Learning the Universe with Machine Learning: Steps to Open the Pandora Box

Shirley Ho

Flatiron Institute, Center for Computational Astrophysics

Siyu He (Flatiron/CMU), Yin Li (Berkeley), Yu Feng (Berkeley), Wei Chen (FaceBook), Siamak Ravanbakhsh(UBC), Barnabas Poczos (CMU), Junier Oliver, Jeff Schneider (CMU)
Layne Price (Amazon), Sebastian Fromenteau (UNAM)

Inverse problem working group meeting

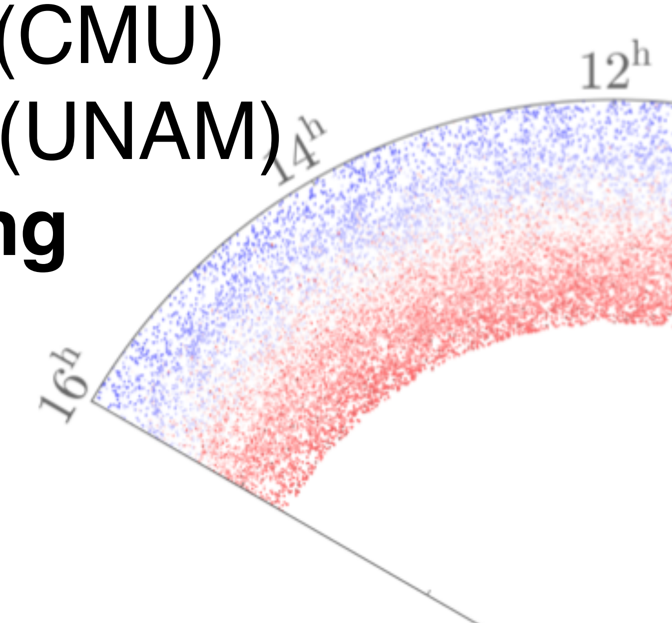

Our Universe as we know it ..

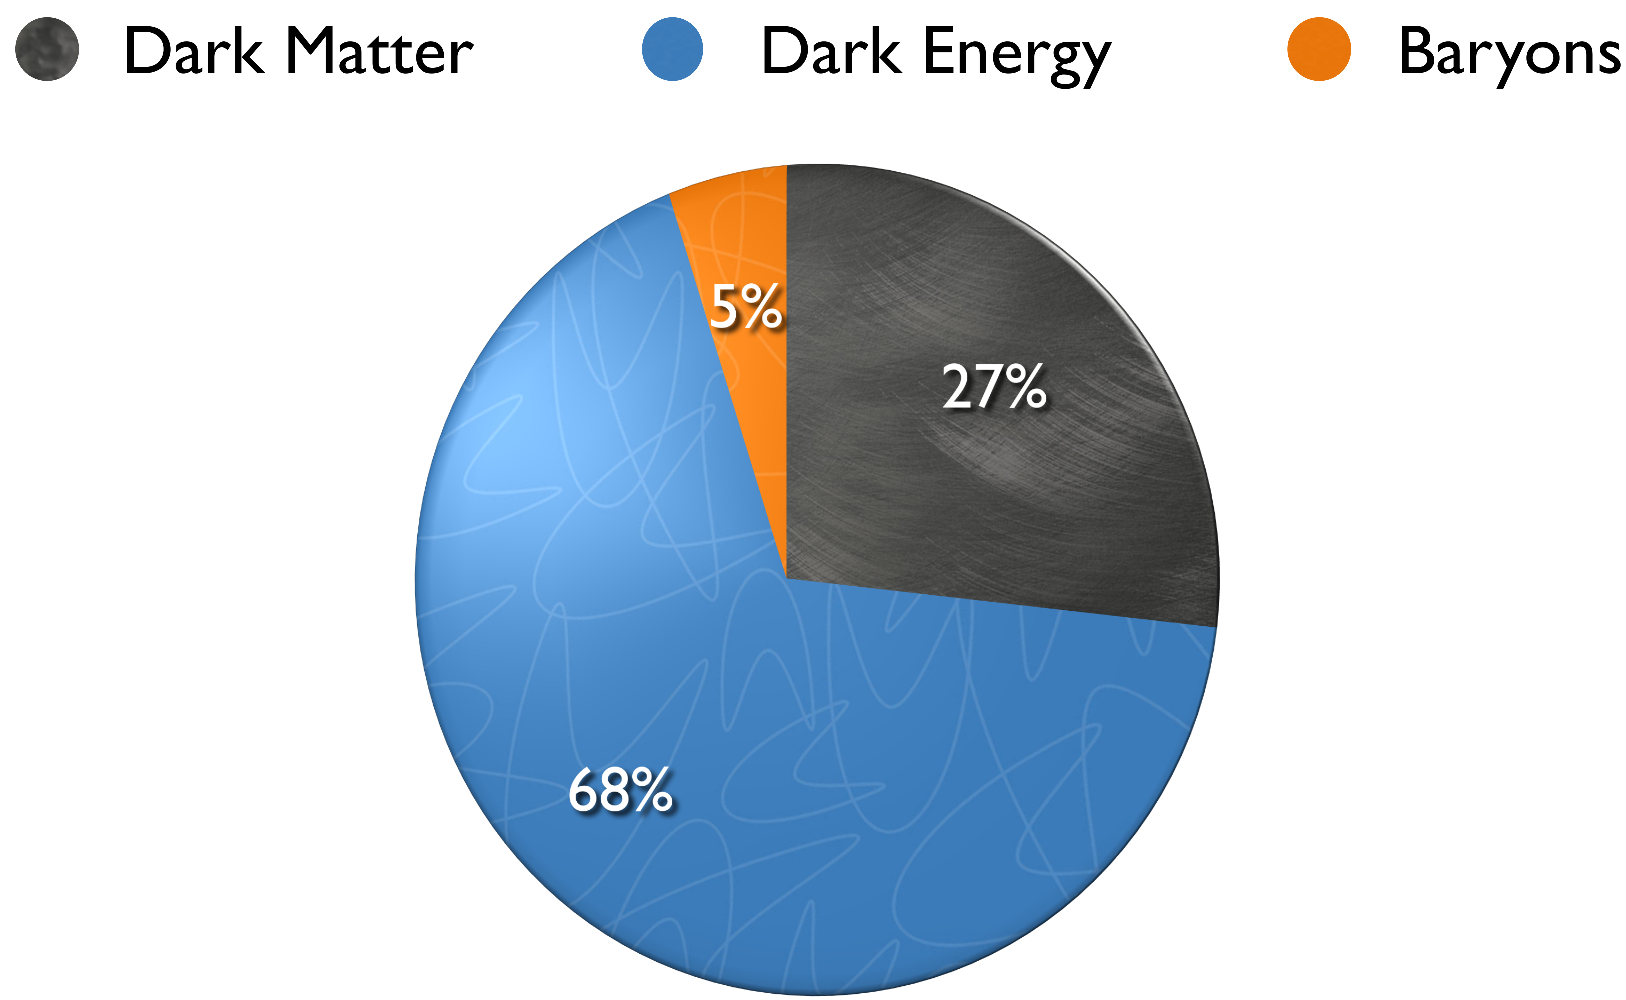

Our Universe as we know it ..

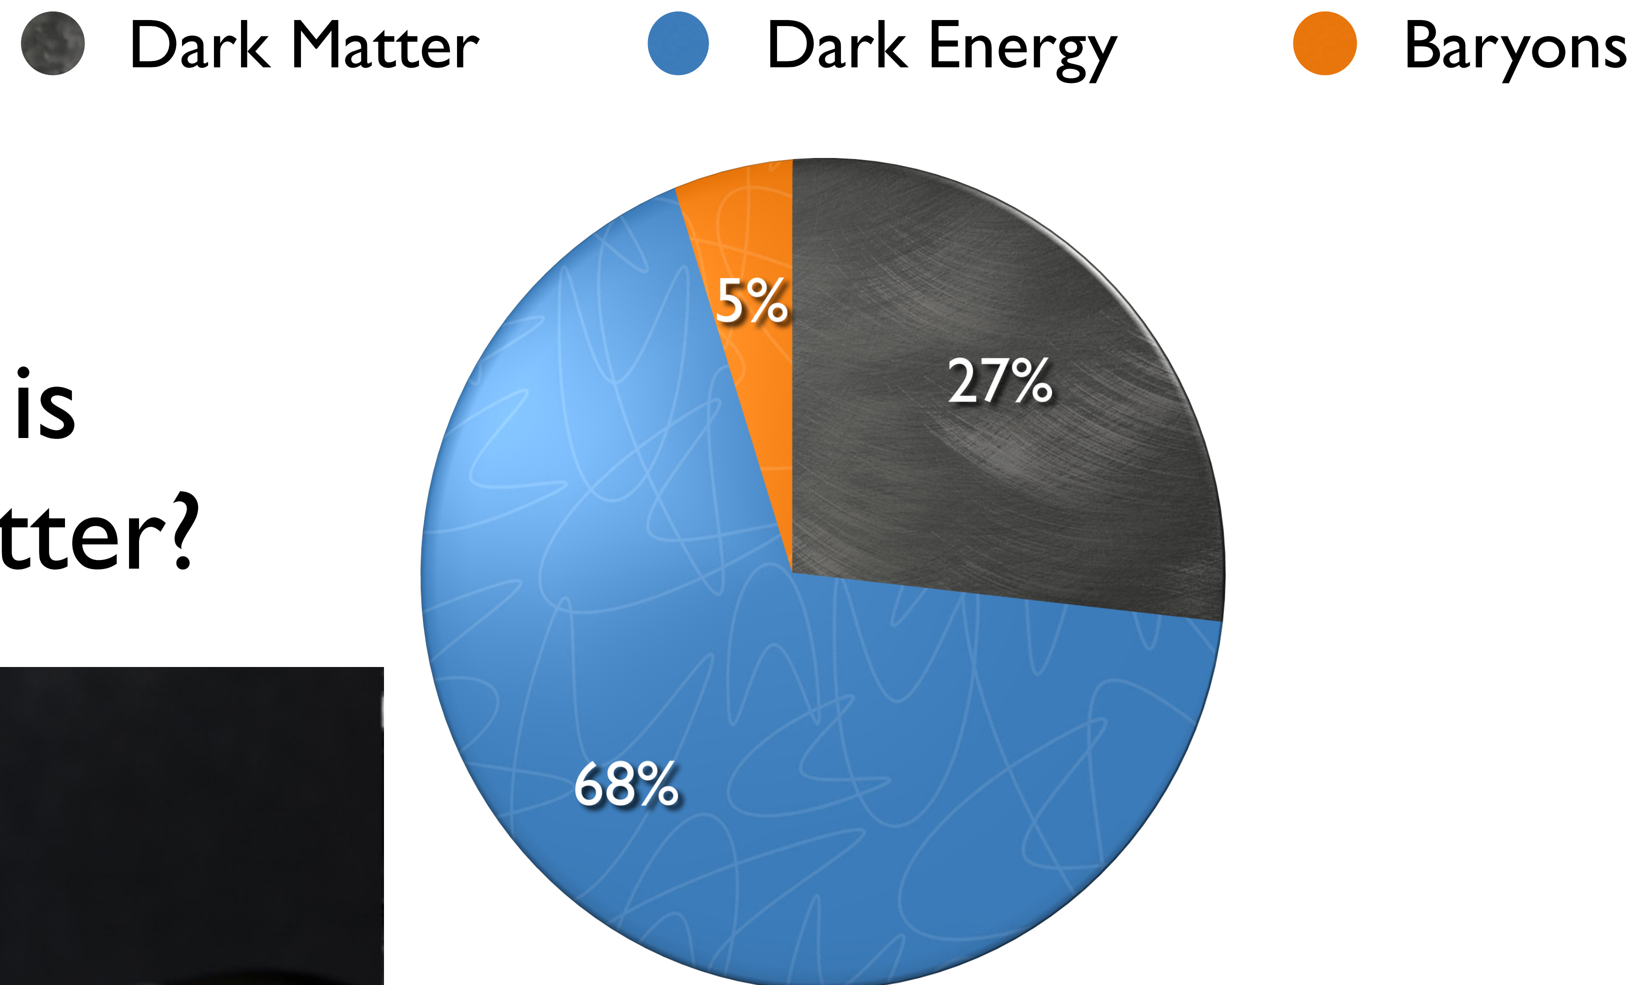

What is Dark Matter?

In case you're wondering, dark matter and dark energy are not Star Trek concepts – they're real forms of energy and matter; at least that's what most astrophysicists claim. Dark matter is a kind of matter hypothesized in astronomy and cosmology to account for gravitational effects that appear to be the result of invisible mass. The problem with it is that it cannot be directly seen with telescopes, and it neither emits nor absorbs light or other electromagnetic radiation at any significant level.

Our Universe

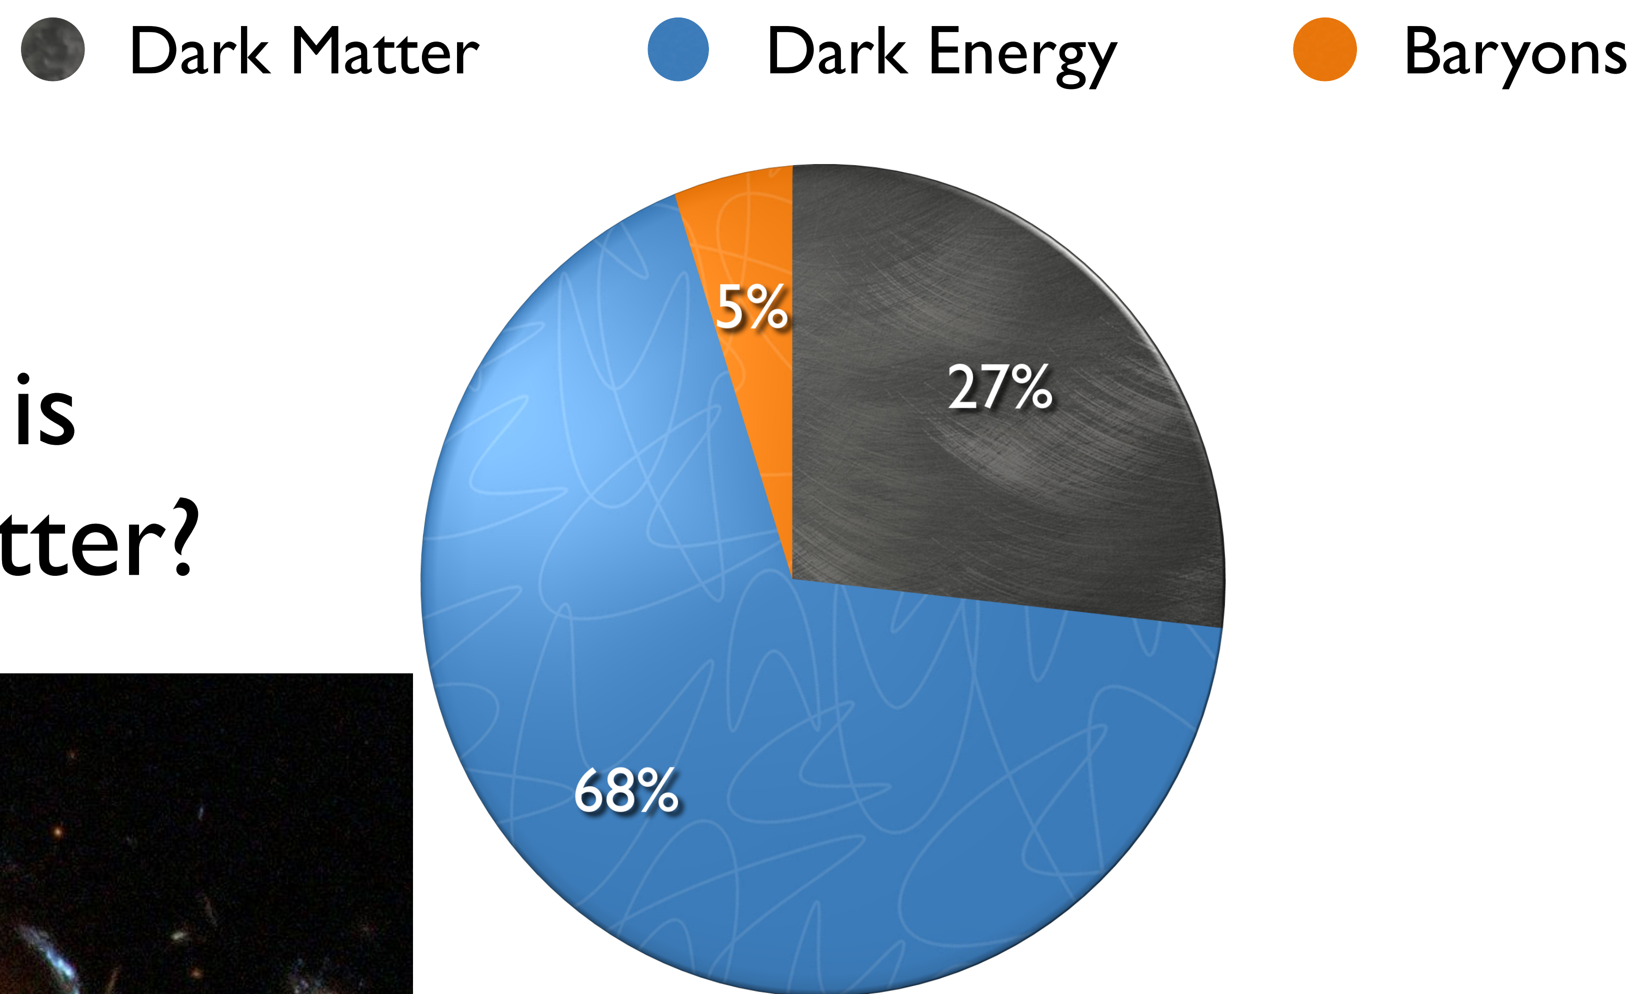

What is Dark Matter?

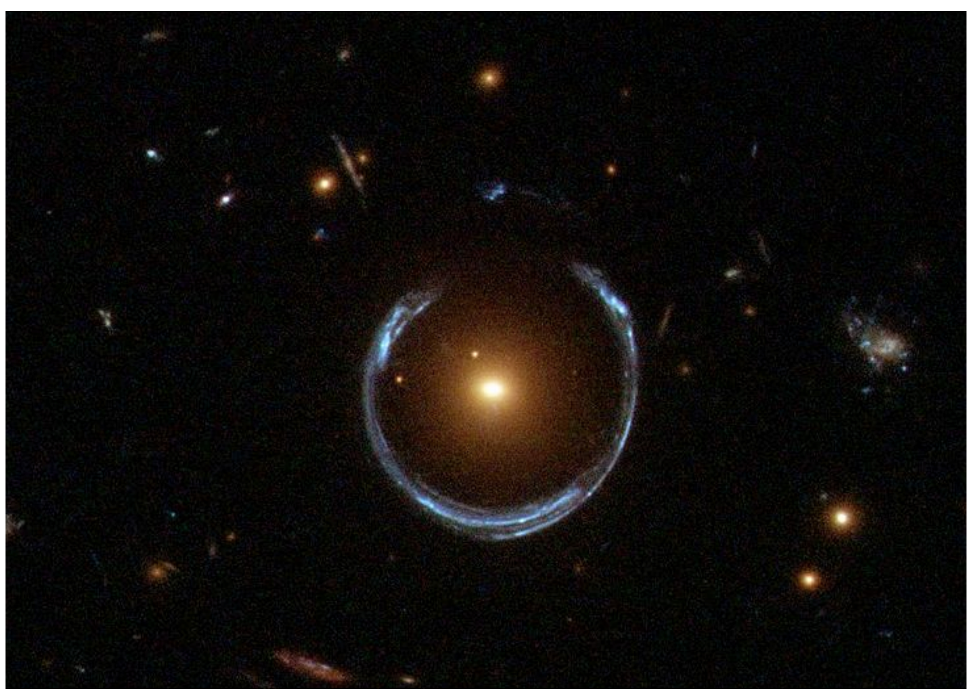

Maybe WIMP?
LHC is looking for this, but
maybe best bet is in cosmology?

Our Universe

What is
Dark Matter?

