

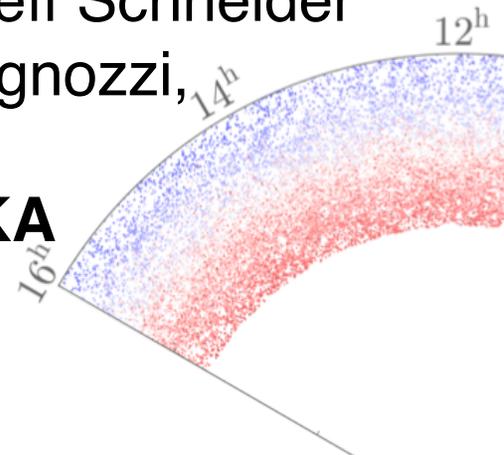
# Learning Science with Machine Learning: Opening the Pandora Box

**Shirley Ho**

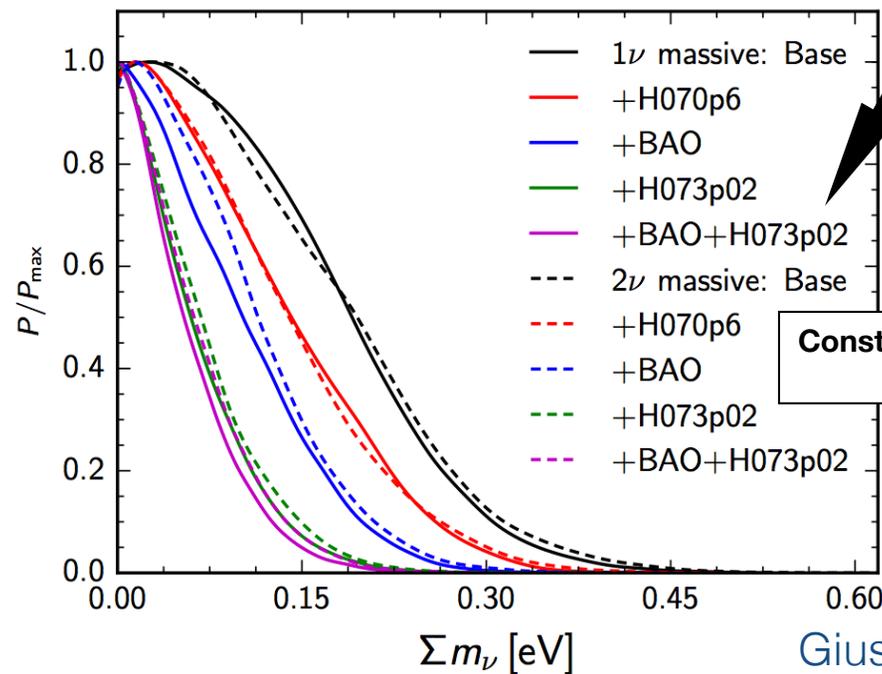
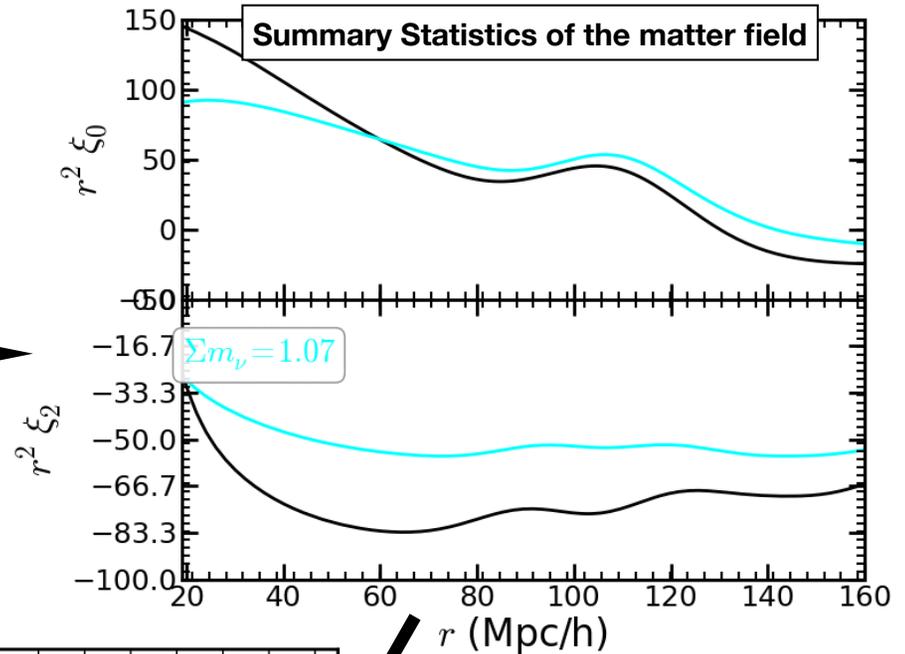
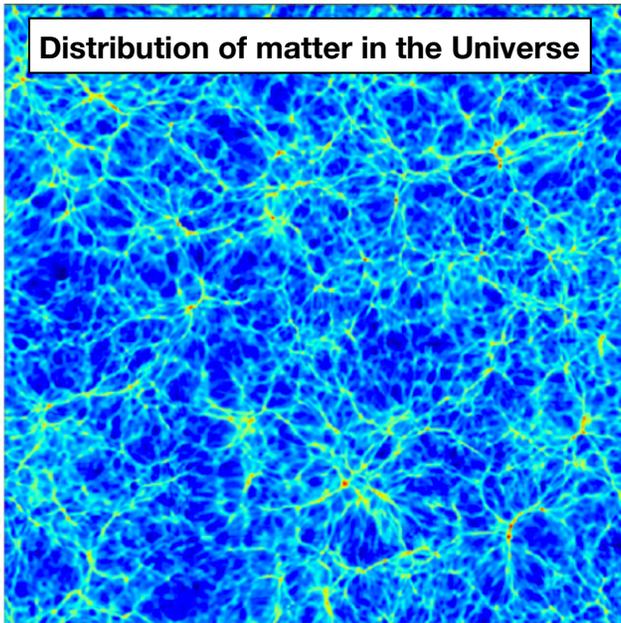
CMU/Berkeley-> Flatiron Institute/Princeton University

**Siyu He (Flatiron/CMU), Yin Li (Berkeley), Yu Feng (Berkeley),**  
Siamak Ravanbakhsh, Barnabas Poczos, Elena Giusarma,  
Emmanuel Schaan, Simone Ferraro, Junier Oliver, Jeff Schneider  
Layne Price, Sebastian Fromenteau, Sunny Vagnozzi,  
Katie Freese, Kam-Biu Luk

**Artificial Intelligence with CERN and SKA**  
**Alan Turing Institute, 2018**

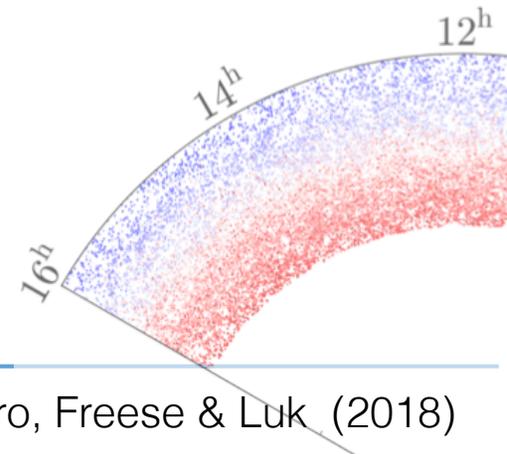
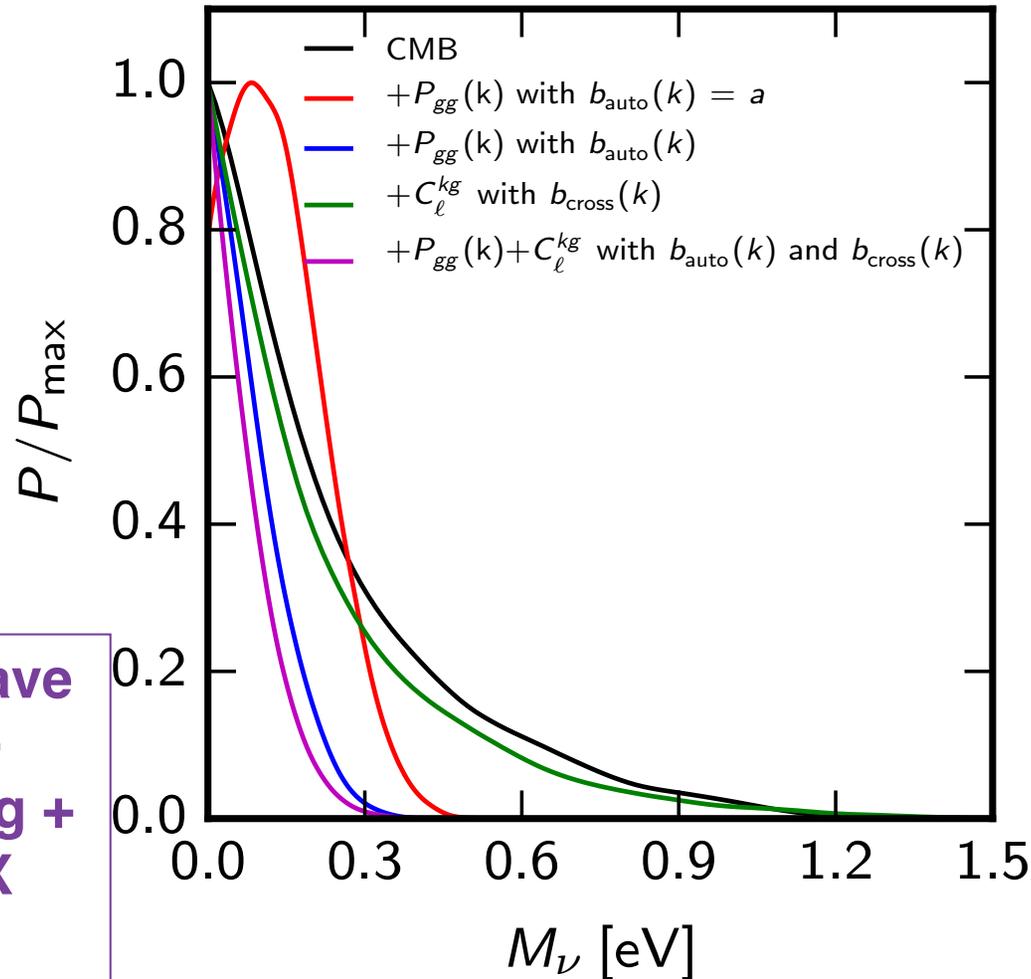


# Traditional Cosmology analysis



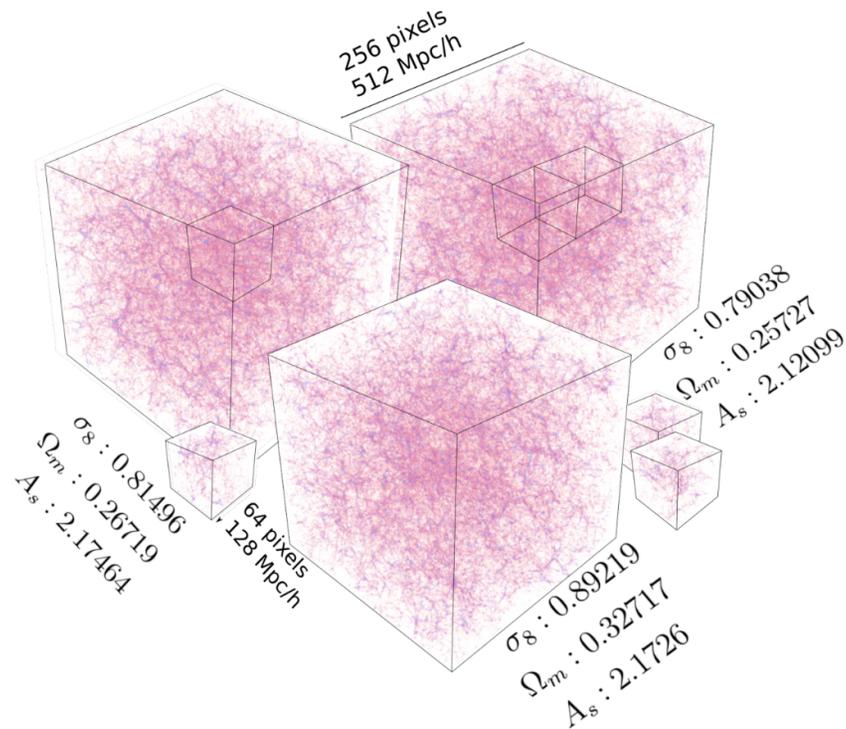
Constraints on important parameters such as neutrino masses

