De-noising Graph Super Resolution with Diffusion Models and transformers

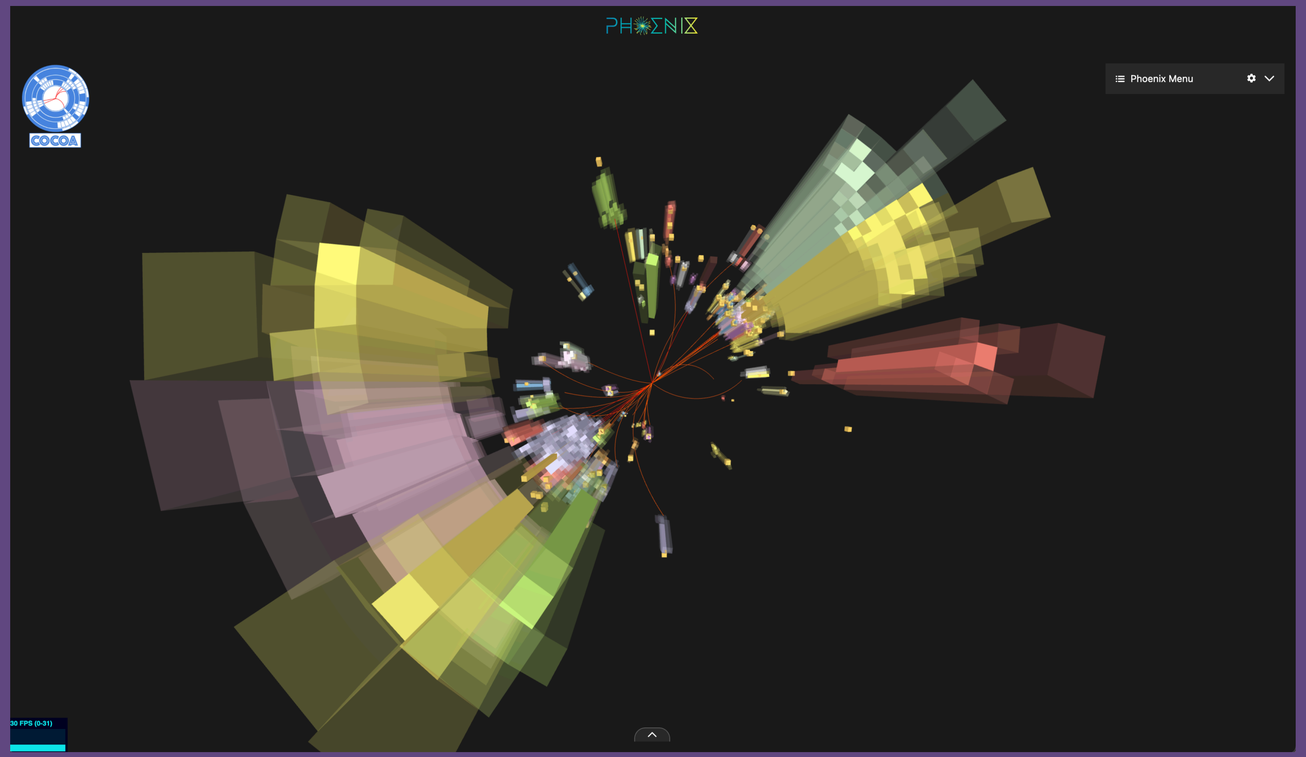

ML4Jets
06 November, 2023

Nilotpal Kakati, *Etienne Dreyer, Eilam Gross*

(nilotpal.kakati@cern.ch)

מכון ויצמן למדע
WEIZMANN INSTITUTE OF SCIENCE

Super Resolution with Diffusion

- Slightly less popular cousin of “text to image with Diffusion”
  - ➡ Still quite popular in CV
  - ➡ Not really studied in Particle Physics

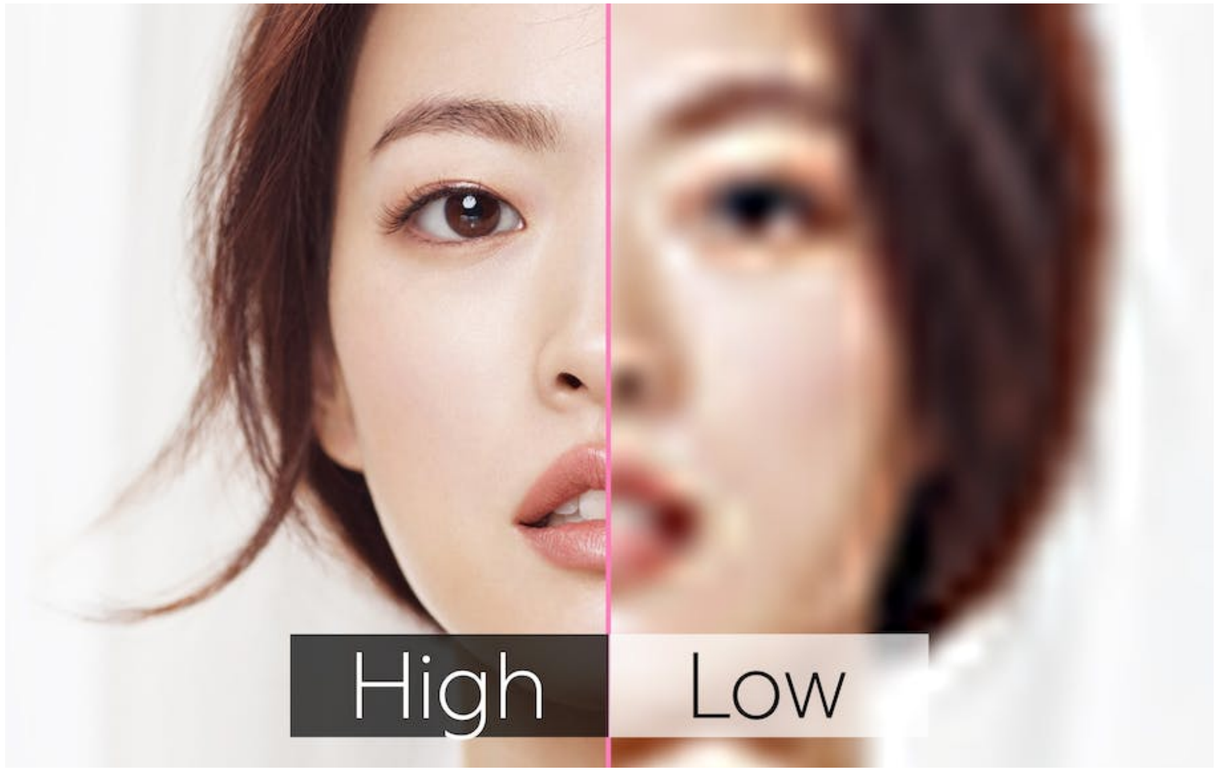

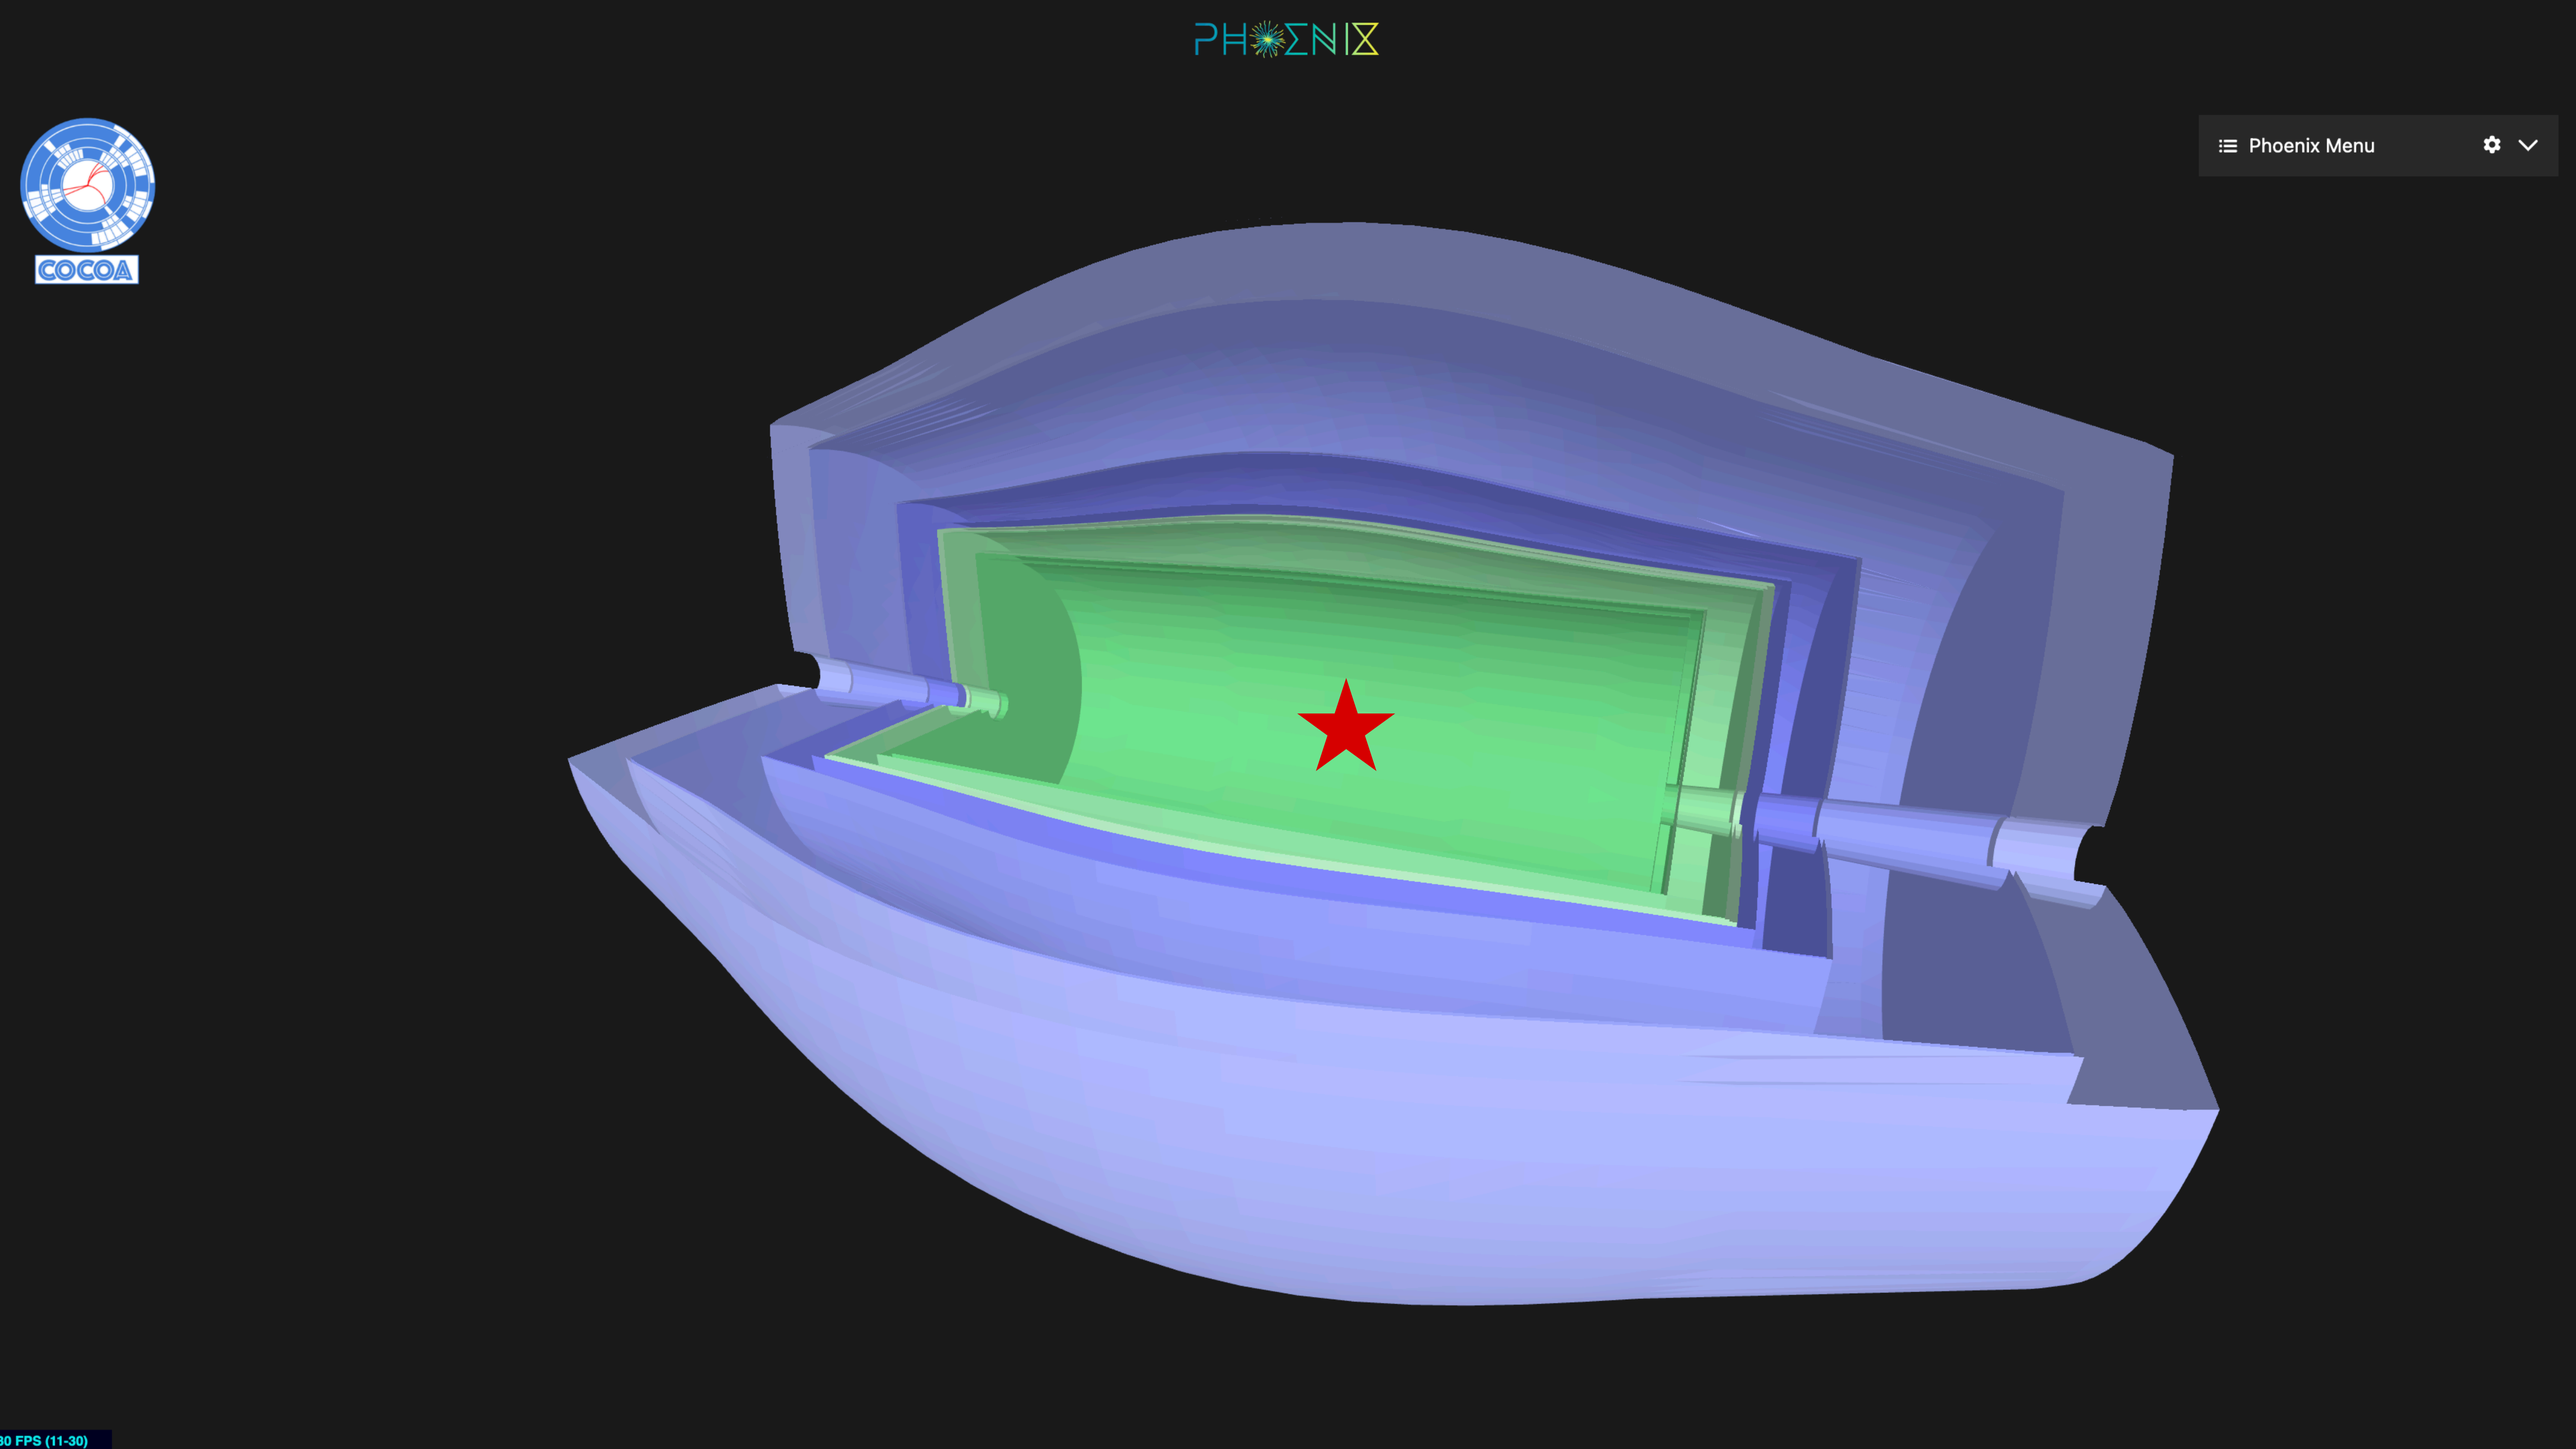
PHOENIX
COCOA
Phoenix Menu
30 FPS (11-30)