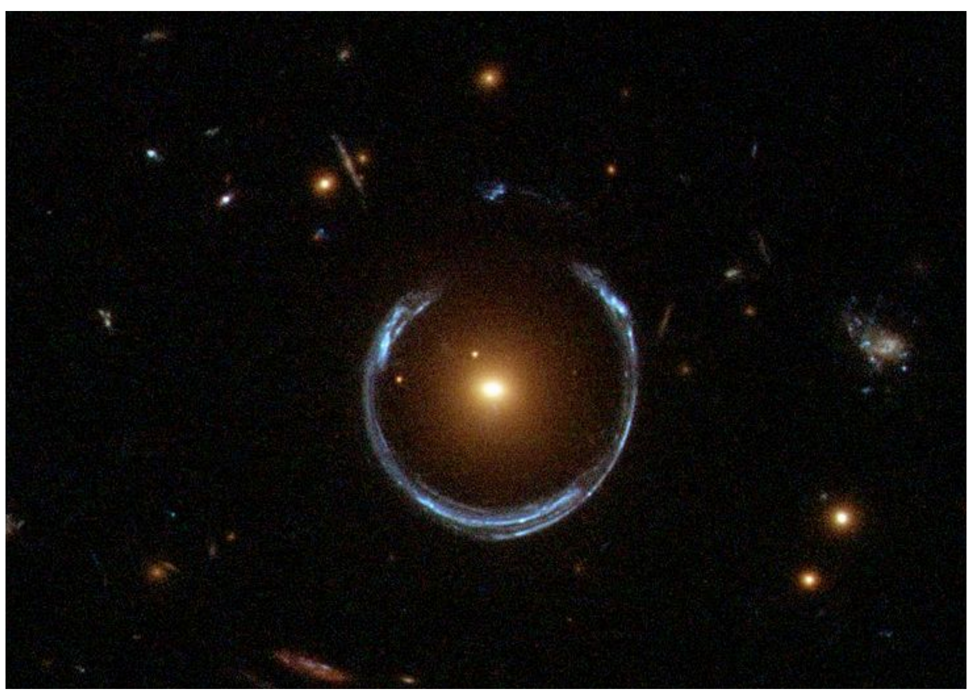

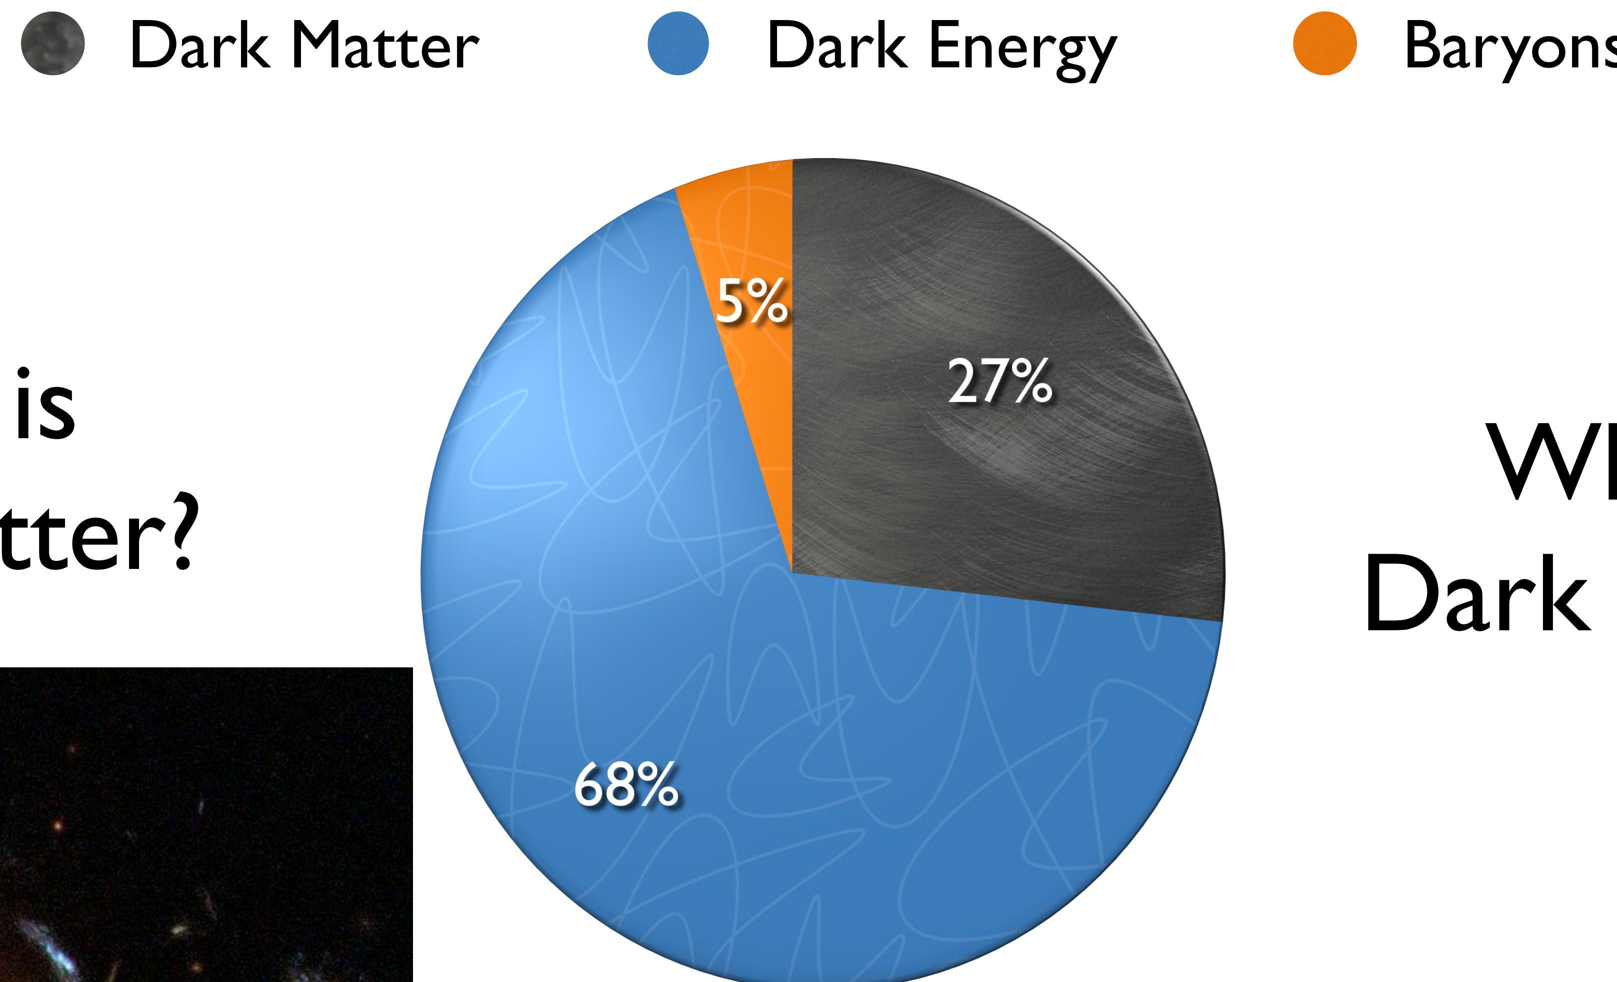

What is
Dark Energy?

Our Universe

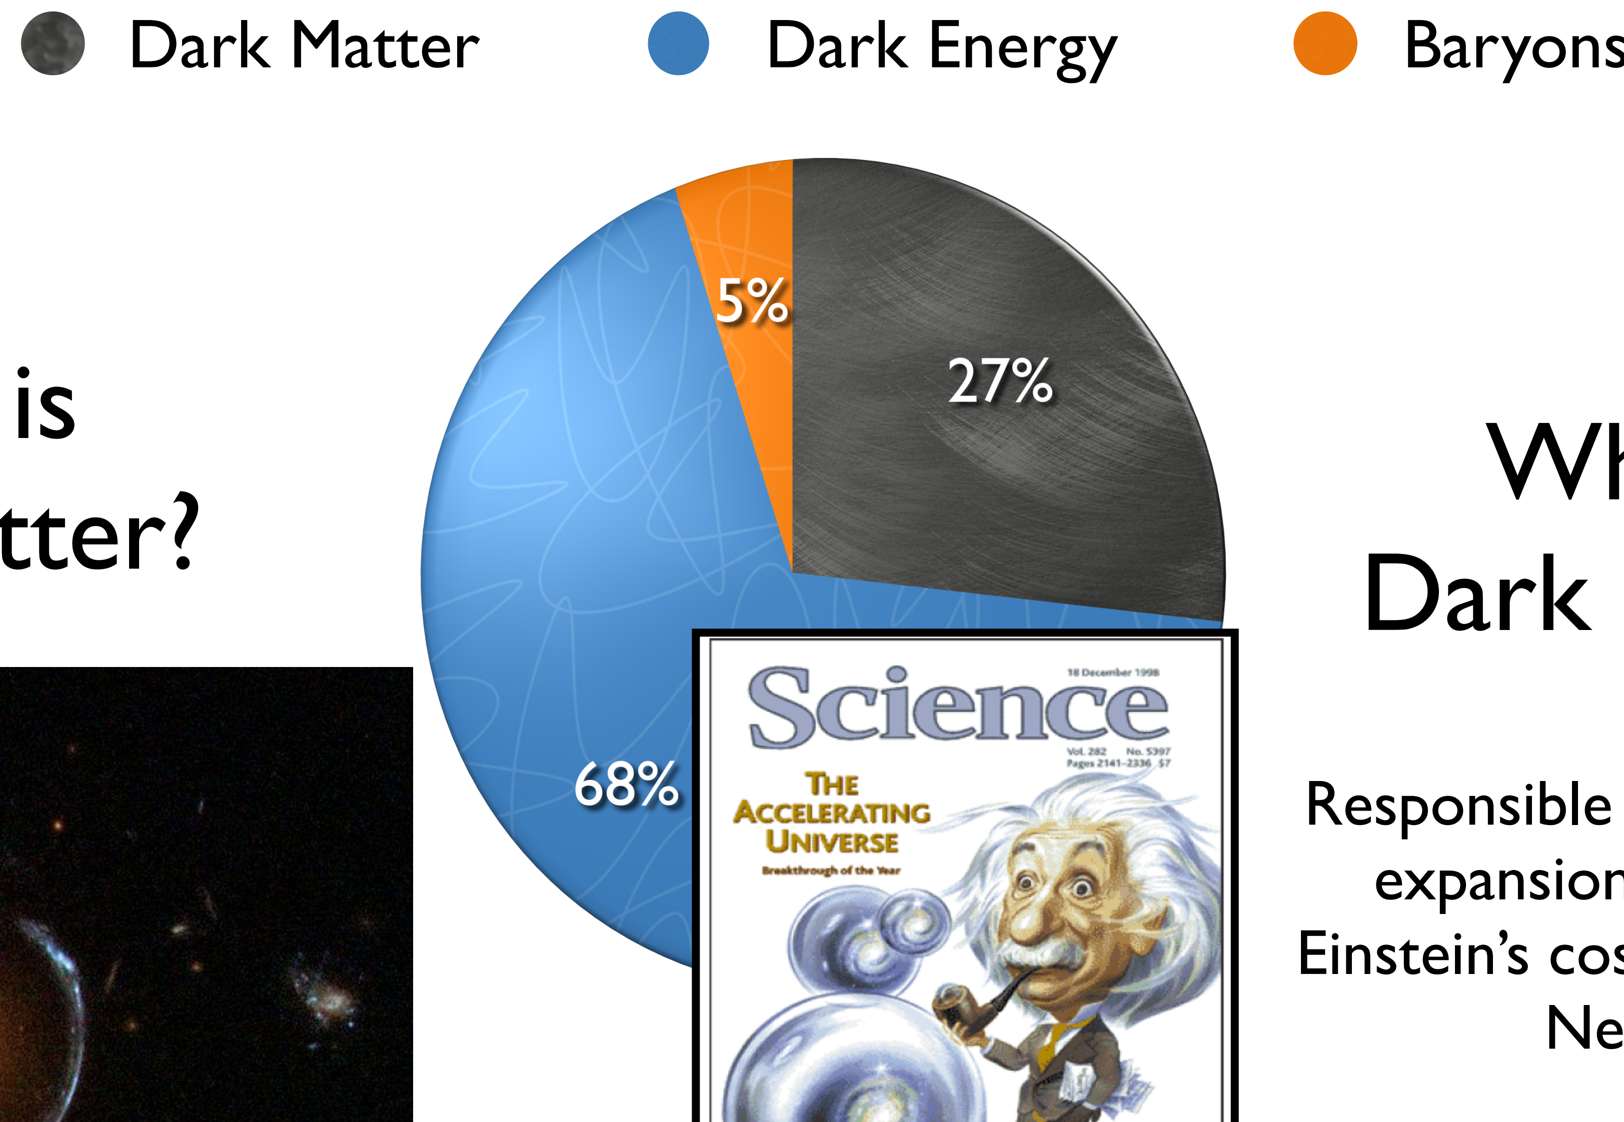

What is Dark Matter?

What is Dark Energy?

Responsible for accelerating the expansion of the Universe.
Einstein's cosmological constant?
New Physics?

Meanwhile, dark energy is a hypothetical form of energy which permeates all of space and tends to accelerate the expansion of the universe. Basically, ever since the 1990s, observations have revealed that the Universe is expanding at an accelerating rate. This baffled researchers; ok, it's clear that it expands, but why is it expanding faster? If anything, it should expand slower, due to all the gravitational attraction. Well, dark energy is the most accepted hypothesis to explain the observations since the 1990s indicating that the universe is expanding at an accelerating rate. The evidence for dark energy is indirect, just like with dark matter. Dark energy is thought to be very homogeneous, not very dense and have a negative pressure (acting repulsively) in order to explain the observed acceleration of the expansion of the universe.

Our Universe

Dark Matter · Dark Energy · Baryons

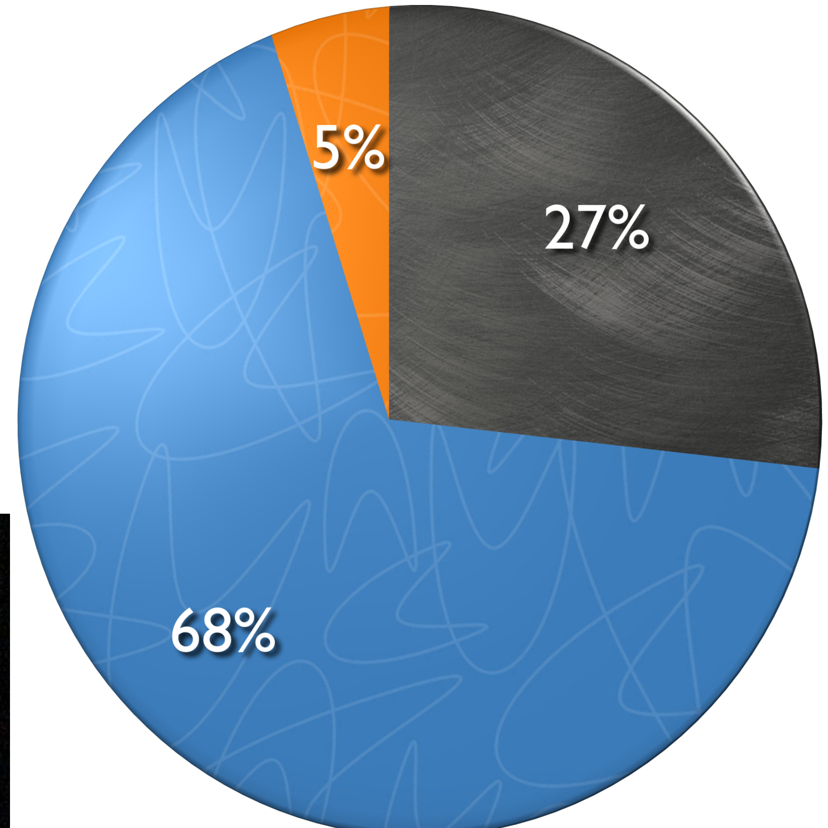

What is Dark Matter?

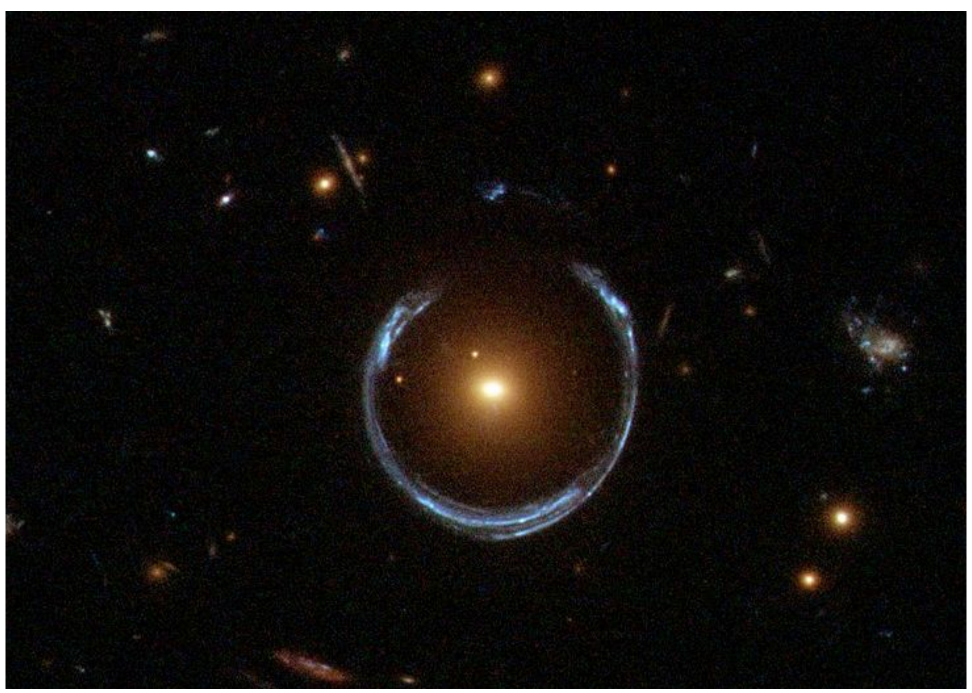

What is Dark Energy?