# Standard Cosmological Analysis: Use summary statistics to constrain physical parameters such as neutrino masses

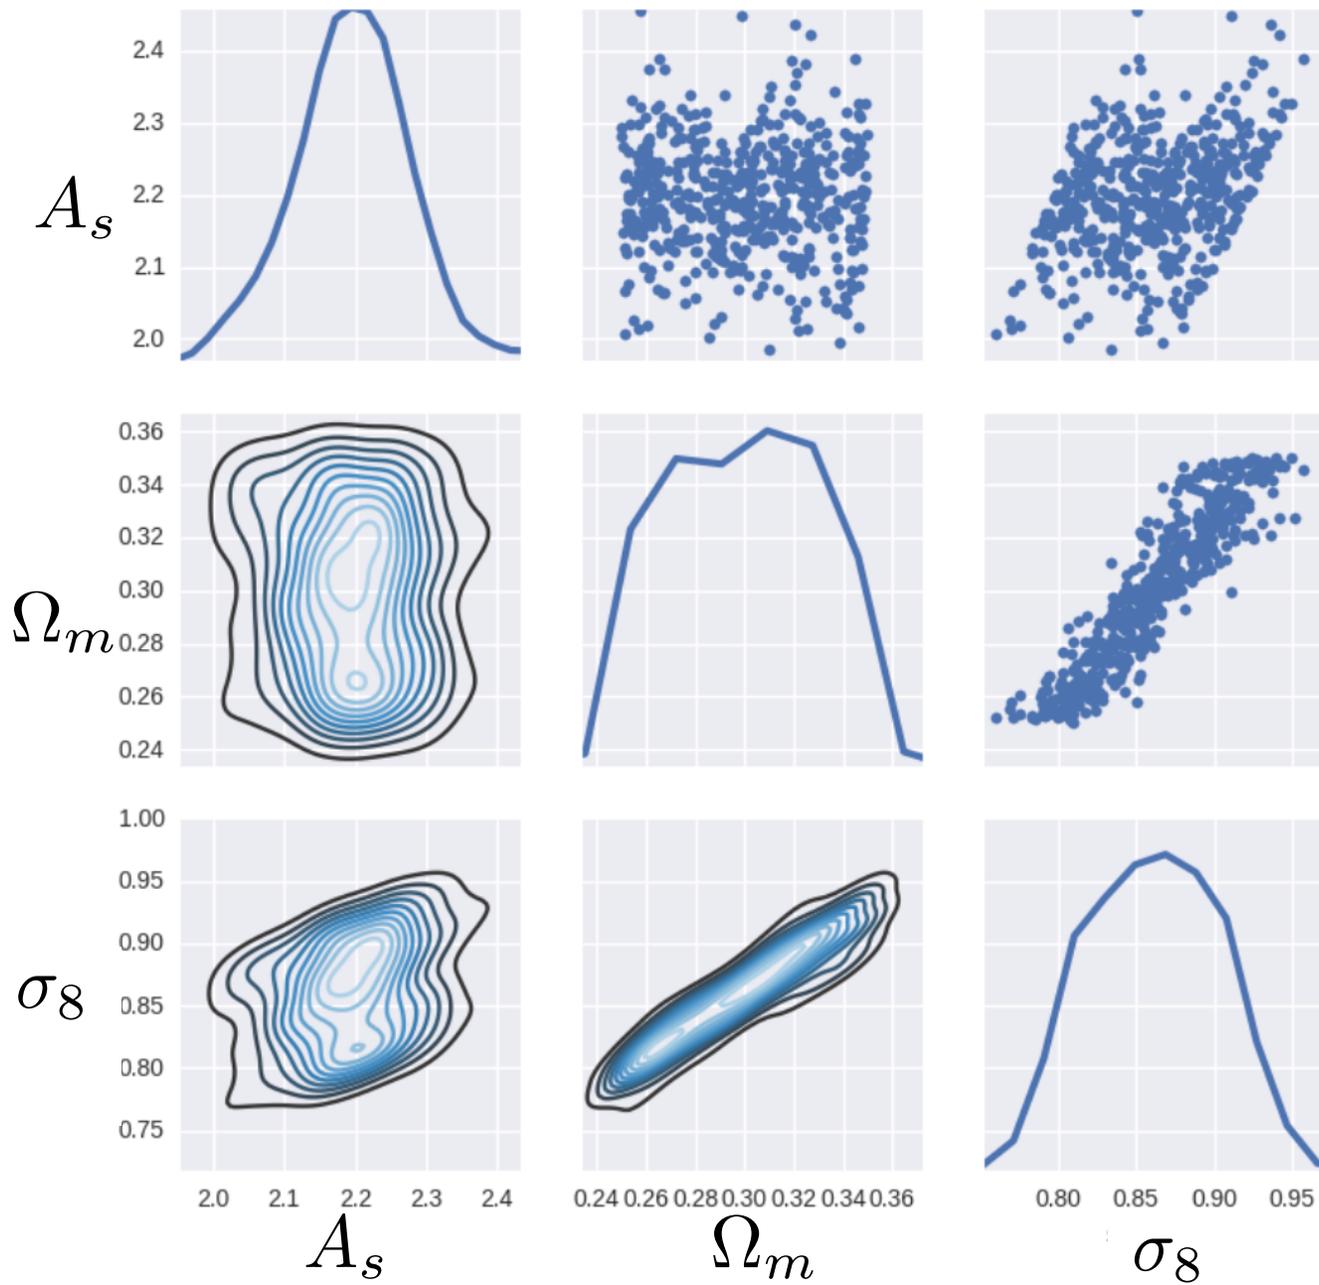


Cosmic Microwave  
Background+  
galaxy clustering +  
cmb lensing X  
galaxy

# Can we use Machine Learning to extract more information from the following datasets?



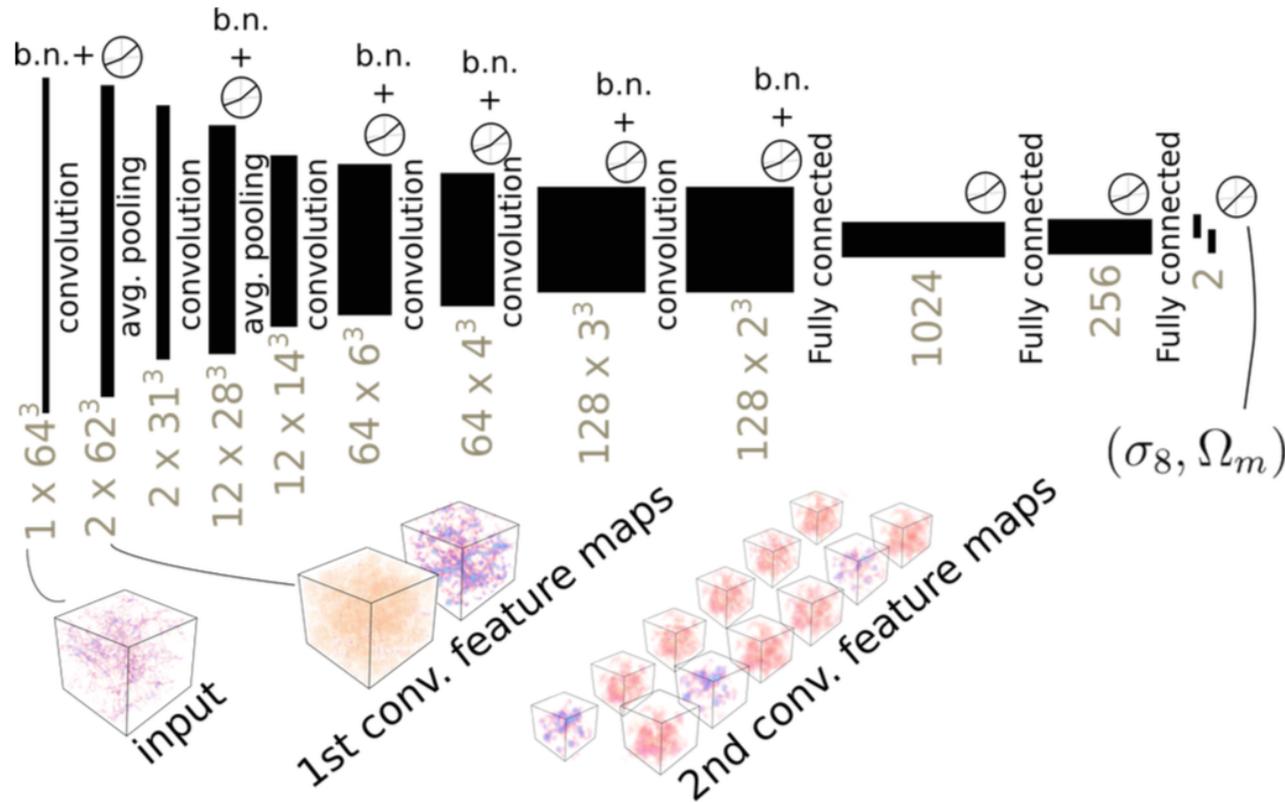
*Figure 1.* Dark matter distribution in three cubes produced using different sets of parameters. Each cube is divided into small sub-cubes for training and prediction. Note that although cubes in this figure are produced using very different cosmological parameters in our constrained sampled set, the effect is not visually discernible.



Ravanbakhsh, Oliver, Price, **Ho**, Schendier & Poczós ICML 2016

# Deep learning algorithm: Convolutional Neural Net (Conv-Net)

Ravanbakhsh, Oliver, Price, **Ho**, Schendier & Poczos **ICML** 2016



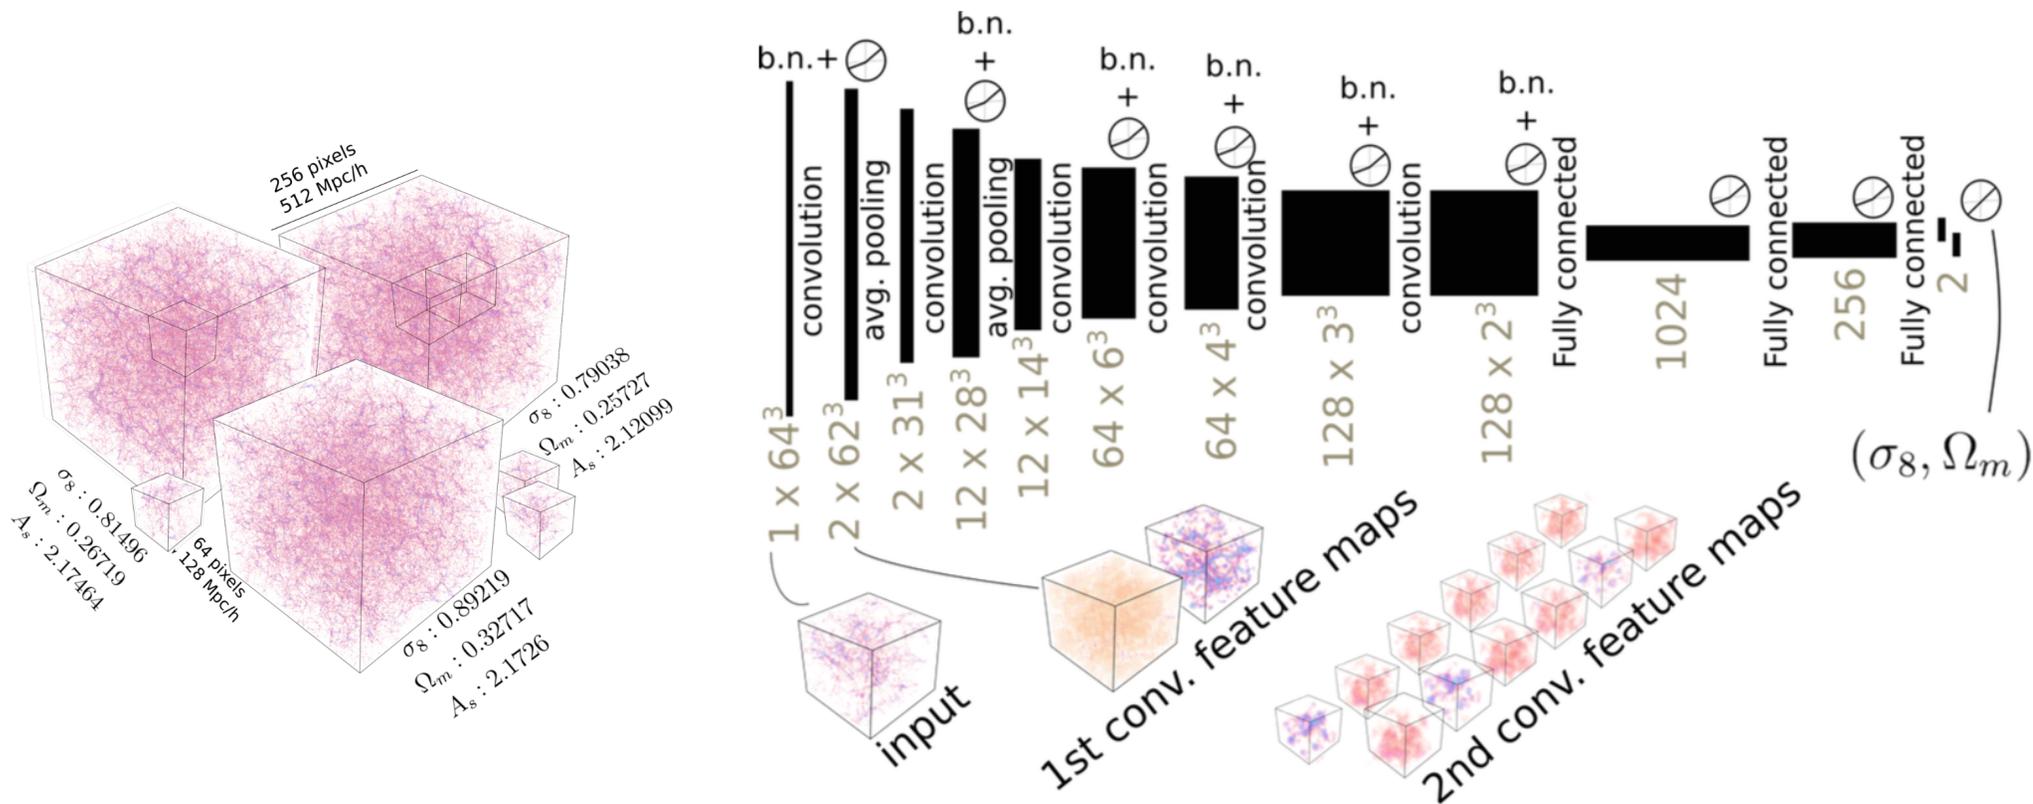
*Figure 6.* The architecture of our 3D conv-net. The model has six convolutional and 3 fully connected layers. The first two convolutional layers are followed by average pooling. All layers, except the final layer, use leaky rectified linear units, and all the convolutional layers use batch-normalization (b.n.).

