics Network (DCIGN)




- ▶ You have a task to accomplish, which can be represented as a smooth function from your inputs to the answer you want
 - ▶ Train an algorithm to learn an approximation of the optimal solution function (Machine Learning)

A mostly complete chart of Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org

- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probabilistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool


Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Deep Feed Forward (DFF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)




Auto Encoder (AE)




Variational AE (VAE)




Denoising AE (DAE)




Sparse AE (SAE)




Markov Chain (MC)




Hopfield Network (HN)




Boltzmann Machine (BM)




Restricted BM (RBM)




Deep Belief Network (DBN)




Deep Convolutional Network (DCN)



Deconvolutional Network (DN)



Deep Convolutional Inverse Graphics Network (DCIGN)




- ▶ You have a task to accomplish, which can be represented as a smooth function from your inputs to the answer you want
 - ▶ Train an algorithm to learn an approximation of the optimal solution function (Machine Learning)
- ▶ NNs are the best ML solution on the market *today*

A mostly complete chart of Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org

- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probabilistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool


Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Deep Feed Forward (DFF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)




Auto Encoder (AE)




Variational AE (VAE)




Denoising AE (DAE)




Sparse AE (SAE)




Markov Chain (MC)




Hopfield Network (HN)




Boltzmann Machine (BM)




Restricted BM (RBM)




Deep Belief Network (DBN)




Deep Convolutional Network (DCN)



Deconvolutional Network (DN)