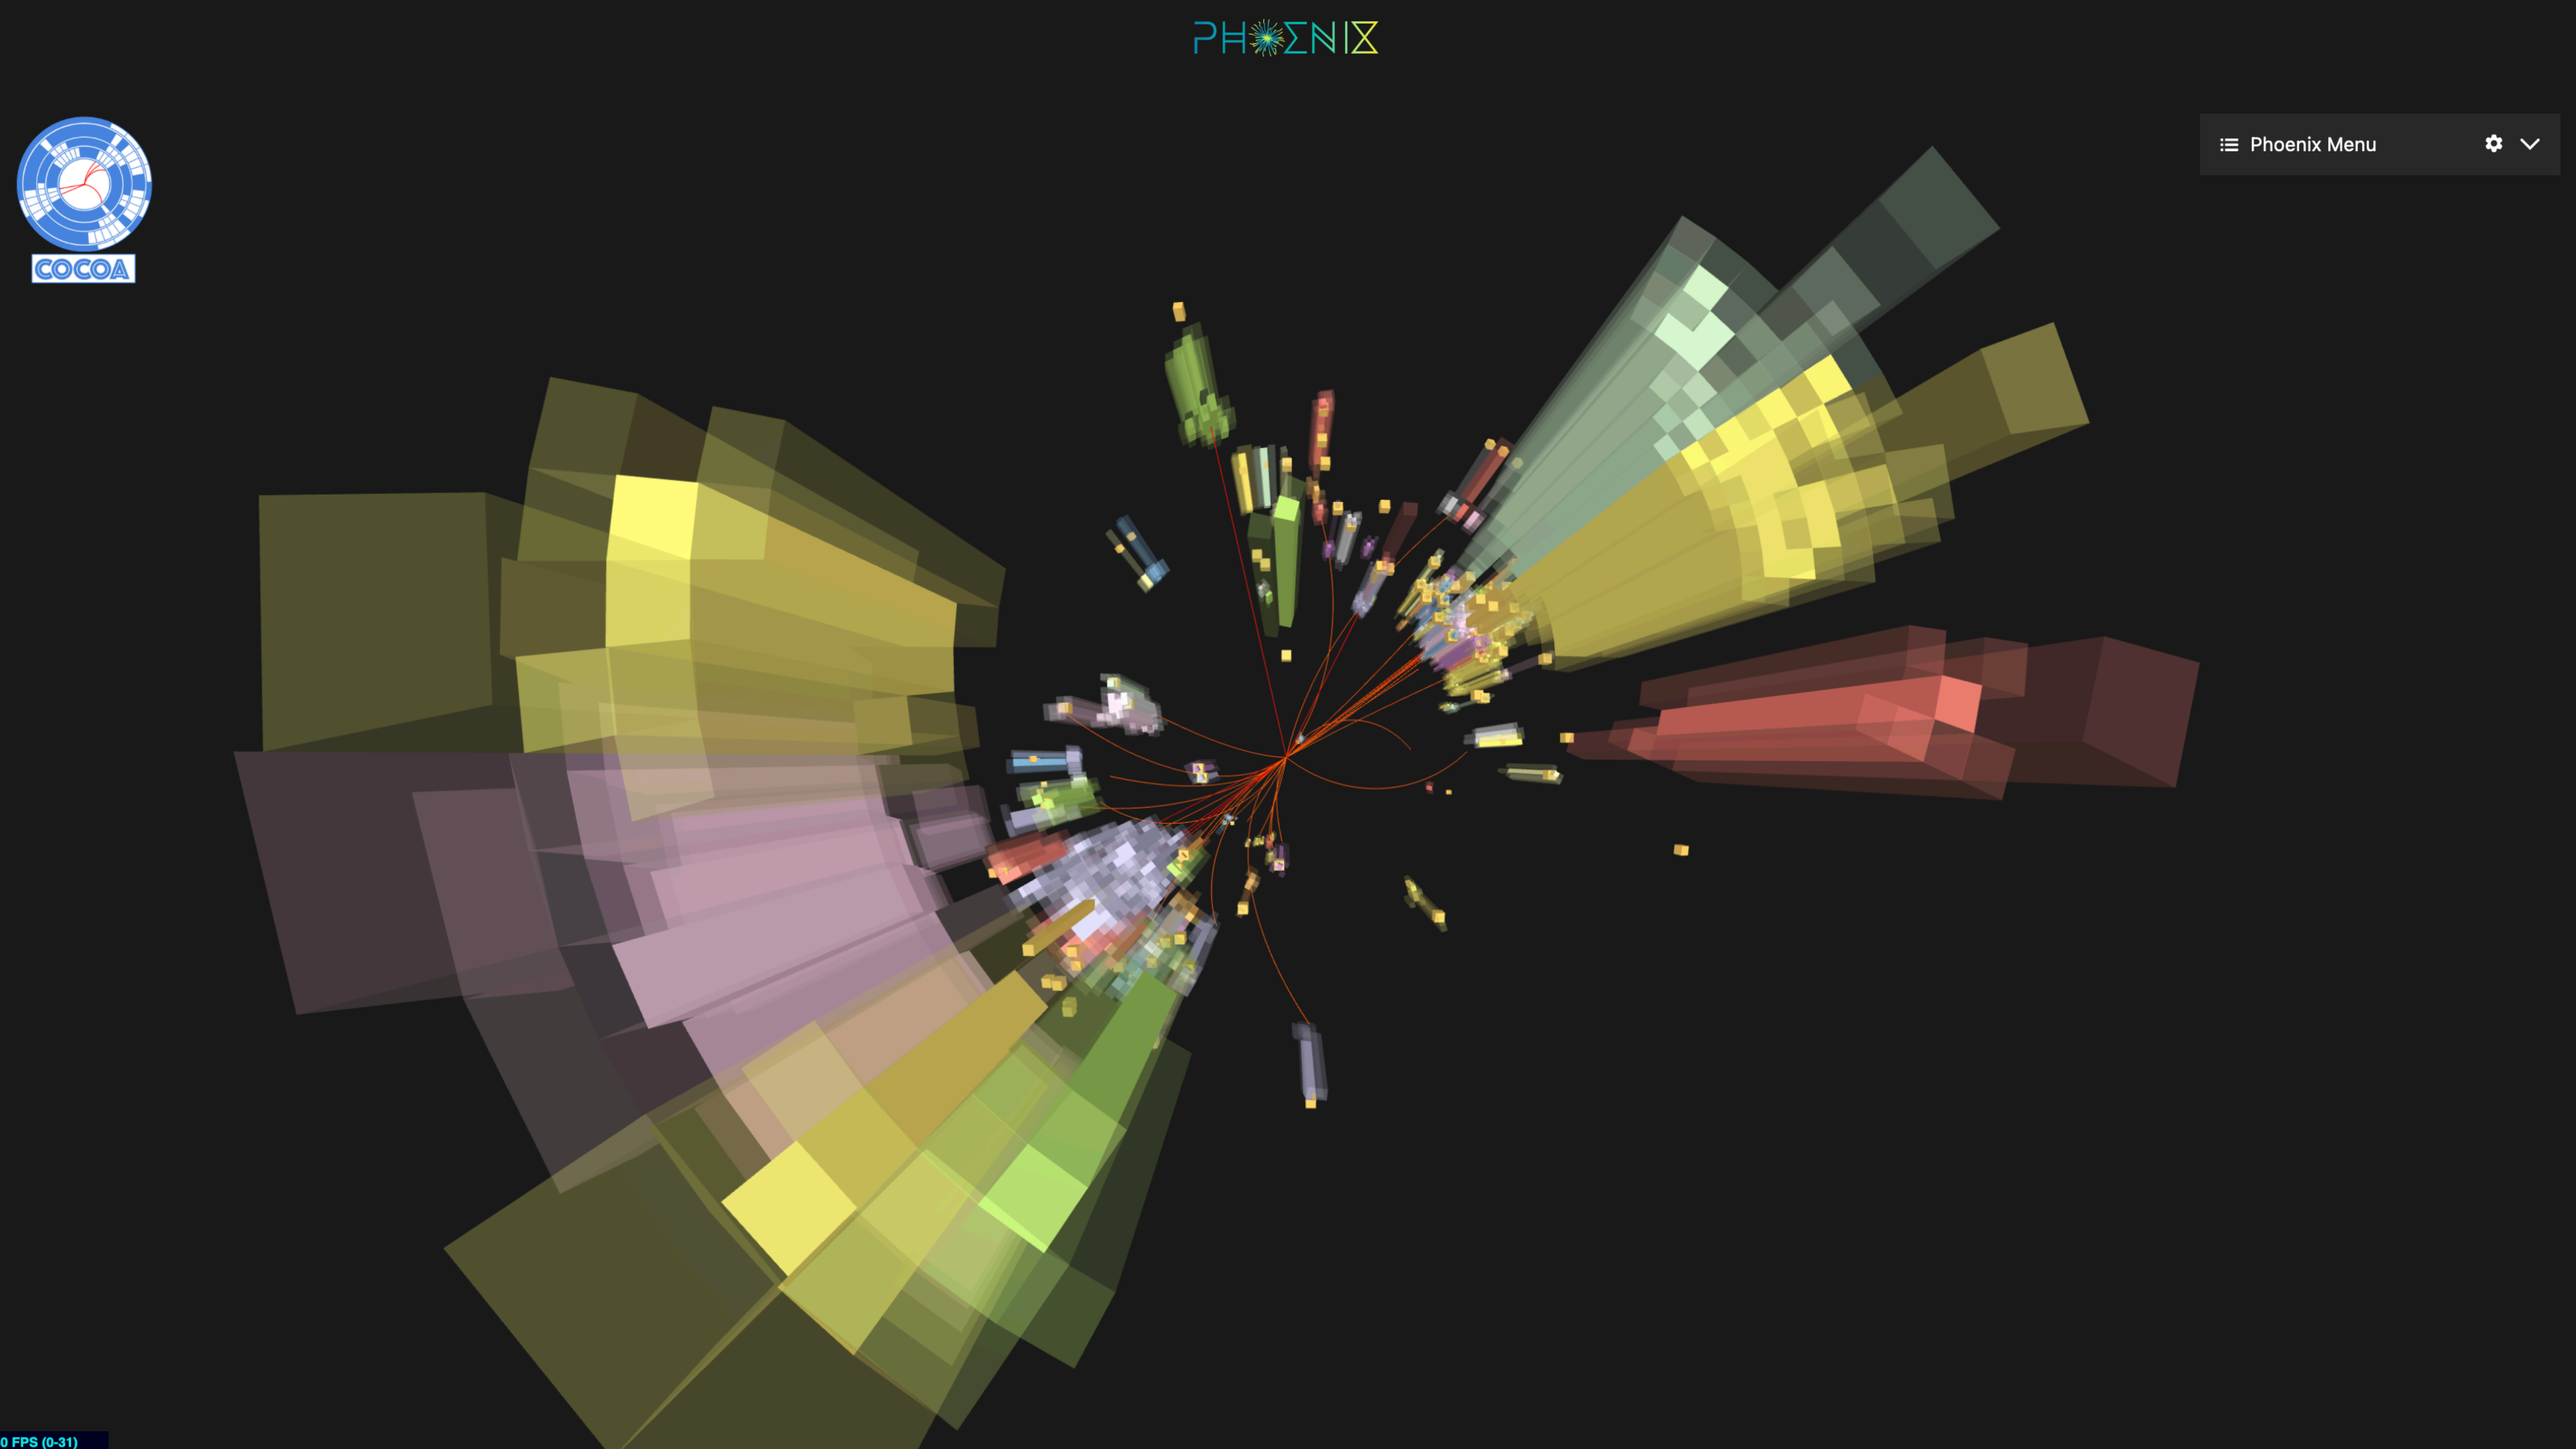
PHOENIX
COCOA
Phoenix Menu
0 FPS (0-31)

Why Super Resolution in Particle Physics?