# Deep learning algorithm: Convolutional Neural Net (Conv-Net)

- Convolutional Neural Net performs non-linear convolution on the input volume of pixels (the dataset in this case) and
- output a volume of scores (of different dimension).
- Trying to learn underlying **mapping** between input and output
- The scores in this case are the cosmological parameters.
- The non-linear convolution are trained with gradient descent so that the output scores are consistent with the training sample

# Training, Validation and Testing with $O(100)$ s of simulations

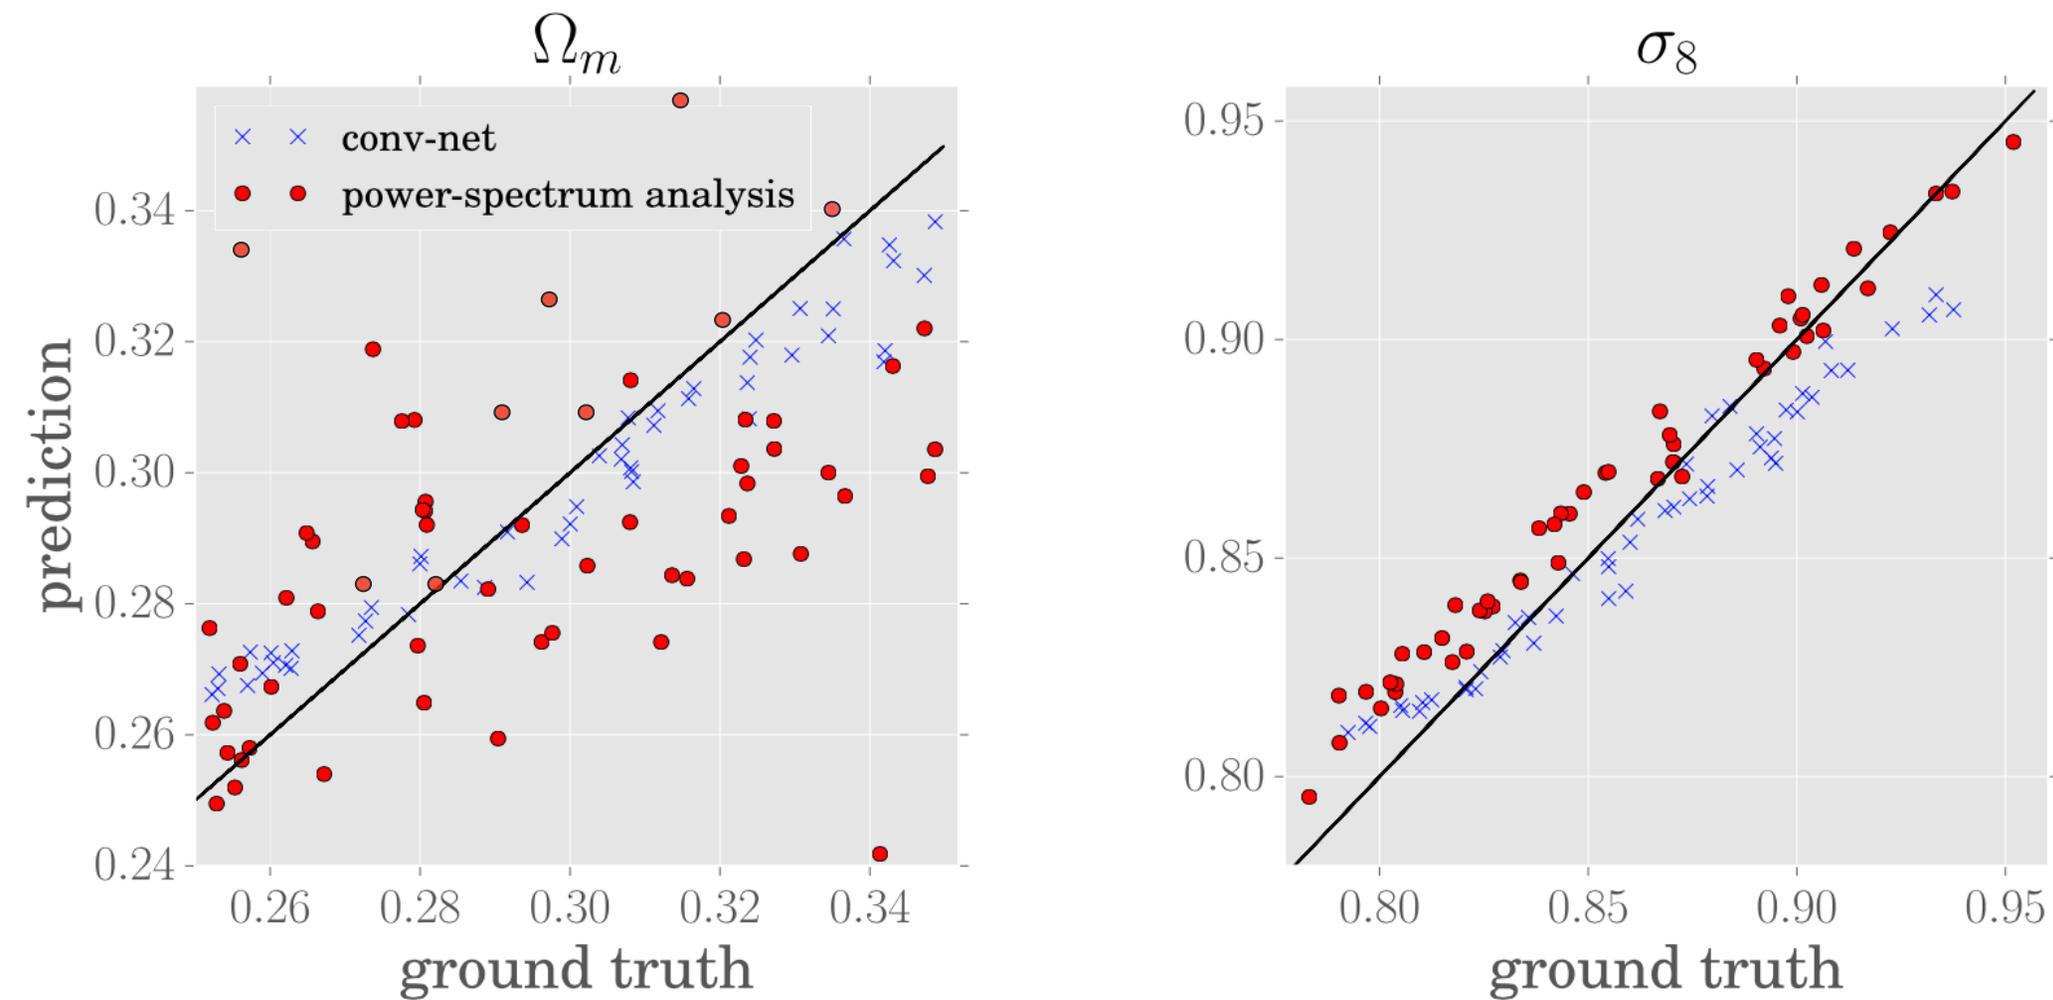
Ravanbakhsh, Oliver, Price, **Ho**, Schendier & Poczos ICML 2016



**Training:** Input N-body simulations with known cosmological parameters to train the ConvNet

**Validation:** Input next set of simulations with known cosmological parameters to fine tune the hidden parameters in ConvNet (eg. Number of layers)

**Test:** Input N-body simulations with unknown cosmological parameters and predict with ConvNet



*Figure 2.* Prediction and ground truth of  $\Omega_m$  and  $\sigma_8$  using 3D conv-net and analysis of the power-spectrum on 50 test cube instances.

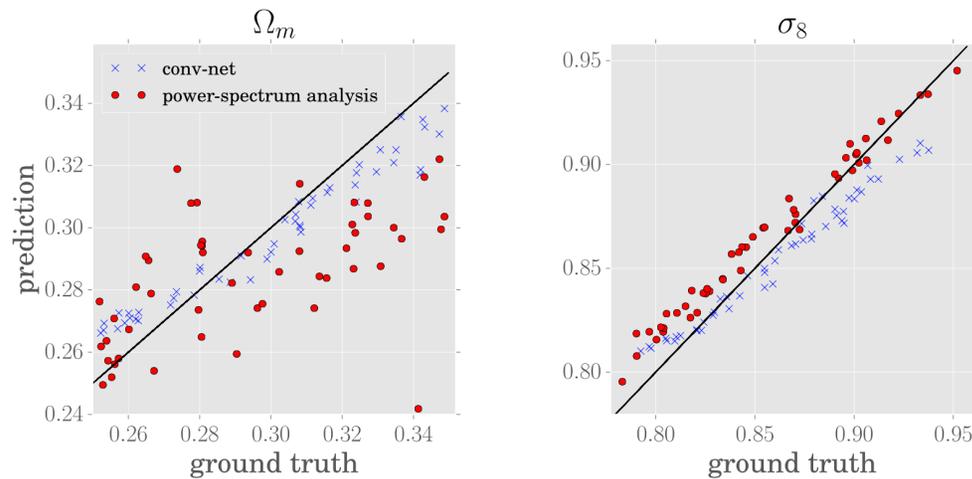
# Now as scientists, we have lots of questions...

- As an astrophysicist, do we understand where the extra information is coming in from?
- Can we get a correct estimate of the error ?
- Can we interpret the model learnt in Machine Learning?
- Can we compress the model learnt into physical laws?
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

# Now as scientists, we have lots of questions...

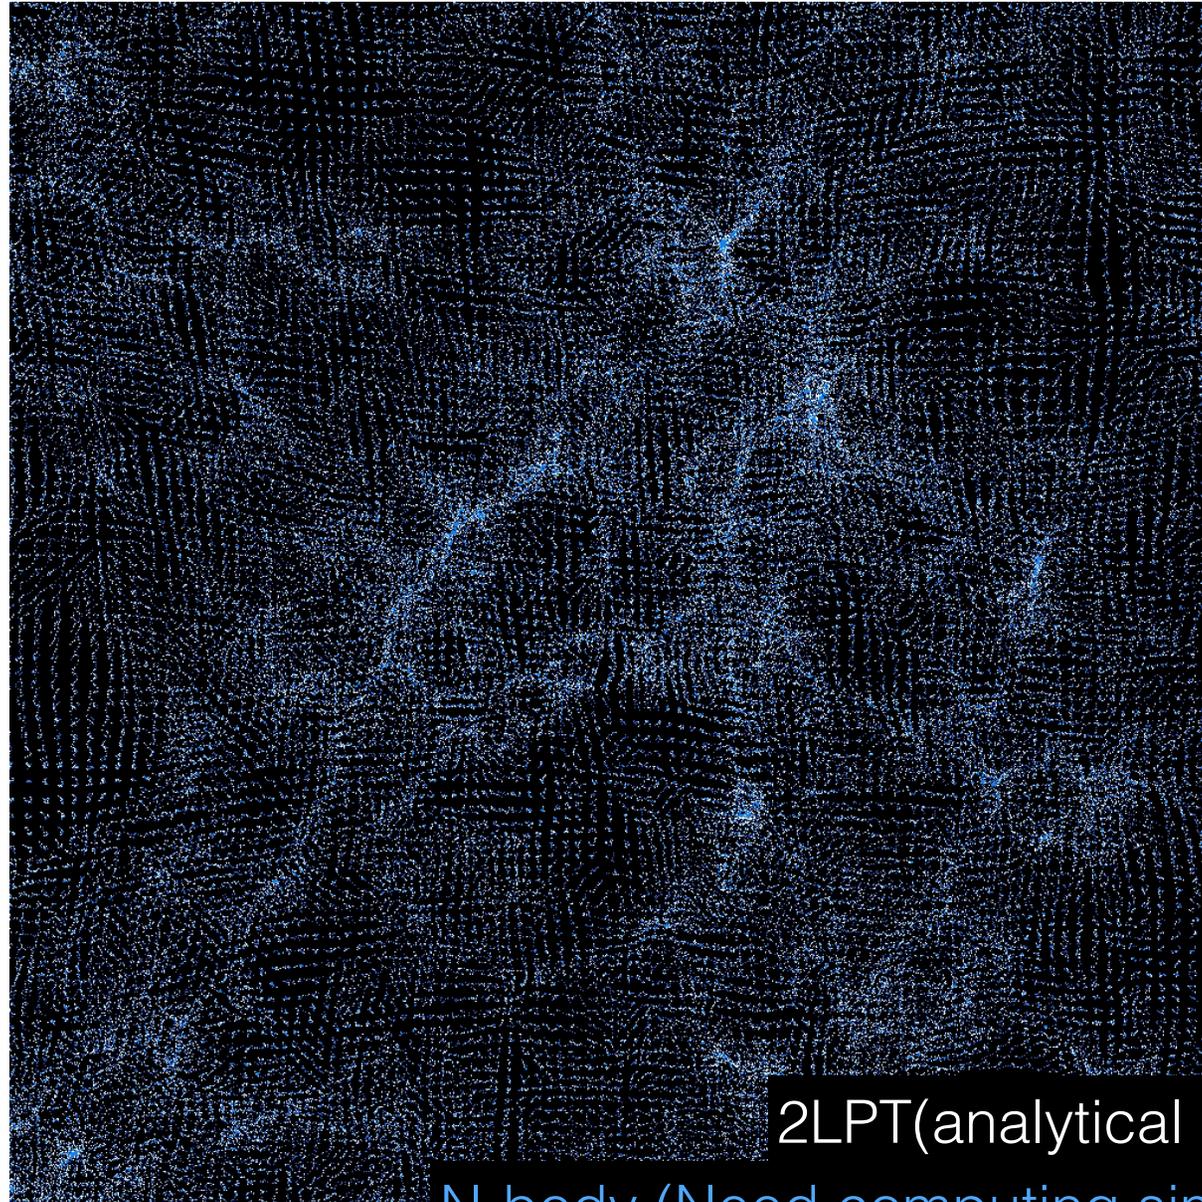
- As an astrophysicist, do we understand where the extra information is coming in from?
- Can we get a correct estimate of the error ?
- **Can we interpret the model learnt in Machine Learning?**
- Can we compress the model learnt into physical laws?
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

# Can we interpret what the model is learning?



- We think the difference between the traditional method and what machine learning is doing may come from the difference between what we can model analytically and what is the full information.
- So we made the following experiment, which is geared to learn the difference between **analytical modeling** and **the full information** in the density field.