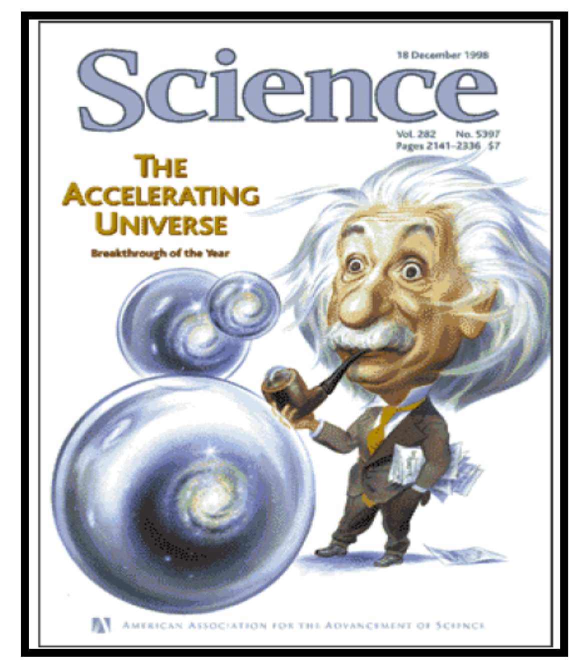

Current Cosmology analysis

Distribution of matter in the Universe

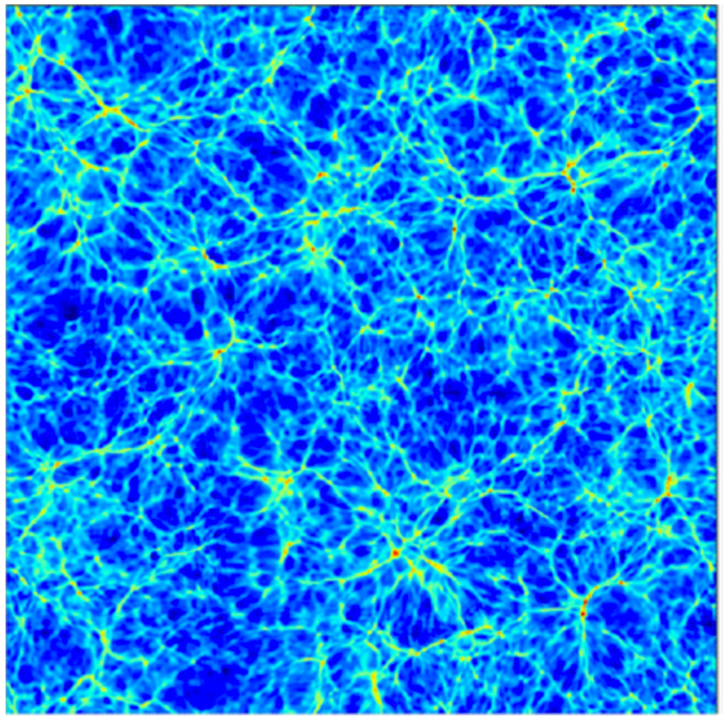

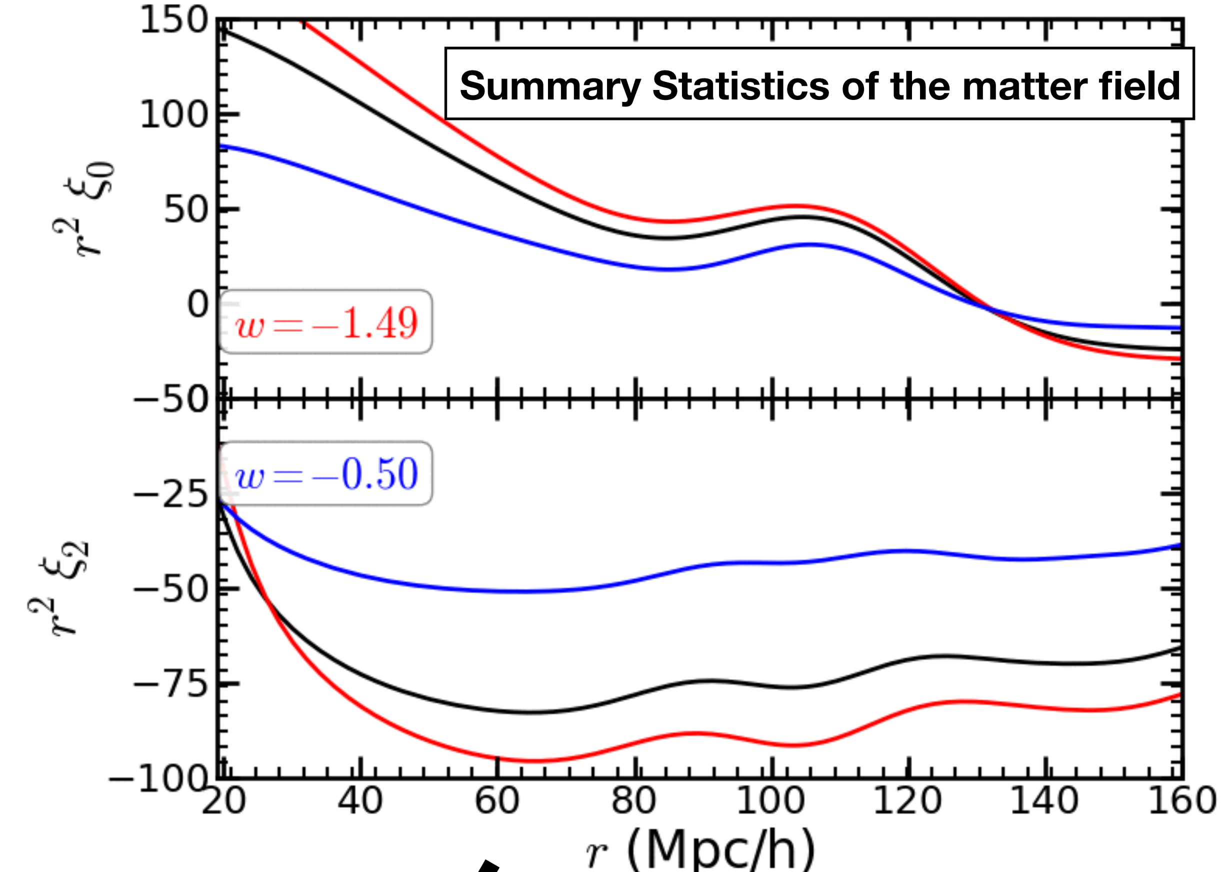

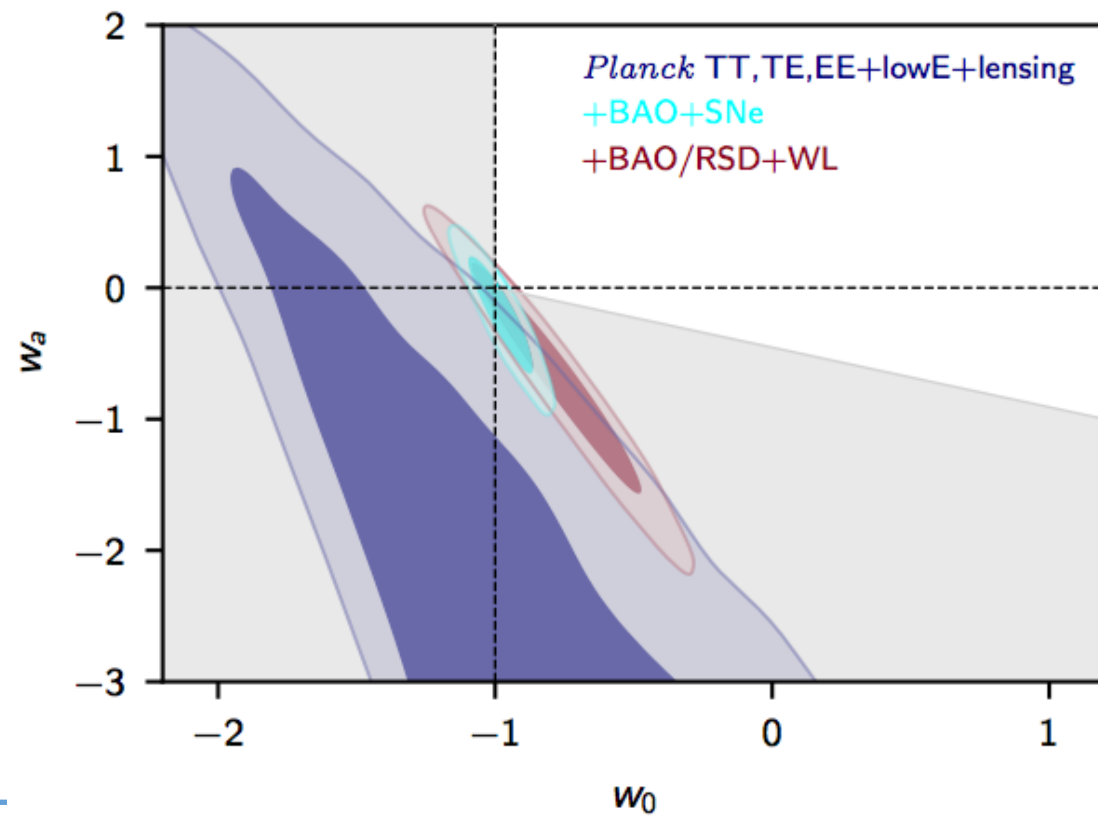

Constraints on Dark Energy

Planck Collaboration paper I 2018

Current Cosmology analysis

Distribution of matter in the Universe

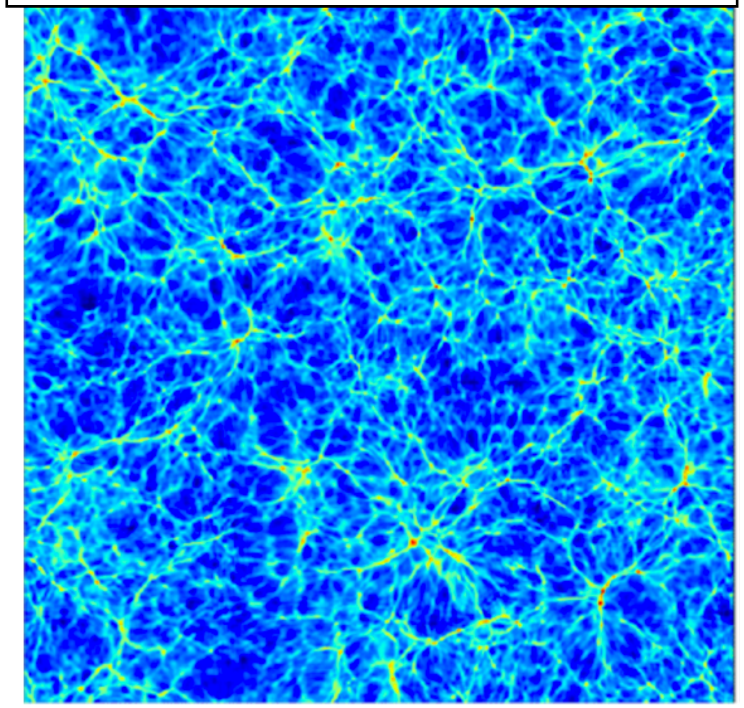

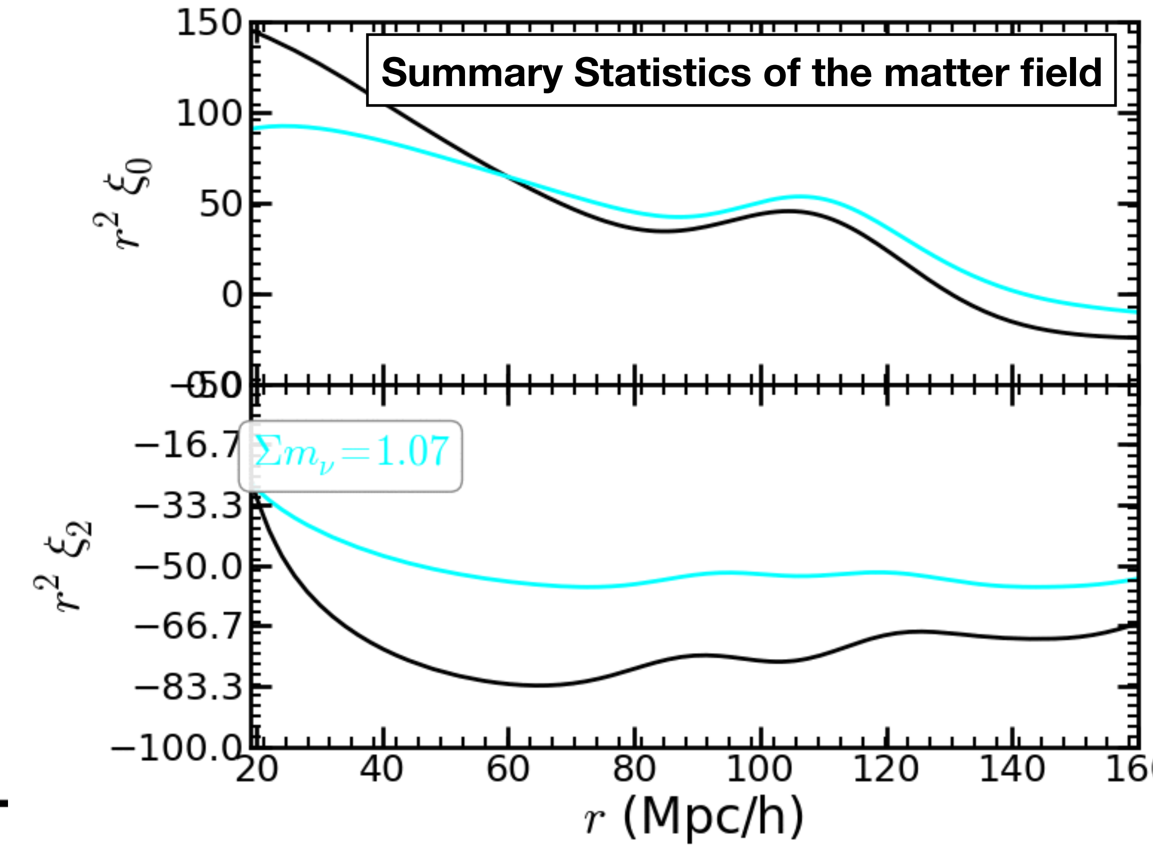

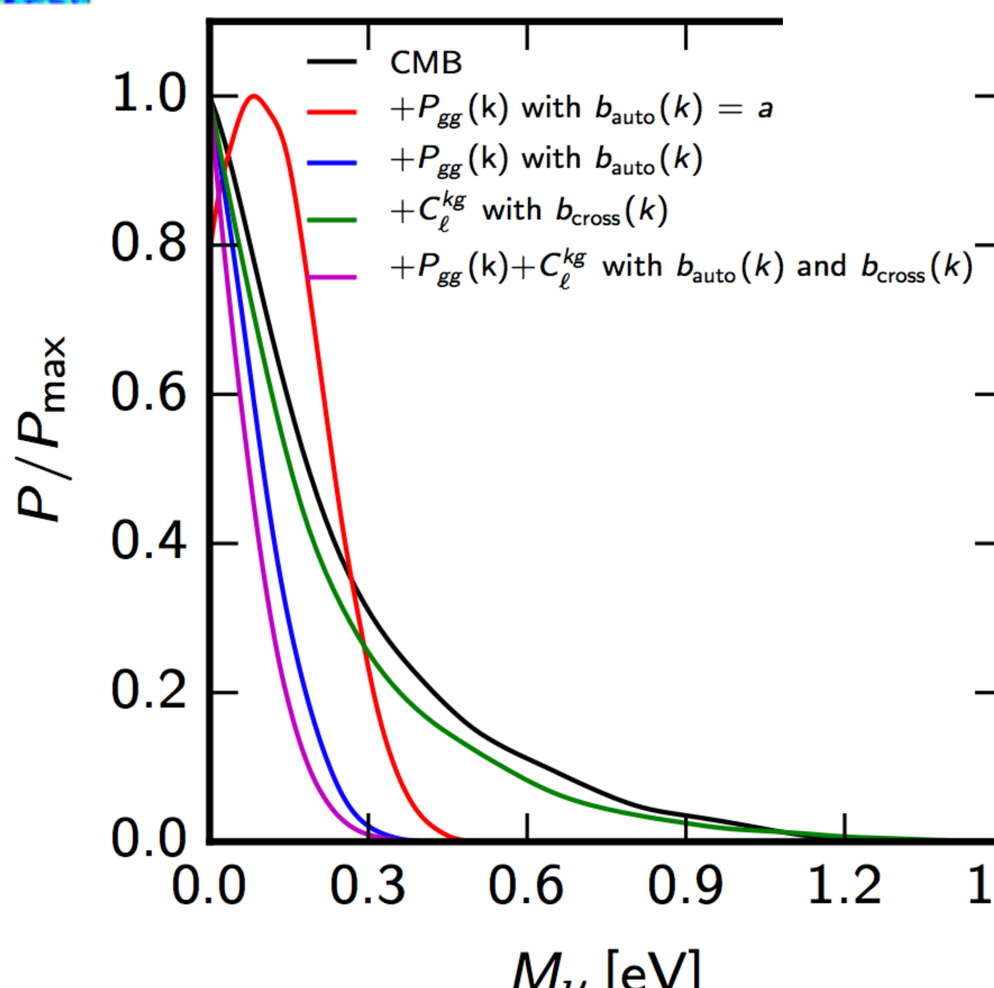

Constraints on neutrino masses

Giusarma,Vagnozzi et al. 2018

Can we do better than
what we have done before?

Mastering the Game of Go without Human Knowledge

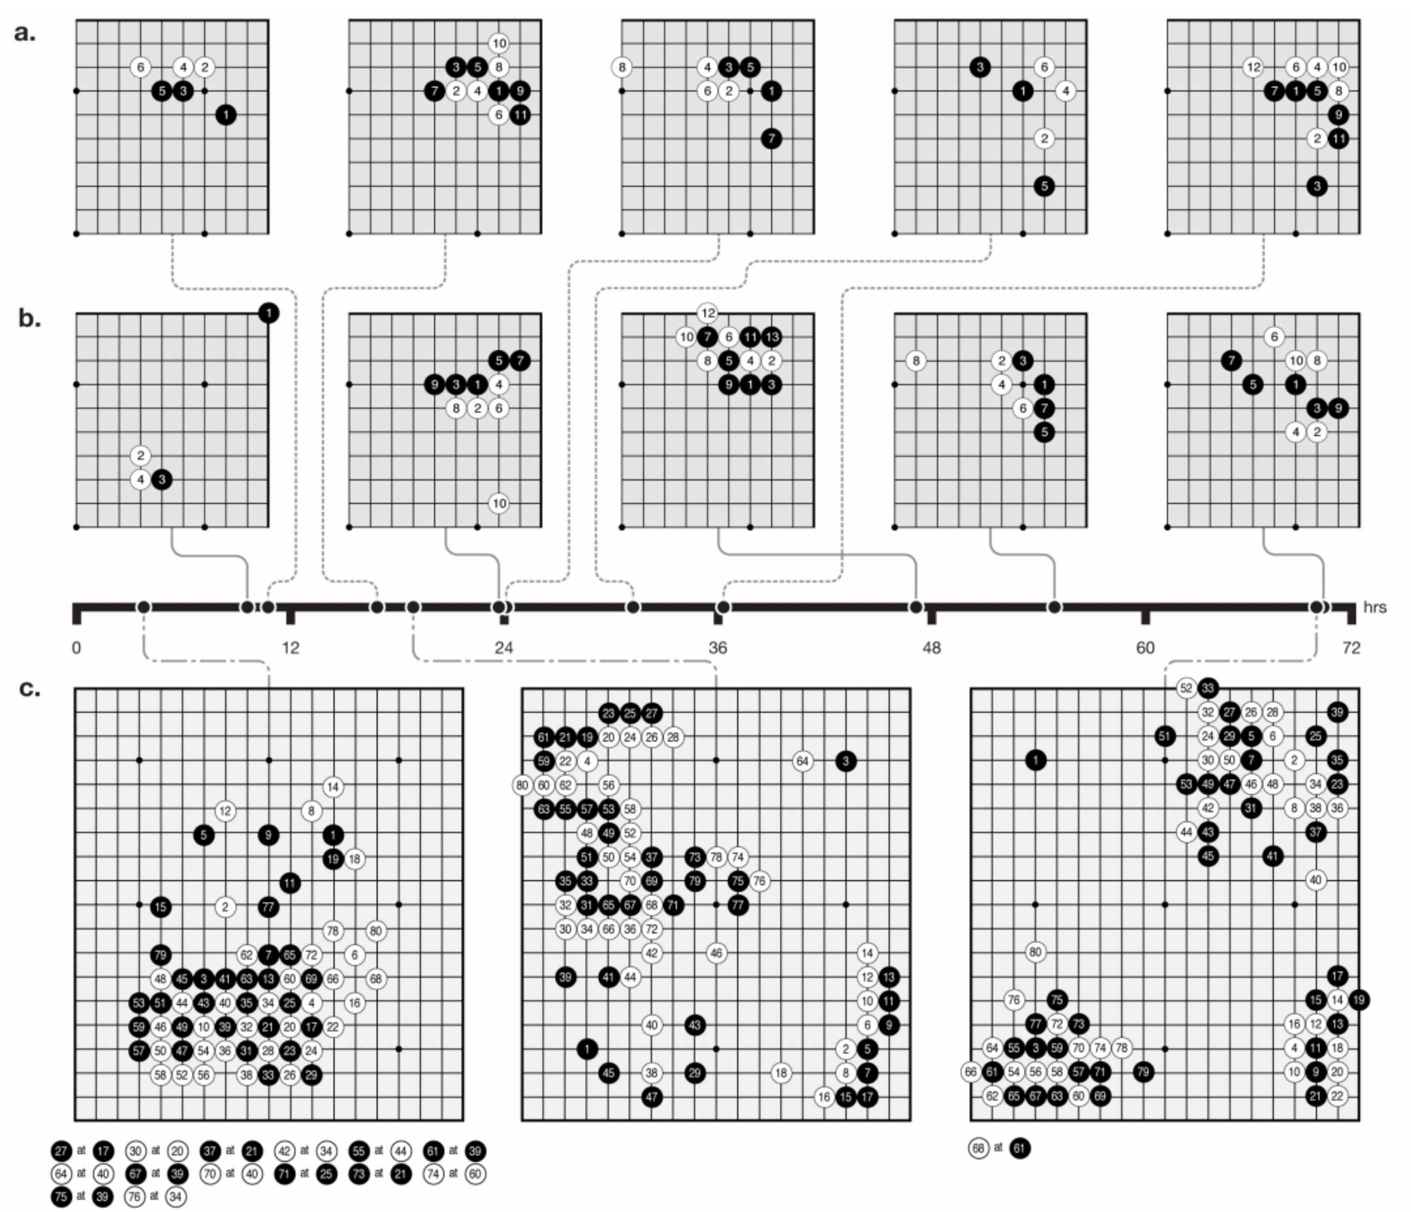

Five human joseki (common corner sequences)
discovered by AlphaGo during training.
Silver, Schrittwieser, Simonyan Nature 2016