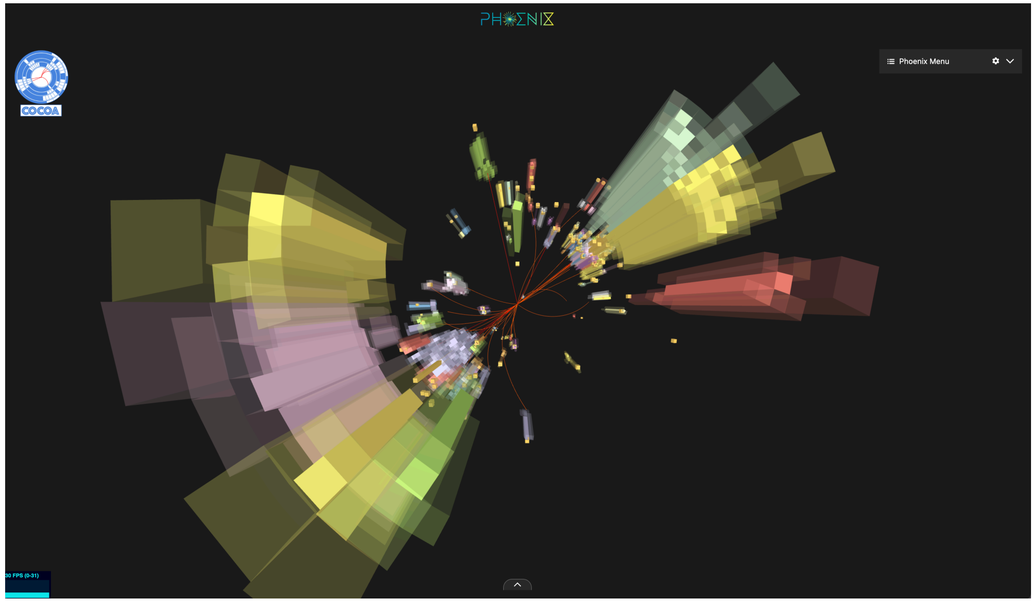

Why Super Resolution in Particle Physics?

- Reconstruction quality depends on the detector granularity
  - ➡ More granular -> better reconstruction
  - ➡ Granularity puts a cap on theoretical reconstruction capability

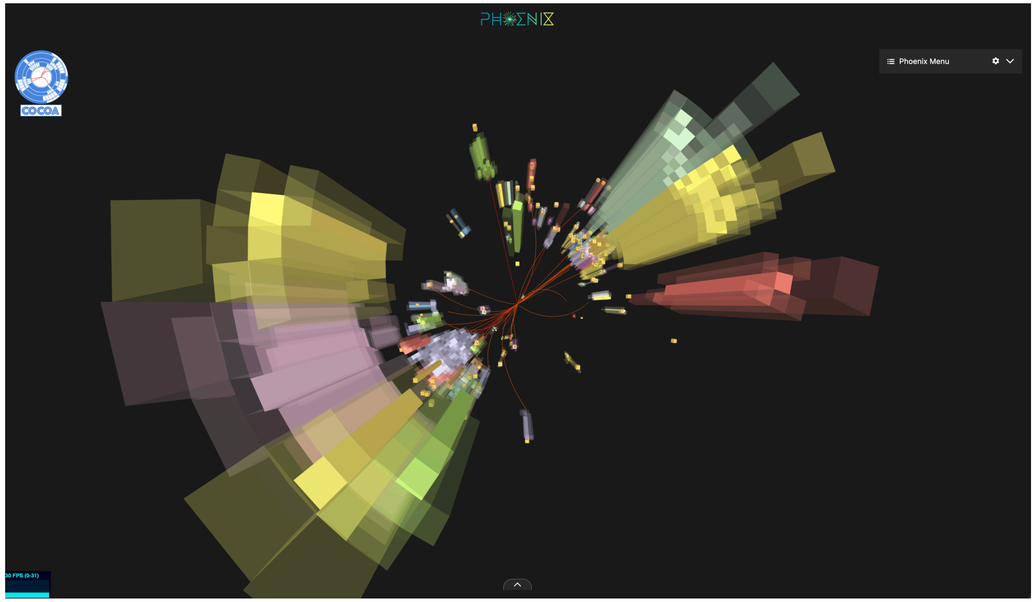

Why Super Resolution in Particle Physics?

- Reconstruction quality depends on the detector granularity
  - More granular -> better reconstruction
  - Granularity puts a cap on theoretical reconstruction capability

- High granularity detectors (simulations) are very expensive!
  - Increasing resolution in post can be a solution!

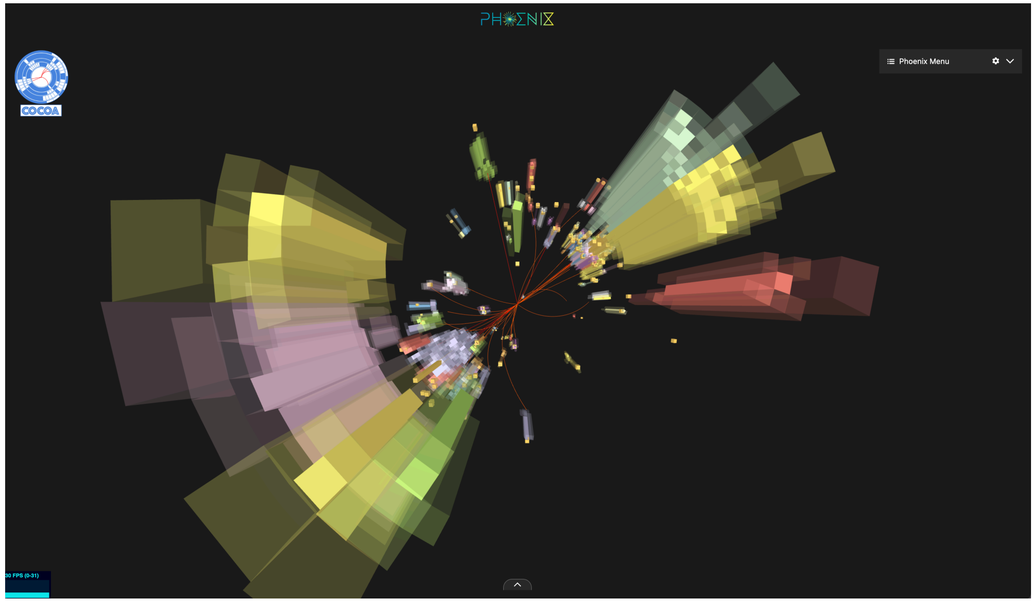

Why Super Resolution in Particle Physics?

- Reconstruction quality depends on the detector granularity
  - ➡ More granular -> better reconstruction
  - ➡ Granularity puts a cap on theoretical reconstruction capability

- High granularity detectors (simulations) are very expensive!
  - ➡ Increasing resolution in post can be a solution!

- Graph super resolution is not a common problem in general
  - ➡ Graphs are very natural in Particle Physics
  - ➡ Hence Graph Super resolution

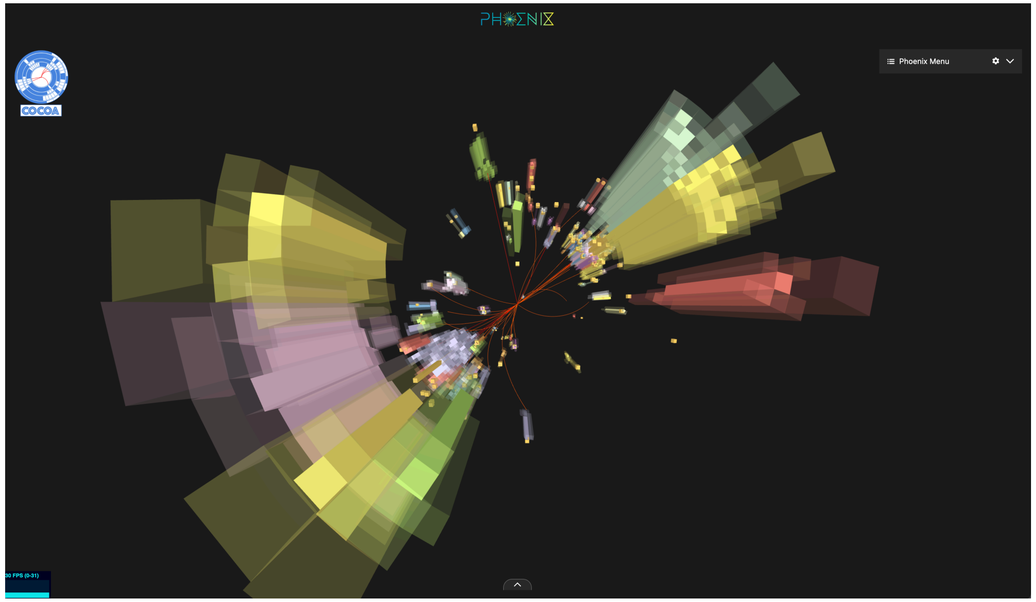

The SetUp

- COCOA mod (*https://iopscience.iop.org/article/10.1088/2632-2153/acf186/pdf*)

- Shooting single electron as a starting point

The SetUp

- COCOA mod (*https://iopscience.iop.org/article/10.1088/2632-2153/acf186/pdf*)

- Shooting single electron as a starting point

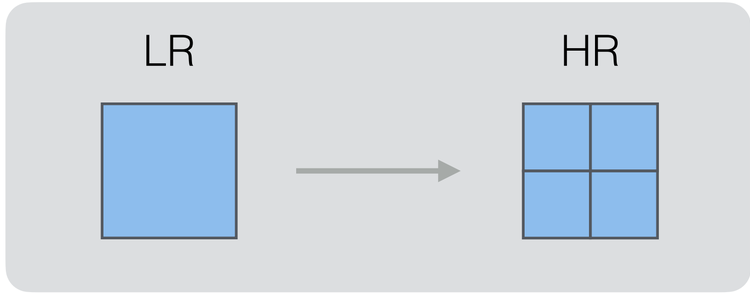

The SetUp

- COCOA mod (*https://iopscience.iop.org/article/10.1088/2632-2153/acf186/pdf*)

- Shooting single electron as a starting point

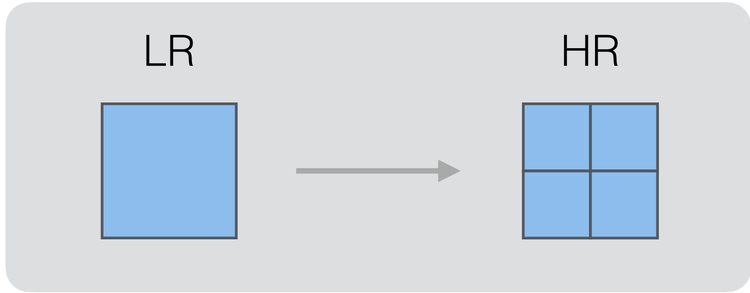

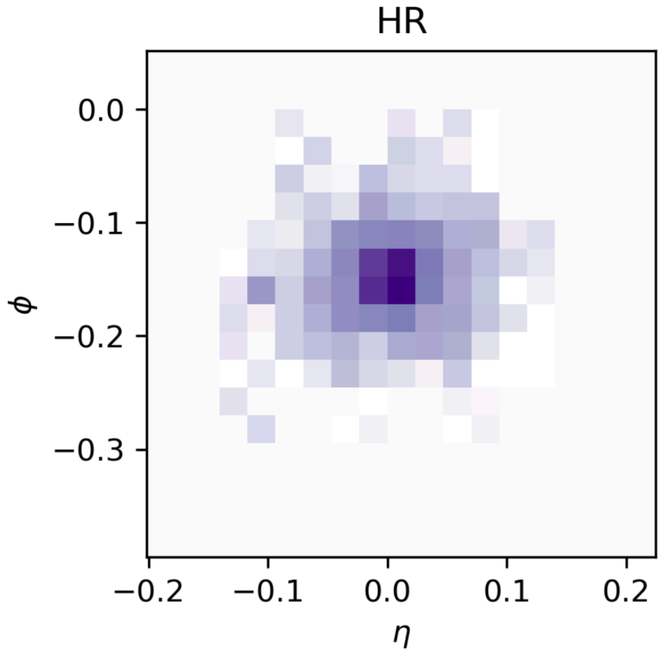

The SetUp

- COCOA mod (*https://iopscience.iop.org/article/10.1088/2632-2153/acf186/pdf*)

- Shooting single electron as a starting point

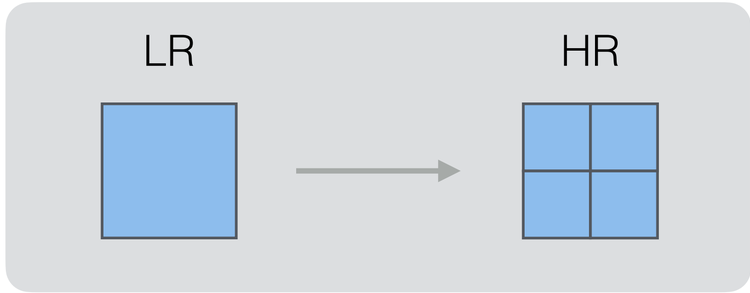

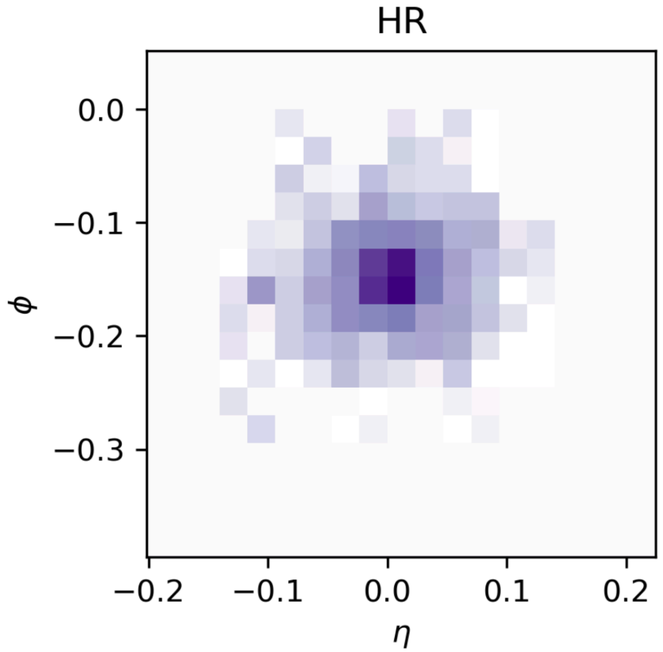

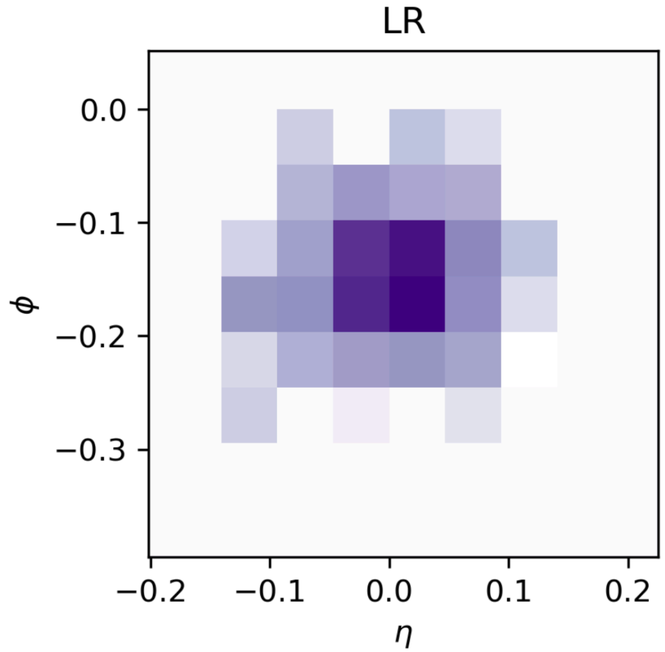

The SetUp

- COCOA mod (*https://iopscience.iop.org/article/10.1088/2632-2153/acf186/pdf*)
- Shooting single electron as a starting point