# Analytical physics (2nd order Lagrangian Perturbation Theory) vs Computer Simulations (N-body)

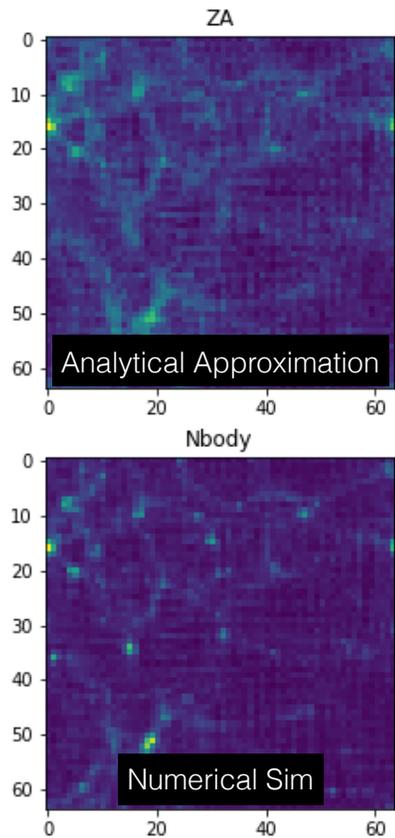


2LPT(analytical physics): White

N-body (Need computing simulations): Blue

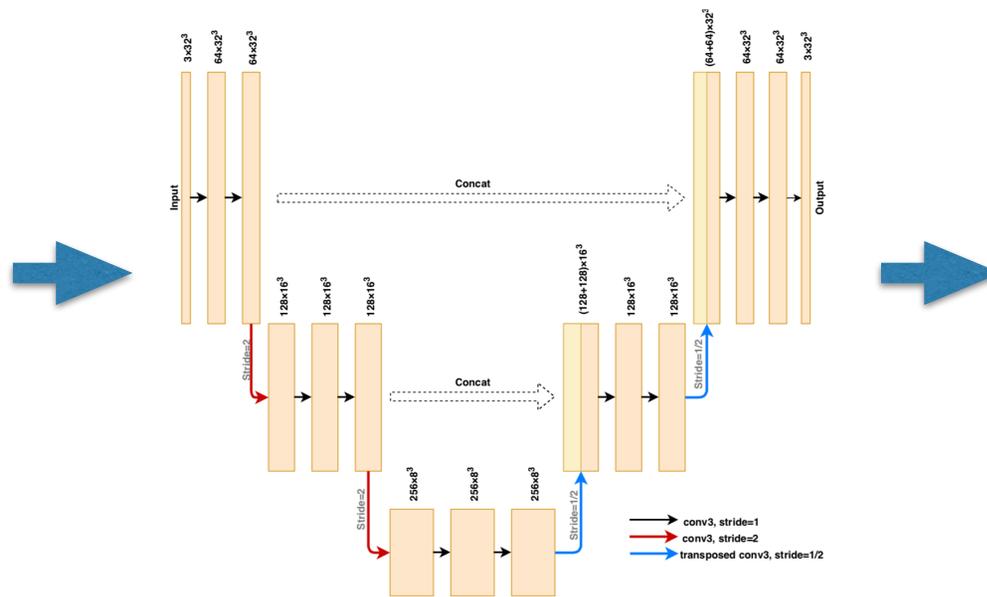
# Predicting from Analytical approximated (Zeldovich Approximation) fields to numerically simulated (FastPM) fields

## Training



## Machine learning [UNET]

Slight variant to Residual NN



10,000 pairs of [Analytical, Sim] boxes  
For training

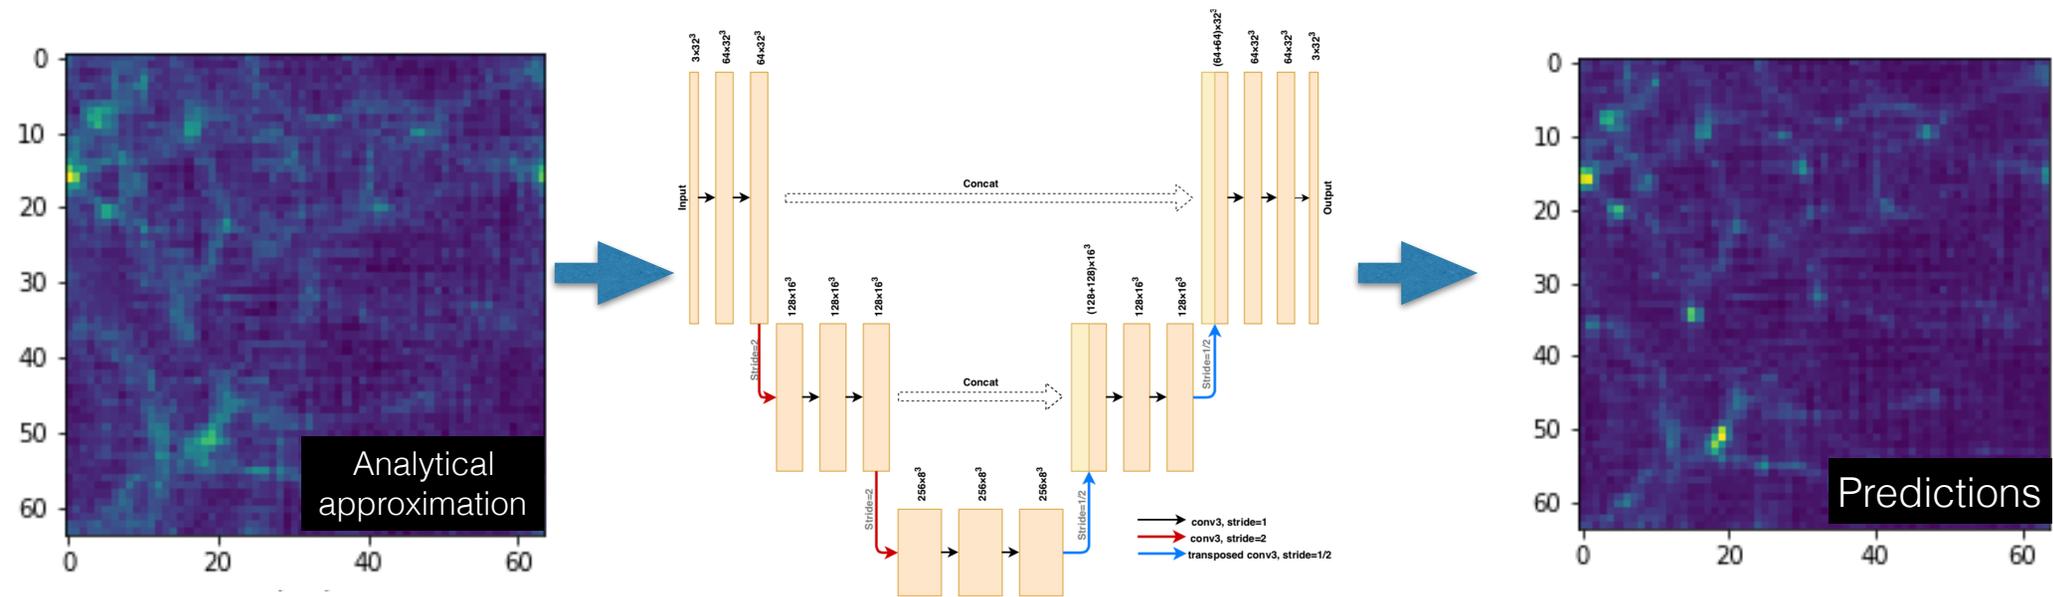
Predicting from  
Analytical approximated (Zeldovich approximation) fields  
to numerically simulated (FastPM) fields

Input

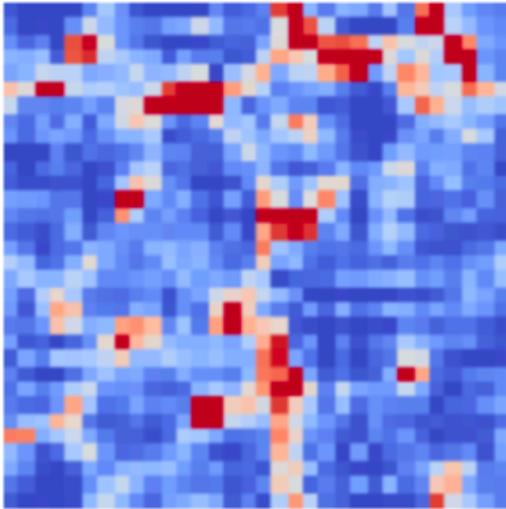
Machine learning [UNET]

Slight variant to Residual NN

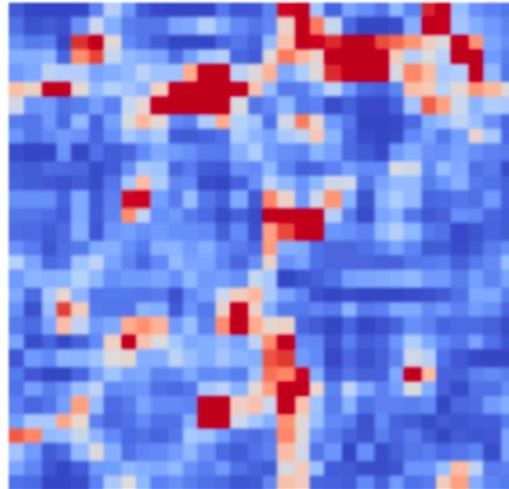
Prediction



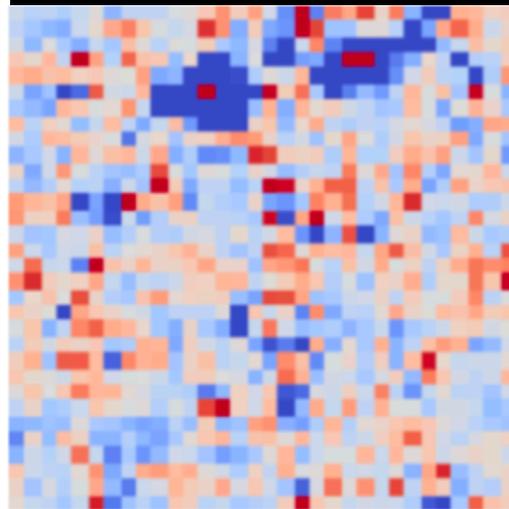
Truth/numerical simulations



Best analytical



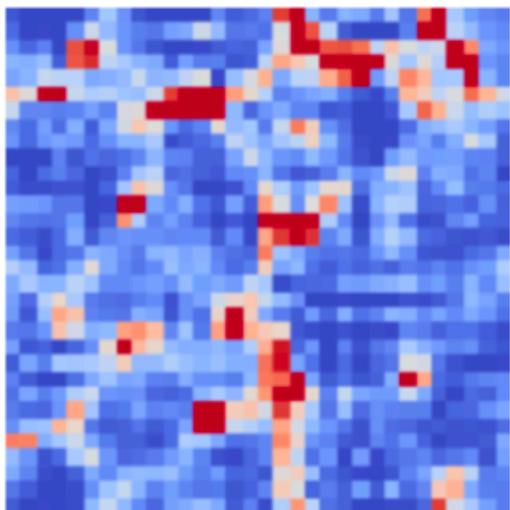
Numerical - Analytical  
residuals



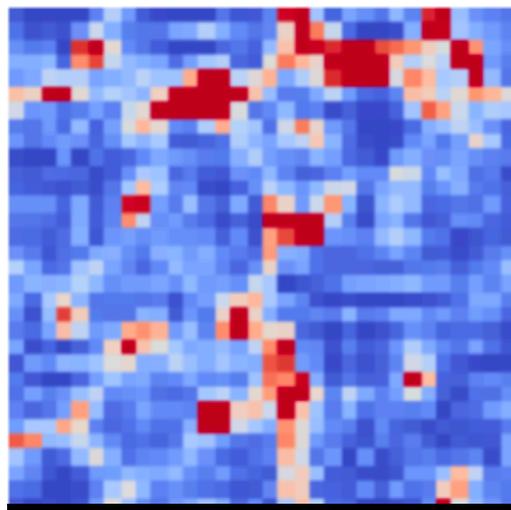
**Residuals**

## Density field comparisons

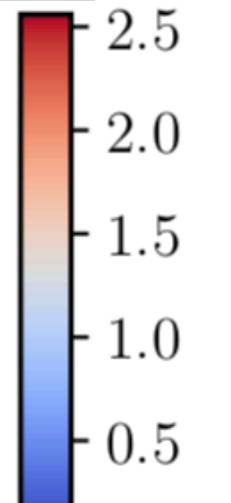
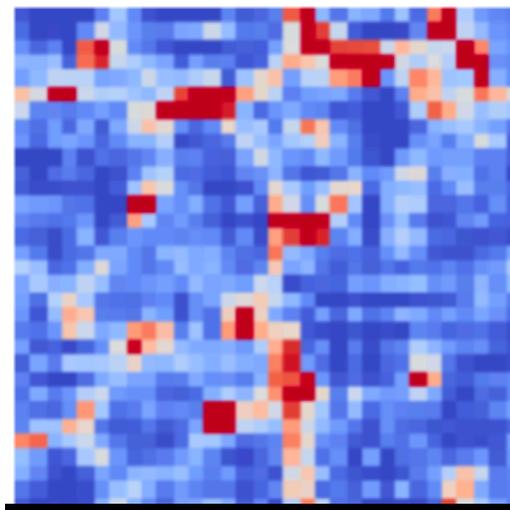
Truth/numerical simulations



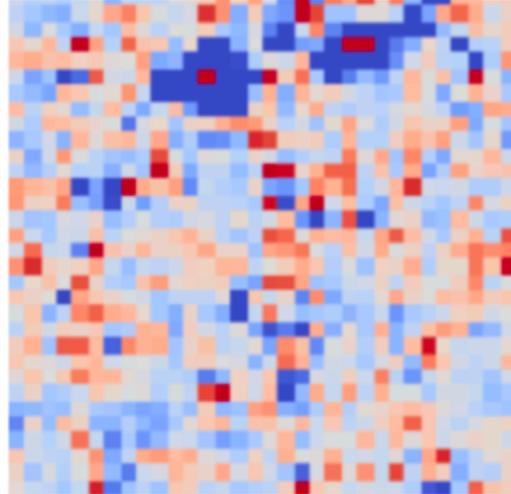
Best analytical



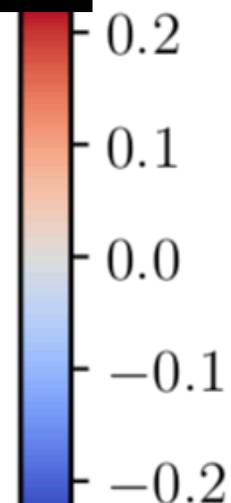
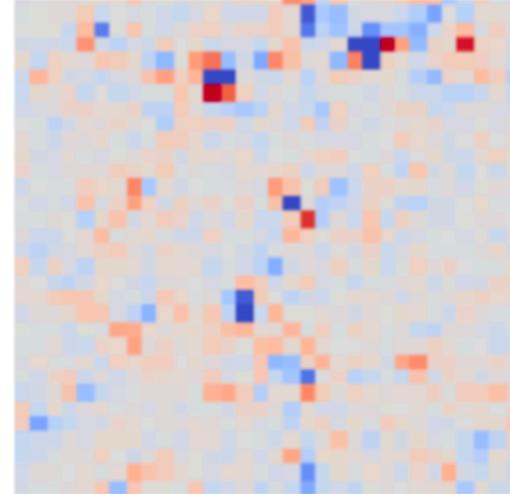
Machine Learning Predictions