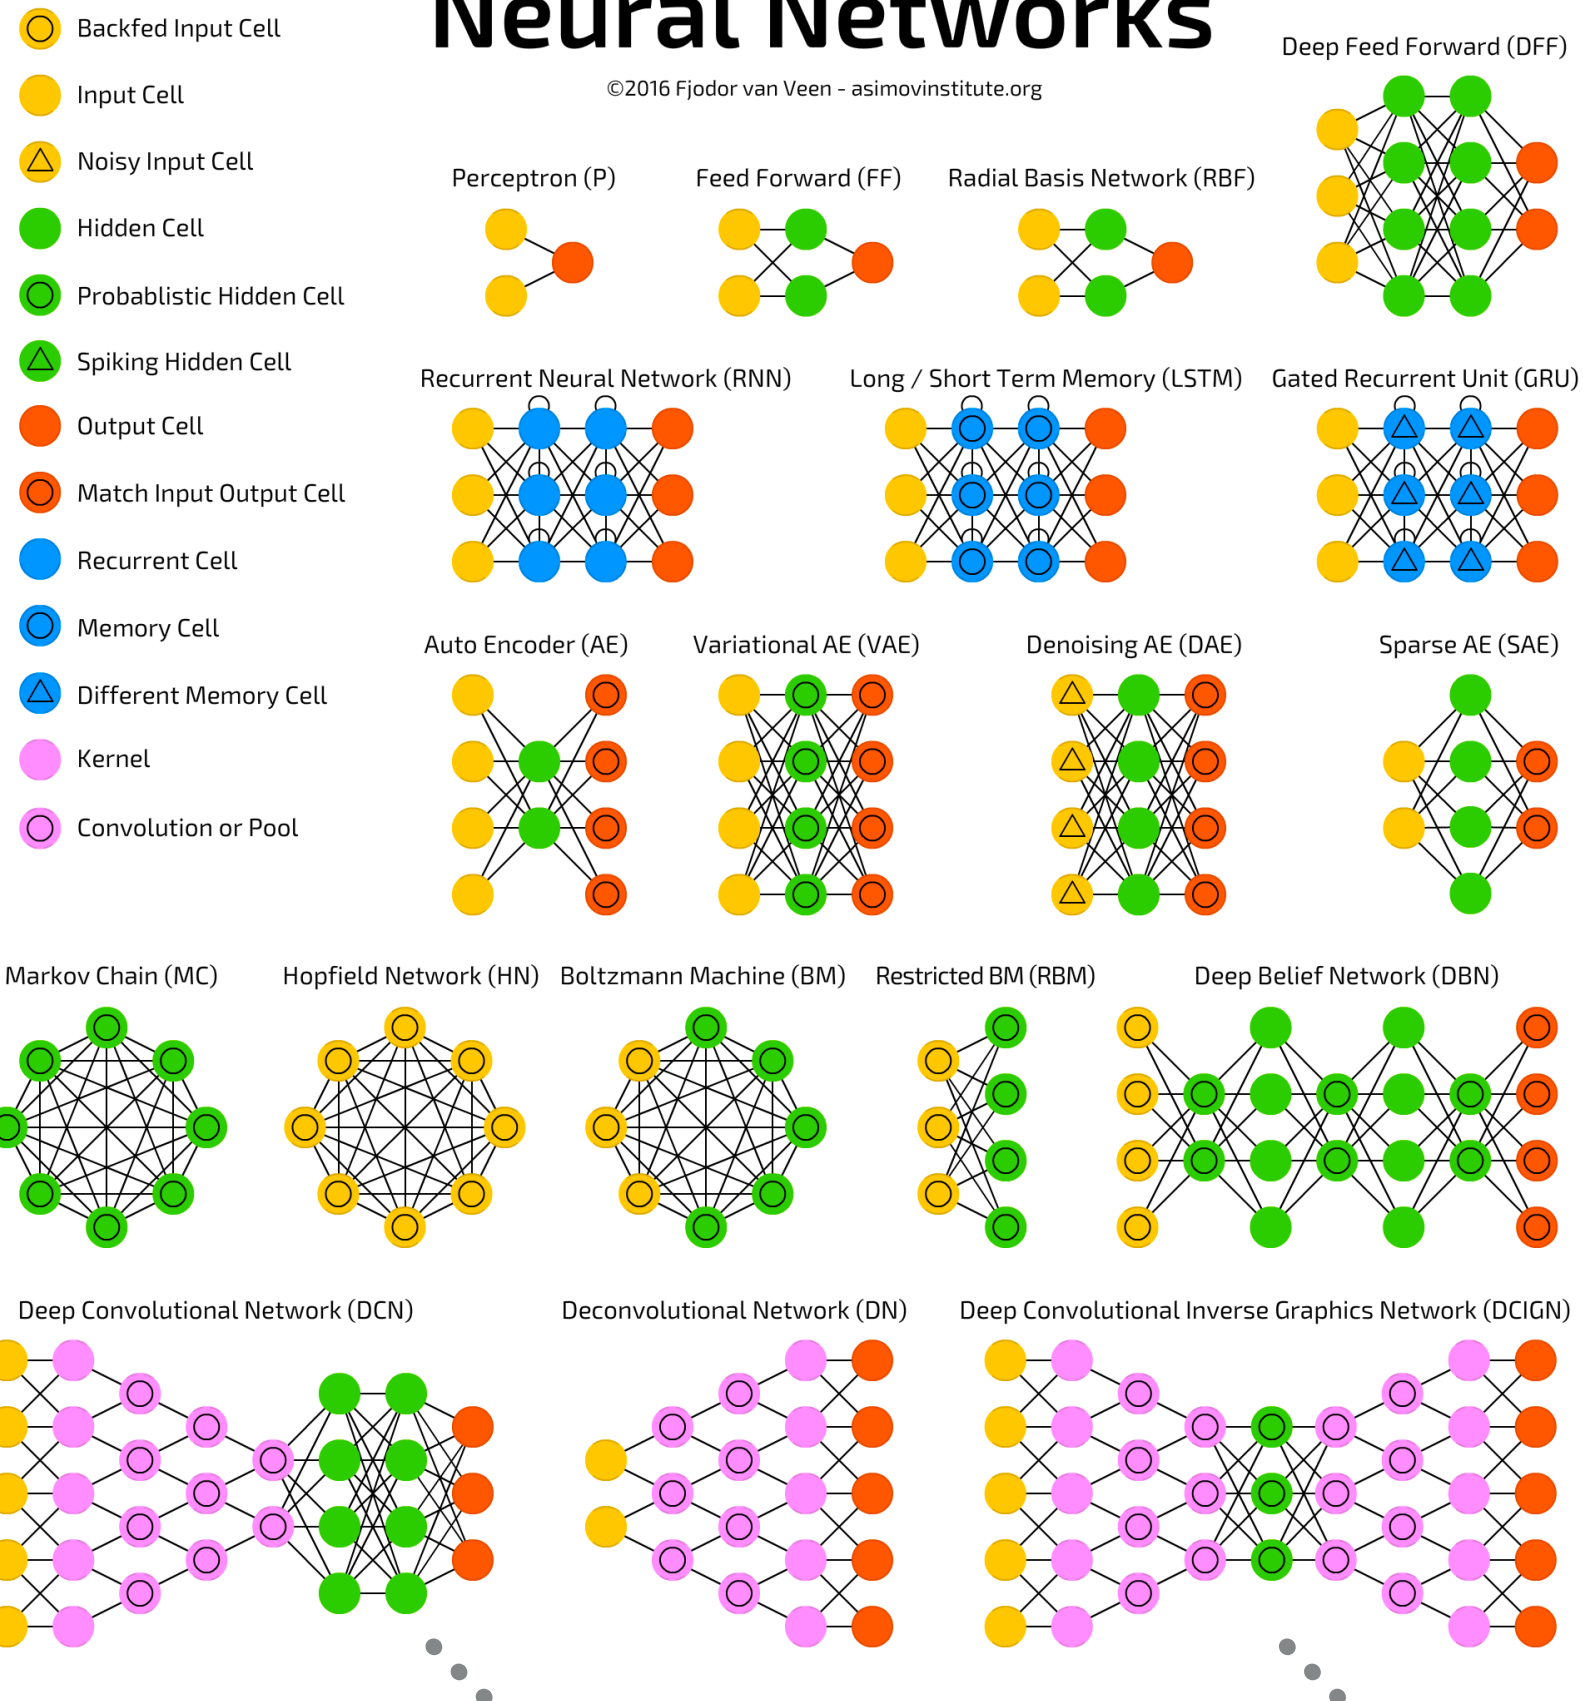
Deep Convolutional Inverse Graphics Network (DCIGN)