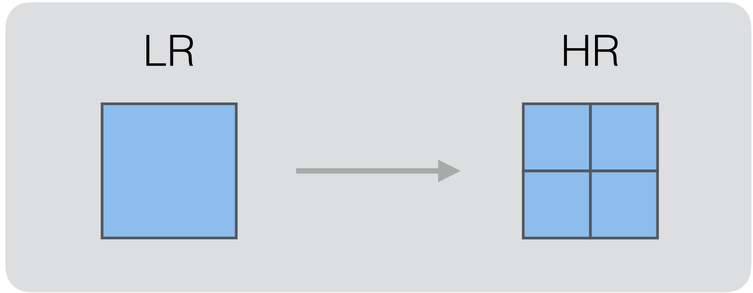

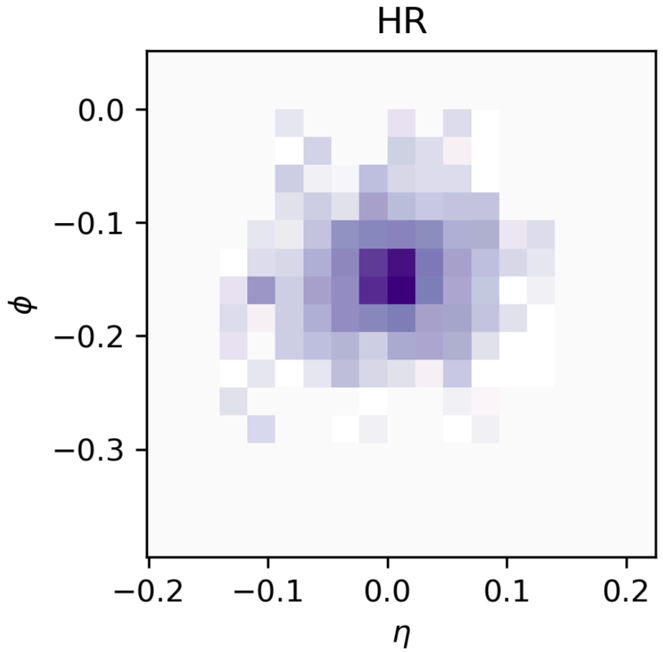

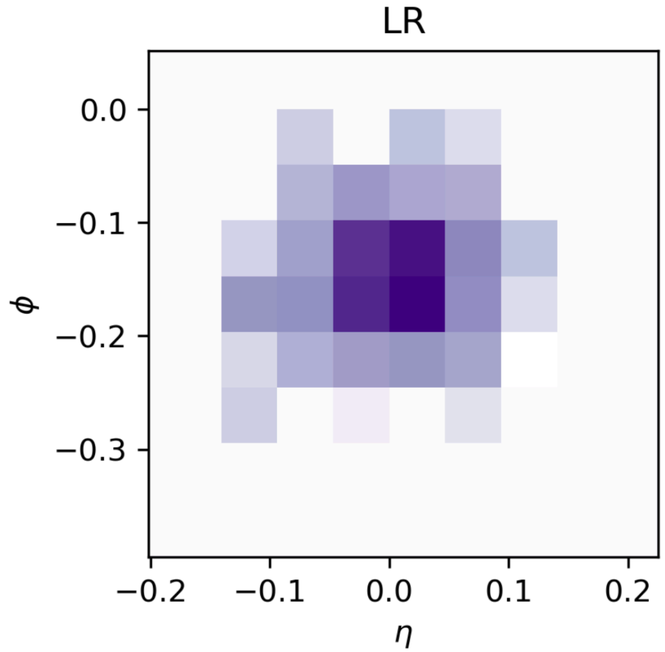

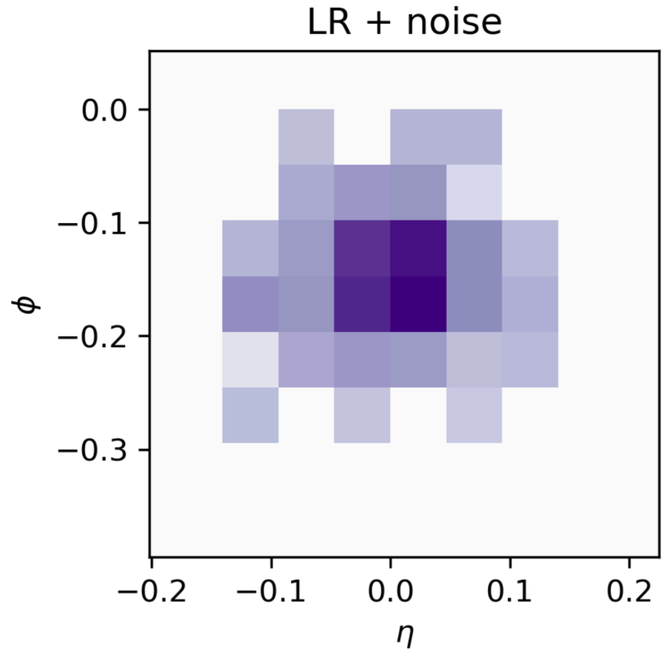

Diffusion set up

- Inspired by the SR3 paper

Image Super-Resolution via Iterative Refinement (*https://arxiv.org/pdf/2104.07636.pdf*)

Diffusion set up

- Inspired by the SR3 paper

Image Super-Resolution via Iterative Refinement (*https://arxiv.org/pdf/2104.07636.pdf*)

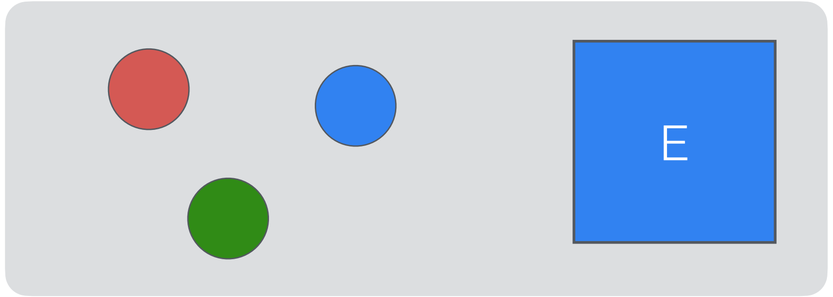

Diffusion set up

- Inspired by the SR3 paper

Image Super-Resolution via Iterative Refinement (*https://arxiv.org/pdf/2104.07636.pdf*)

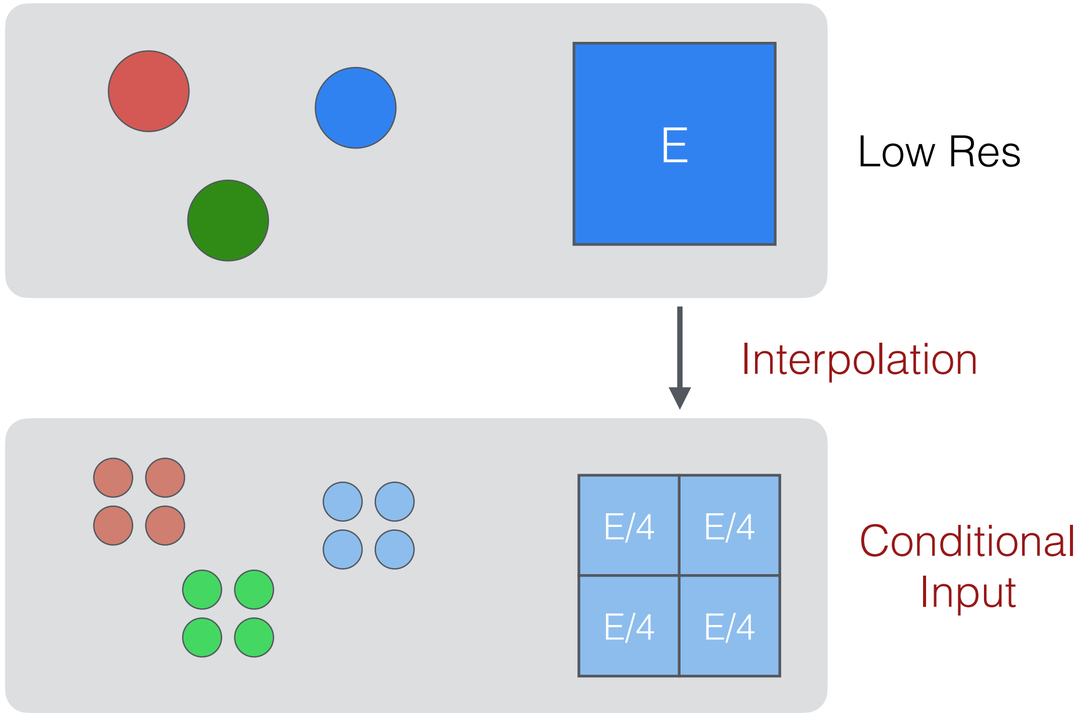

Diffusion set up

- Inspired by the SR3 paper

Image Super-Resolution via Iterative Refinement (*https://arxiv.org/pdf/2104.07636.pdf*)