Numerical - Analytical residuals



Numerical - Machine Learning residuals



## Density field comparisons

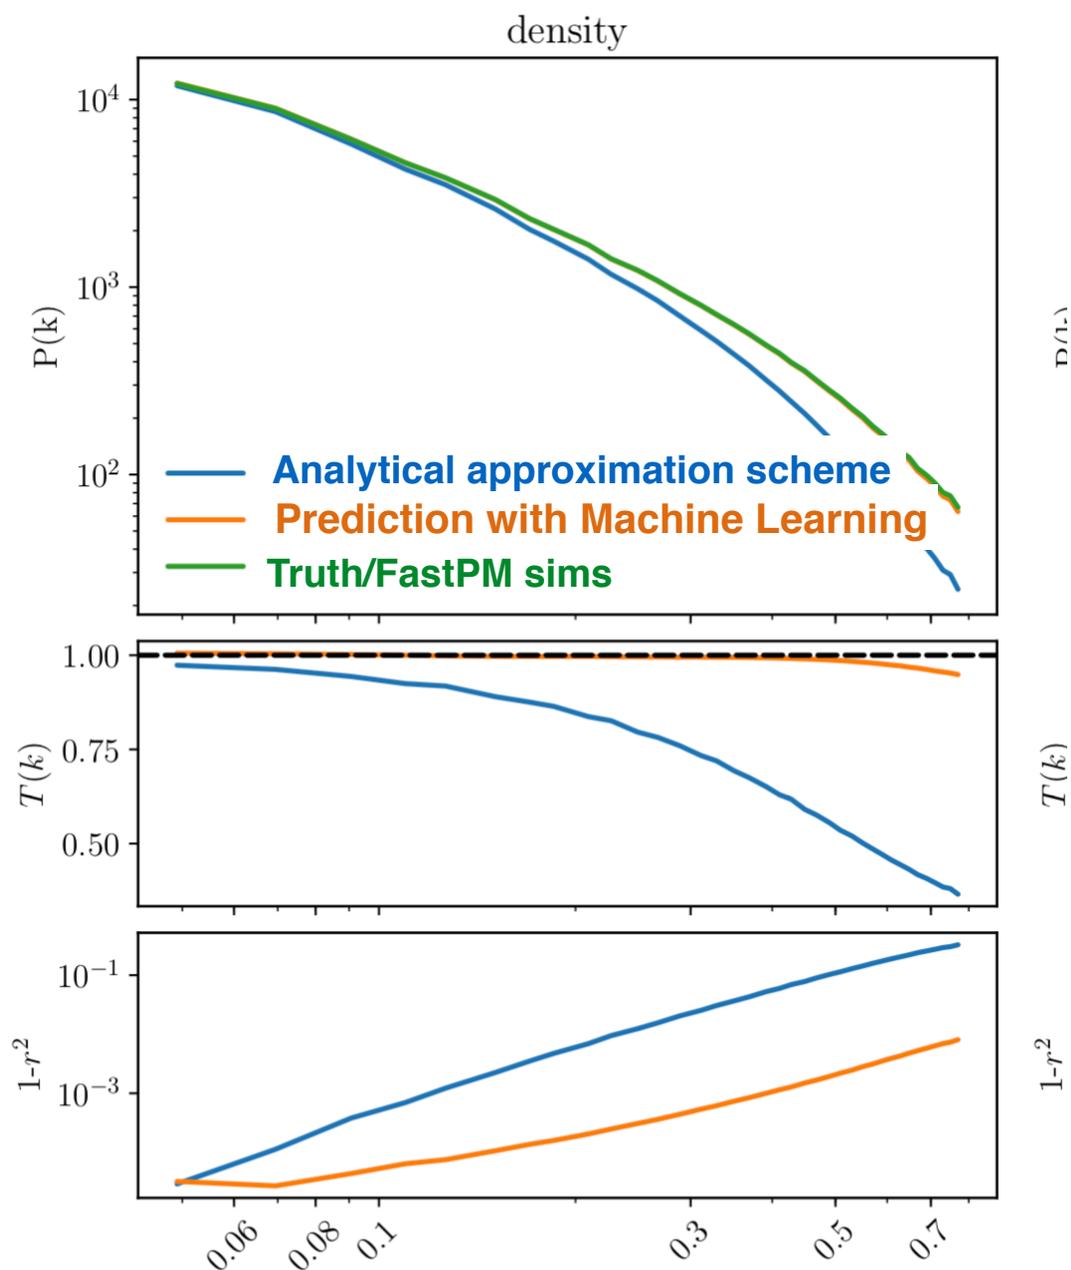
Checking the following:

- 1) the average power-spectrum of 1000 sims, and
- 2) ratios to the true power-spectrum ( $T(k)$ ), and
- 3) The cross-correlation coefficients.

The simulations can be predicted in  $O(1)$  minutes post training and validation.

$$T(k) = \frac{P_{\text{pred}}(k)}{P_{\text{true}}(k)}$$

$$r(k) = \frac{P_{\text{pred} \times \text{true}}(k)}{\sqrt{P_{\text{pred}}(k)P_{\text{true}}(k)}}$$



(a) Results from the density field

# Foray into understanding what the heck the Model is learning

- We first train a network with [ZA, N-body] pairs, and make prediction using ZA inputs. And we have seen that the predictions are pretty good.
- Then we analyze what the network has learned by decomposing the **input** into different **Fourier modes** and look at the **predicted power-spectra** of these modes.
- Different Fourier modes in the following form:

$$\psi(\hat{x}) = A_{\hat{k}_i} \hat{k}_i \cos(\vec{k}_i \cdot \vec{x})$$

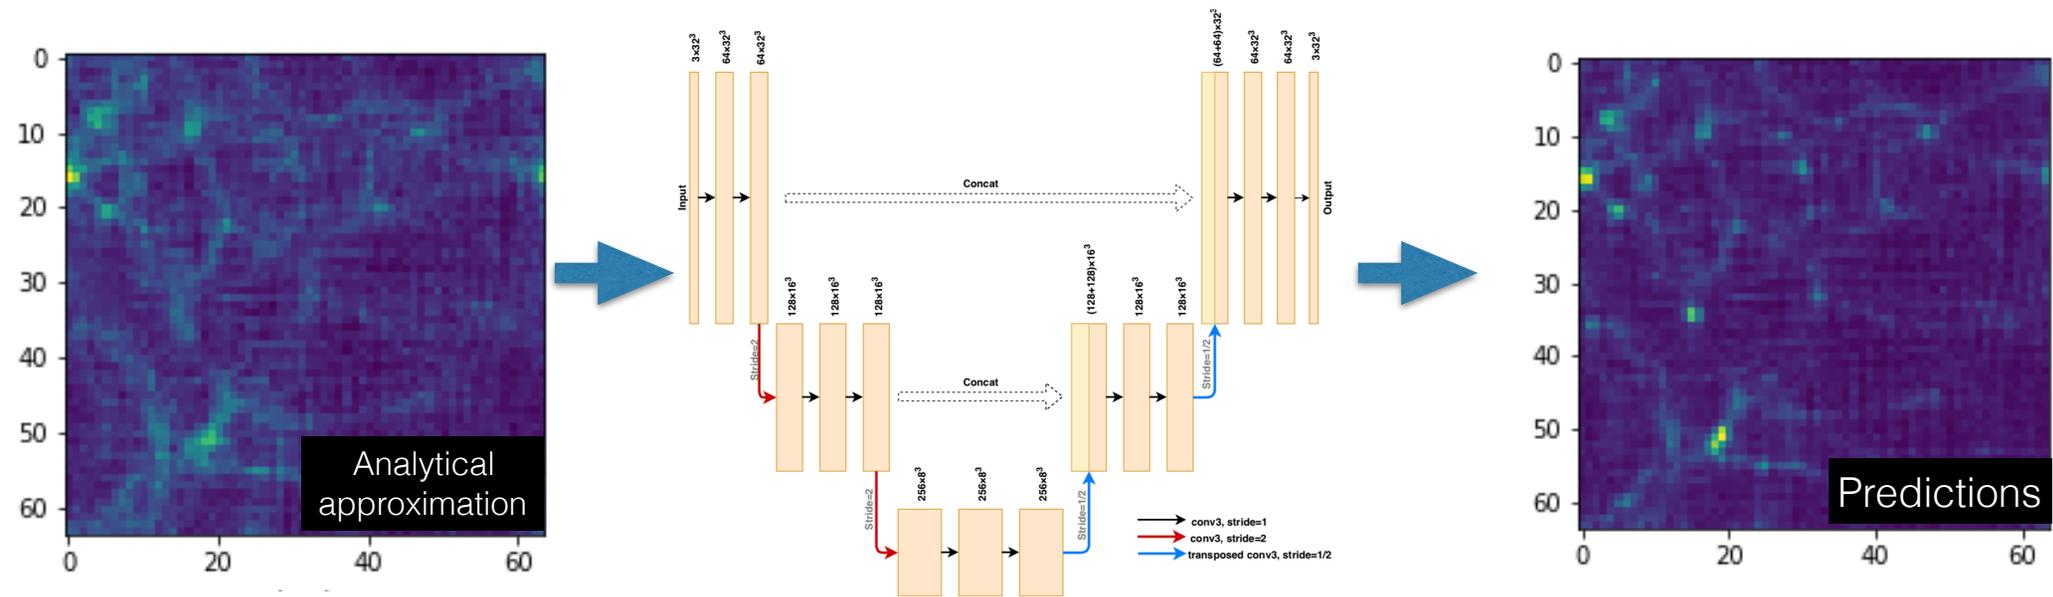
# Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

## UNET

Input

Slight variant to Residual NN

Prediction





# Input mode: A Plane Wave

**What happen if we have power only one scale?**

Preliminary results. Please wait for our publication.

The transfer function shows that the U-Net model captures quite well at the dominate scale, which indicates the U-Net model is able to capture scale information. The U-Net model also captures the other modes of FastPM that are two orders smaller than the dominant mode and come from the numerical artifact of FastPM simulations.

# Interrogating the learned model

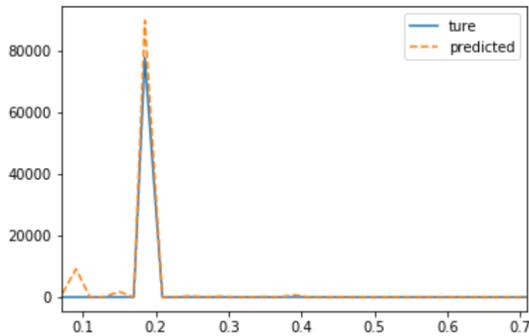
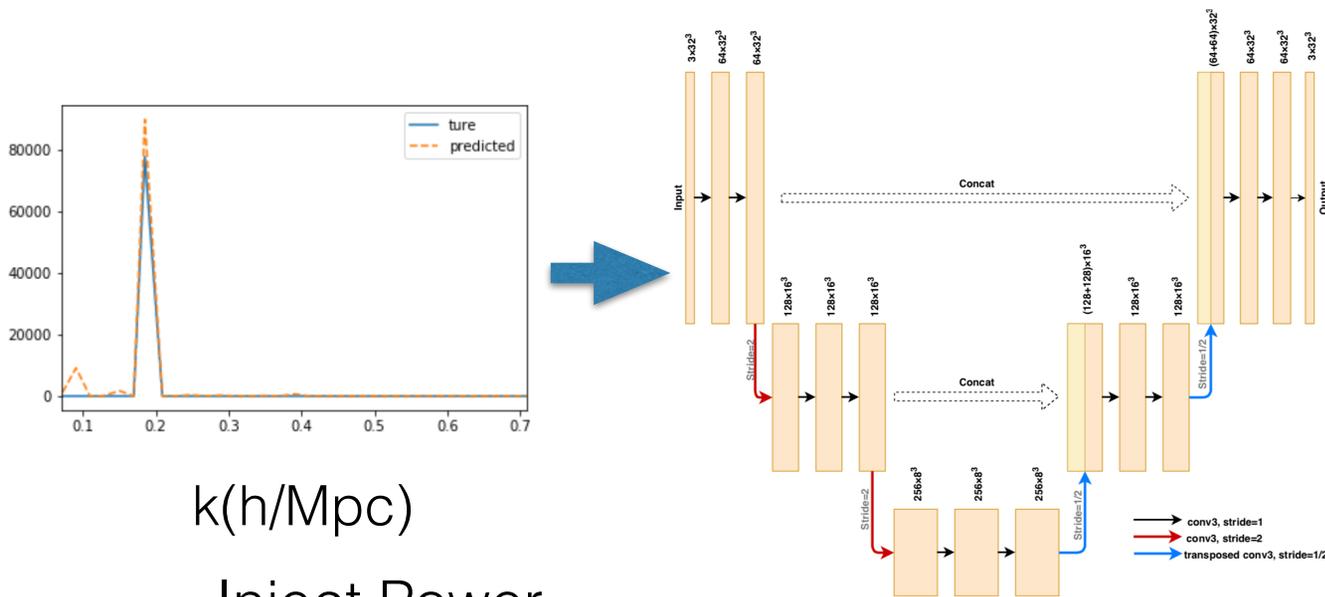
Is Rotational Invariance learnt by the model?