Machine learning in image recognition

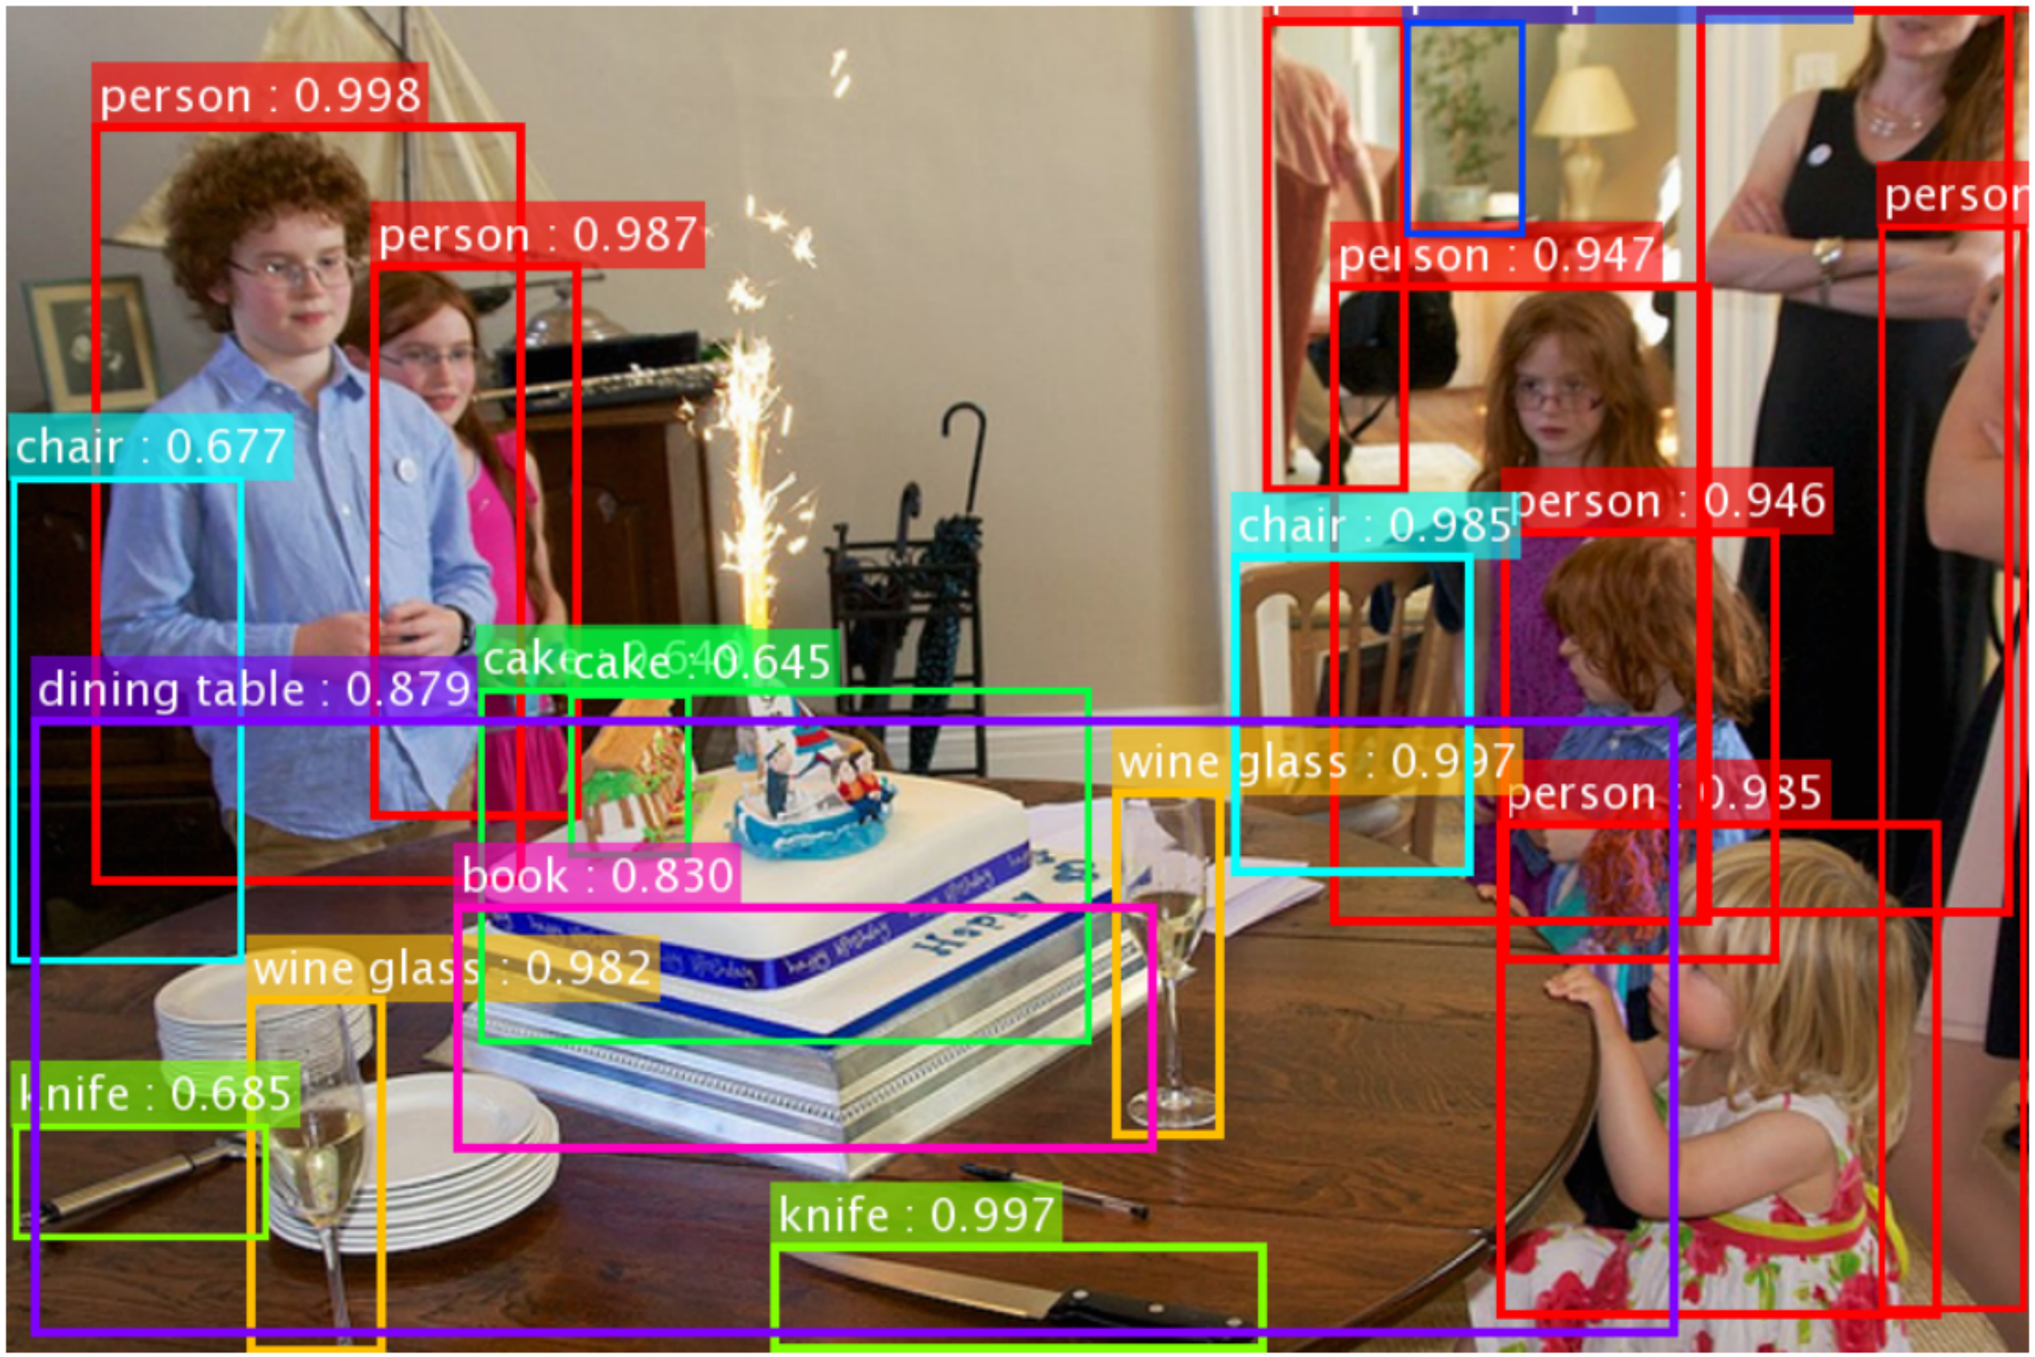

ResNet's object detection result on Common Object in Context

Kaimin He (Facebook Research, now Microsoft Research Asia) et al. 2016

Can machine learning help us understand the Universe?

Can we use Machine Learning to help us understand the Universe?
Extracting more information from the astronomical dataset

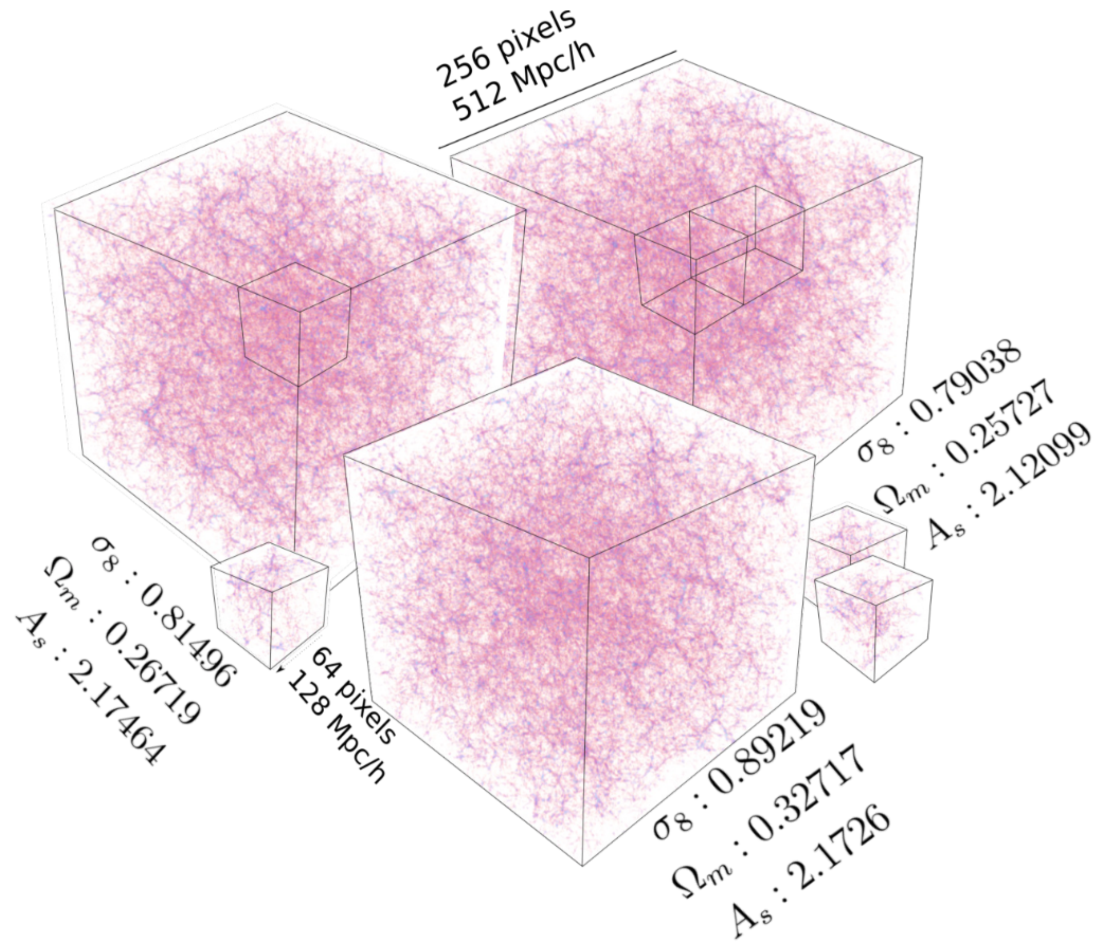

Figure 1. Dark matter distribution in three cubes produced using different sets of parameters. Each cube is divided into small sub-cubes for training and prediction. Note that although cubes in this figure are produced using very different cosmological parameters in our constrained sampled set, the effect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

Ravanbakhsh, Oliver, Price, **Ho**, Schendier & Poczos
International Conference of Machine Learning 2016

Can we use Machine Learning to help us understand the Universe? Introducing our machine learning network (Convolutional Neural Net)

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

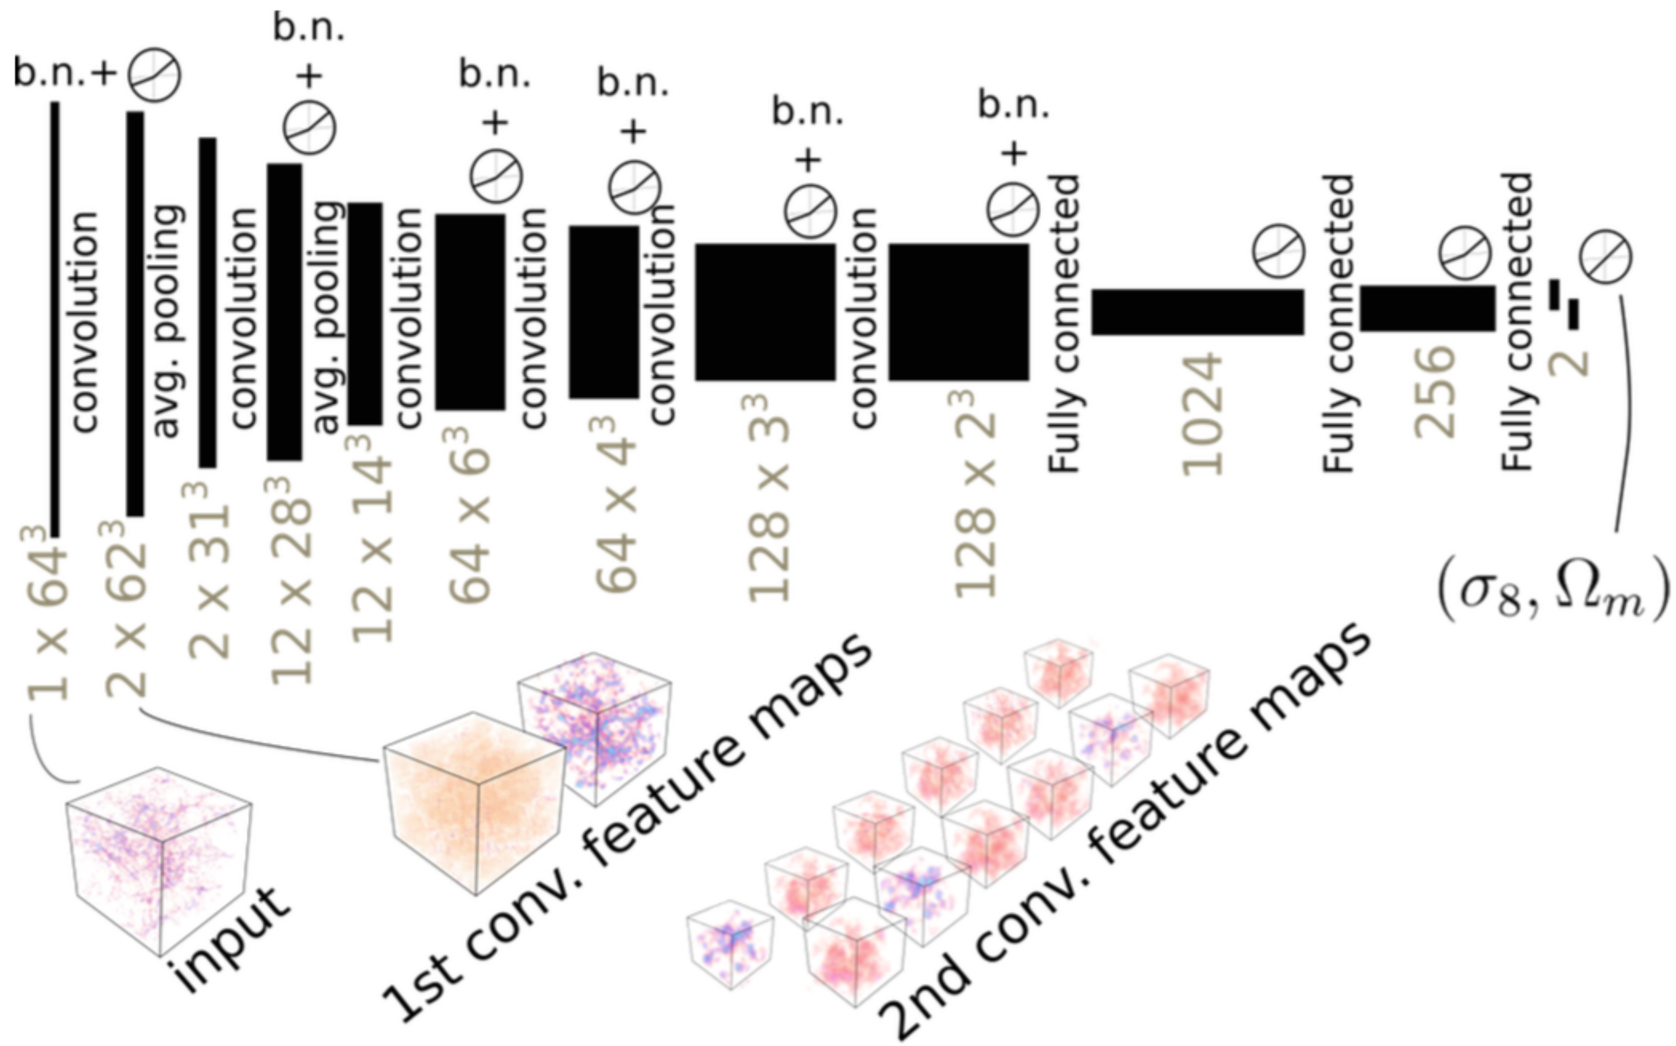

Figure 6. The architecture of our 3D conv-net. The model has six convolutional and 3 fully connected layers. The first two convolutional layers are followed by average pooling. All layers, except the final layer, use leaky rectified linear units, and all the convolutional layers use batch-normalization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

Can we use Machine Learning to help us understand the Universe? Training, Validation and Test

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

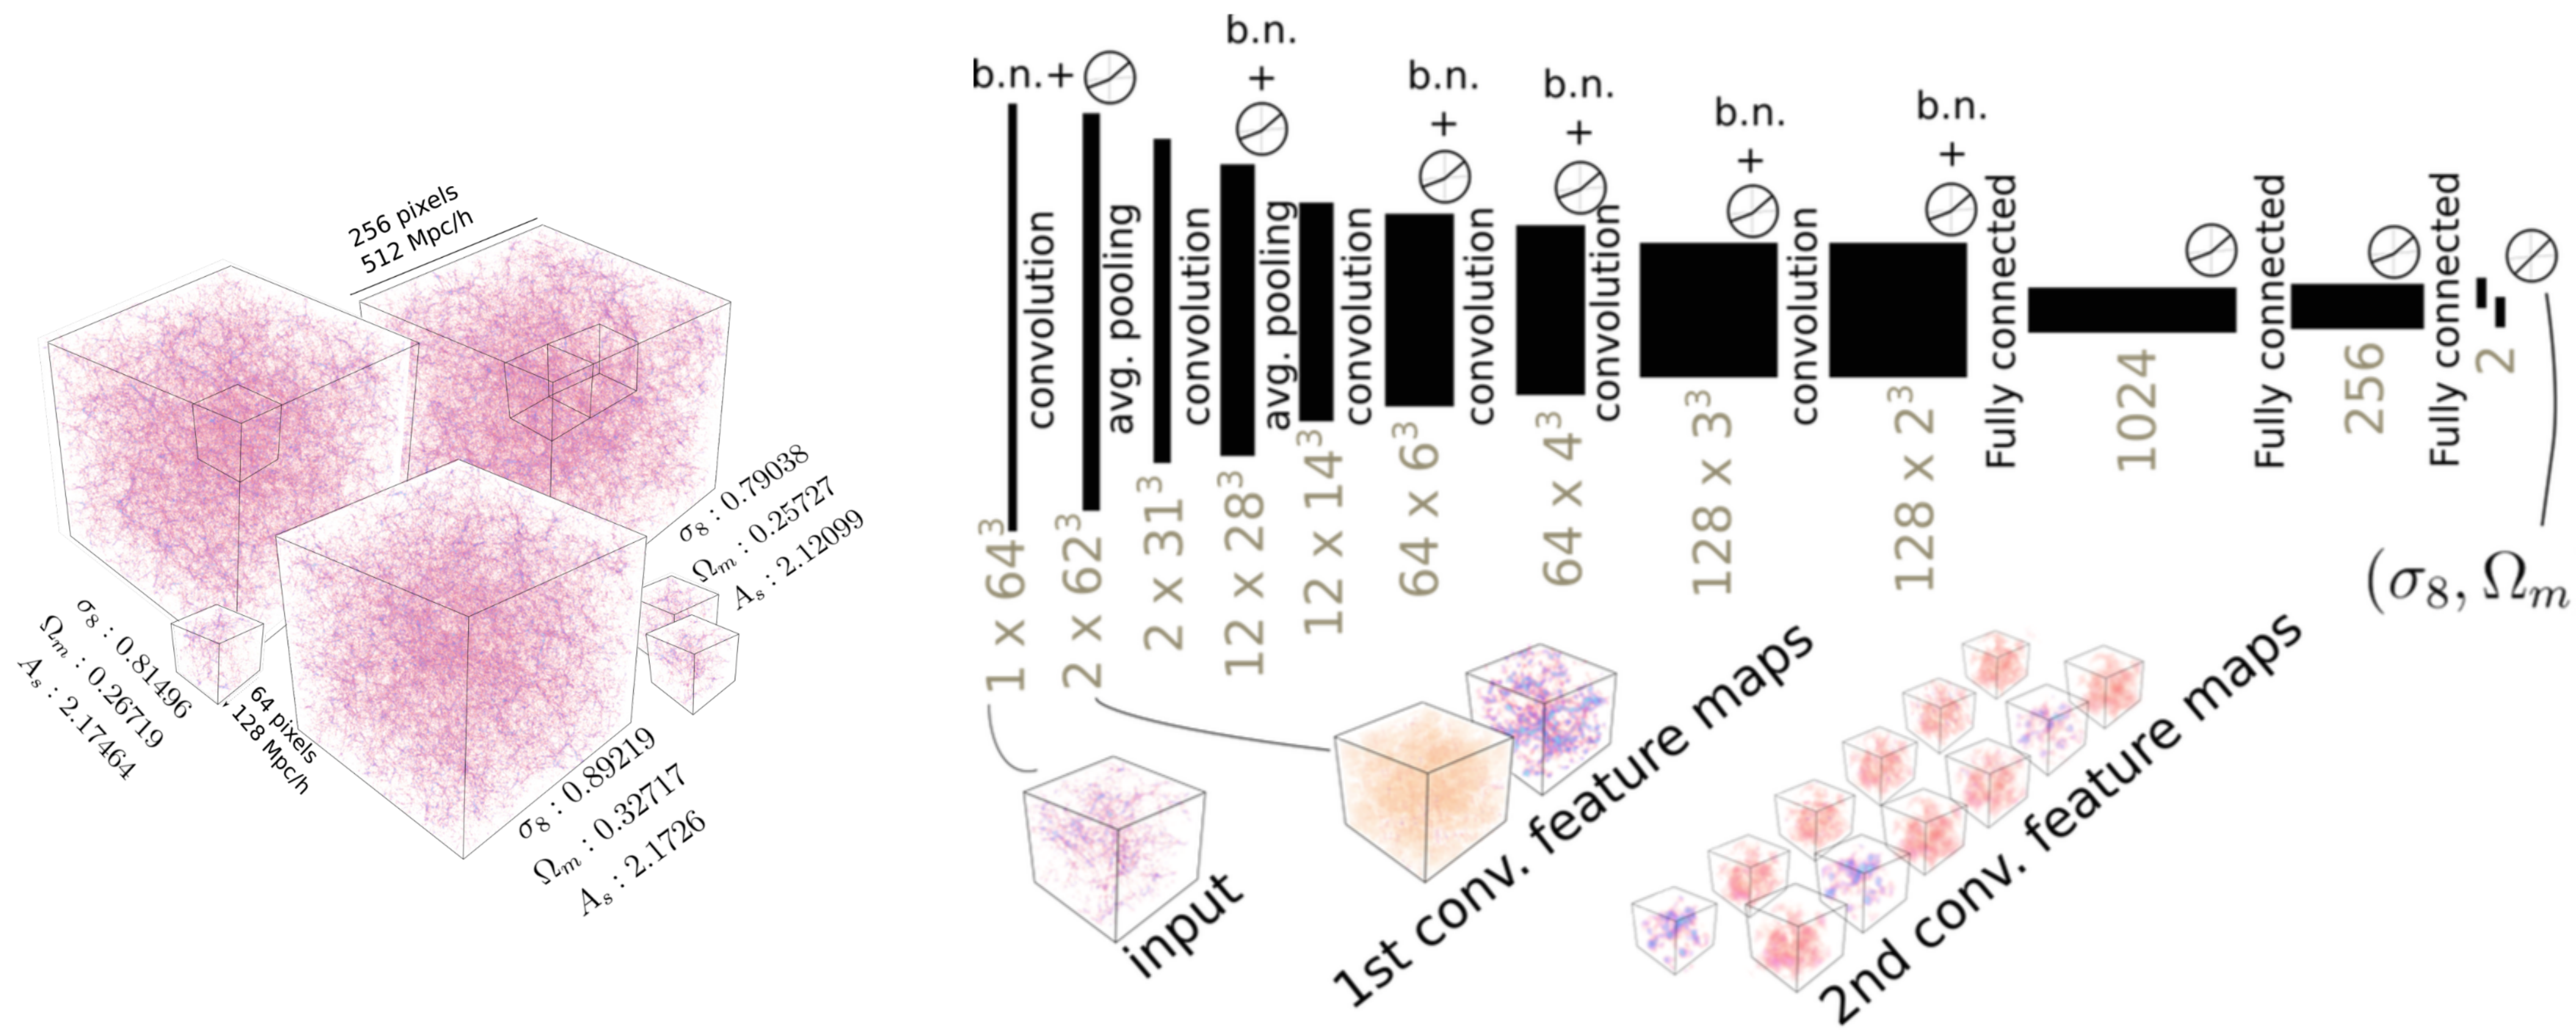

Training: Input N-body simulations with known cosmological parameters to train the ConvNet

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

Can we use Machine Learning to help us understand the Universe? It achieves higher accuracies than our traditional method.