- ▶ You have a task to accomplish, which can be represented as a smooth function from your inputs to the answer you want
 - ▶ Train an algorithm to learn an approximation of the optimal solution function (Machine Learning)
- ▶ NNs are the best ML solution on the market *today*
 - ▶ Each node performs a math operation on the input

A mostly complete chart of Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org



- ▶ You have a task to accomplish, which can be represented as a smooth function from your inputs to the answer you want
 - ▶ Train an algorithm to learn an approximation of the optimal solution function (Machine Learning)
- ▶ NNs are the best ML solution on the market *today*
 - ▶ Each node performs a math operation on the input
 - ▶ Edges represent the flow of nodes' inputs & outputs

A mostly complete chart of
Neural Networks





©2016 Fjodor van Veen - asimovinstitute.org



$$\ell_i^{k-1}$$


$$\ell_i^{k-1}$$

- ▶ Classic feed forward architecture with some modifications responsible for revolutions in computer vision, language processing, etc.

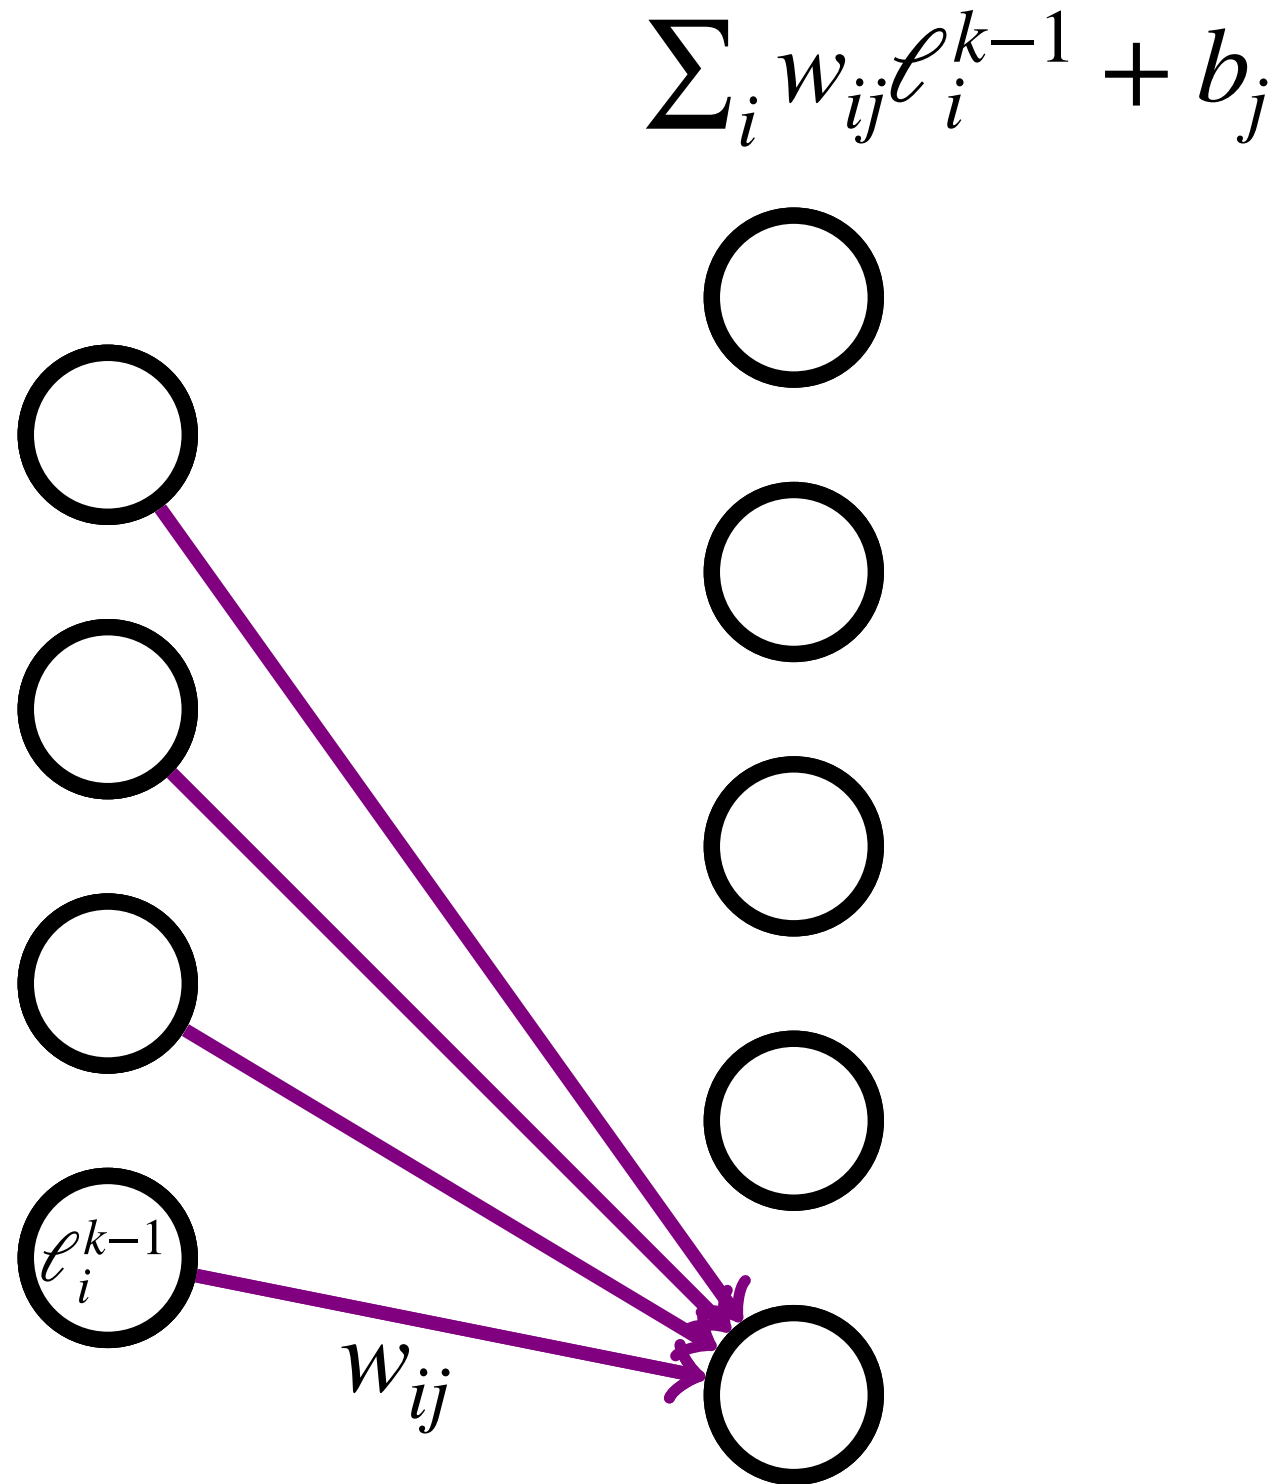


$$\ell_i^{k-1}$$

- ▶ Classic feed forward architecture with some modifications responsible for revolutions in computer vision, language processing, etc.
- ▶ Each input multiplied by a weight