UNET

Input

Slight variant to Residual NN

Prediction

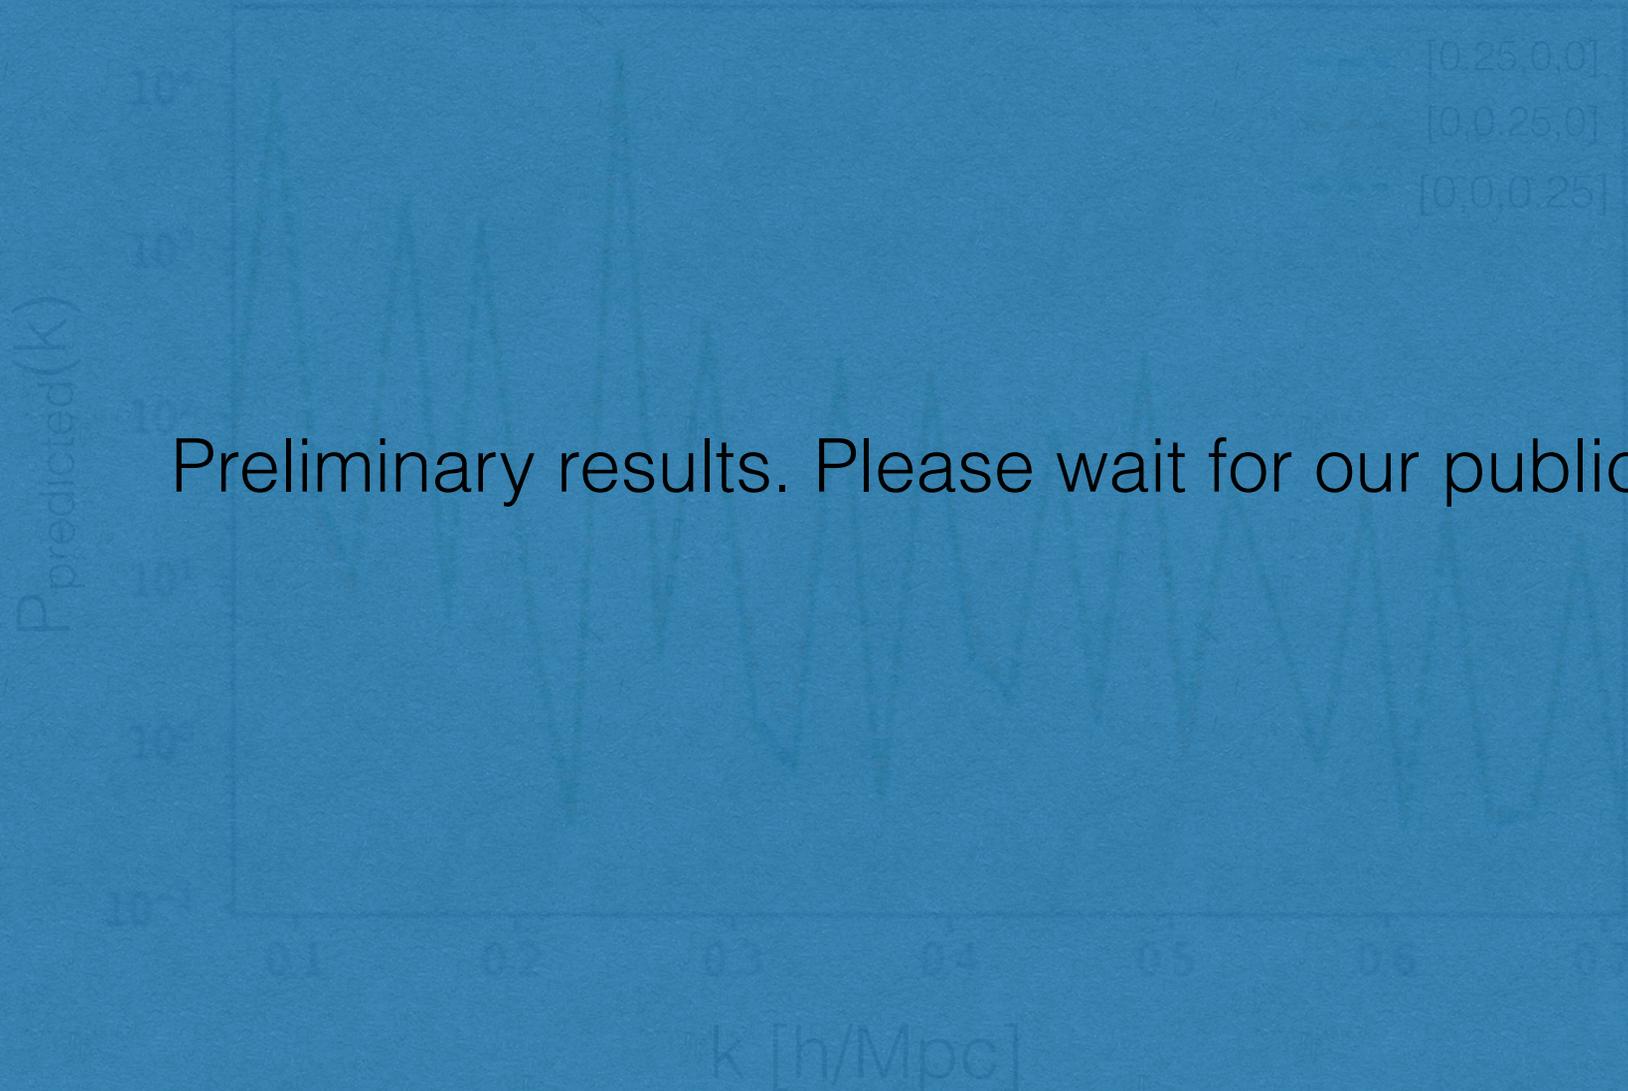


$k$  (h/Mpc)

Inject Power

at  $[k_x, k_y, k_z] = [0.25, 0, 0]$   
 $= [0, 0.25, 0]$   
 $= [0, 0, 0.25]$

# Is Rotational Invariance learnt by the model?

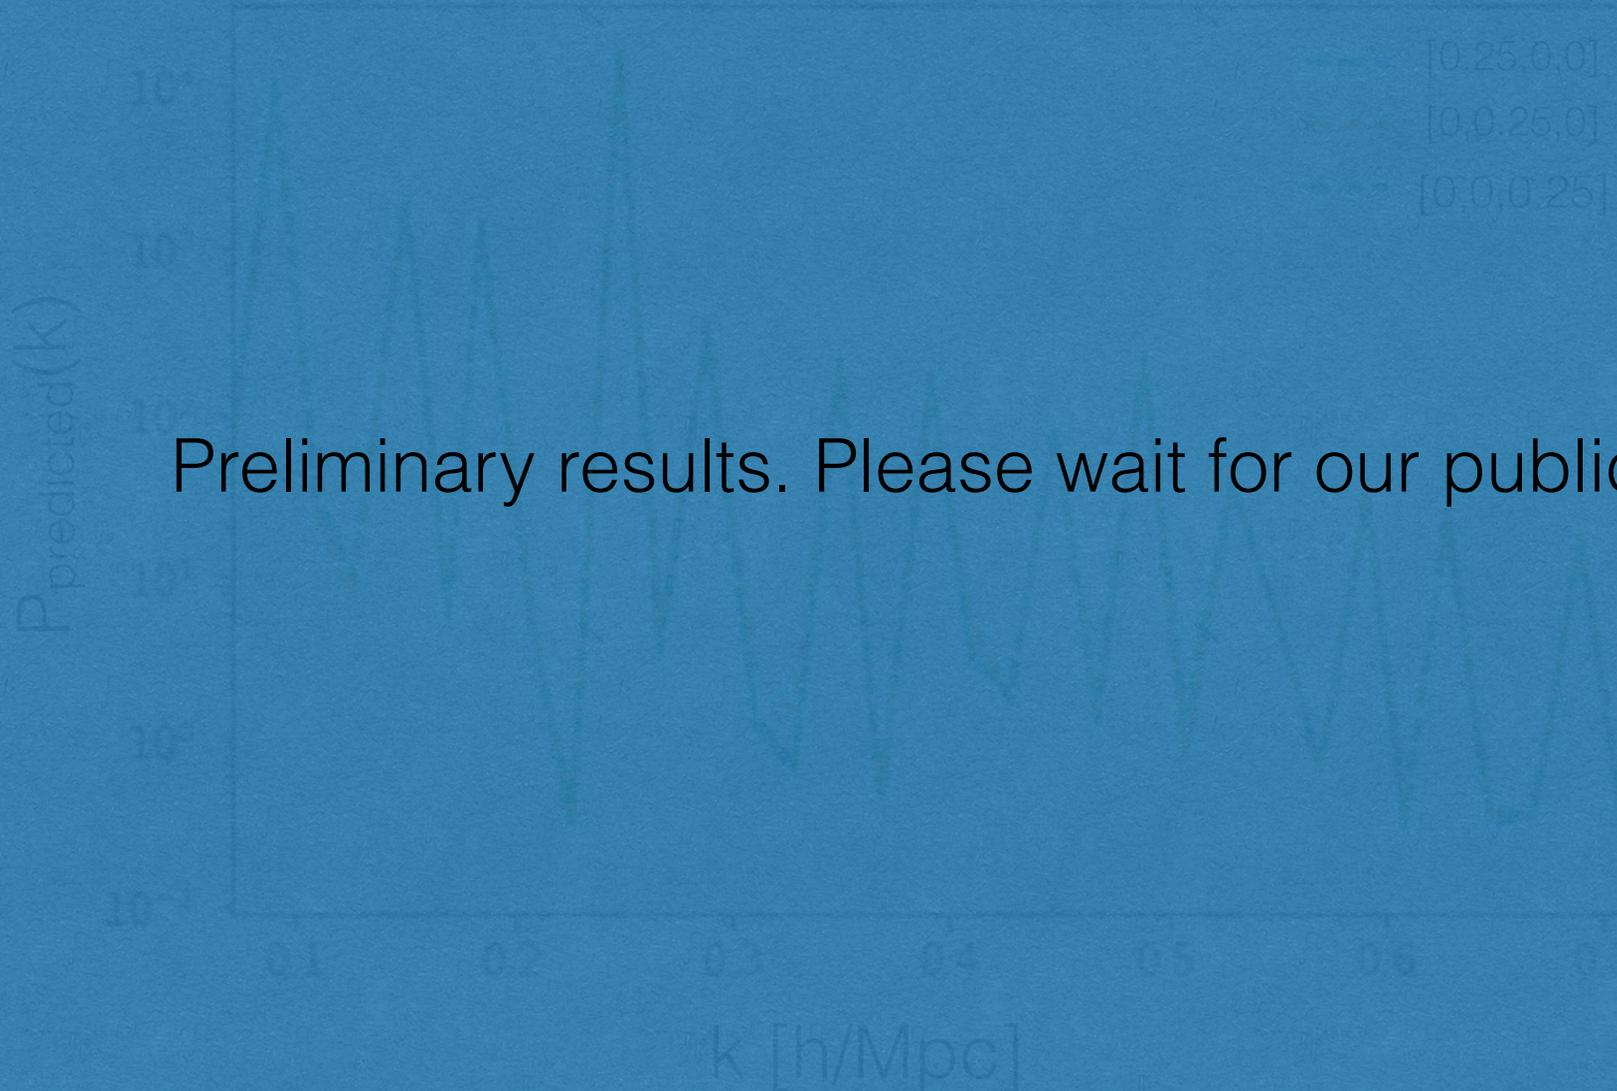


Preliminary results. Please wait for our publication.

# Is Rotational Invariance learnt by the model?

Yes, *predicted power is the similar no matter which orientation:*

rotational invariance is learnt!



Preliminary results. Please wait for our publication.

# Interrogating the learned model

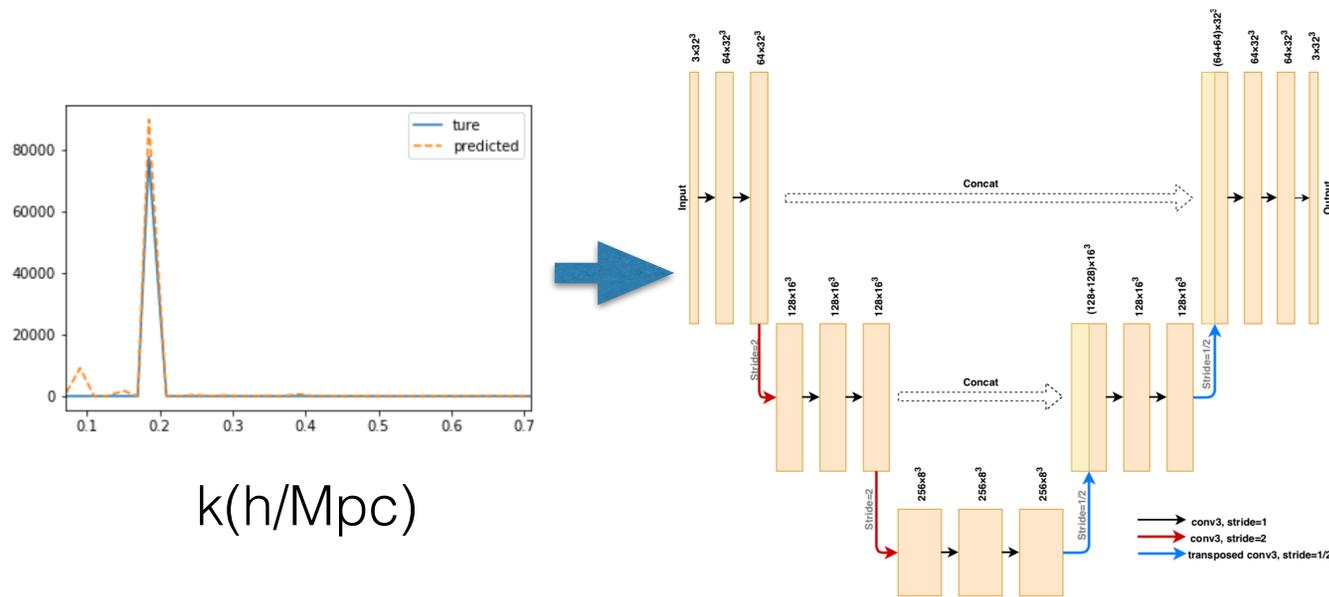
What happens if we change the phase of the input mode?