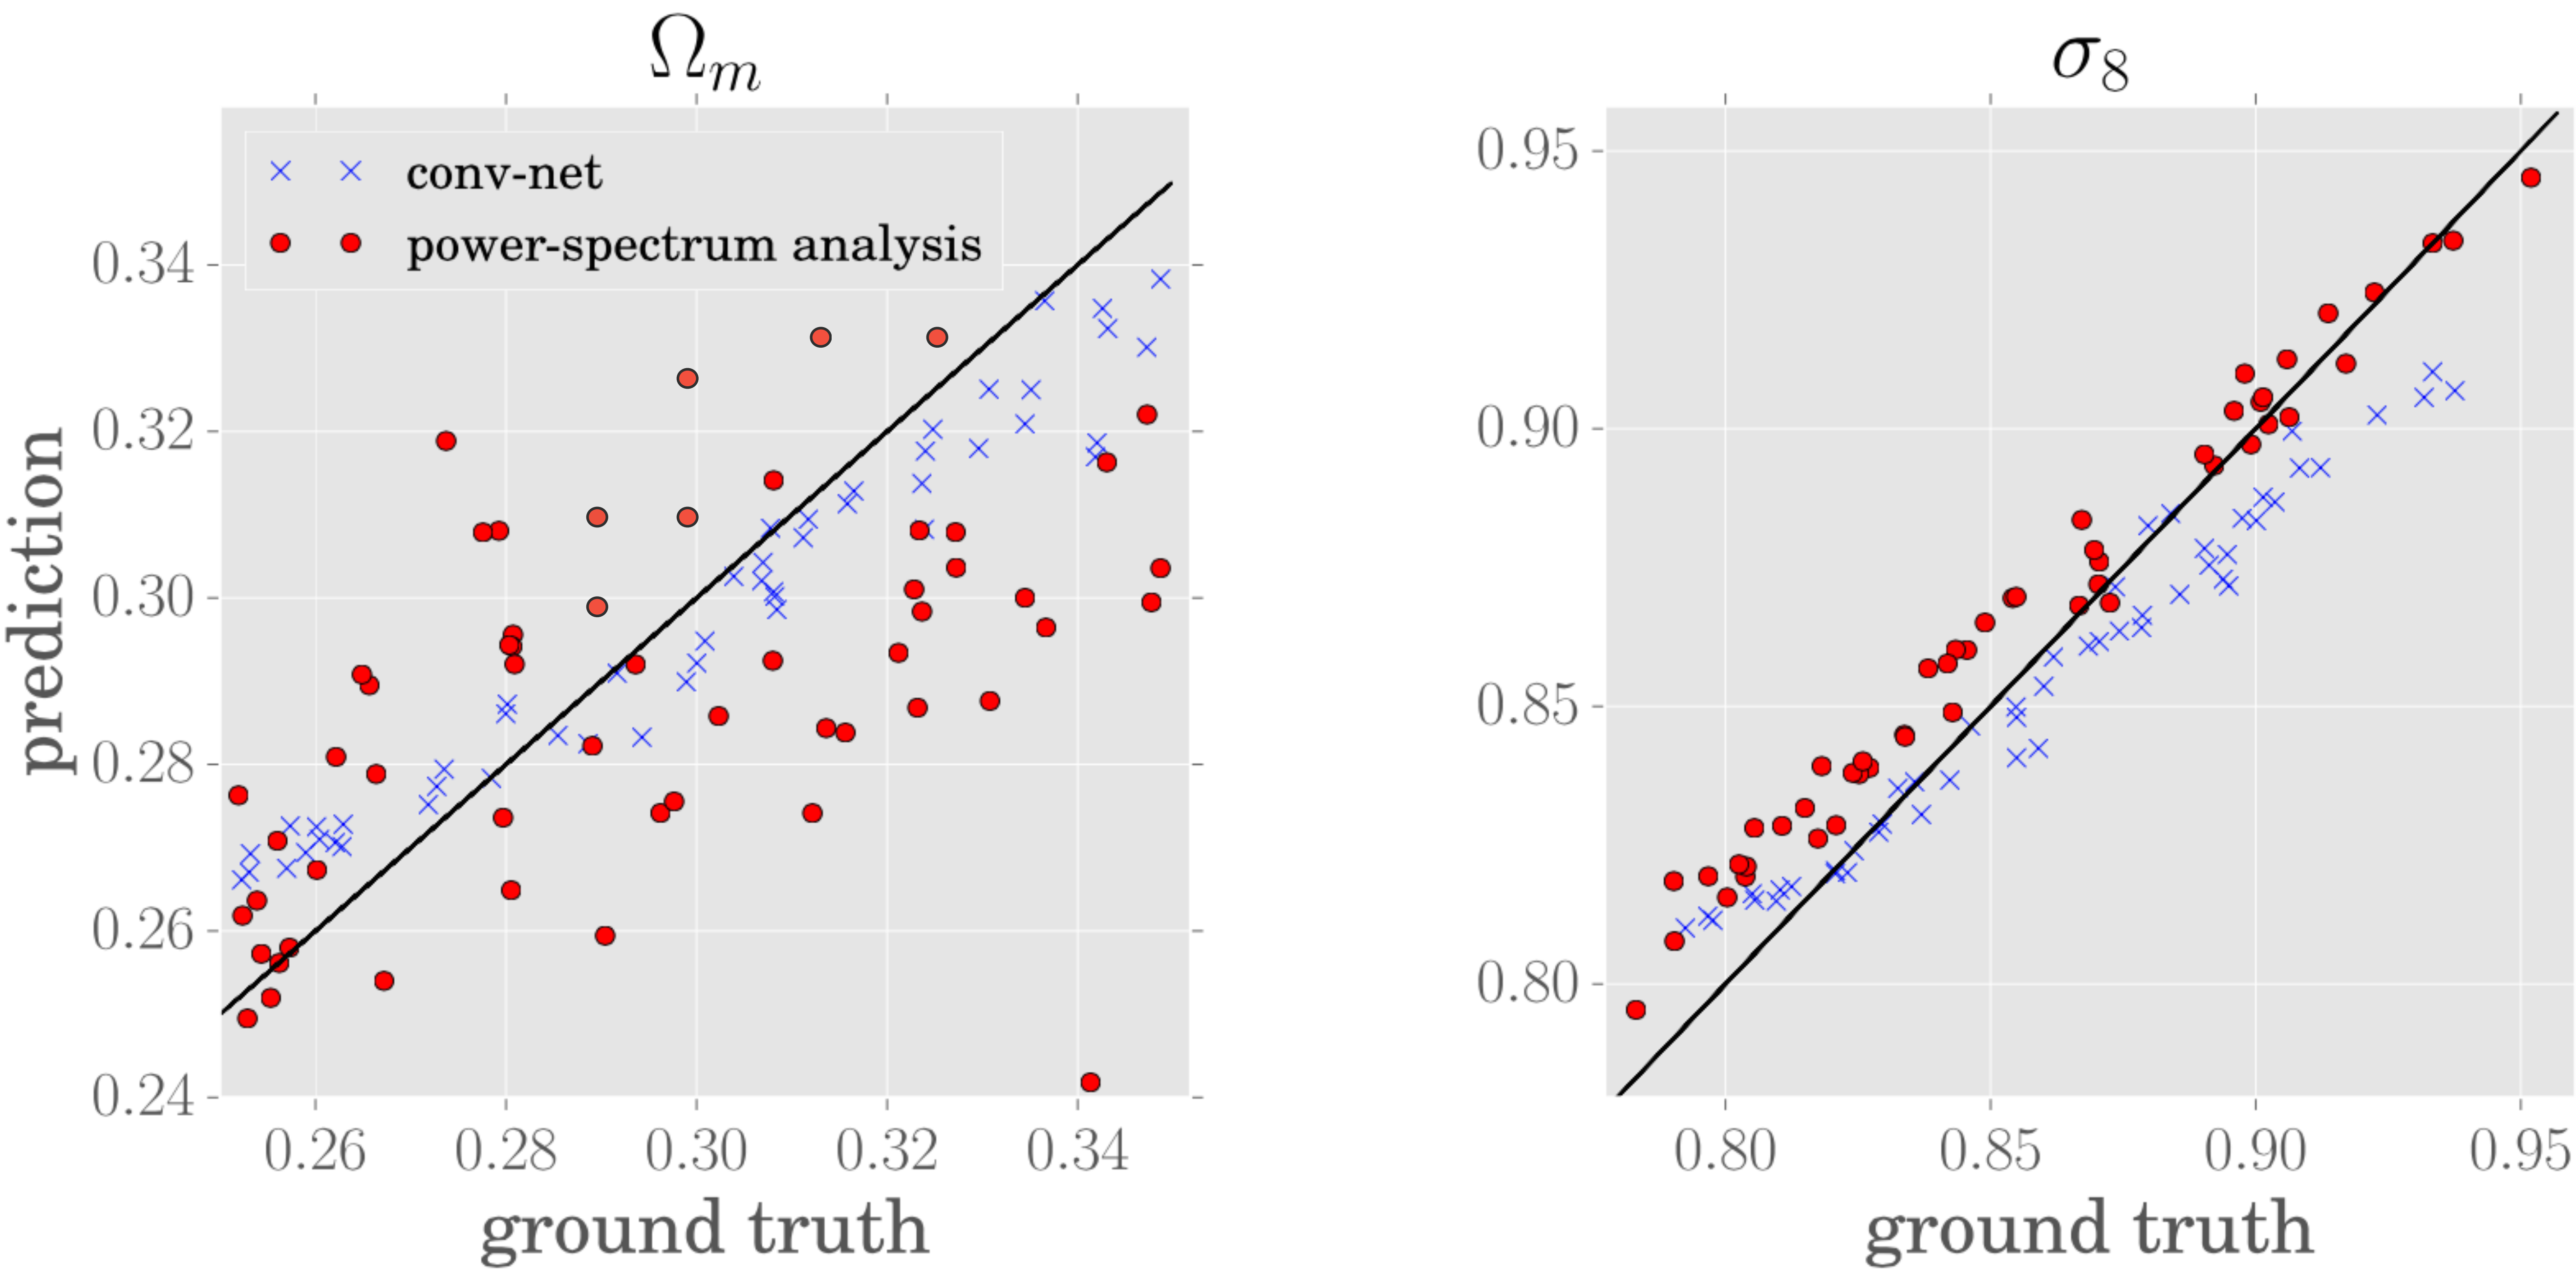

Ravanbakhsh, Oliver, Price, **Ho**, Schendier & Poczos **ICML** 2016

Now as scientists, we have lots of questions: AKA: Outline for the remaining of the talk

- Can we get a correct estimate of the error ?
    - See He, Ravanbaksh & Ho *International Conference for Learning Representations 2018*
- Can we interpret the model learnt in Machine Learning?
- What is the model learning?
- 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: AKA: Outline for the remaining of the talk

- Can we get a correct estimate of the error ?
    - See He, Ravanbaksh & Ho *International Conference for Learning Representations 2018*
- **Can we interpret the model learnt in Machine Learning?**
- What is the model learning?
- 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?

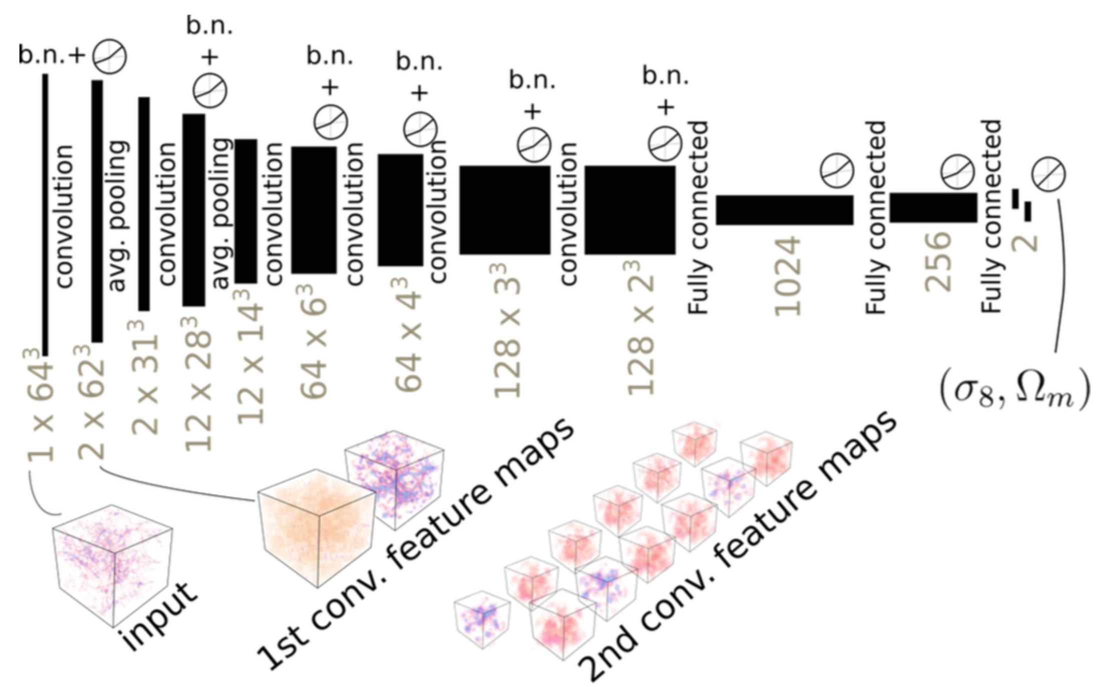

- We now design the following experiment, to learn the difference between **analytical modeling** and **the full information** in the density field.

- In other words: Can we understand what gravity does to billions of dark matter particles over many years, without using computer to simulate the physical laws step by step?

- Or: Can we use machine learning to skip the simulation of a complex physical system?

Can we use machine learning to simulate a physical system?

- There are recent work that tries to simulate simpler physical systems. See the next video by my collaborators.

- They are able to simulate Kepler's law quite well up to ~1000 time steps.

Battaglia, Pascanu, Lai, et al. NeurIPS 2016

Can we use machine learning to simulate the Universe ?

Using Machine Learning to simulate the Universe: The Setup of the Experiment

Inputs

Machine Learning model

Outputs

Analytical approximation of the non-linear evolution of the Universe

Using Machine Learning to simulate the Universe: The Setup of the Experiment

Inputs

Machine Learning model

Outputs

Positions and velocities of all particles, evolved under gravity after X years

Using Machine Learning to simulate the Universe: The Setup of the Experiment

Inputs

Machine Learning model

Outputs

Instead of using numerical simulations of newton's laws for all the particles, with smart algorithms to run really fast.
We will attempt to use machine learning to "learn"/ interpolate from a large number of pre-run simulations.
We call these "training data".

From Analytical approximated (Zeldovich approximation) fields to numerically simulated (FastPM) fields

Inputs

Machine Learning model

Outputs

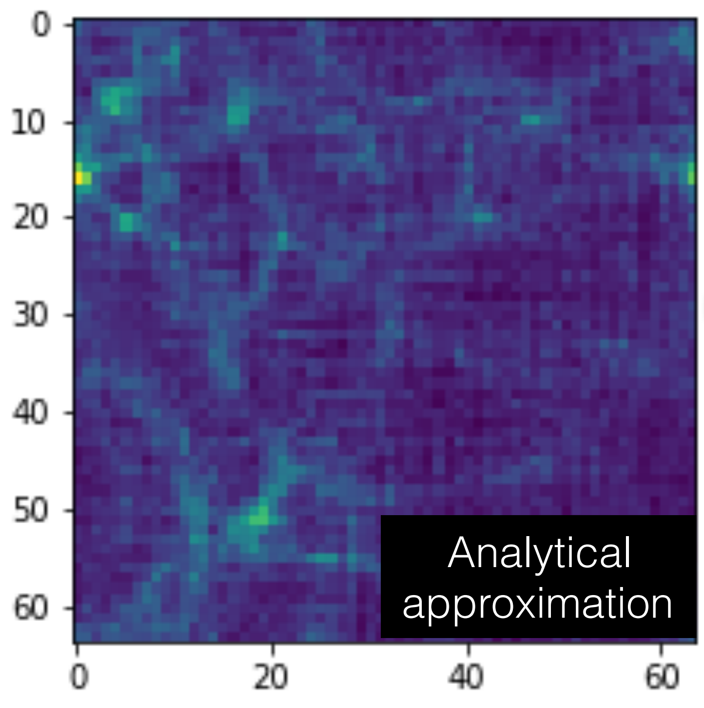

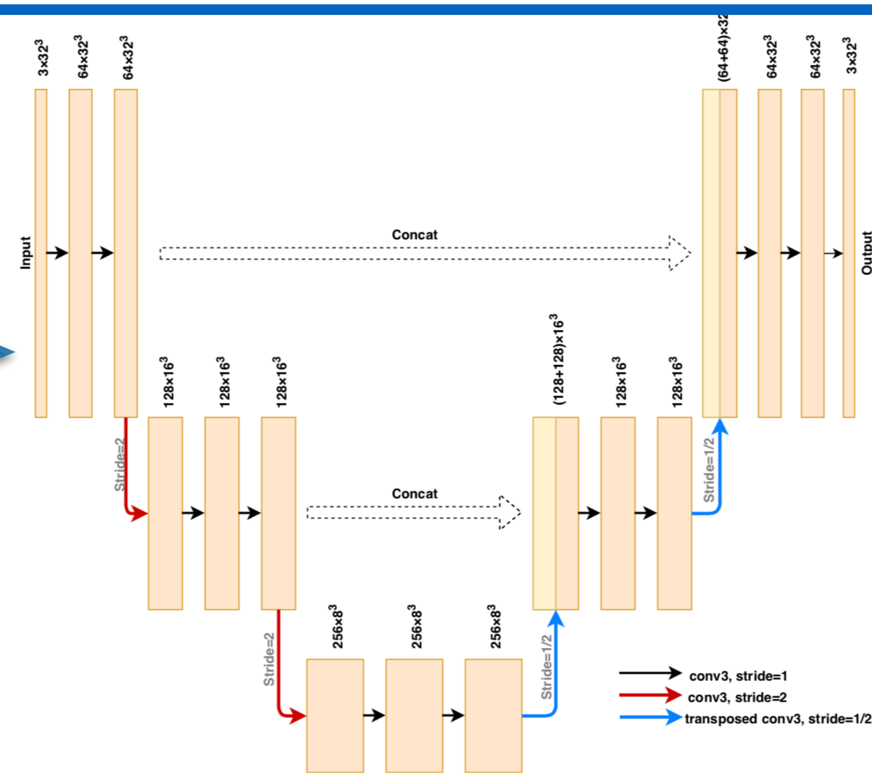

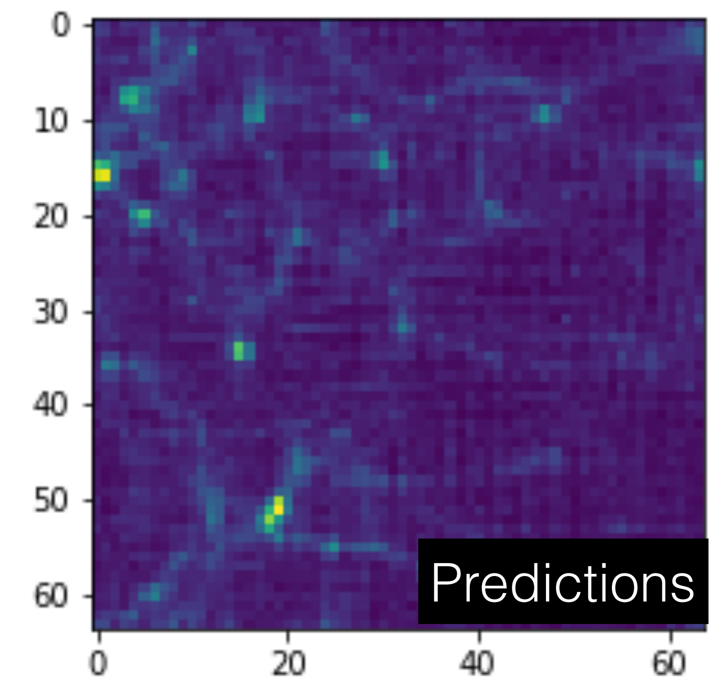

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

From Analytical approximated (Zeldovich approximation) fields to numerically simulated (FastPM) fields: Training

Training

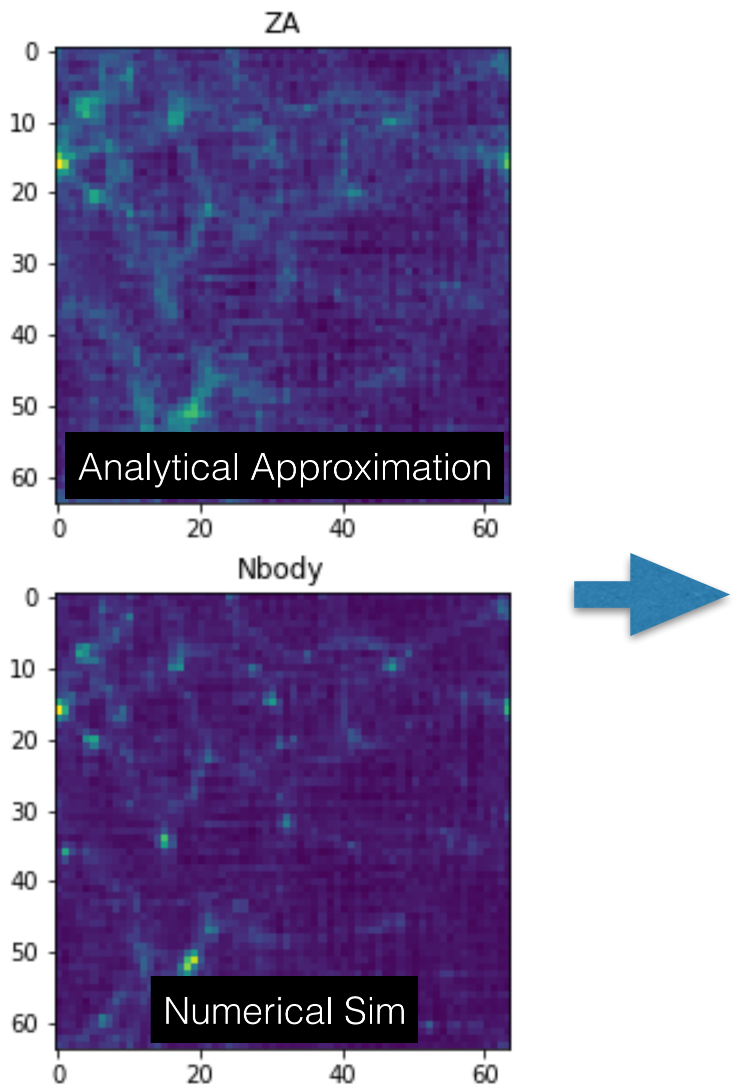

8,000 pairs of
[Analytical, Sim] boxes
For training

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

From Analytical approximated (Zeldovich approximation) fields to numerically simulated (FastPM) fields: Model