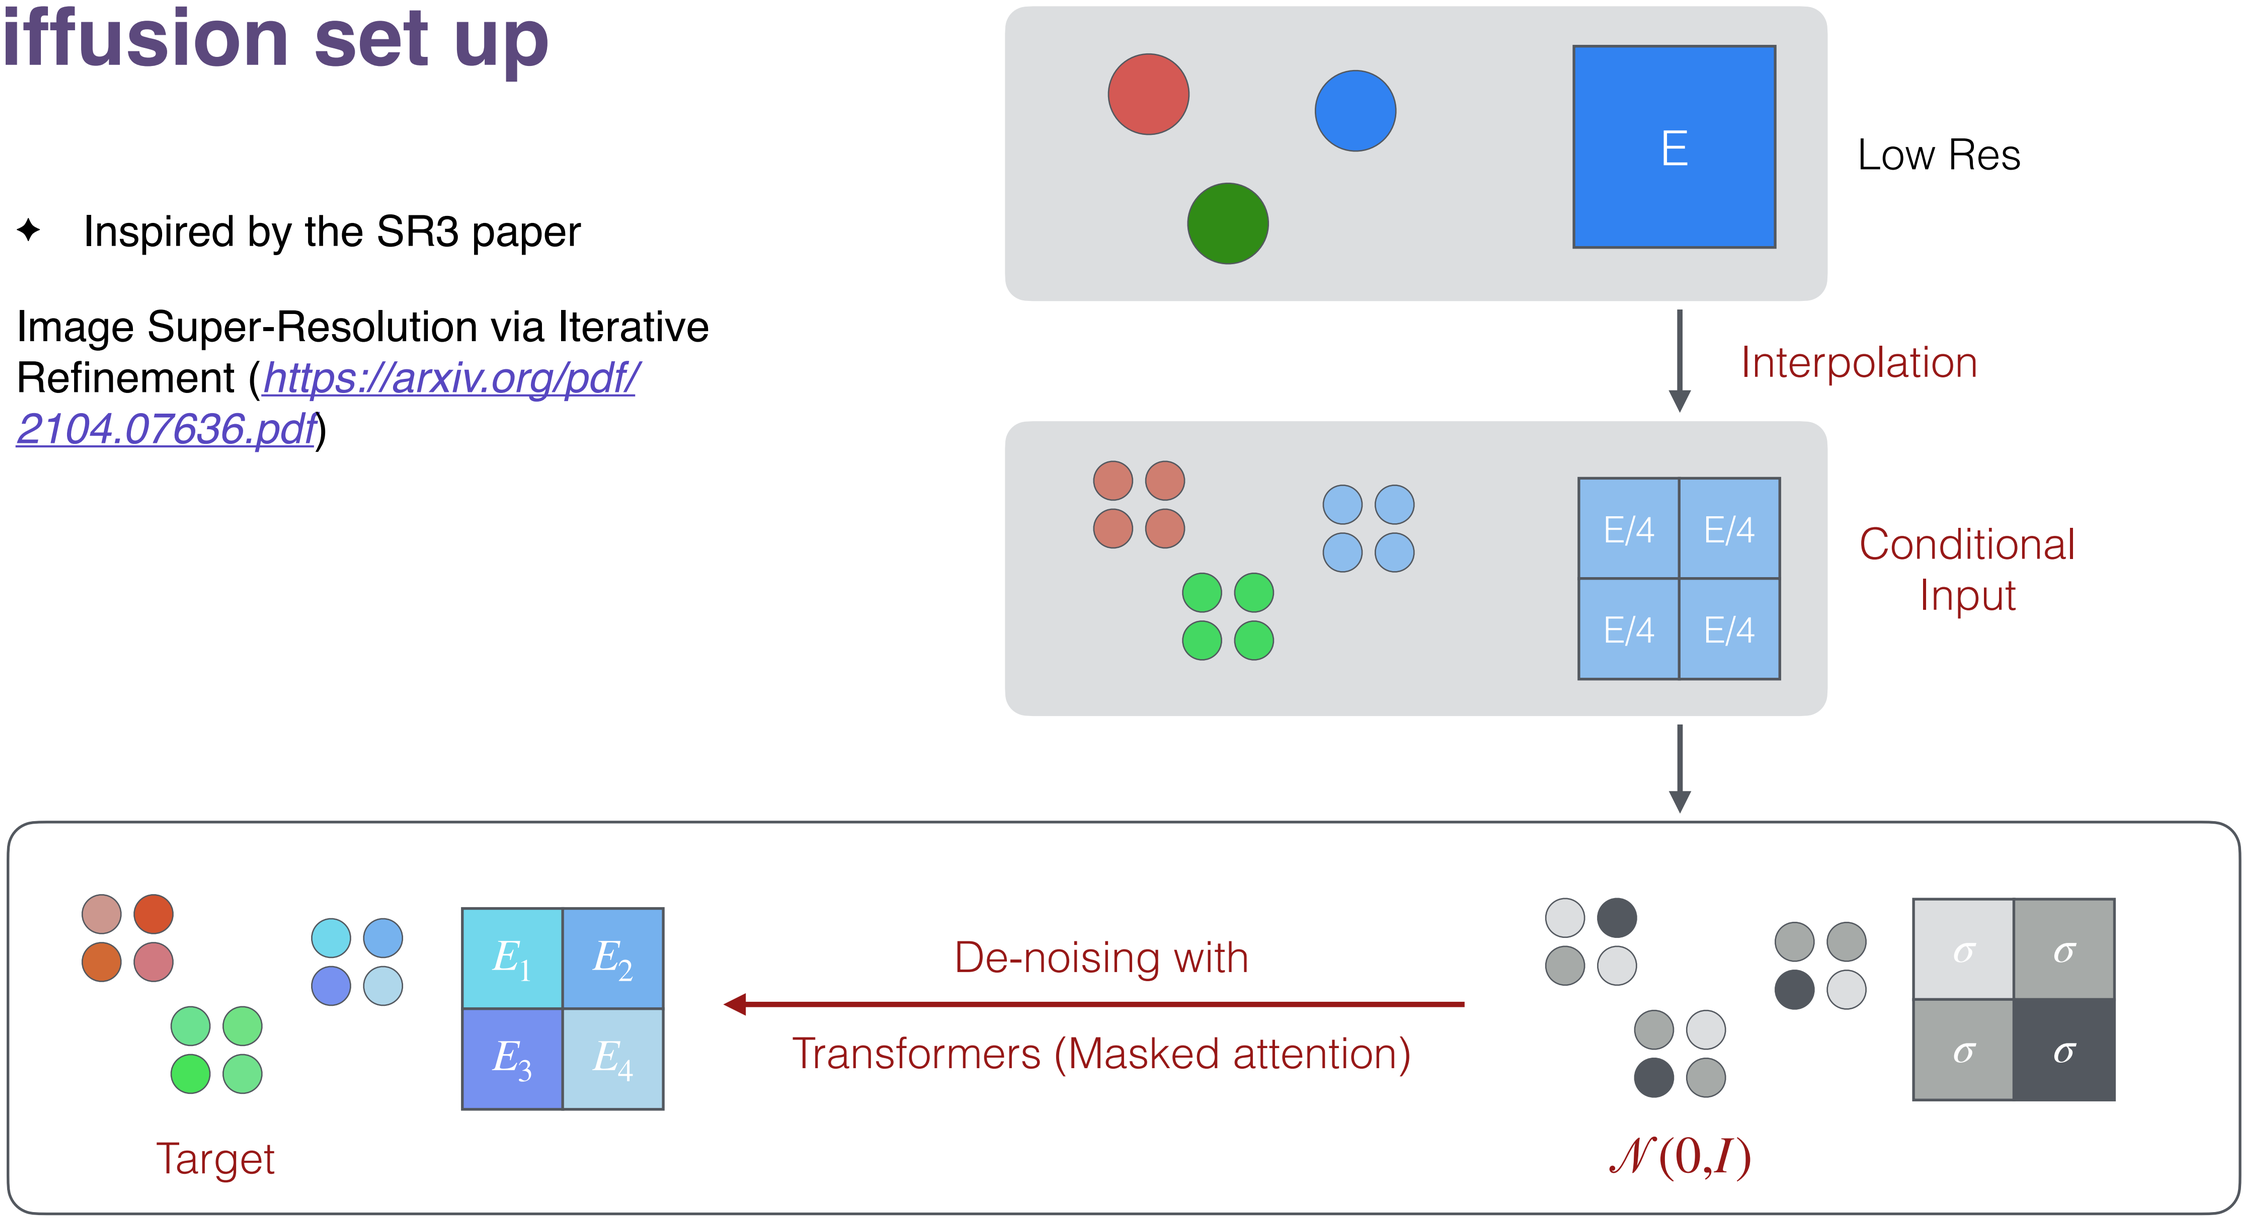

Sneak peak into the results

LR HR (target)