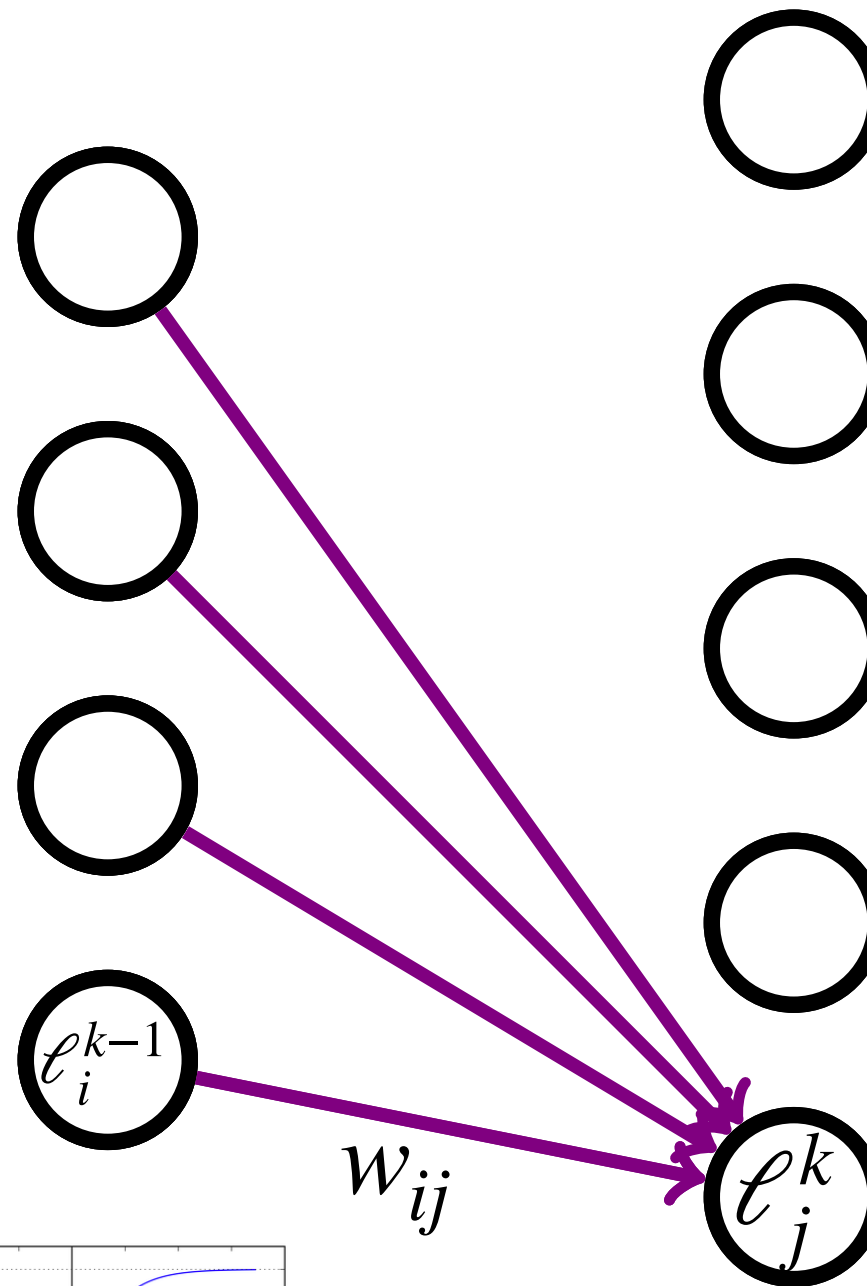


- ▶ Classic feed forward architecture with some modifications responsible for revolutions in computer vision, language processing, etc.
- ▶ Each input multiplied by a weight
- ▶ Weighted values are summed, bias is added





- ▶ Classic feed forward architecture with some modifications responsible for revolutions in computer vision, language processing, etc.
- ▶ Each input multiplied by a weight
- ▶ Weighted values are summed, bias is added
- ▶ Nonlinear activation function is applied

$$\ell_j^k = \phi \left(\sum_i w_{ij} \ell_i^{k-1} + b_j \right)$$





- ▶ Classic feed forward architecture with some modifications responsible for revolutions in computer vision, language processing, etc.
- ▶ Each input multiplied by a weight
- ▶ Weighted values are summed, bias is added
- ▶ Nonlinear activation function is applied


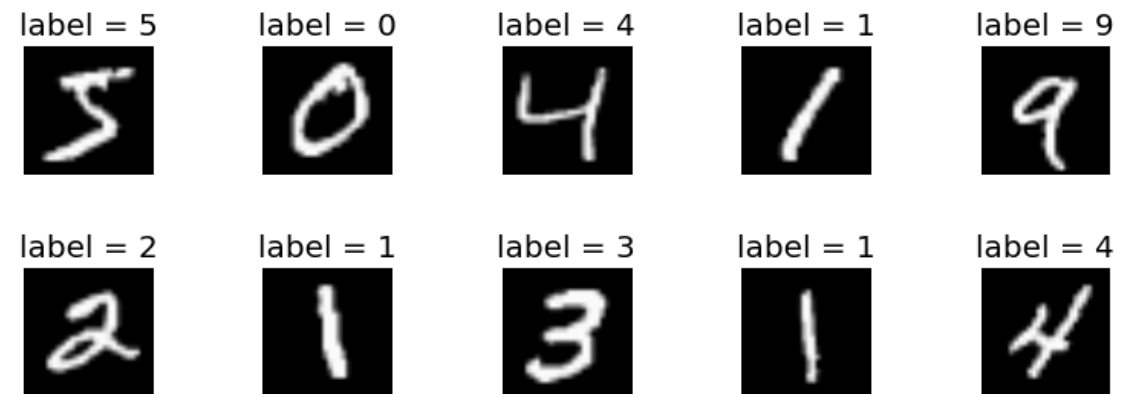
$$\ell_j^k = \phi \left(\sum_i w_{ij} \ell_i^{k-1} + b_j \right)$$





A sufficiently "wide" neural network can approximate any function!





- ▶ A network is trained by specifying inputs, targets, and a loss function
 - ▶ Target is what the network should learn for that input, can be a "truth" label (supervised) or the input itself (unsupervised)
 - ▶ Loss function quantifies how many mistakes the network makes
- ▶ Training is the minimization of the loss function by varying the network parameters



- ▶ A network is trained by specifying inputs, targets, and a loss function
 - ▶ Target is what the network should learn for that input, can be a "truth" label (supervised) or the input itself (unsupervised)
 - ▶ Loss function quantifies how many mistakes the network makes
- ▶ Training is the minimization of the loss function by varying the network parameters