UNET

Input

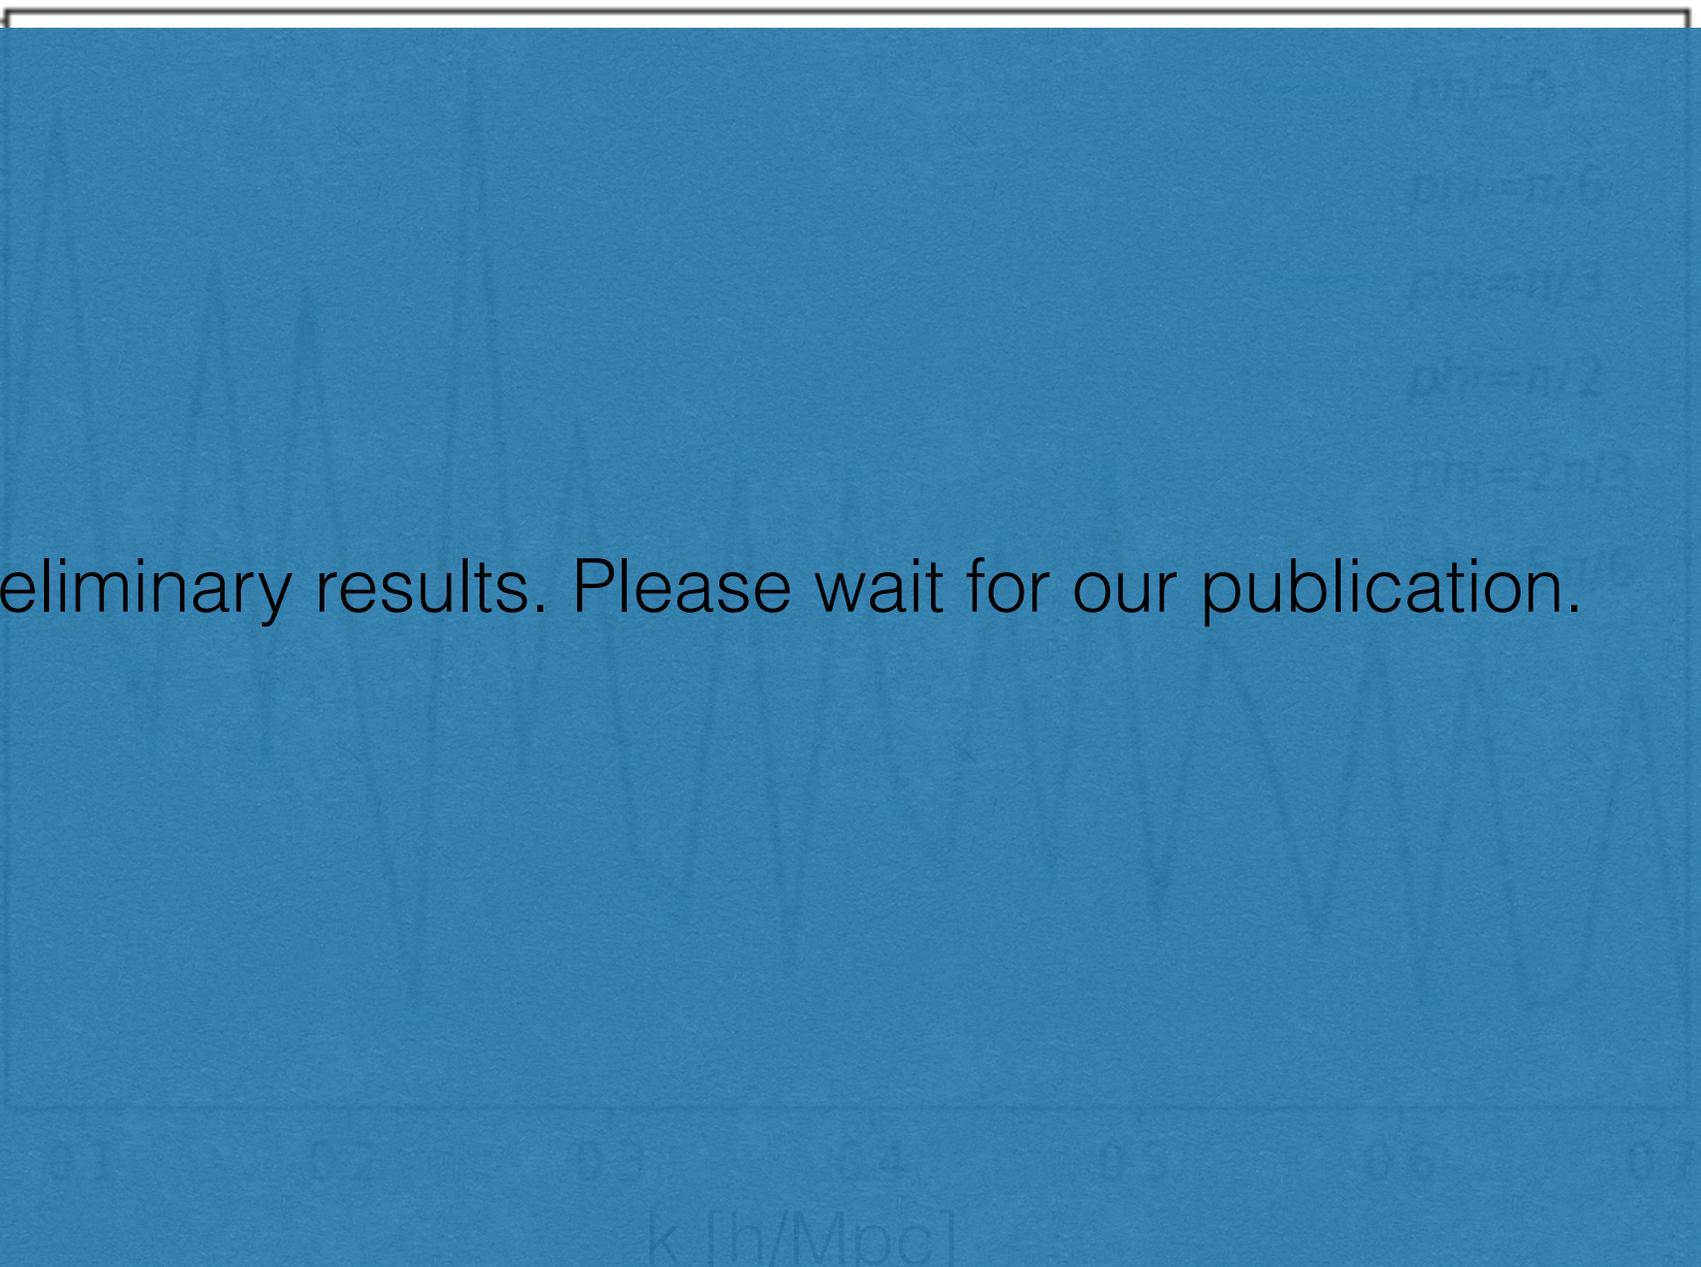
Slight variant to Residual NN

Prediction



Inject Power  
At same k,  
but different phases

# What happens if we change the phase of the input mode?



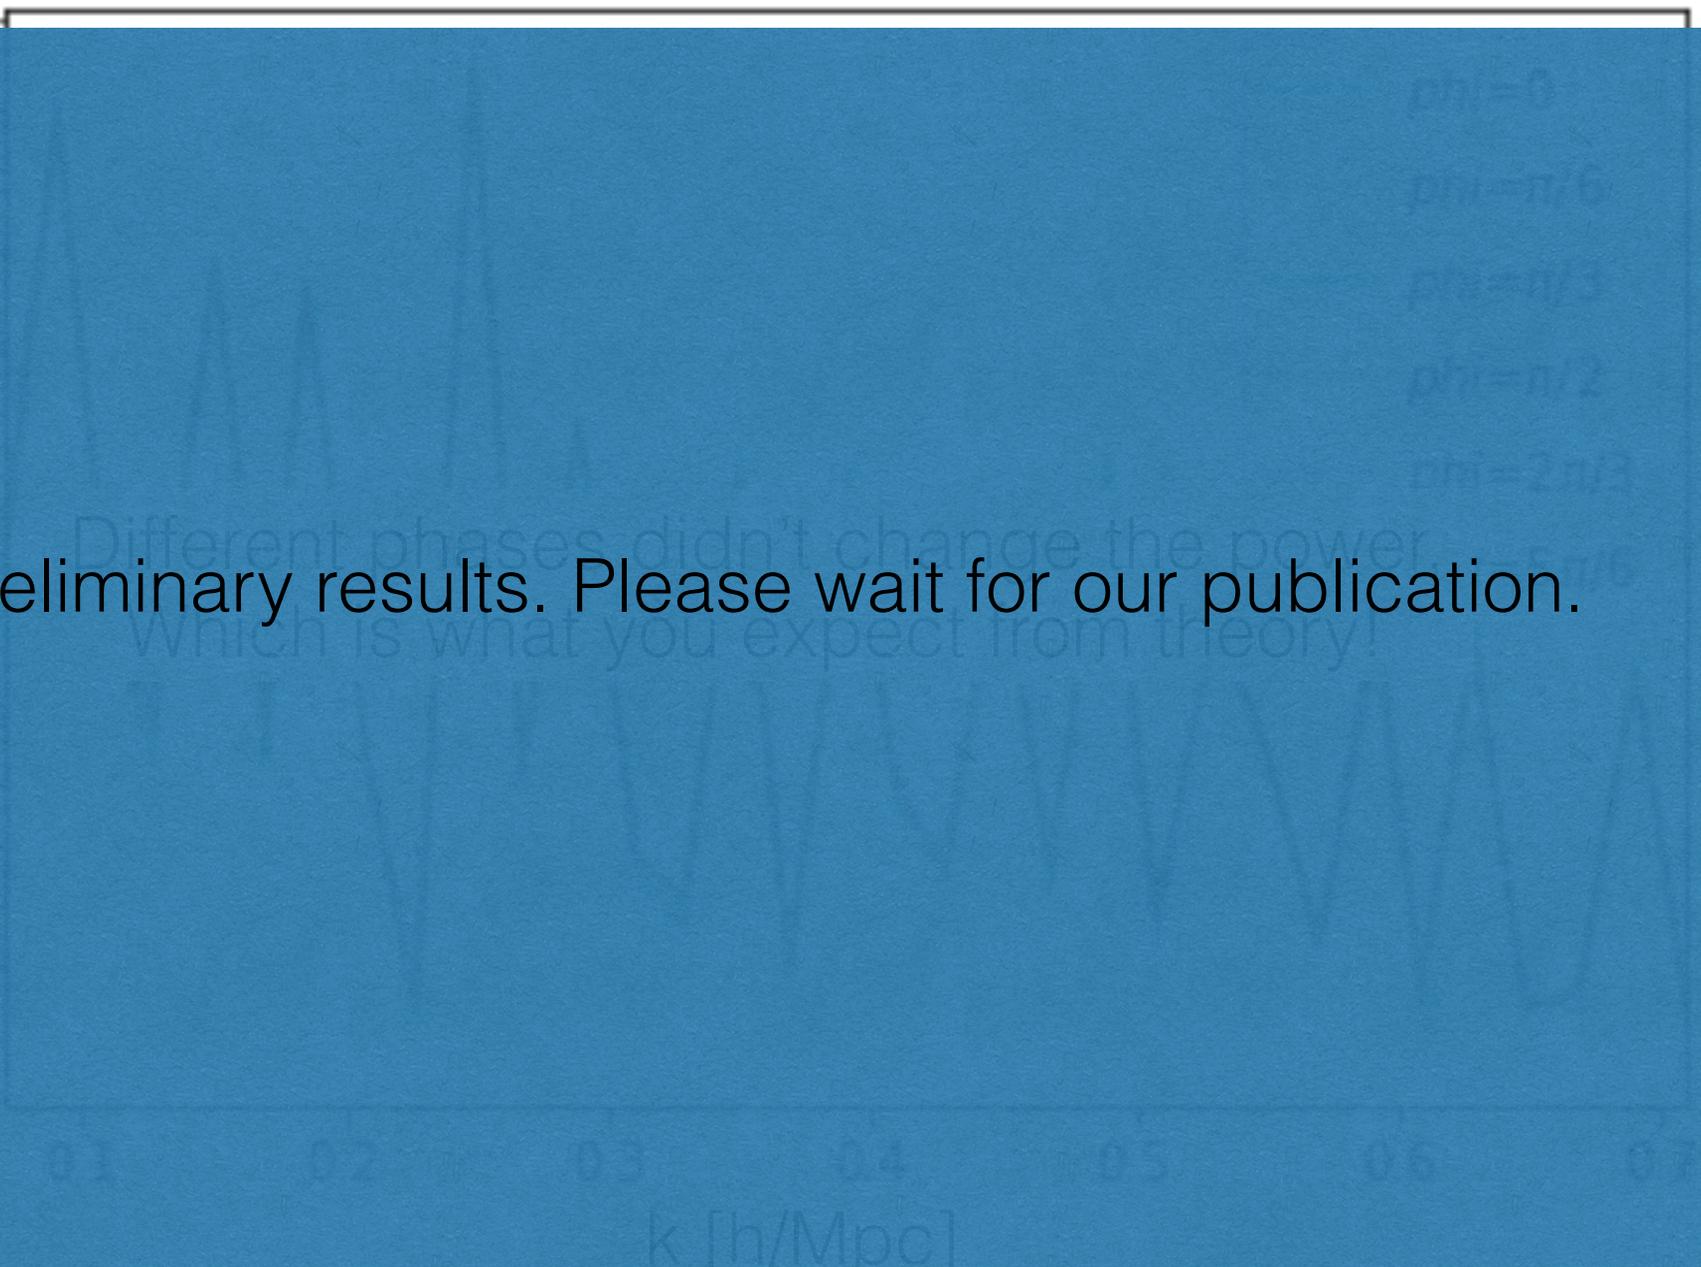
Predicted(k)

$k$  [h/Mpc]

$\phi=0$   
 $\phi=\pi/6$   
 $\phi=\pi/3$   
 $\phi=\pi/2$   
 $\phi=2\pi/3$   
 $\phi=\pi$

Preliminary results. Please wait for our publication.

# What happens if we change the phase of the input mode?



Predicted(k)

Different phases didn't change the power, which is what you expect from theory!

$\phi=0$   
 $\phi=\pi/6$   
 $\phi=\pi/3$   
 $\phi=\pi/2$   
 $\phi=2\pi/3$

$k$  [h/Mpc]

Preliminary results. Please wait for our publication.

It seems like physics are being  
learned by the model...