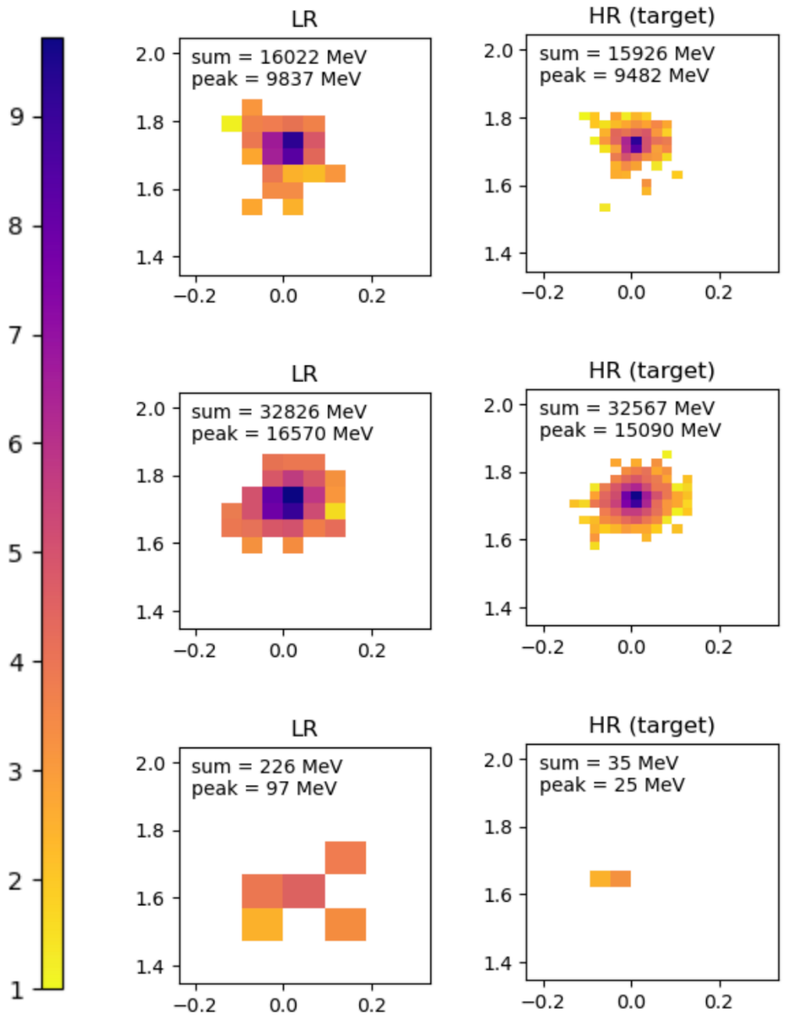

Sneak peak into the results

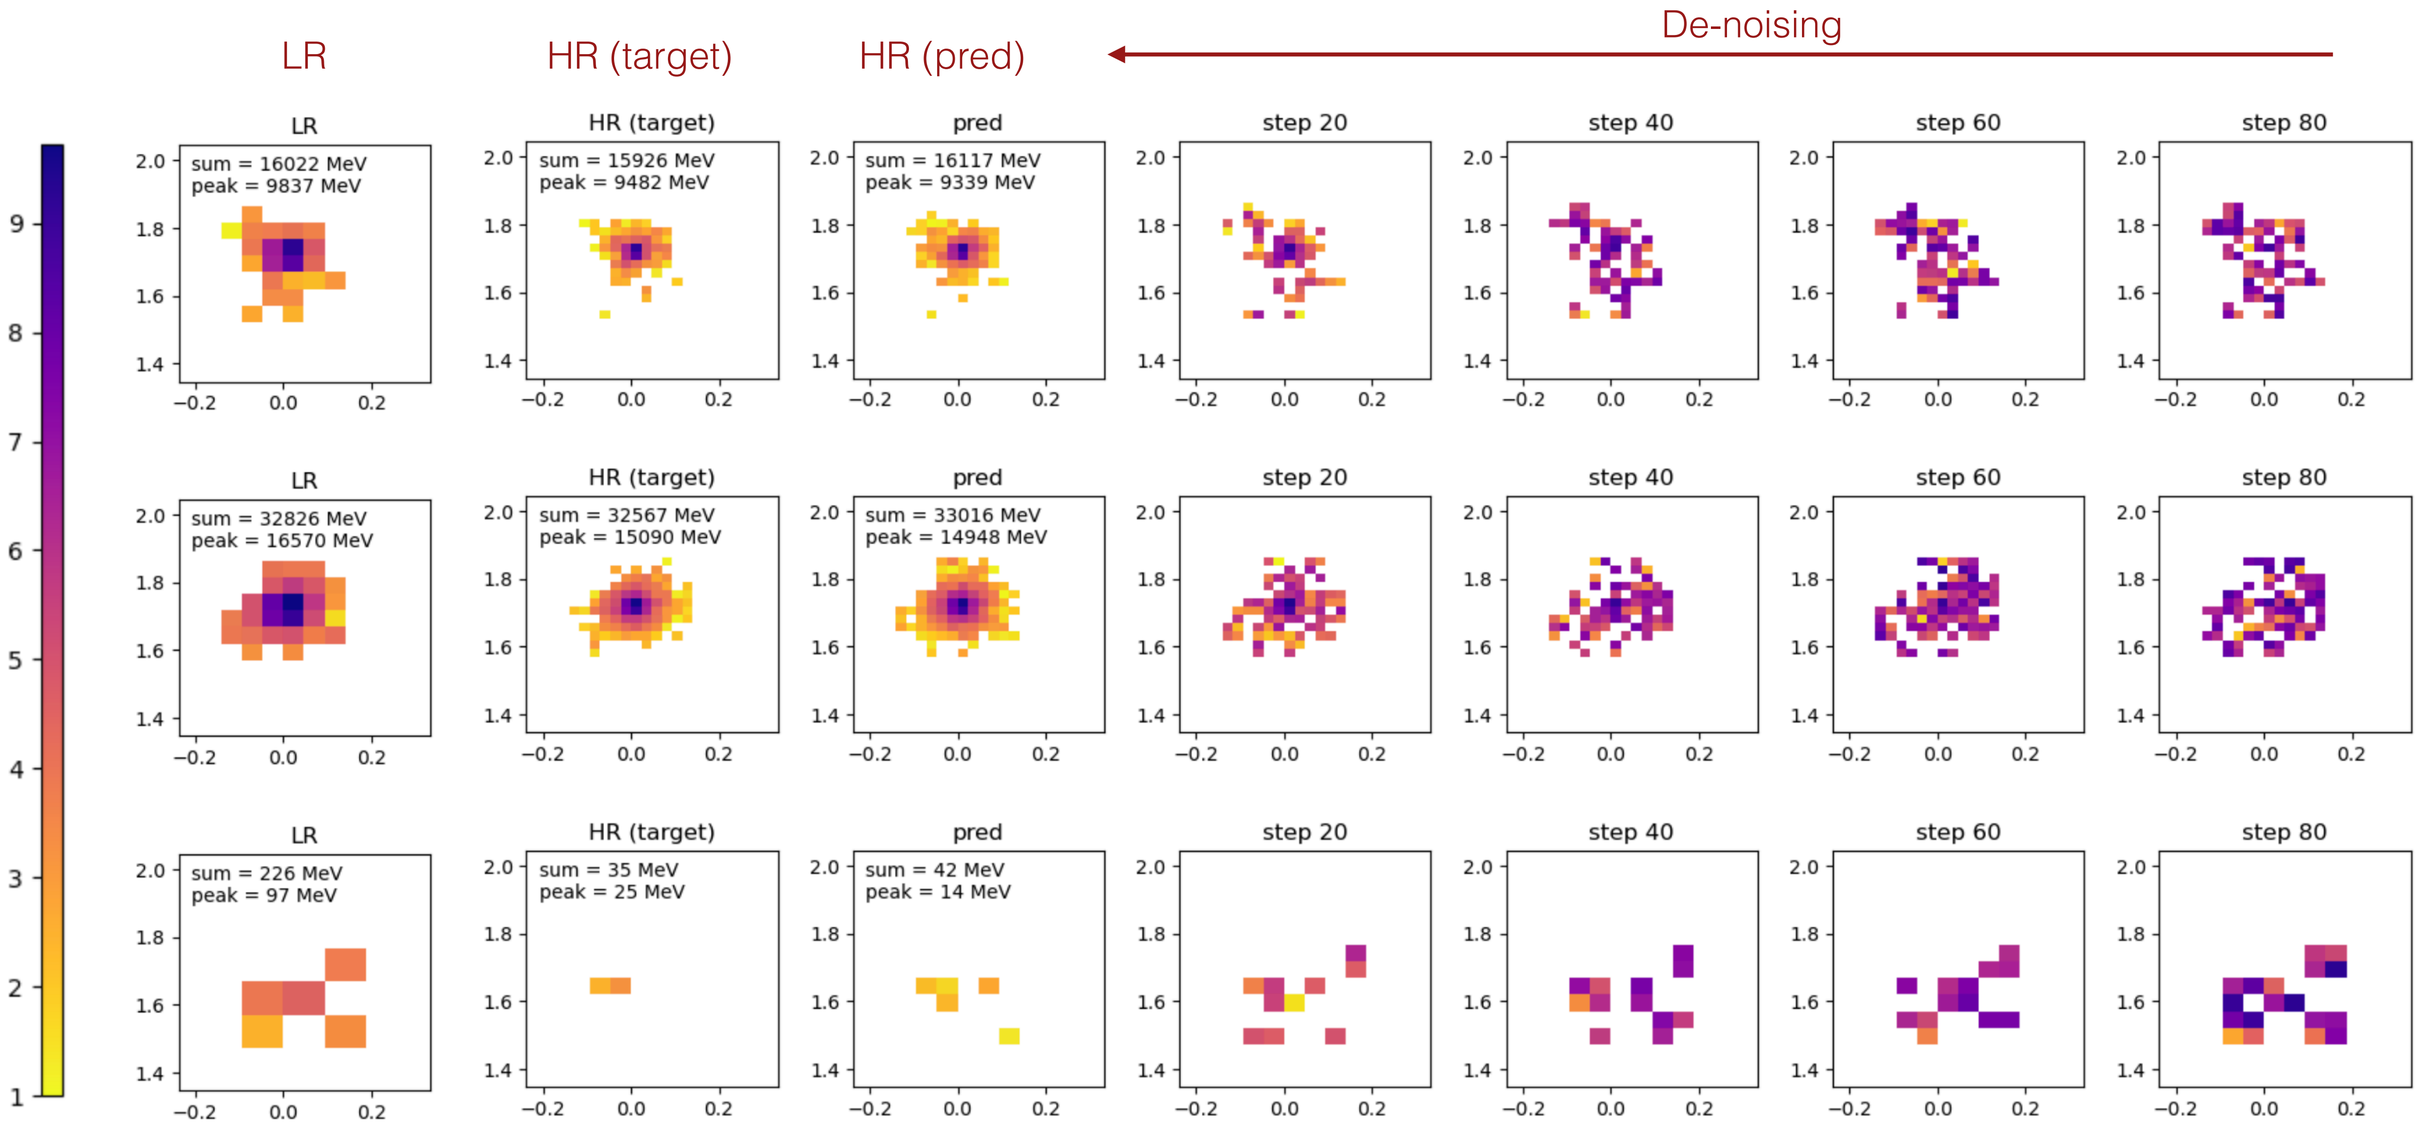

Sneak peak into the results

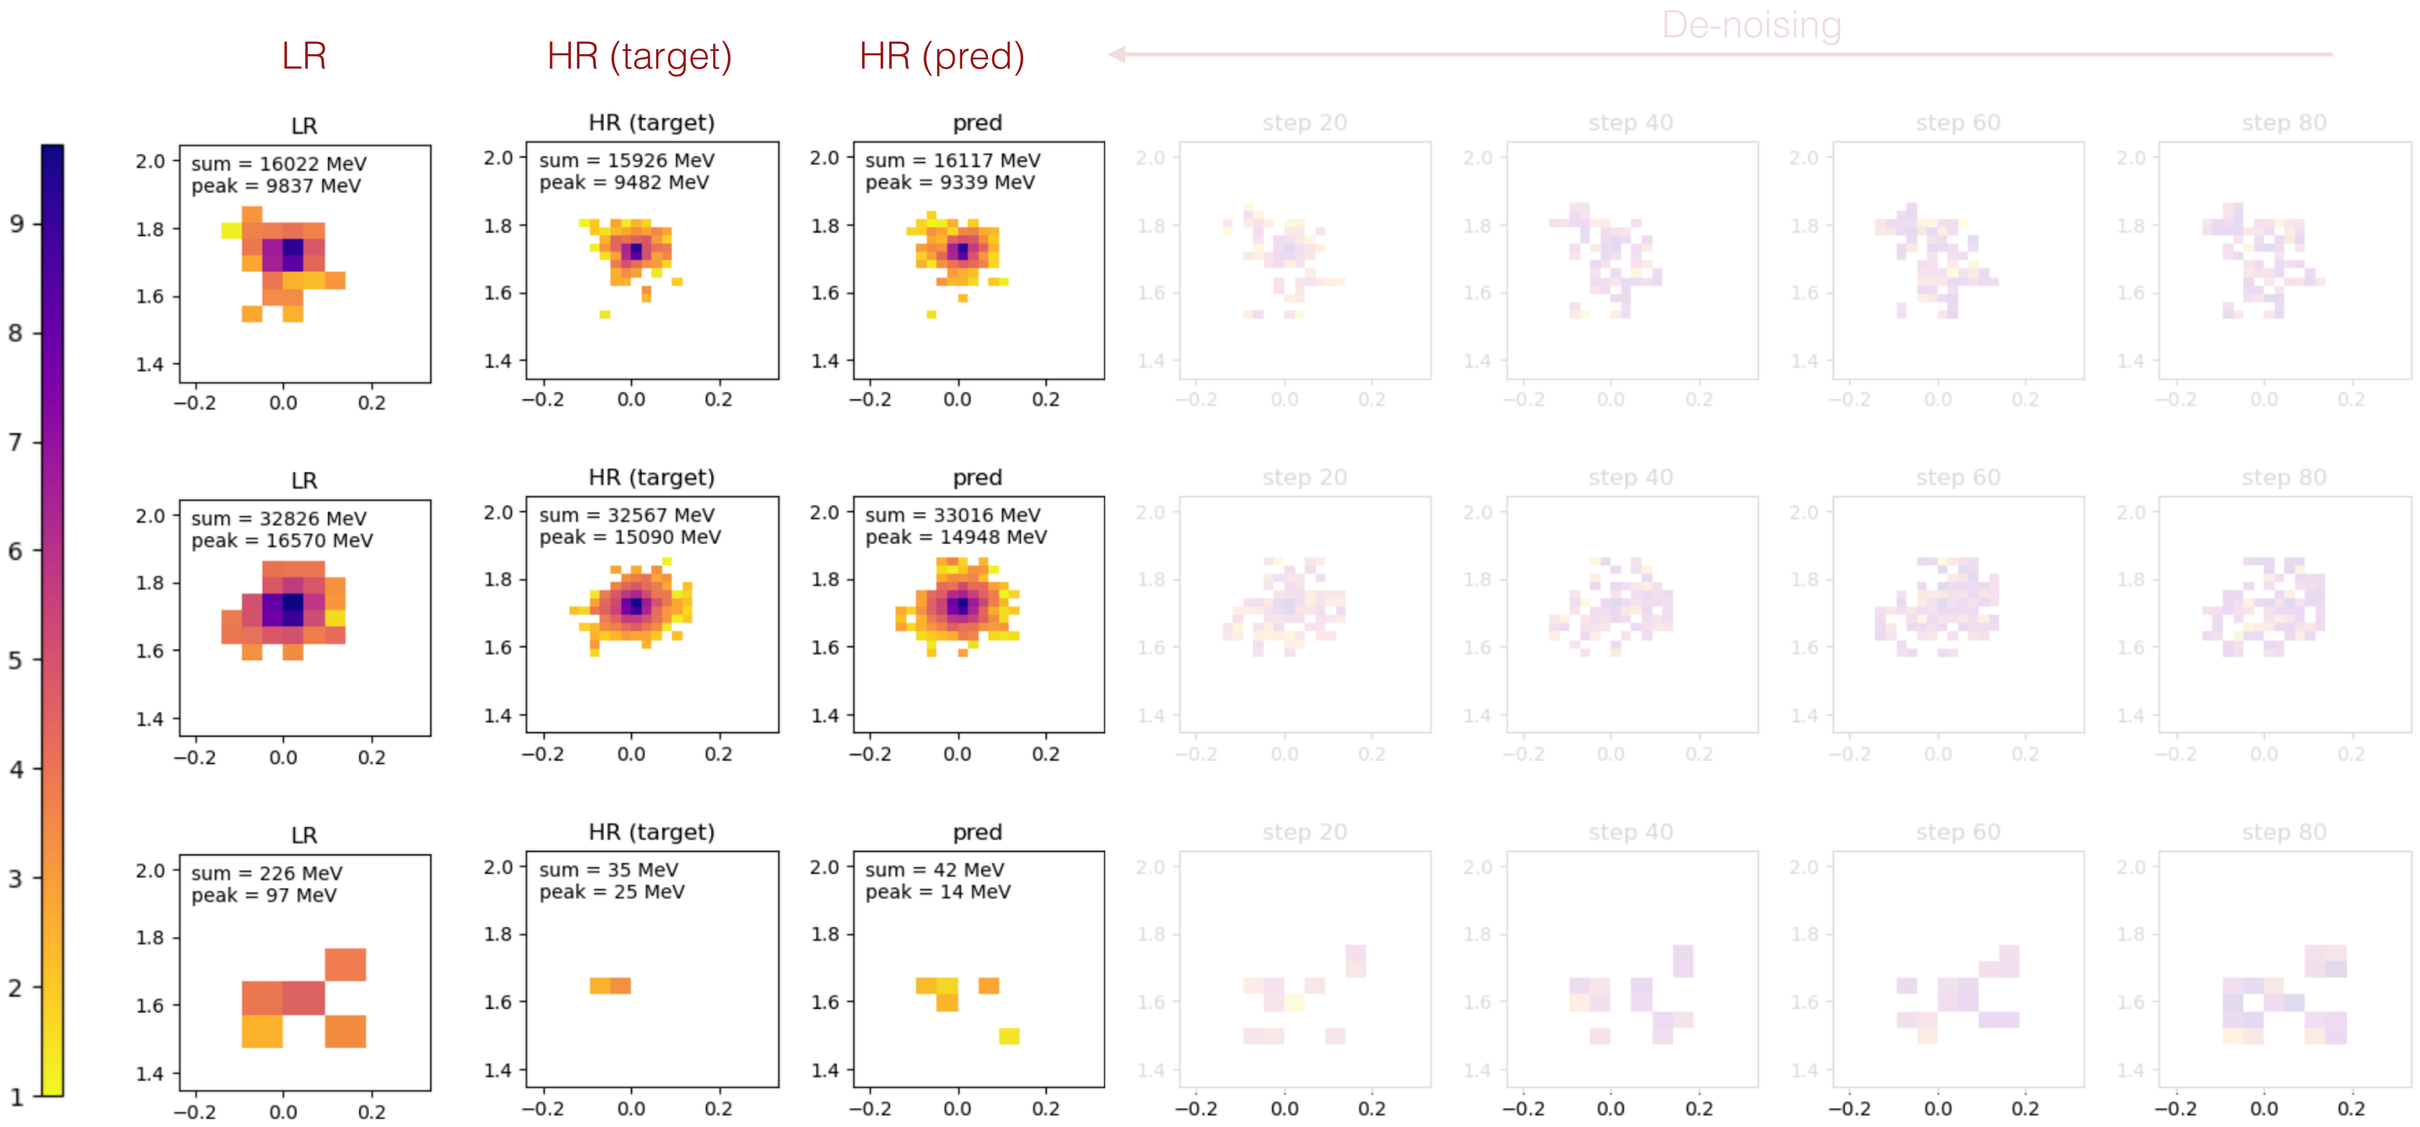

De-noising

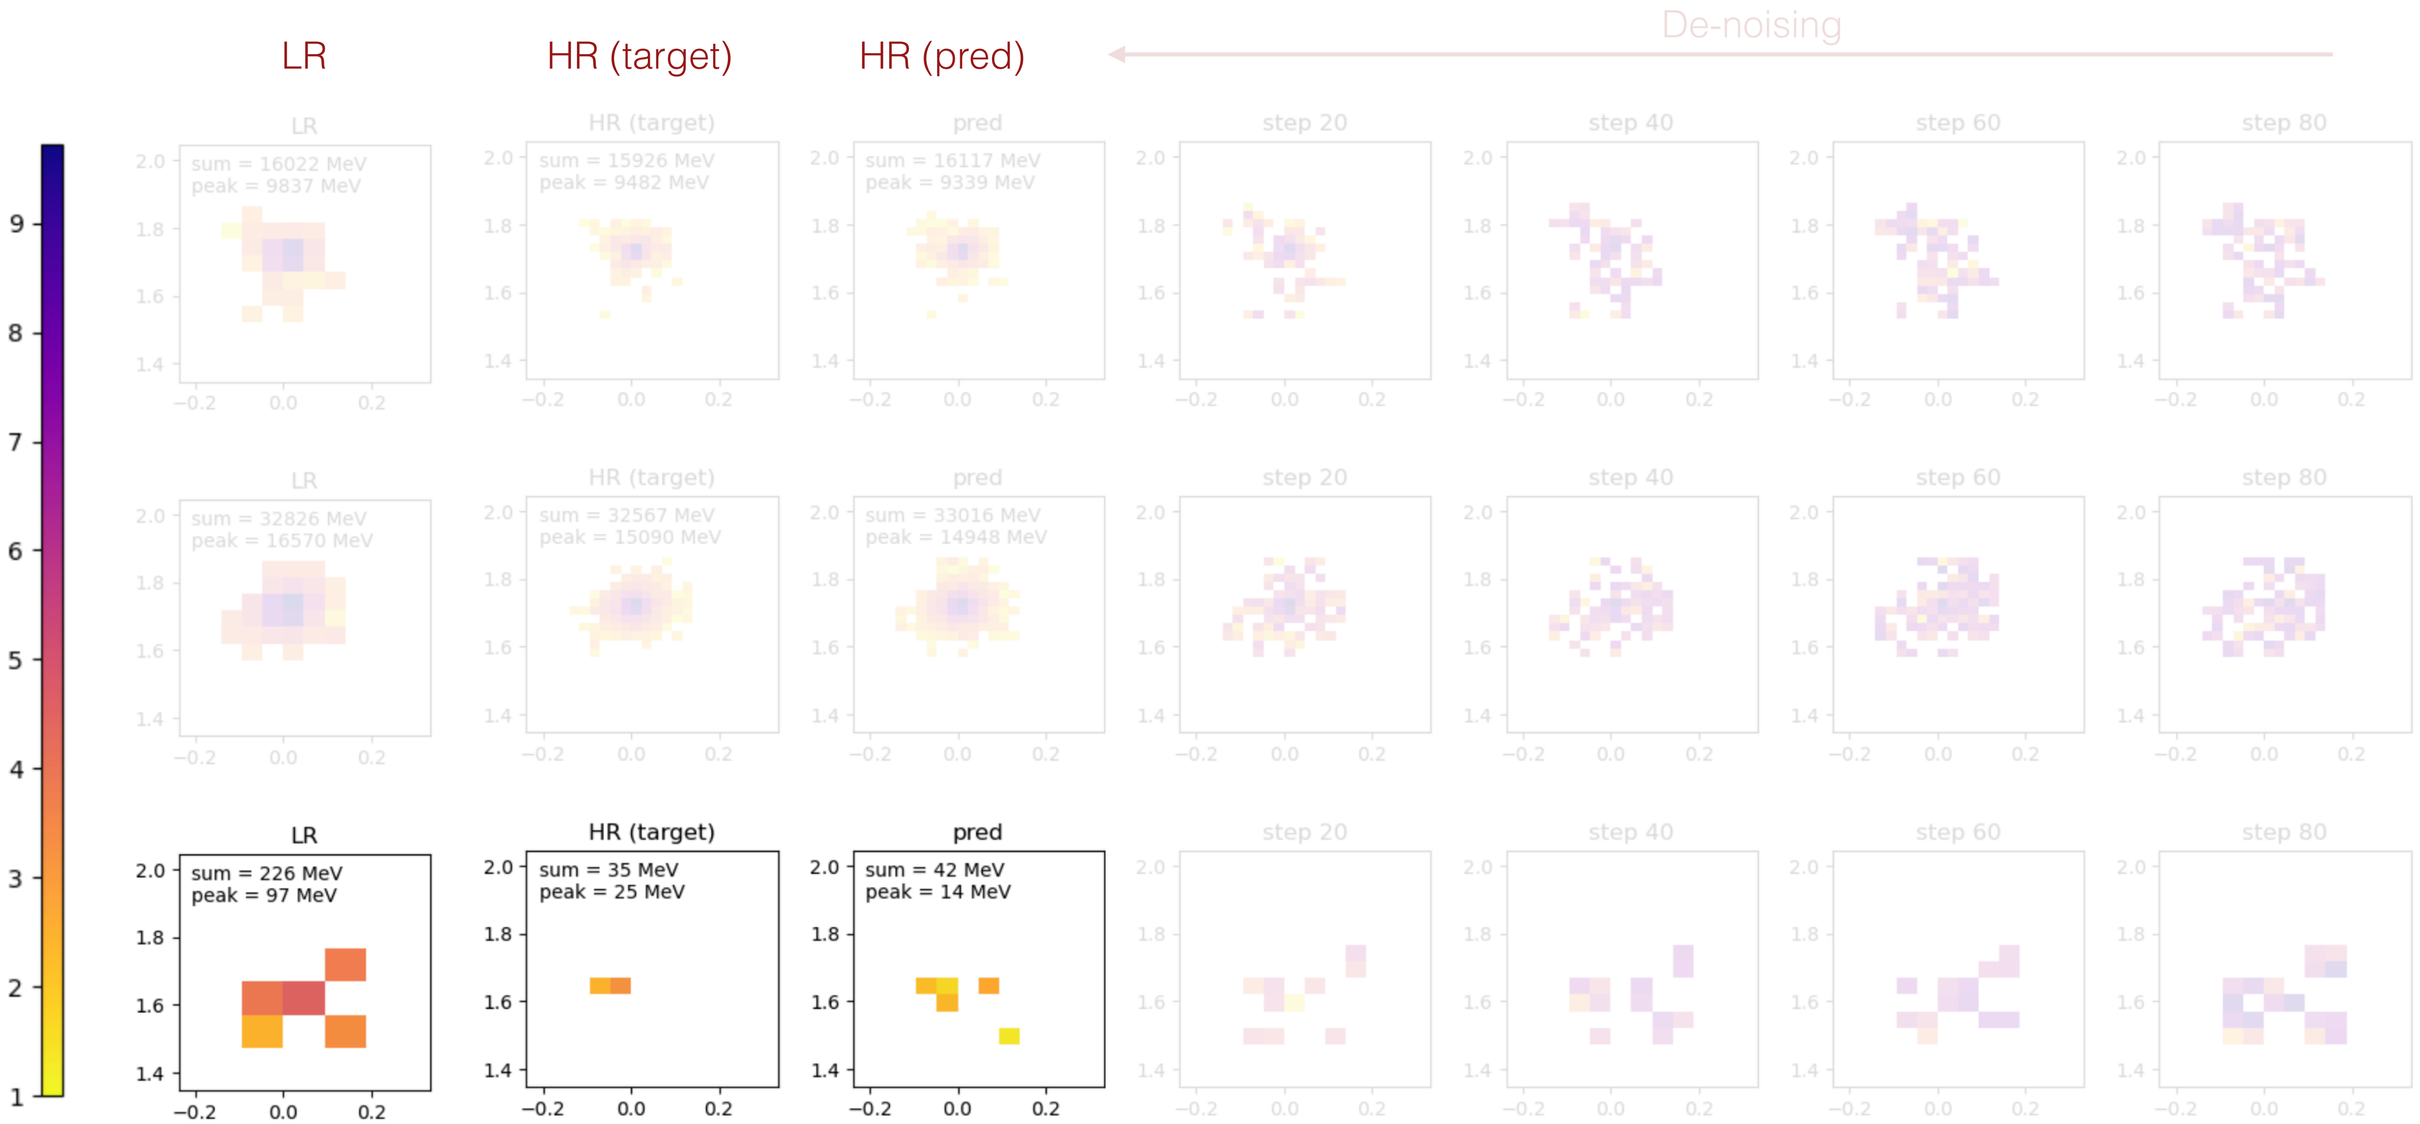
LR
HR (target)
HR (pred)
De-noising
LR
sum = 16022 MeV
peak = 9837 MeV
HR (target)
sum = 15926 MeV
peak = 9482 MeV
pred
sum = 16117 MeV
peak = 9339 MeV
step 20
step 40
step 60
step 80
LR
sum = 32826 MeV
peak = 16570 MeV
HR (target)
sum = 32567 MeV
peak = 15090 MeV
pred
sum = 33016 MeV
peak = 14948 MeV
step 20
step 40
step 60
step 80
LR
sum = 226 MeV
peak = 97 MeV
HR (target)
sum = 35 MeV
peak = 25 MeV
pred
sum = 42 MeV
peak = 14 MeV
step 20
step 40
step 60
step 80

A more Interesting case!