- ▶ Using $g \rightarrow bb$ jets as a proxy in **double muon** tagged jet sample
- ▶ Associated data/MC uncertainty 3-5%

- The first four SIP values for selected tracks ordered in decreasing SIP;
- For each τ -axis we consider the first two SIP values for their respective associated tracks ordered in decreasing SIP, to further discriminate against single b quark and light flavor jets from QCD when one or both SV are not reconstructed due to IVF inefficiencies;
- The measured IP significance in the plane transverse to the beam axis, 2D SIP, of the first two tracks (first track) that raises the SV invariant mass above the bottom (charm) threshold of 5.2 (1.5) GeV;
- The number of SV associated to the jet;
- The significance of the 2D distance between the primary vertex and the secondary vertex, flight distance, for the SV with the smallest 3D flight distance uncertainty, for each of the two τ -axes;
- The ΔR between the SVs with the smallest 3D flight distance uncertainty and its τ -axis, for each of the two τ -axes;
- The relative pseudorapidity, η_{rel} , of the tracks from all SVs with respect to their τ -axis for the three leading tracks ordered in increasing η_{rel} , for each of the two τ -axes;
- The total SV mass, defined as the total mass of all SVs associated to a given τ -axis, for each of the two τ -axes;
- The ratio of the total SV energy, defined as the total energy of all SVs associated to a given τ -axis, and the total energy of all the tracks associated to the fat jet that are consistent with the primary vertex, for each of the two τ -axes;
- The information related to the two-SV system, the z variable, defined as:

$$z = \Delta R(SV_0, SV_1) \cdot \frac{p_{T,SV_1}}{m(SV_0, SV_1)} \quad (2)$$

where SV_0 and SV_1 are SVs with the smallest 3D flight distance uncertainty. The z variable helps rejecting the $b\bar{b}$ background from gluon splitting relying on the different kinematic properties compared to the $b\bar{b}$ pair from the decay of a massive resonance.







Table 11: Full list of inclusive PF candidate features used as input to the DeepAK8 network

feature
$p_T(PF) / p_T(j)$
$E_{rel}(PF)$
$\Delta\phi(PF, j)$
$\Delta\eta(PF, j)$
$\Delta R(PF, j)$
$\Delta R_m(PF, SV)$
$\Delta R(PF, \text{subject 1})$
$\Delta R(PF, \text{subject 2})$
$w_p(PF)$
f_{HCAL}

Table 10: Full list of charged PF candidate features used as input to the DeepAK8 network

feature	comment
trackEtaRel	BTV
trackPtRatio	BTV
trackPParRatio	BTV
trackSip2dVal	BTV
trackSip2dSig	BTV
trackSip3dVal	BTV
trackSip3dSig	BTV
trackJetDistVal	BTV
$p_T(cPF) / p_T(j)$	
$E_{rel}(cPF)$	
$\Delta\phi(cPF, j)$	
$\Delta\eta(cPF, j)$	
$\Delta R(cPF, j)$	
$\Delta R_m(cPF, SV)$	
$\Delta R(cPF, \text{subject 1})$	
$\Delta R(cPF, \text{subject 2})$	
χ_n^2	
quality	
d_z	
S_z	
d_{xy}	
S_{xy}	
track_dptdpt	track covariance
track_detadeta	track covariance
track_dphidphi	track covariance
track_dxydxy	track covariance
track_dzdz	track covariance
track_dxydz	track covariance
track_dphidxy	track covariance
track_dlambdadz	track covariance

Table 12: Full list of secondary vertex features used as input to the DeepAK8 network

feature
$p_T(SV) / p_T(j)$
$E_{rel}(SV)$
$\Delta\phi(SV, j)$
$\Delta\eta(SV, j)$
$\Delta R(SV, j)$
$p_T(SV)$
m_{SV}
$N_{\text{tracks}}(SV)$
$\chi_n^2(SV)$
$d_{xy}(SV)$
$S_{xy}(SV)$
$d_{3D}(SV)$
$S_{3D}(SV)$
$\cos\theta(SV)$