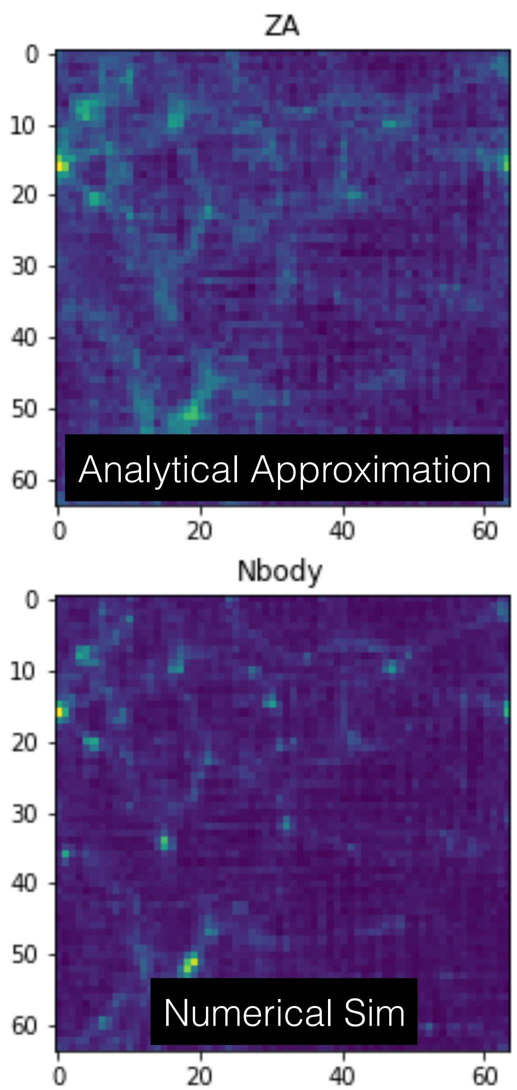

8,000 pairs of
[Analytical, Sim] boxes
For training

Machine learning Model

Slight variant to Residual NN

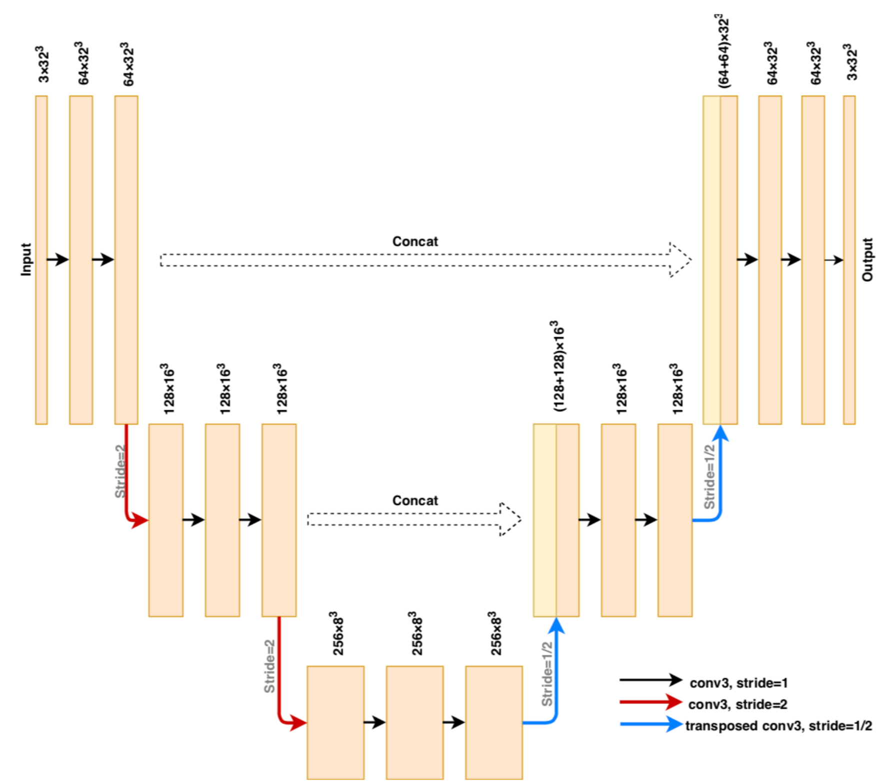

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

From Analytical approximated (Zeldovich approximation) fields to numerically simulated (FastPM) fields: Final setup

Machine learning Model

Slight variant to Residual NN

Input

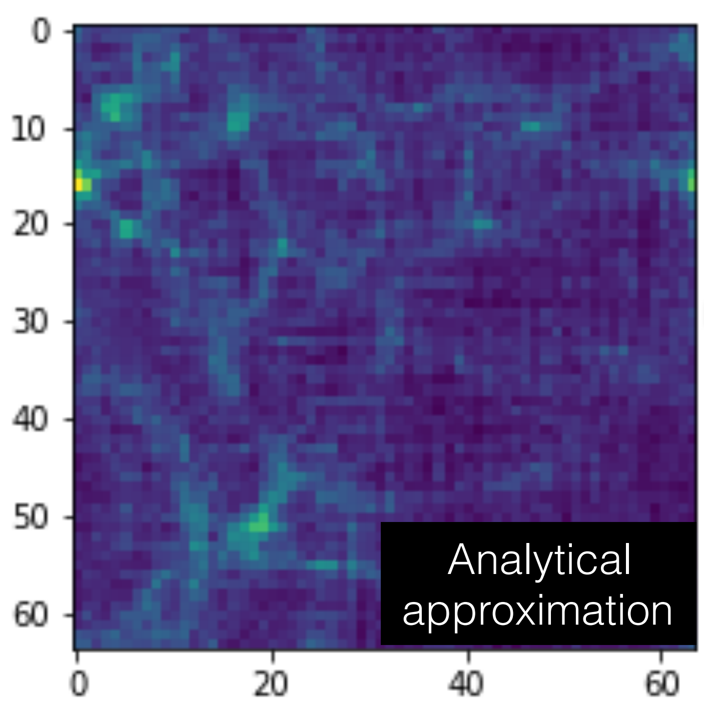

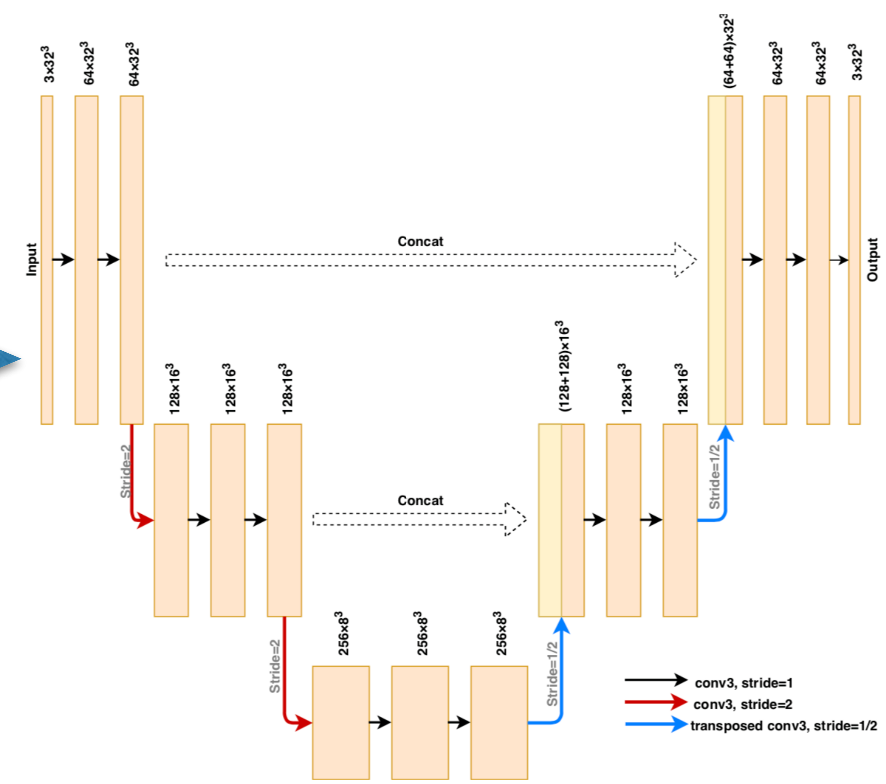

Prediction

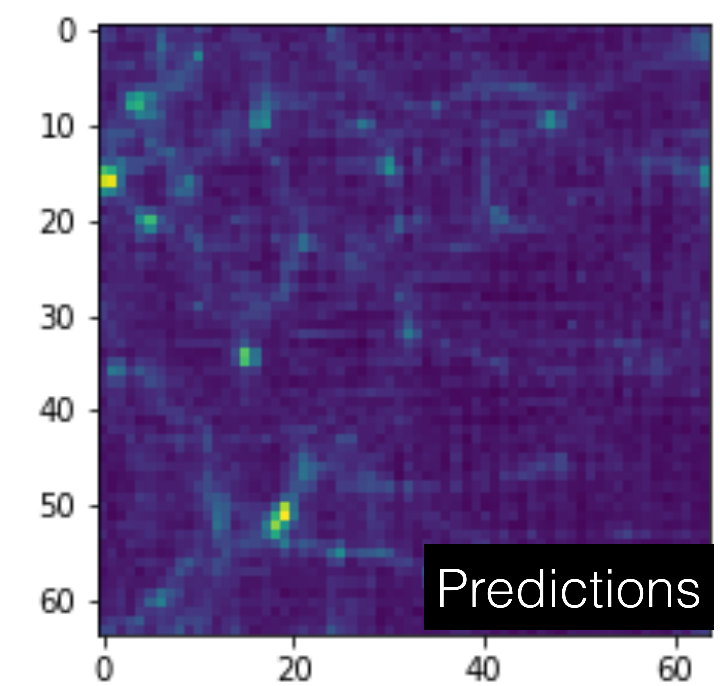

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: How well do we do ?

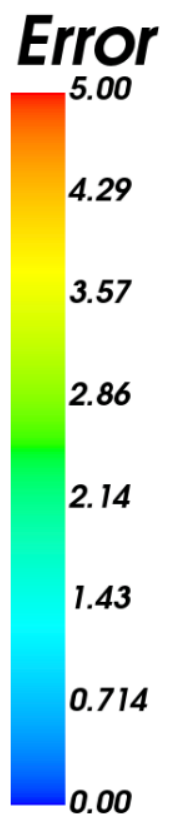

Mpc/h

We will show the errors in displacement field, predicted by
Our benchmark model (2LPT), and our ML model
Displacement field is the difference
between current position to the initial position of the particles

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: How well do we do ?

Benchmark (2LPT) prediction errors

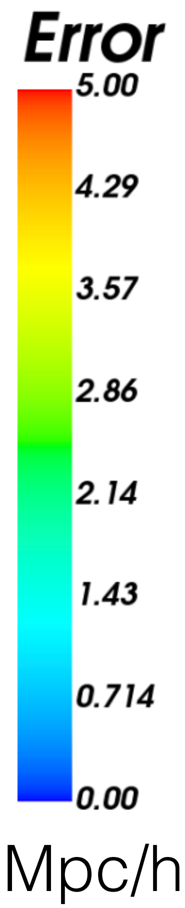

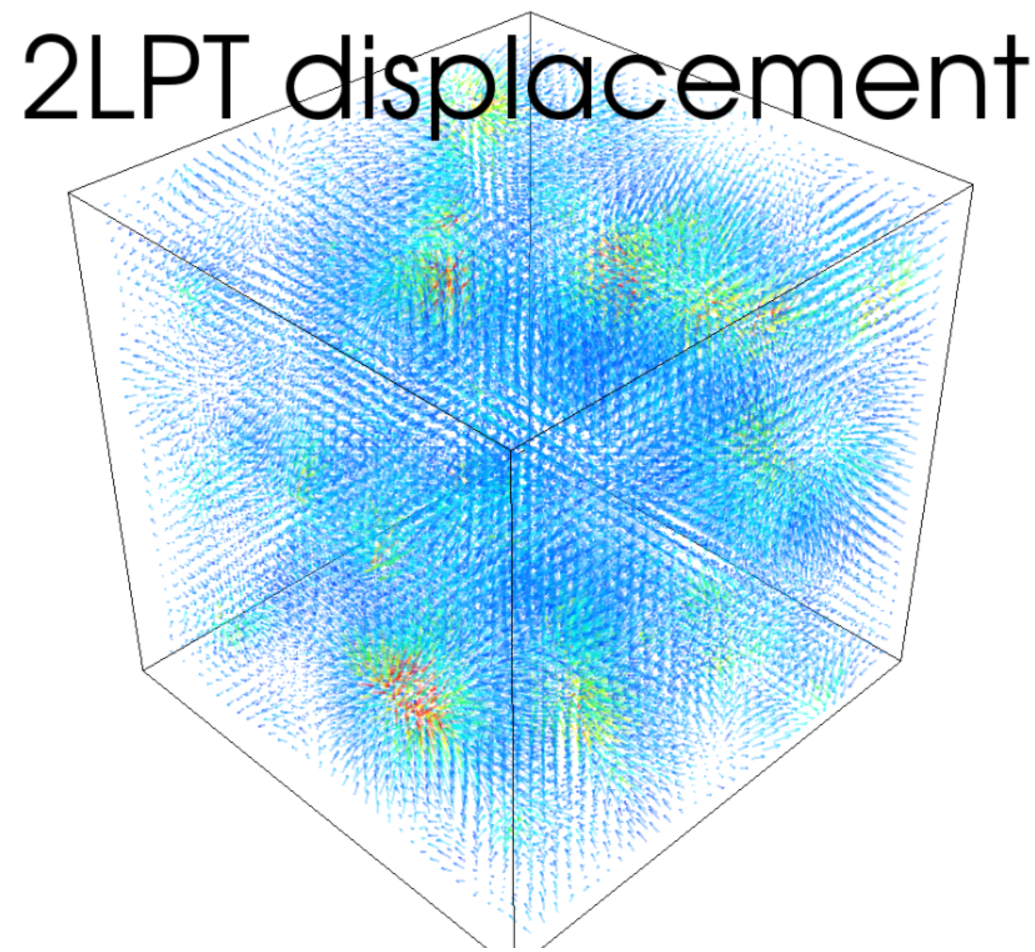

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: How well do we do ?

Machine Learning Model prediction errors

Benchmark (2LPT) prediction errors

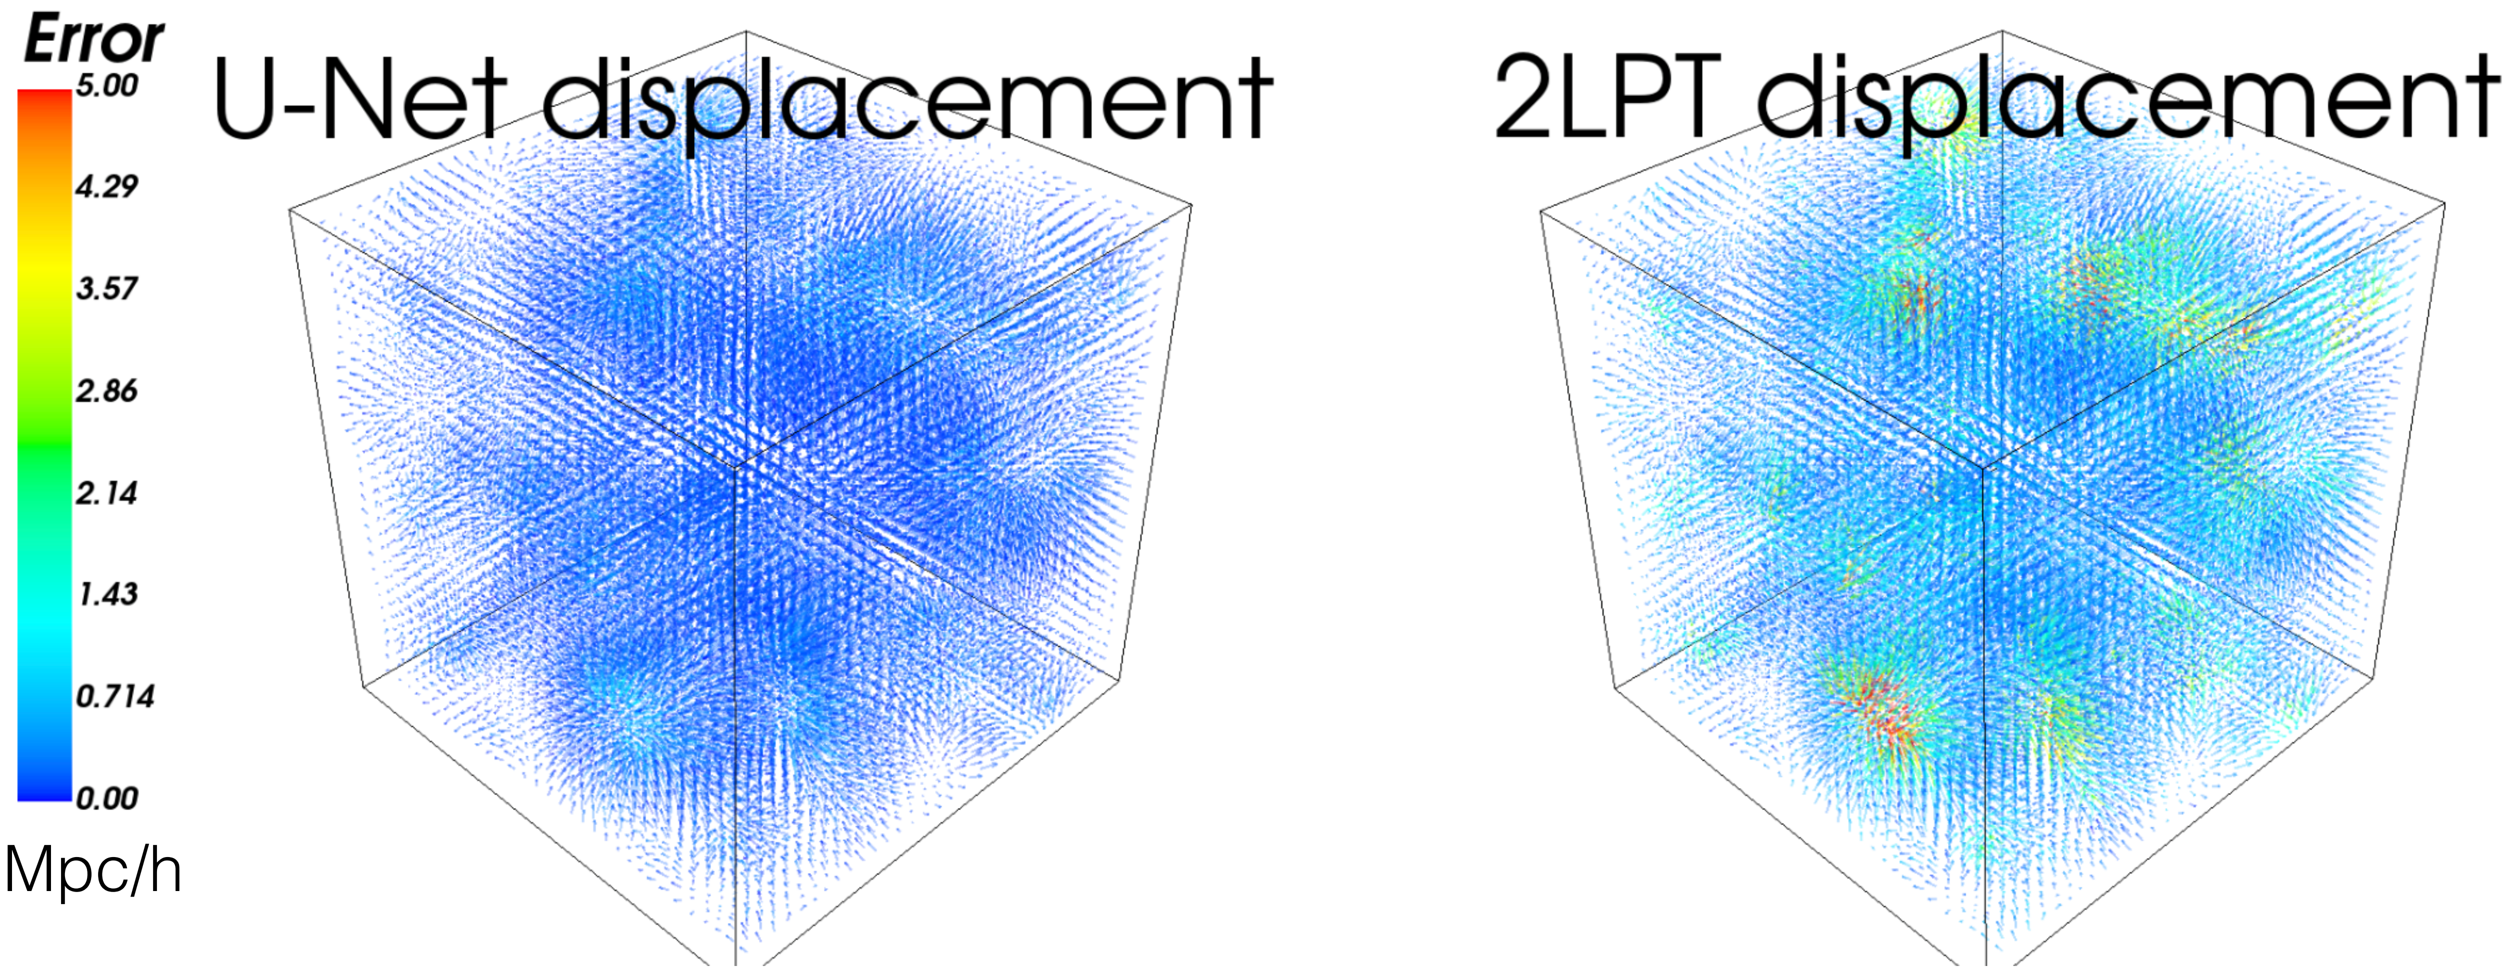

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: How well do we do?