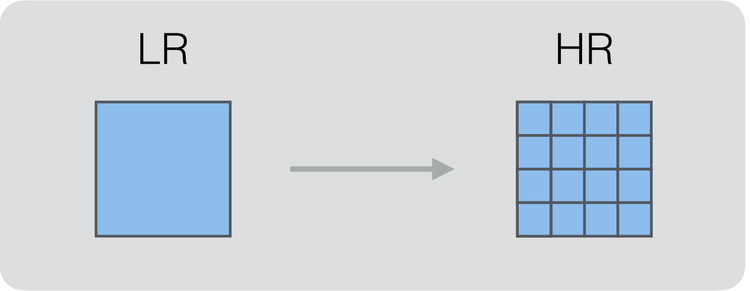

- Multiple particles
- 1-5 particles
- Electrons and photons

Overtraining!
(Couldn't get the full training ready in time)

A more Interesting case!

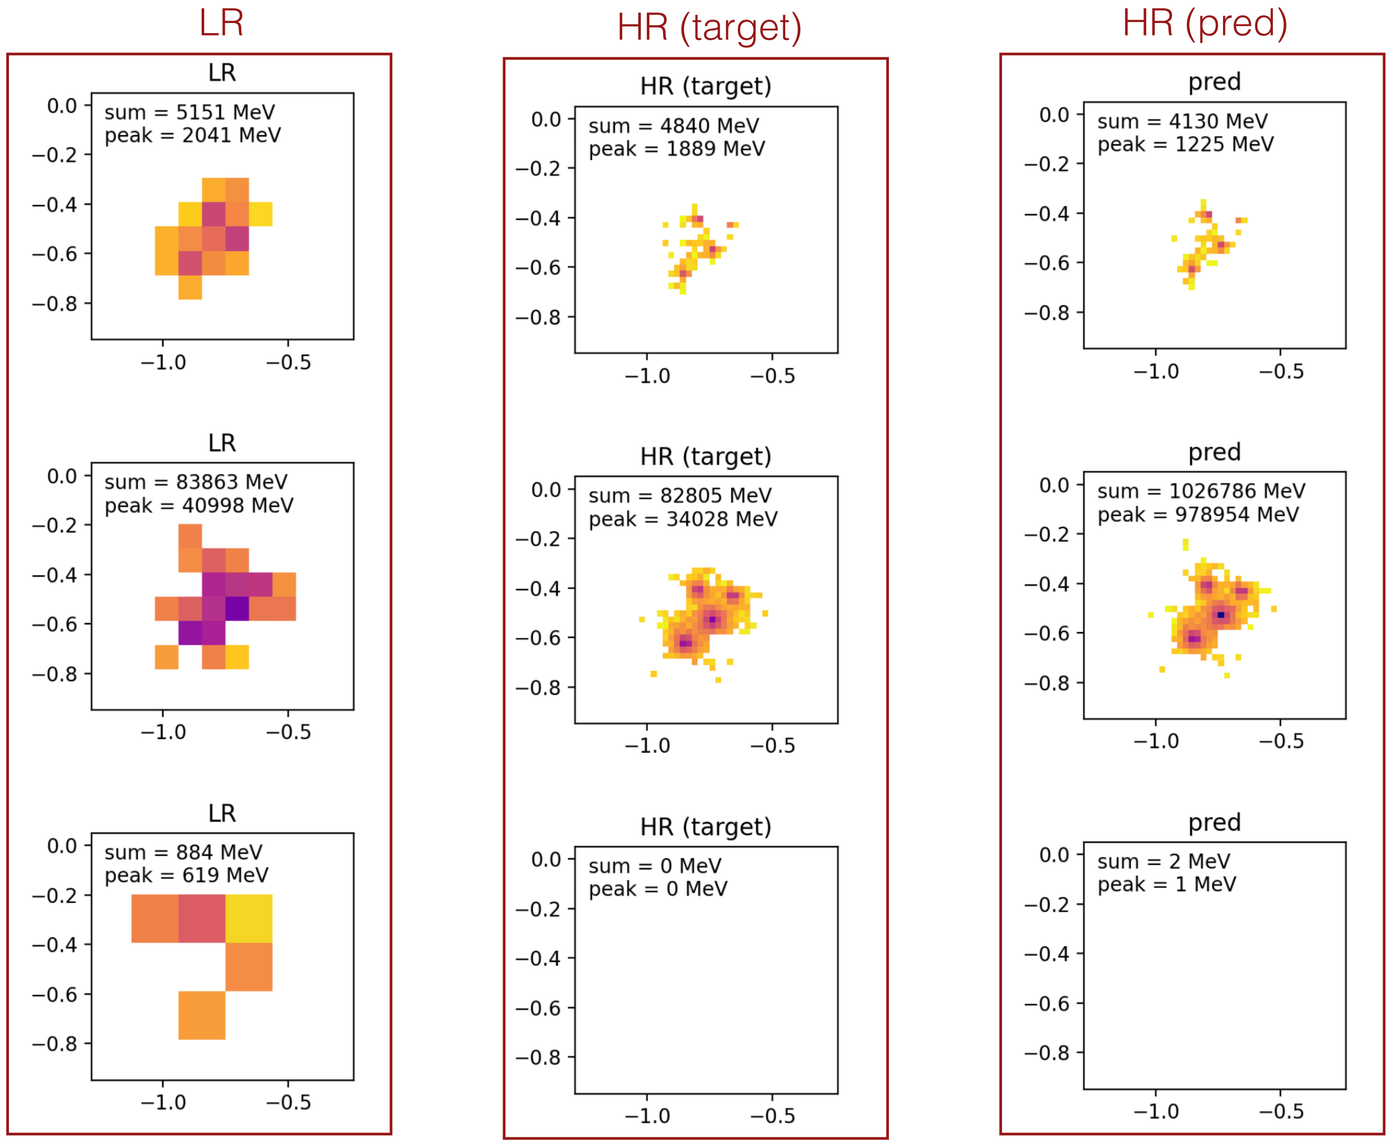

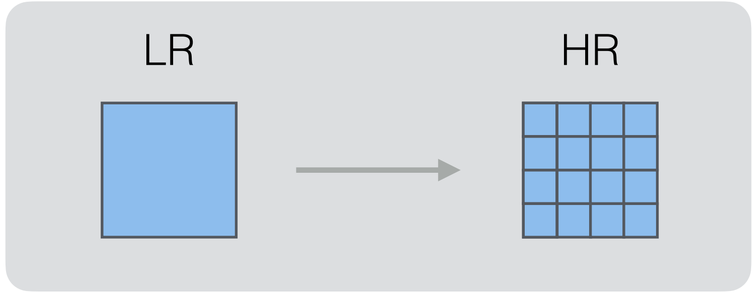

- Multiple particles
- 1-5 particles
- Electrons and photons

Overtraining!
(Couldn't get the full training ready in time)

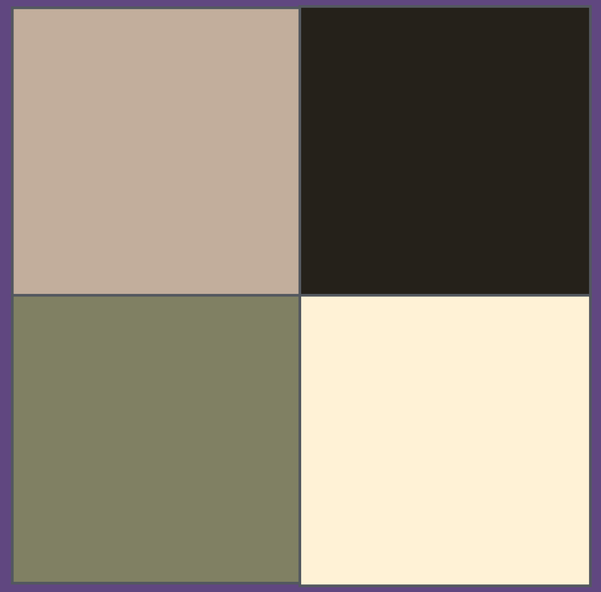

Is it hallucination?

Creating information out of nowhere?

Where does the extra info come from?

Where does the extra info come from?

- From training data!
  - ➡ Energy deposition, by let's say a photon, is not random
  - ➡ Model can learn the HR distribution conditioned on the LR distribution

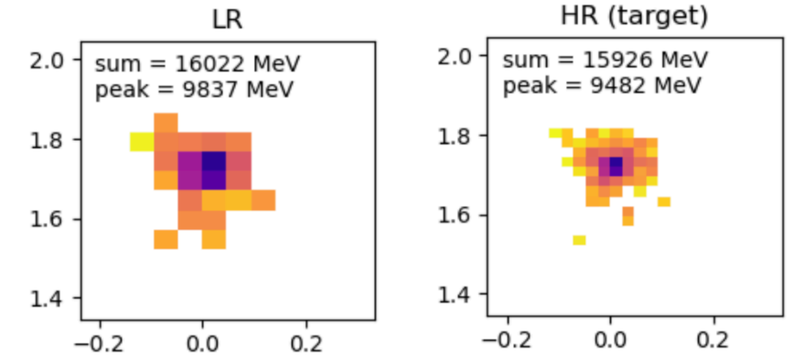

Where does the extra info come from?

- From training data!
  - ➡ Energy deposition, by let's say a photon, is not random
  - ➡ Model can learn the HR distribution conditioned on the LR distribution

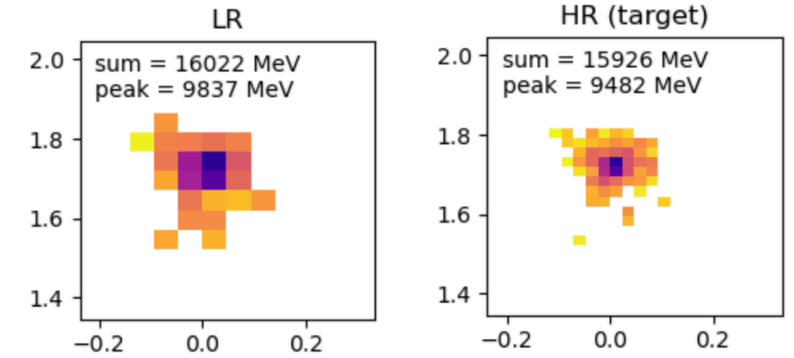

- HR output = Educated estimation of the model based on the patterns learned from the training data

Where does the extra info come from?

- From training data!
  - ➡ Energy deposition, by let's say a photon, is not random
  - ➡ Model can learn the HR distribution conditioned on the LR distribution

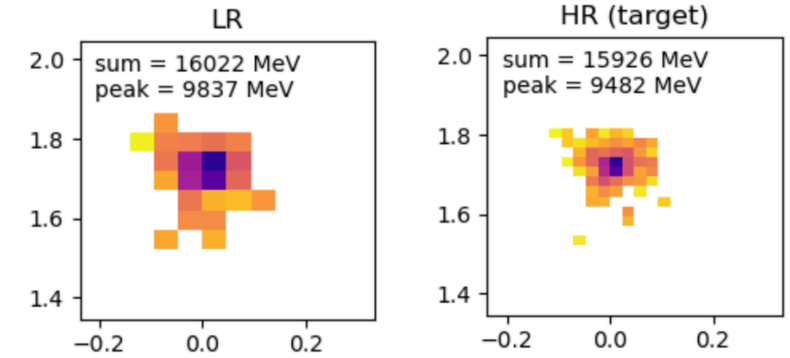

- HR output = Educated estimation of the model based on the patterns learned from the training data
  - ➡ Similar to how SR work in Computer vision

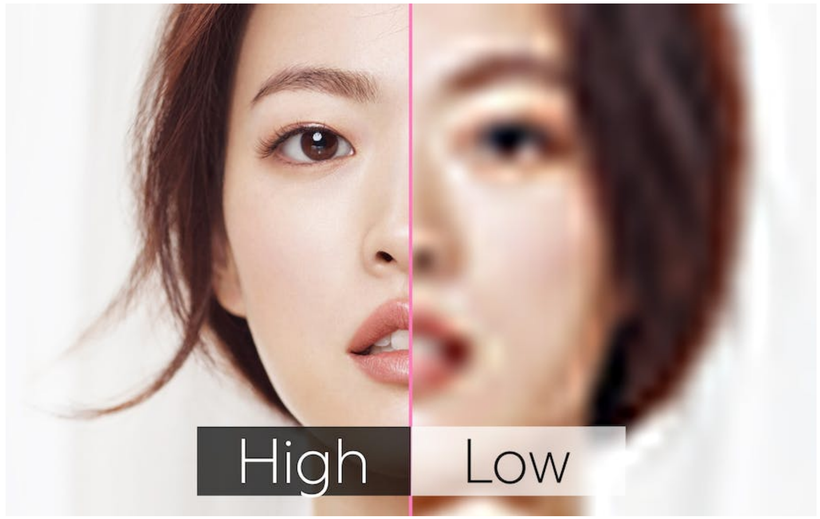

But, how do we still verify we are not hallucinating?