# Now as scientists, we have lots of questions...

- As an astrophysicist, do we understand where the extra information is coming in from?
- Can we get a correct estimate of the error ?
- Can we interpret the model learnt in Machine Learning?
- Can we compress the model learnt into physical laws?
- **More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?**

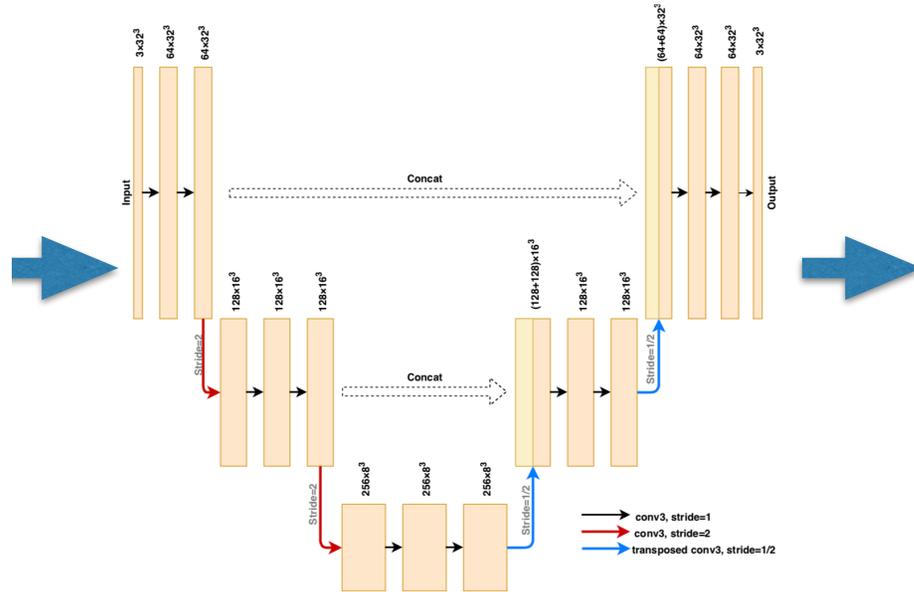
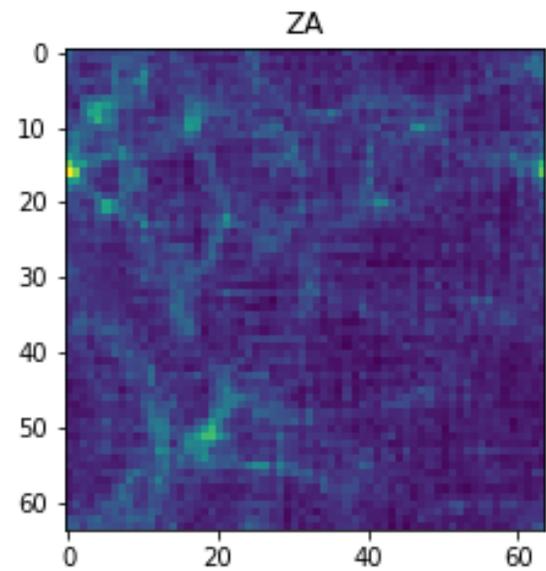
# Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

## UNET

Input

Slight variant to Residual NN

Prediction



ZA maps of

**Different cosmology:**

$$A_s = \{0.2 A_0, 0.8 A_0, 1.2 A_0, 1.8 A_0\}$$

## Experiment:

- 1) We input Analytical approximated field of particles (of one cosmology parameter )
- 2) We predict particle position outputs using ML (or physics)
- 3) Architecture : UNet (a variant of ResNet)
- 4) It works very well (ask me later)
- 5) Question is: What happens if I input a Analytical field with different cosmology ?

## Power-spectrum of Density field

Experiment:

- 1) We input Analytical approximated field of particles (of one cosmology parameter )
- 2) We predict particle position outputs using ML (or physics)
- 3) Architecture : UNet (a variant of ResNet)
- 4) It works very well (ask me later)
- 5) Question is: What happens if I input a Analytical field with different cosmology ?

**Dotted line -> Prediction using ML**

$$T(k) = \frac{P_{\text{pred}}(k)}{P_{\text{true}}(k)}$$

$$r(k) = \frac{P_{\text{pred} \times \text{true}}(k)}{\sqrt{P_{\text{pred}}(k)P_{\text{true}}(k)}}$$

**Dashed Line -> (2LPT) Theoretical predictions**

## Power-spectrum of Density field



## Experiment:

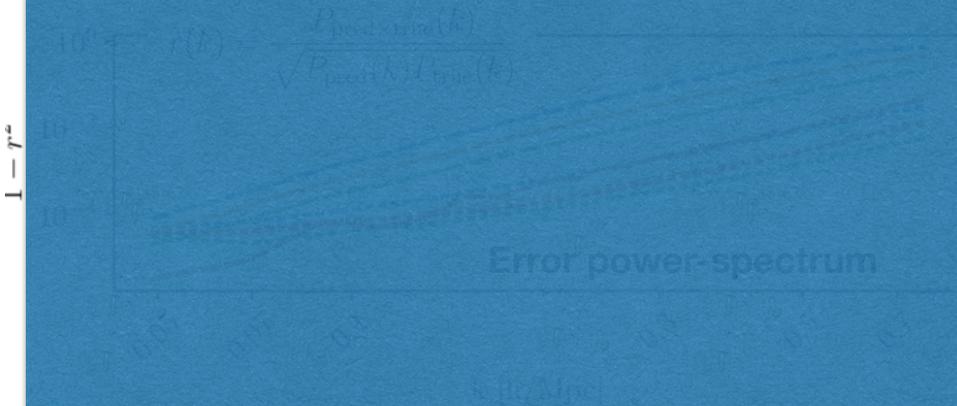
- 1) We input Analytical approximated field of particles (of one cosmology parameter)
- 2) We predict particle position outputs using ML (or physics)
- 3) Architecture : UNet (a variant of ResNet)
- 4) It works very well (ask me later)
- 5) Question is: What happens if I input a Analytical field with different cosmology ?

Dotted line -> Prediction using ML

Preliminary results. Please wait for our publication.



How far are we from the truth?



Error power-spectrum

Dashed Line -> (2LPT) Theoretical predictions

(a) Results from the density field

## Power-spectrum of Density field



## Experiment:

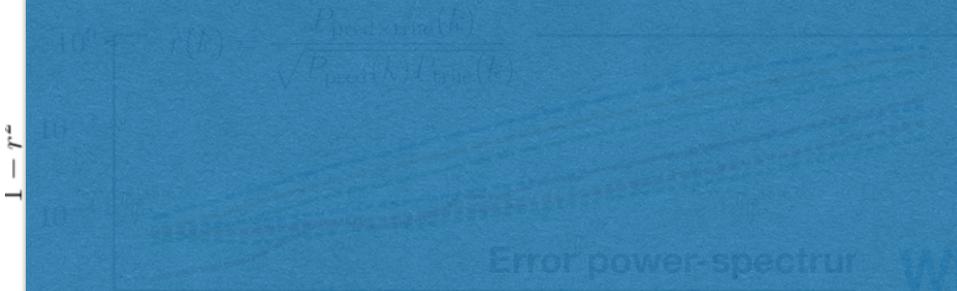
- 1) We input Analytical approximated field of particles (of one cosmology parameter)
- 2) We predict particle position outputs using ML (or physics)
- 3) Architecture : UNet (a variant of ResNet)
- 4) It works very well (ask me later)
- 5) Question is: What happens if I input a Analytical field with different cosmology ?

Dotted line -> Prediction using ML

Preliminary results. Please wait for our publication.



How far are we from the truth?



Error power-spectrum

Dashed Line -> (2LPT) Theoretical predictions

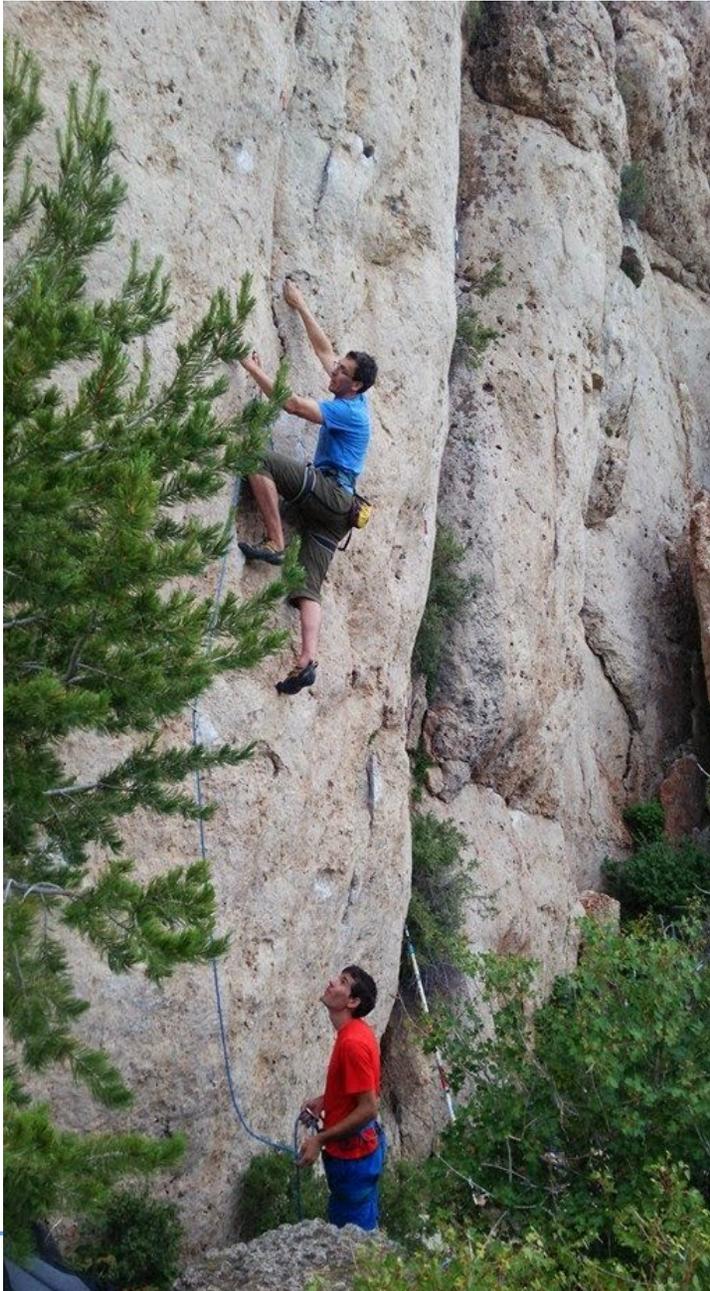
Why can the machine learning algorithm generalize from the one set of cosmology and still predict well for other cosmology? Aka, the test set is not the training set.

(a) Results from the density field

# Let's leave you with questions: Why?

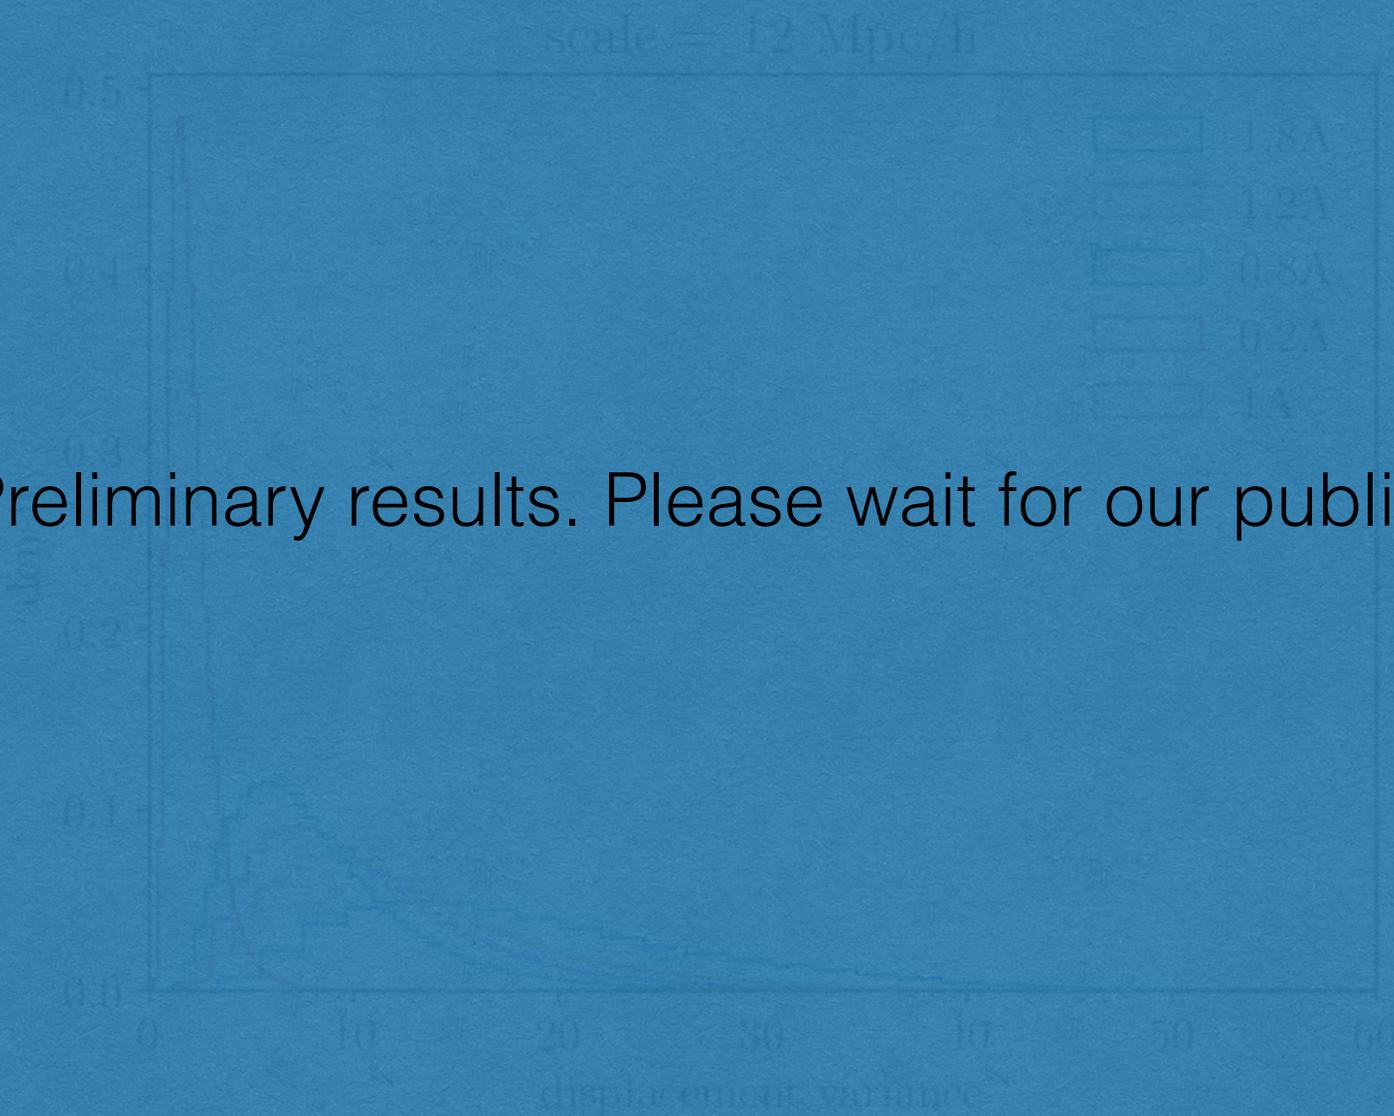
- Is it possible that the model is generalizing rules from the training set that can deal with cosmological inputs with different parameter sets?
- Or maybe the model has seen these parameter sets ?

# My possible climb to fame?



- Understanding Machine Learning?
- Compressing the learned model into physical laws?
- Discover new laws of nature?

# Possible reason ?



Preliminary results. Please wait for our publication.