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.

1000 simulations were predicted in 30 seconds post training and validation.

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

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

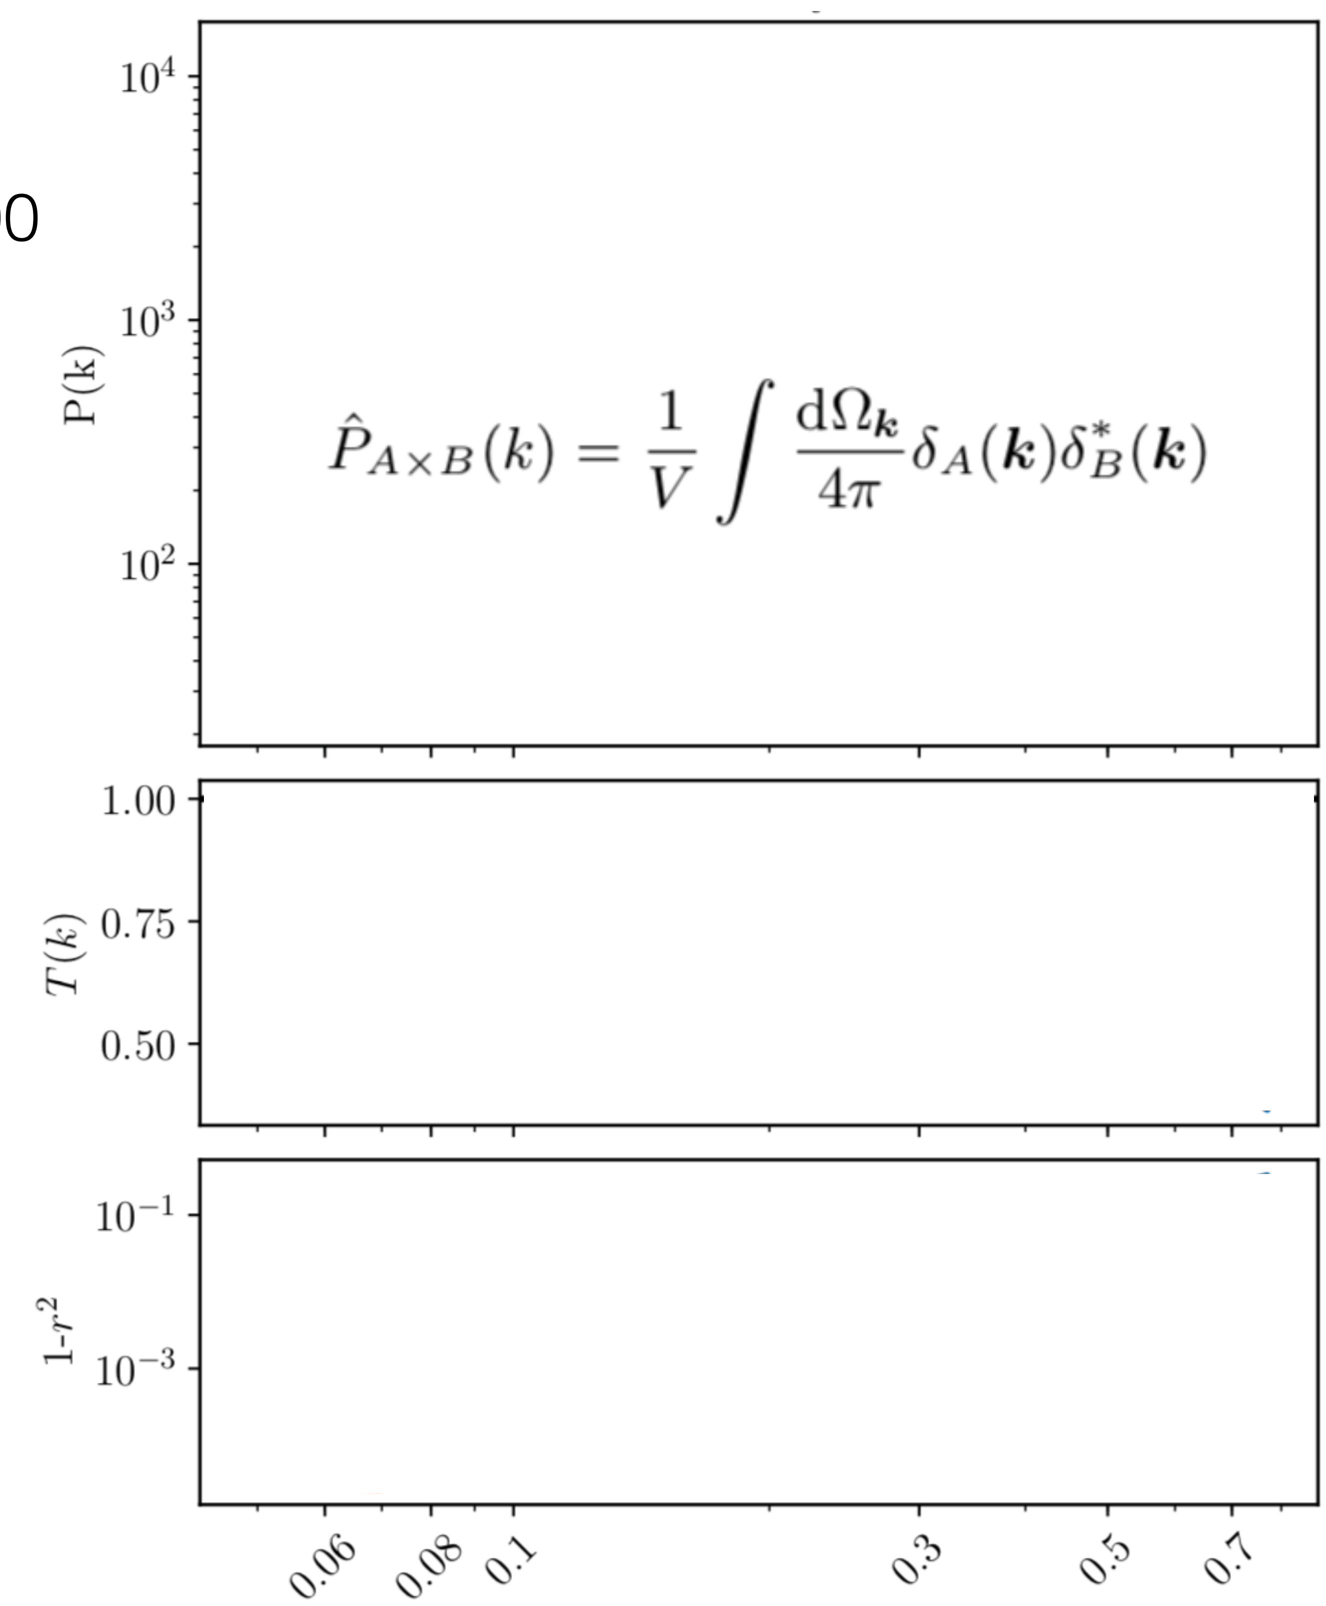

(a) Results from the density field

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: How well do we do?

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.

1000 simulations were predicted in 30 seconds post training and validation.

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

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

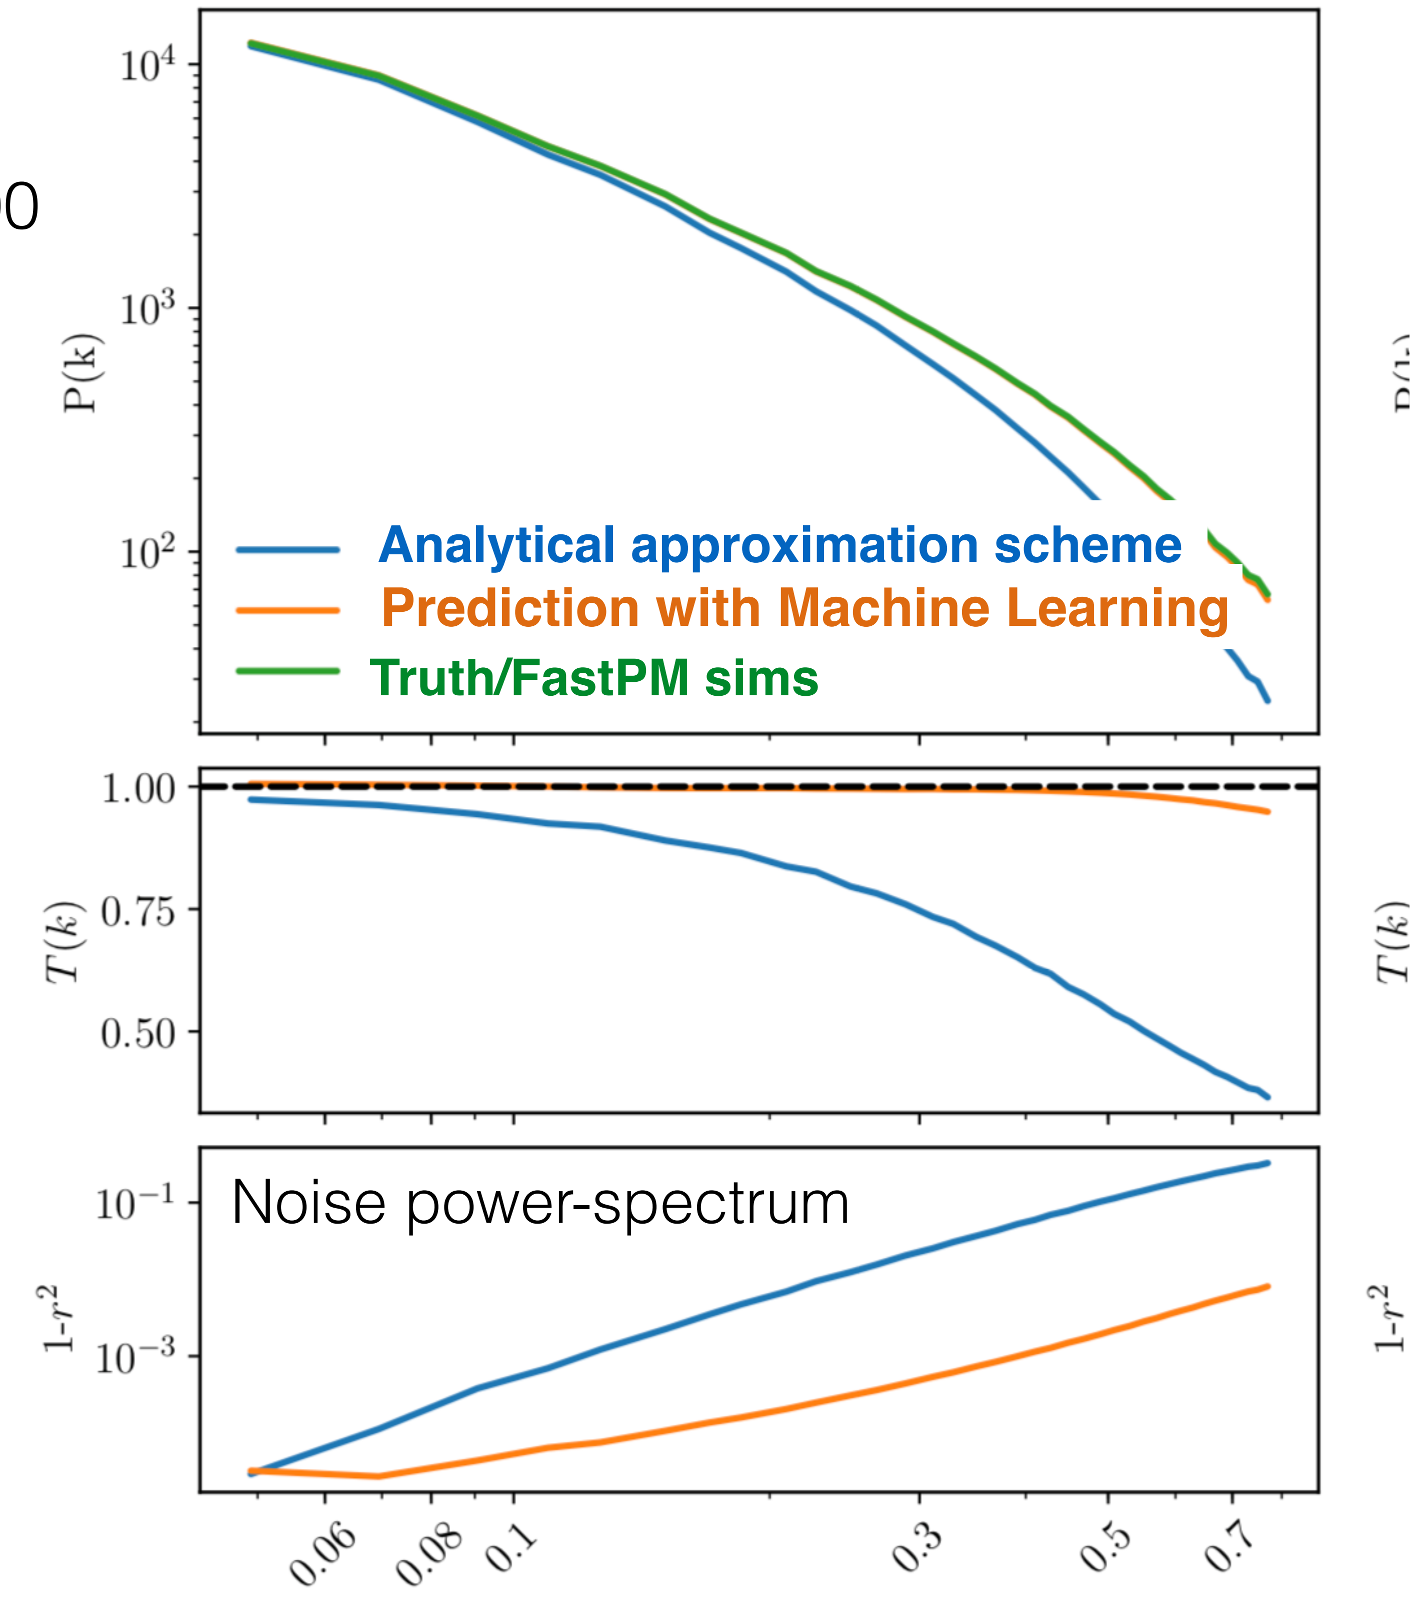

(a) Results from the density field

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Using Machine learning to simulate the Universe: Checking higher order correlation functions

- We checked on the 2-point function, seems like the model is predicting well.
- Then you asked: well, 2-point function is easy, if we have information that is non-gaussian, you want to test more than 2-point function.
- How about 3 point function?

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Projected multipoles of 3 point correlation function

Truth

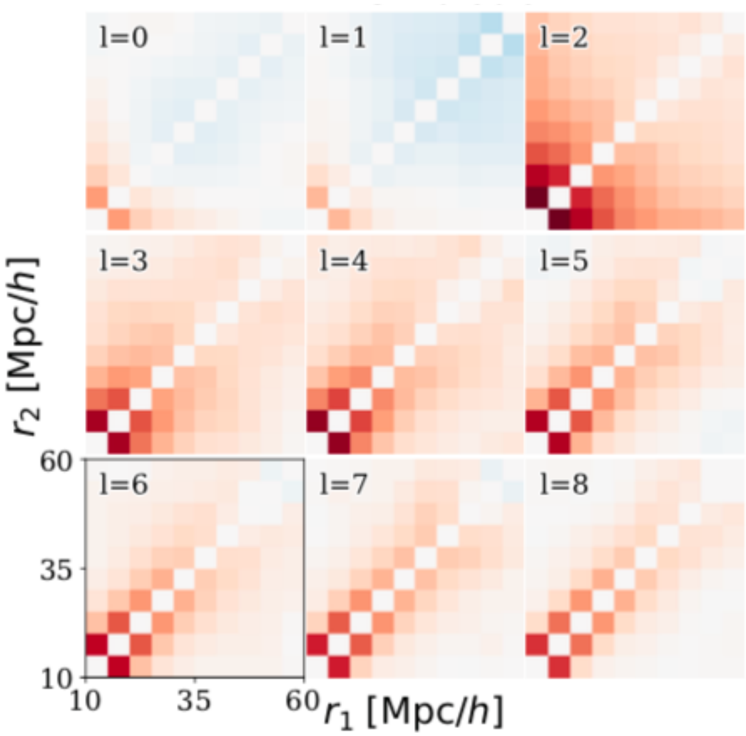

Multipoles generated using nbodykit implementation of 3-point function fast computation
(Hand, Feng et al. 2017; Slepian & Eisenstein 2015)

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Projected multipoles of 3 point correlation function

Truth

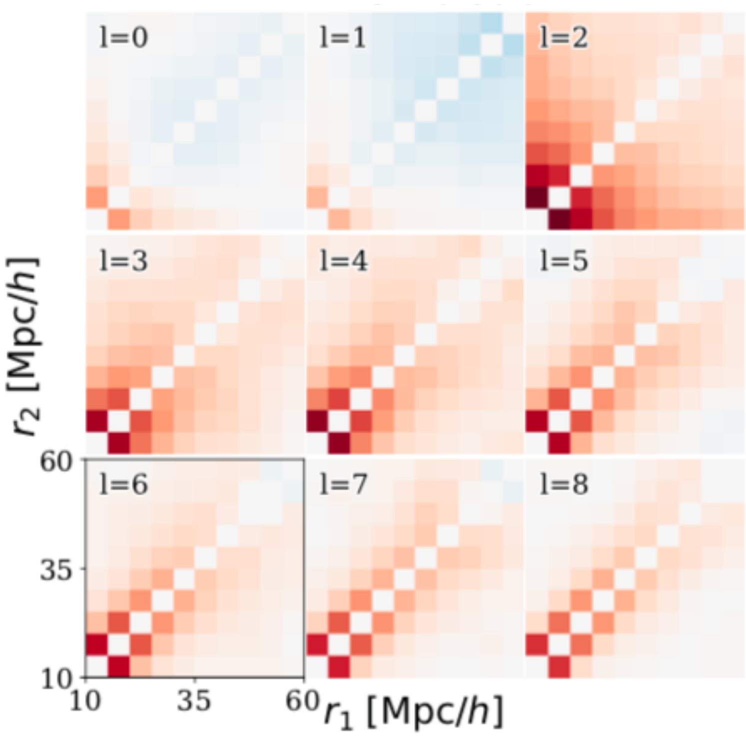

Benchmark

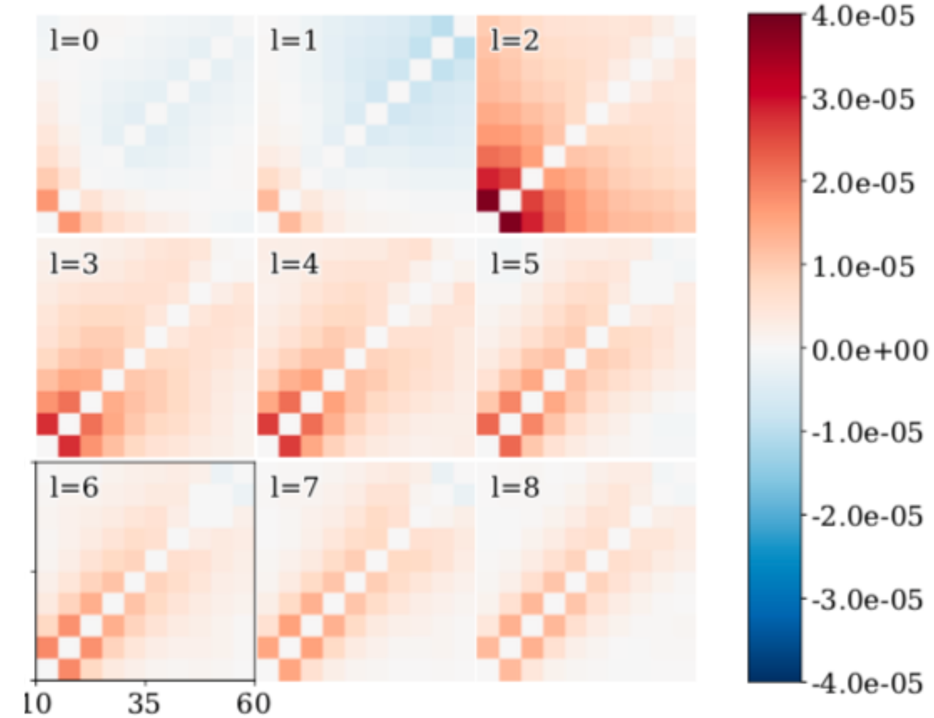

Fractional residuals

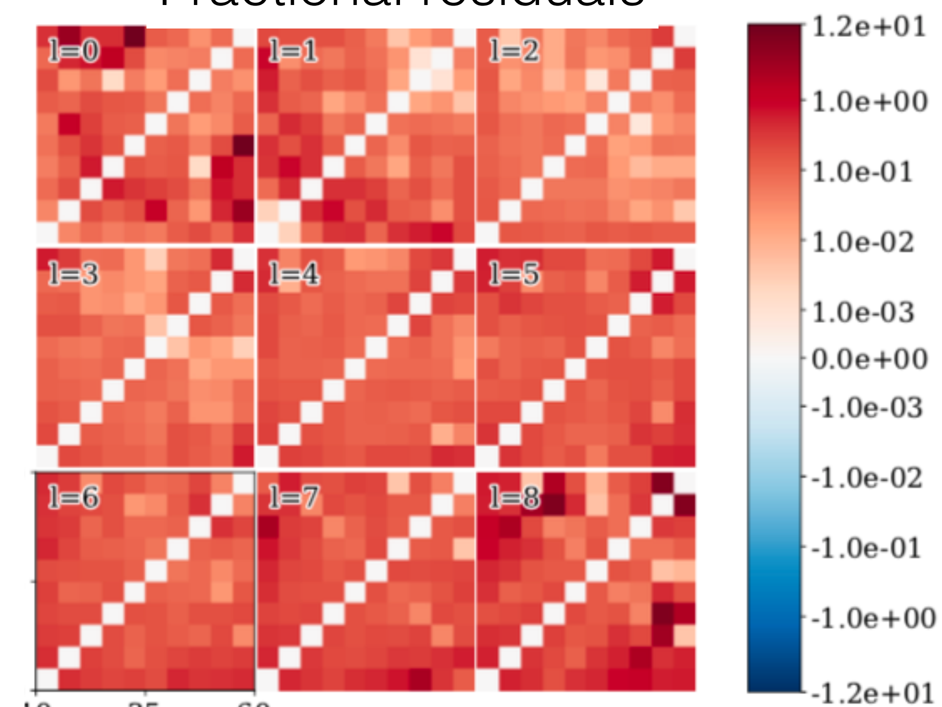

Multipoles generated using nbodykit implementation of 3-point function fast computation (Hand, Feng et al. 2017; Slepian & Eisenstein 2015)