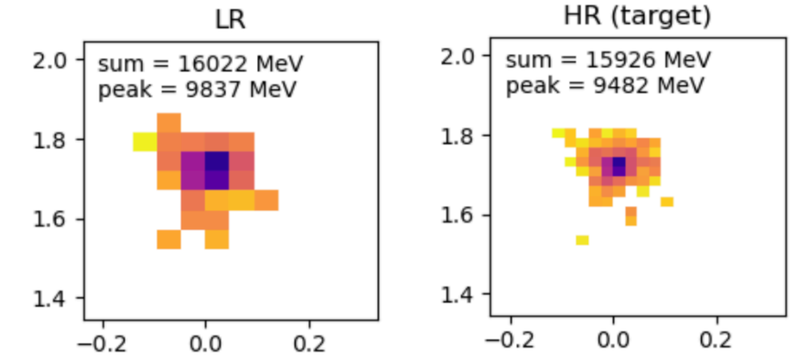

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?

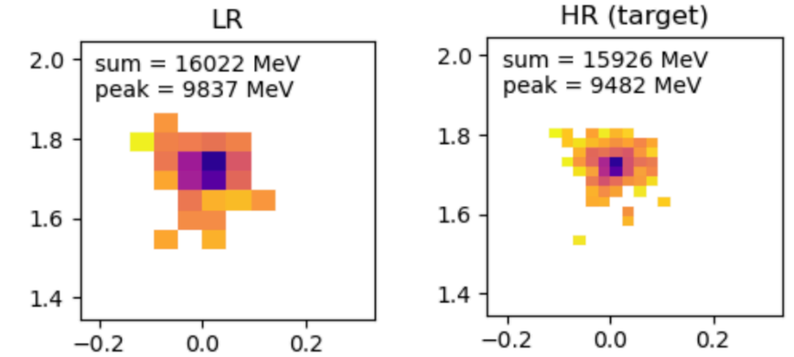

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model

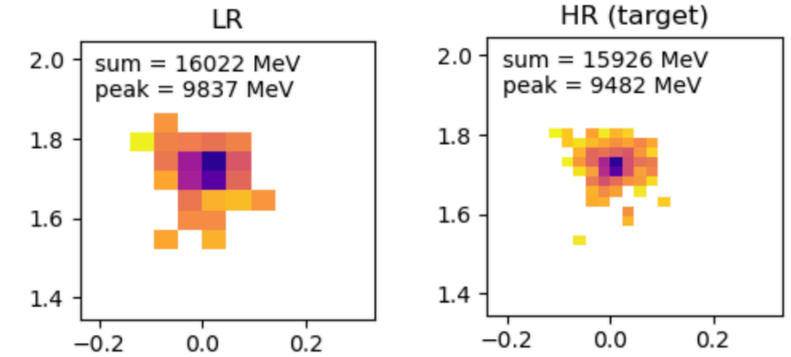

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?

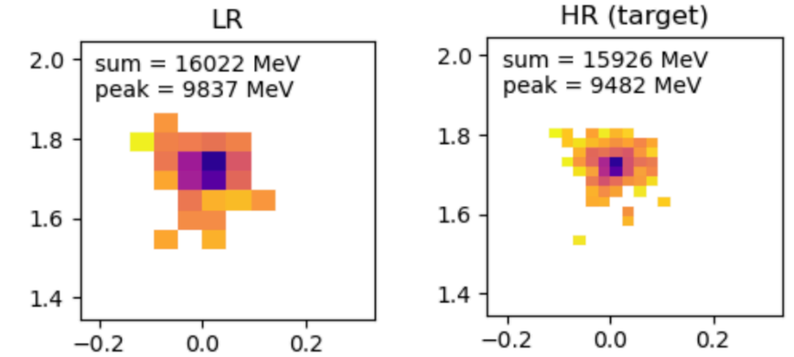

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?
- With simulation,

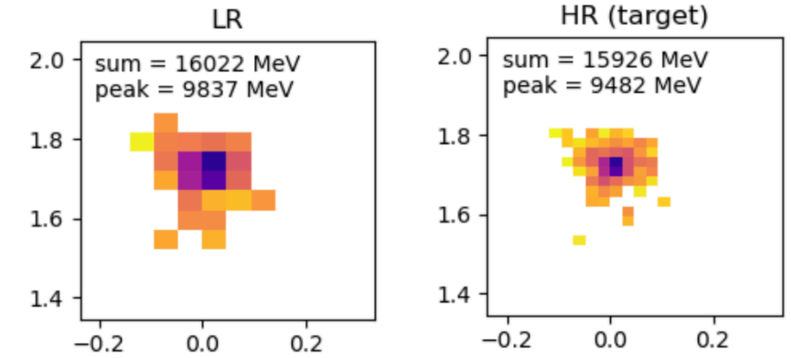

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?
- With simulation,
  - ➡ it’s easier, we can have the truth targets

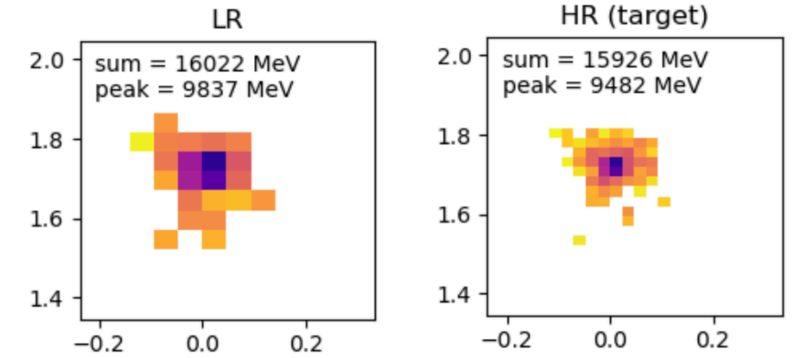

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?
- With simulation,
  - ➡ it’s easier, we can have the truth targets
- With actual data,

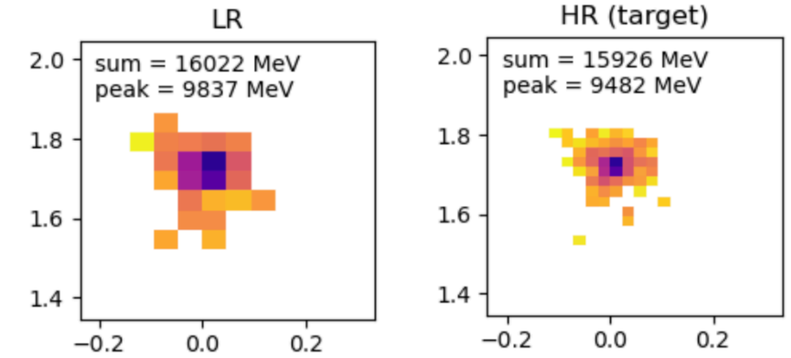

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?
- With simulation,
  - ➡ it’s easier, we can have the truth targets
- With actual data,
  - ➡ Not so easy. (How much we trust our simulations?)

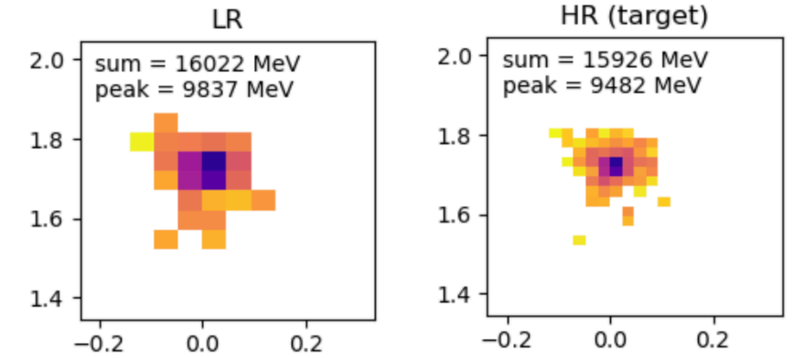

But, how do we still verify we are not hallucinating?

- Qn: The model can predict “realistic looking” outputs, but are they correct?
  - ➡ Ans: Of course, they are not perfect, just like any other ML model
- How do we estimate how good/bad are the HR estimations?
- With simulation,
  - ➡ it’s easier, we can have the truth targets
- With actual data,
  - ➡ Not so easy. (How much we trust our simulations?)
  - ➡ Calibration problem (tricky, but I believe doable)

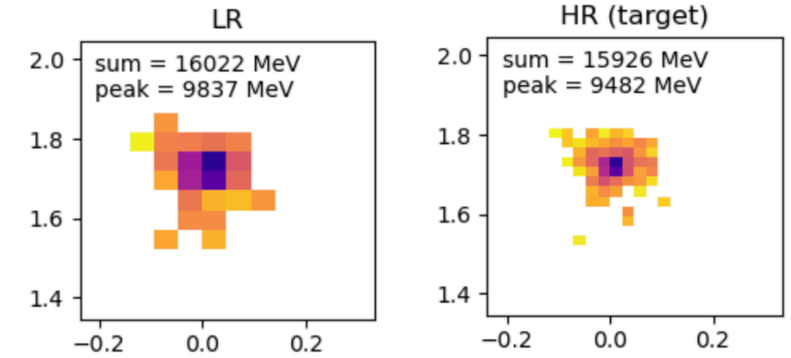

But, most importantly,

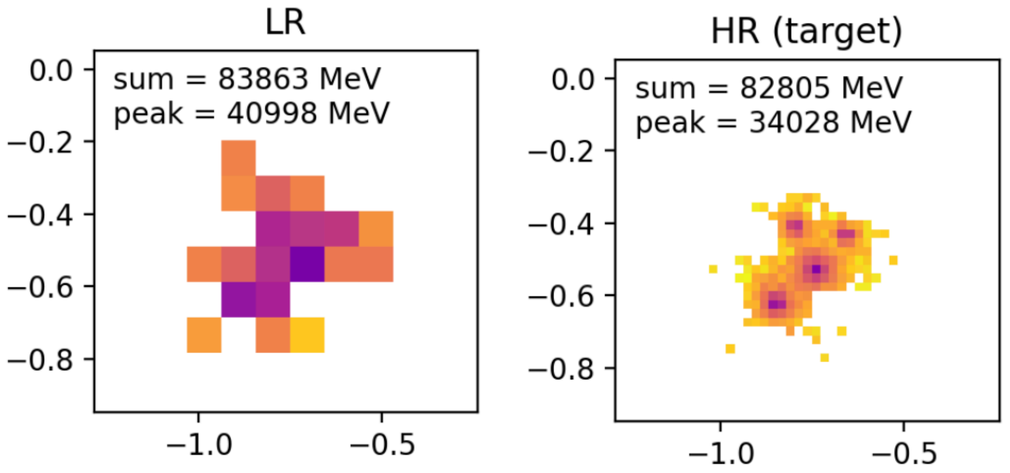

But, most importantly,

- We shouldn’t look at it in isolation

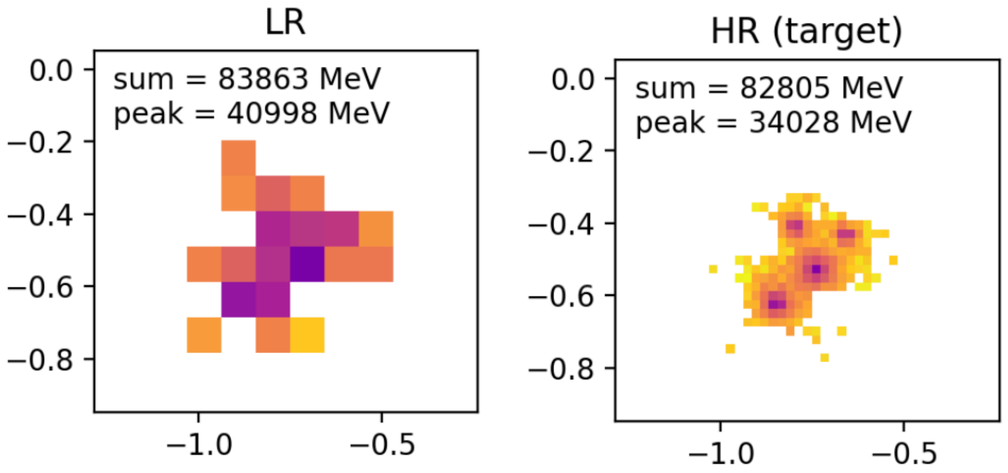

But, most importantly,

- We shouldn't look at it in isolation
- Primary goal -
  - Assist downstream reconstruction task
  - Multi-particle example demonstrates cardinality
  - Plenty more to look at (ongoing study)

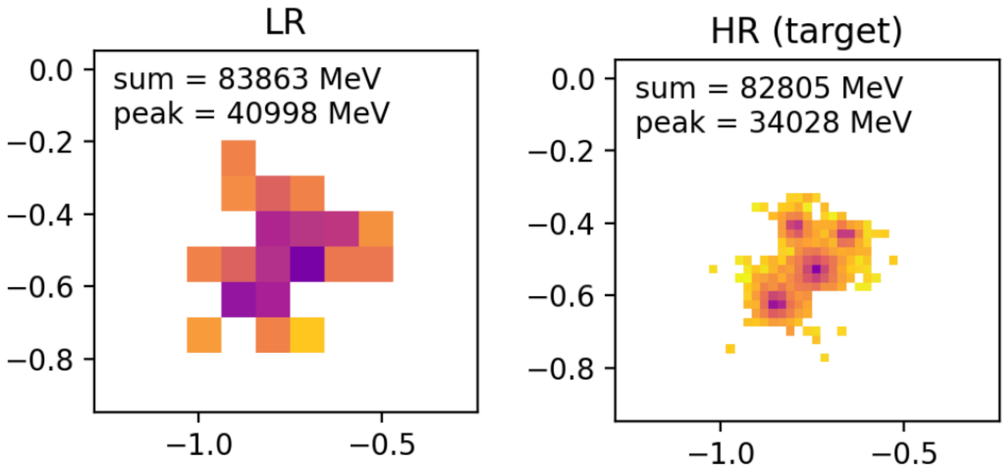

But, most importantly,

- We shouldn't look at it in isolation
- Primary goal -
  - Assist downstream reconstruction task
  - Multi-particle example demonstrates cardinality
  - Plenty more to look at (ongoing study)

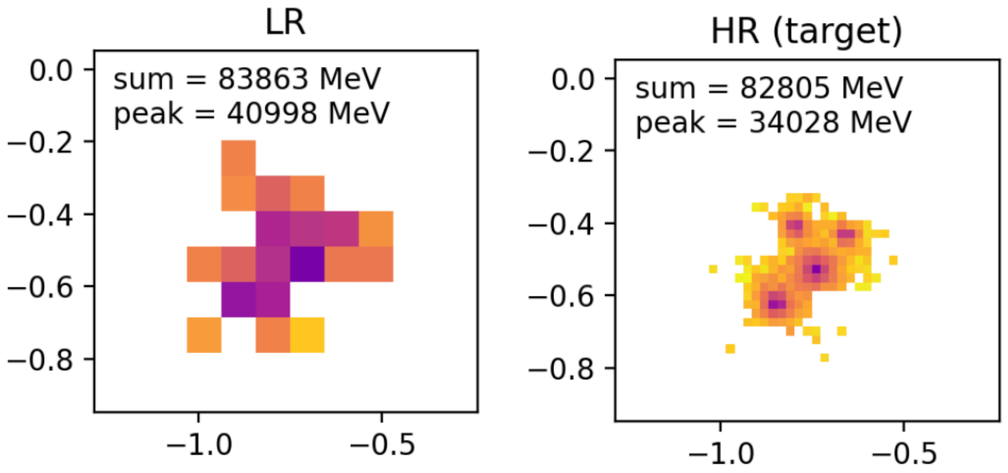

Thanks!