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Projected multipoles of 3 point correlation function

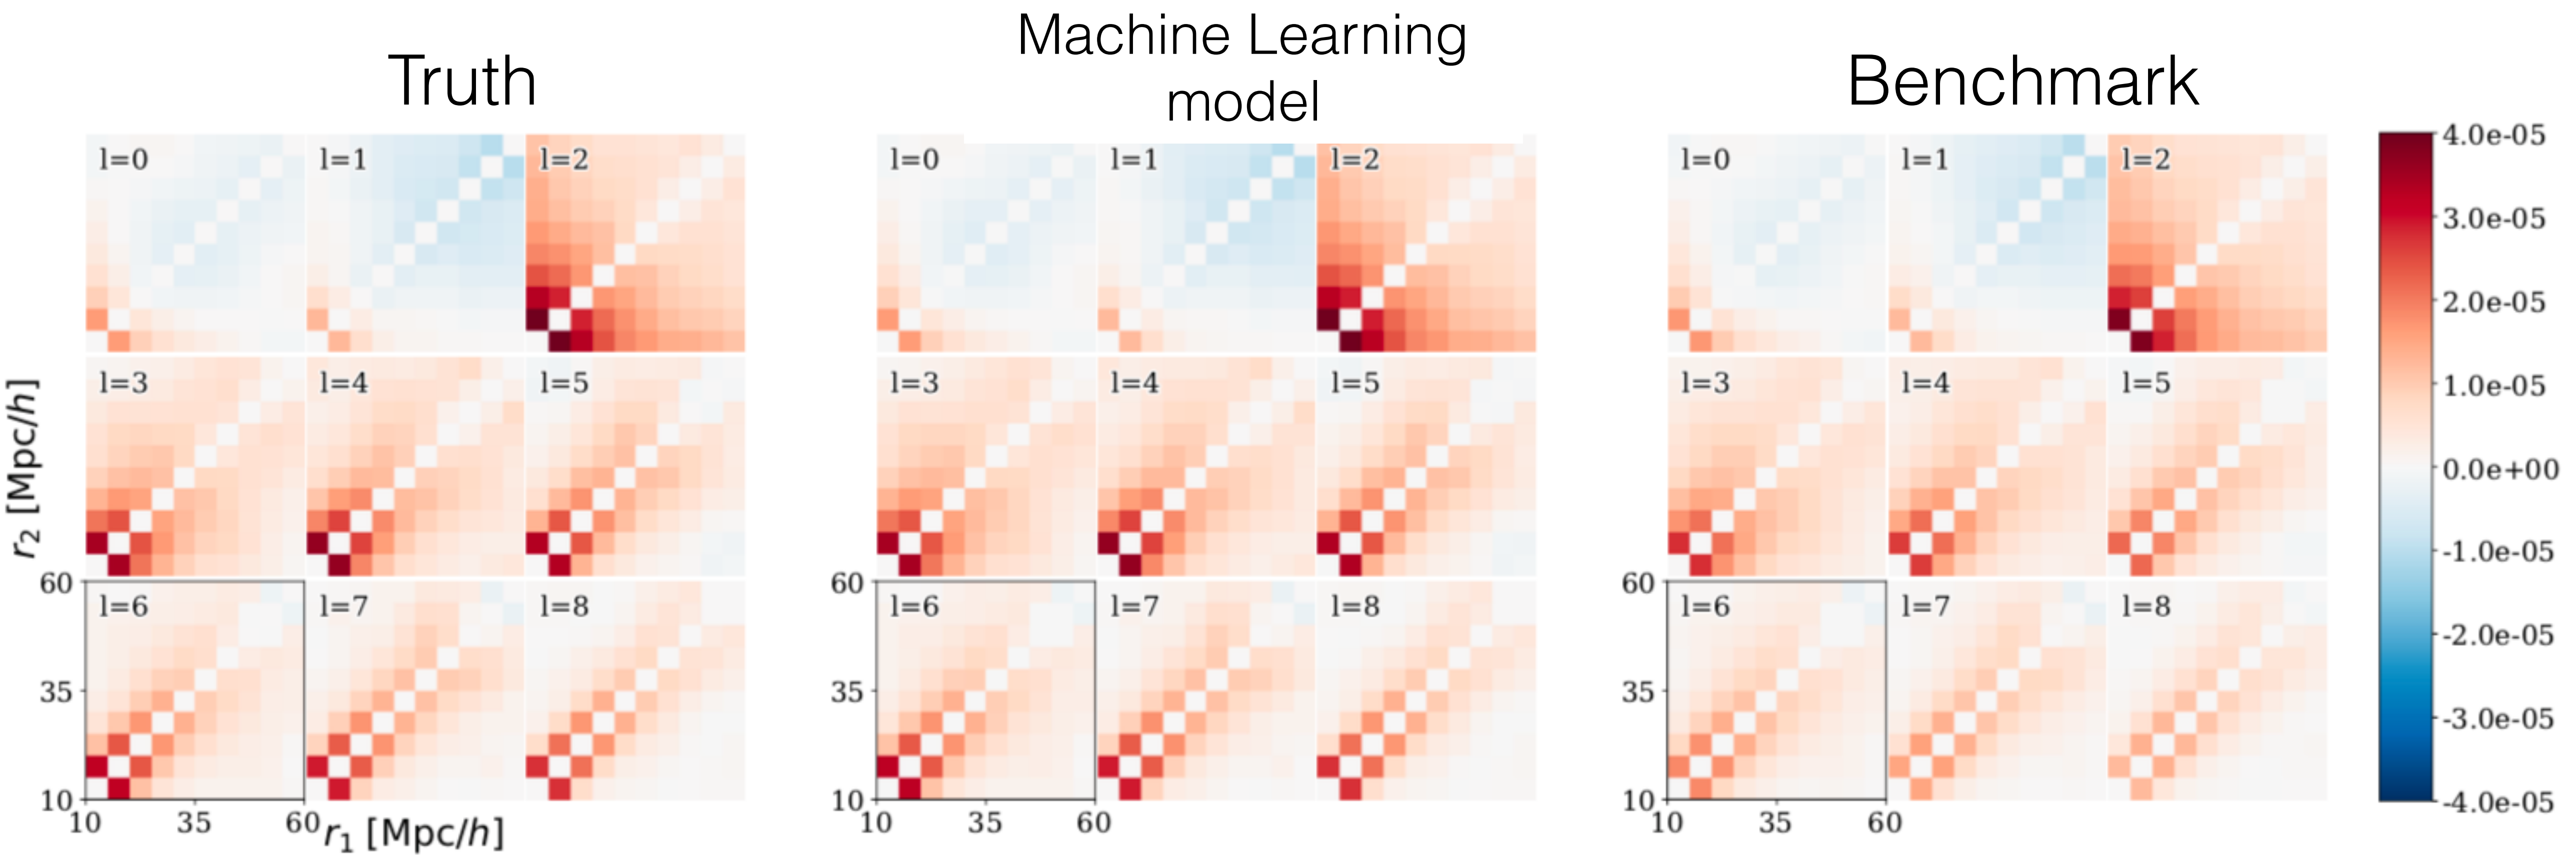

Fractional residuals are ~10 times better

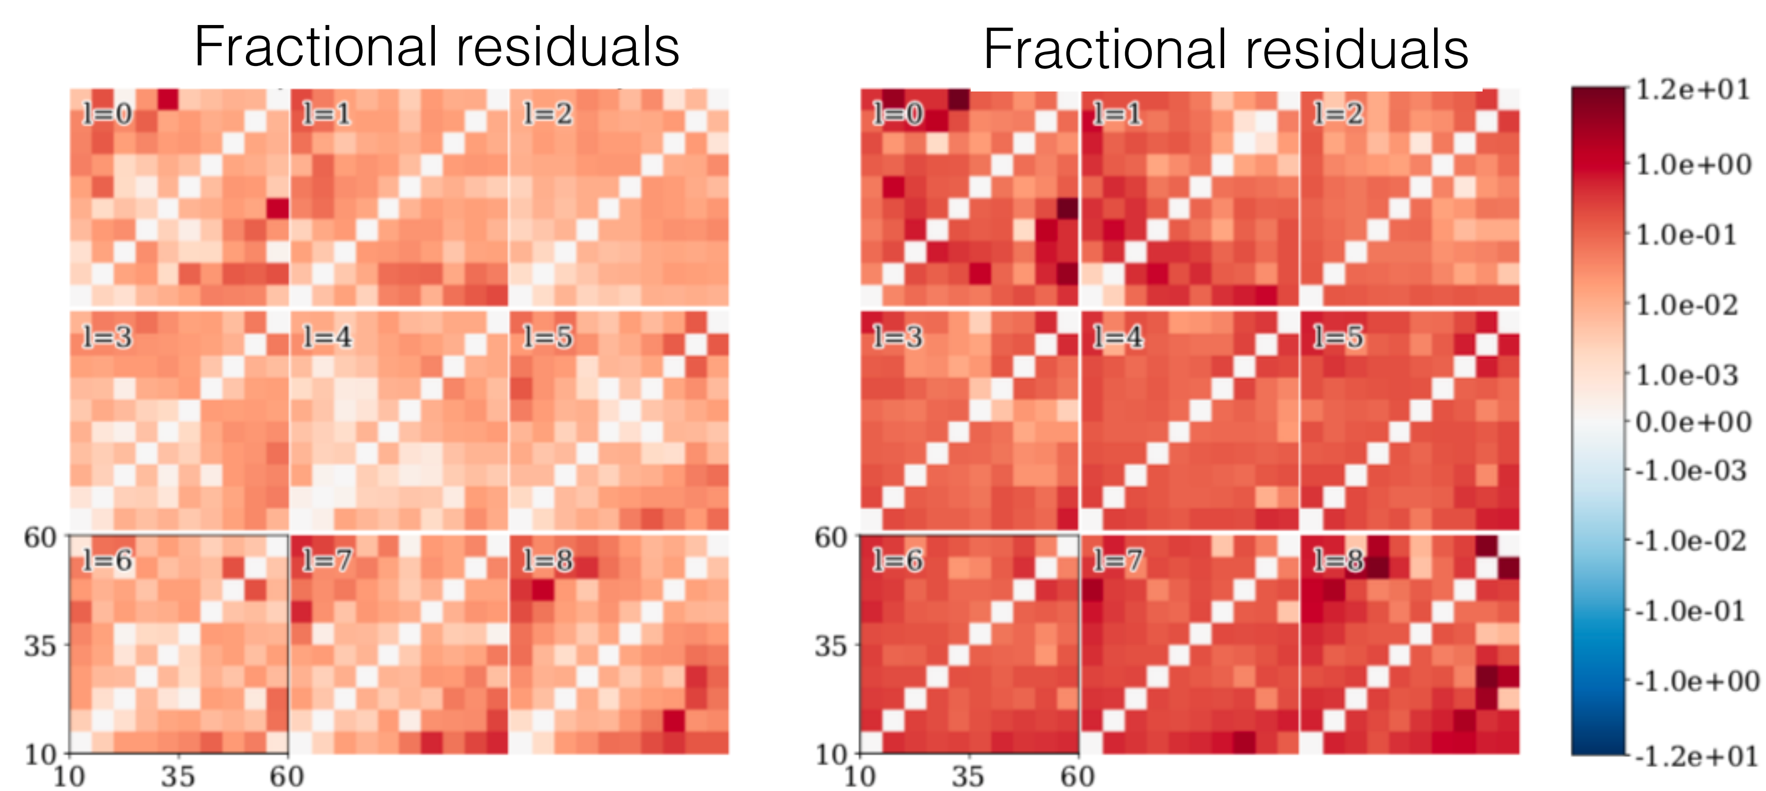

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Now as scientists, we have lots of questions: AKA: Outline for the remaining of the talk

- Can we get a correct estimate of the error ?
    - See He, Ravanbaksh & Ho *International Conference for Learning Representations 2018*
- Can we interpret the model learnt in Machine Learning?
    - **Can we simulate the Universe with Machine Learning -> Yes we can!**
- What is the model learning?
- 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: AKA: Outline for the remaining of the talk

- Can we get a correct estimate of the error ?
  - See He, Ravanbaksh & Ho *International Conference for Learning Representations 2018*
- Can we interpret the model learnt in Machine Learning?
  - Can we simulate the Universe with Machine Learning -> Yes we can!
- **What is the model learning?**
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
  - Use **simple analytical cases** that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.
  - Locating **invariances** in the system
  - Locating where the **information** is coming from
  - … other suggestions are very welcome
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
    - **Use simple analytical cases that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.**
    - Locating invariances in the system
    - Locating where the information is coming from
    - … other suggestions are very welcome
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model I: Simple analytical cases

- 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:

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

From Analytical approximated fields to numerically simulated fields

Input

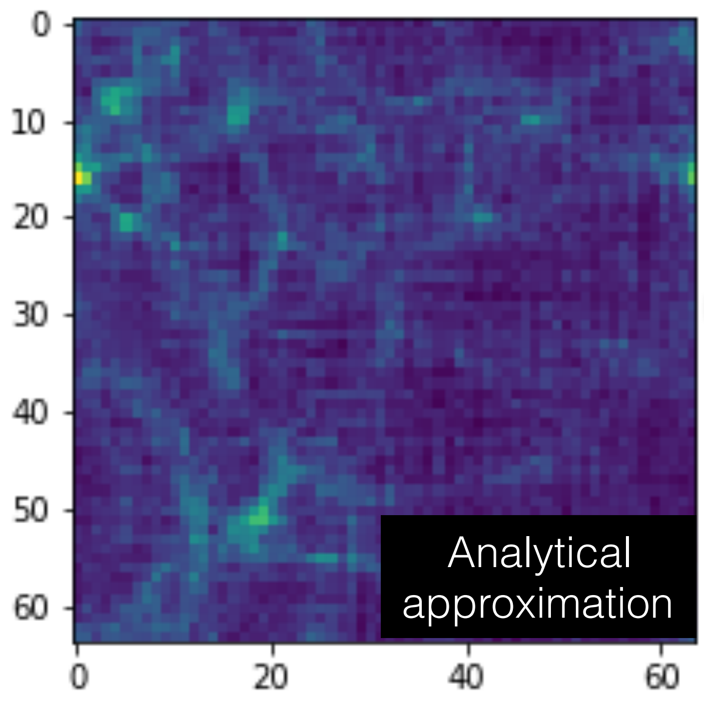

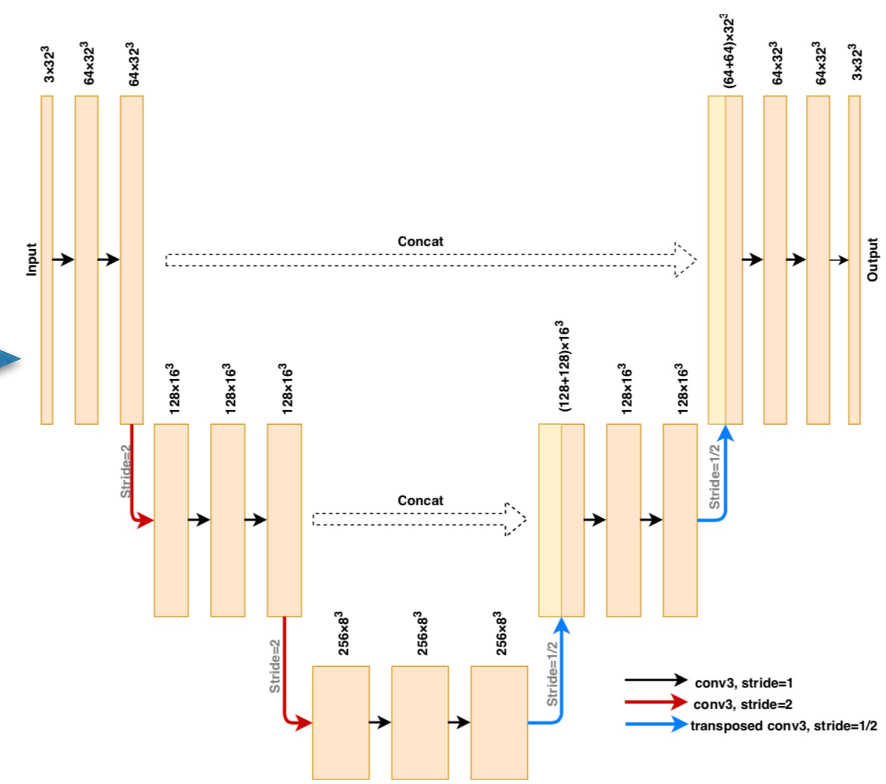

Prediction

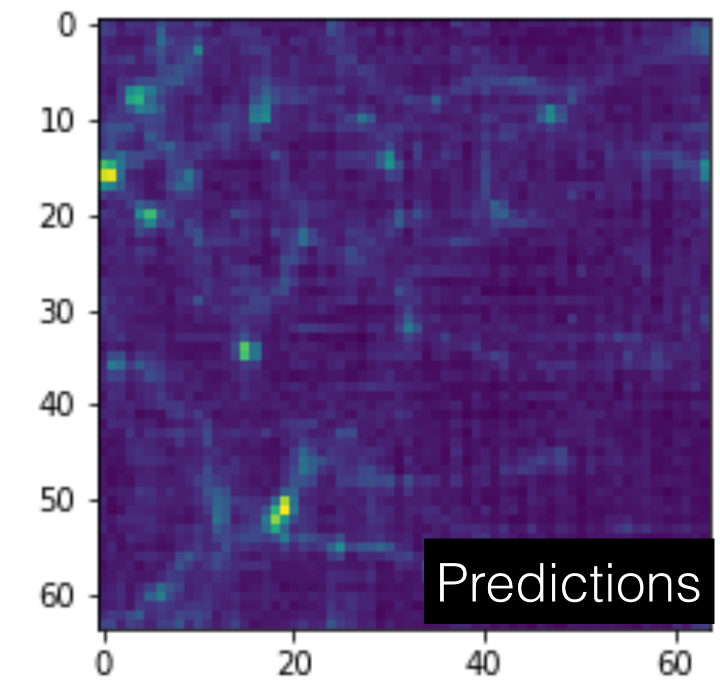

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Interrogating the machine learning model I:

What happen if we have power at only one scale?

Input

Prediction

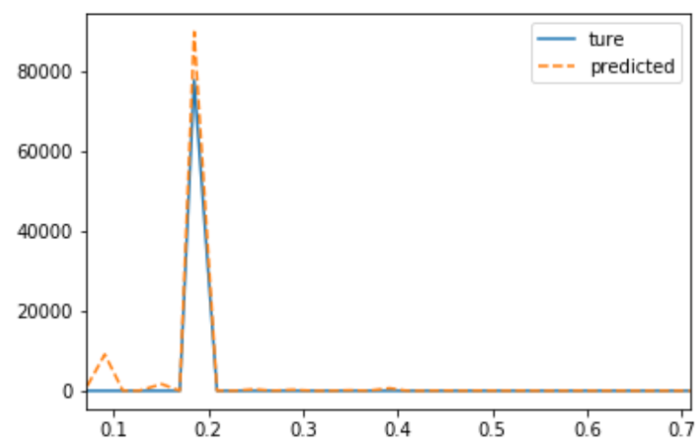

k(h/Mpc)

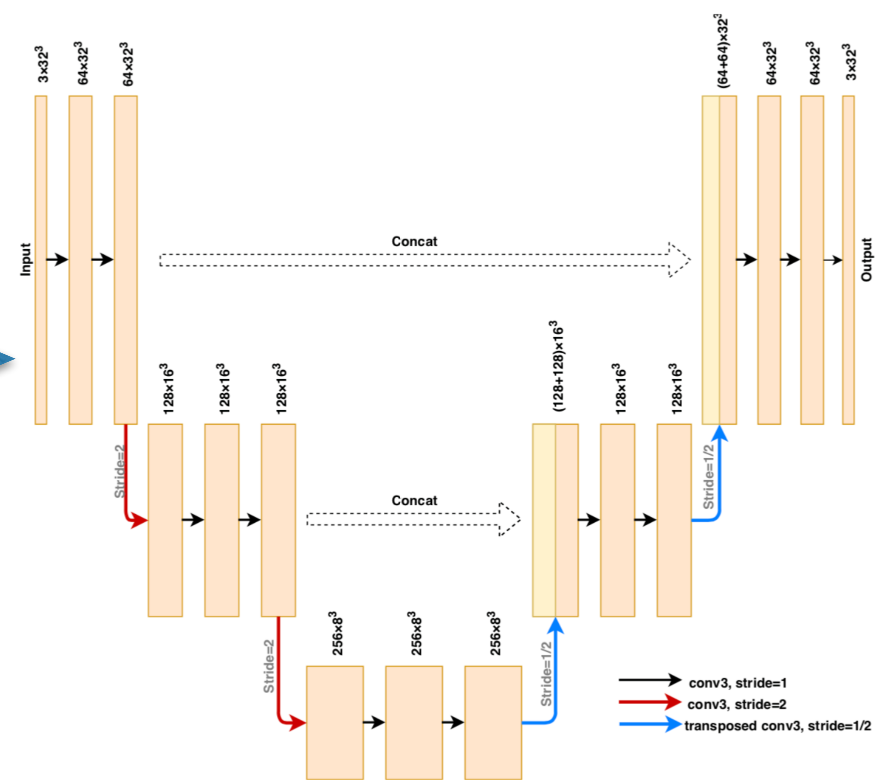

Inject Power
at one scale as input

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I:

What happen if we have power at only one scale?

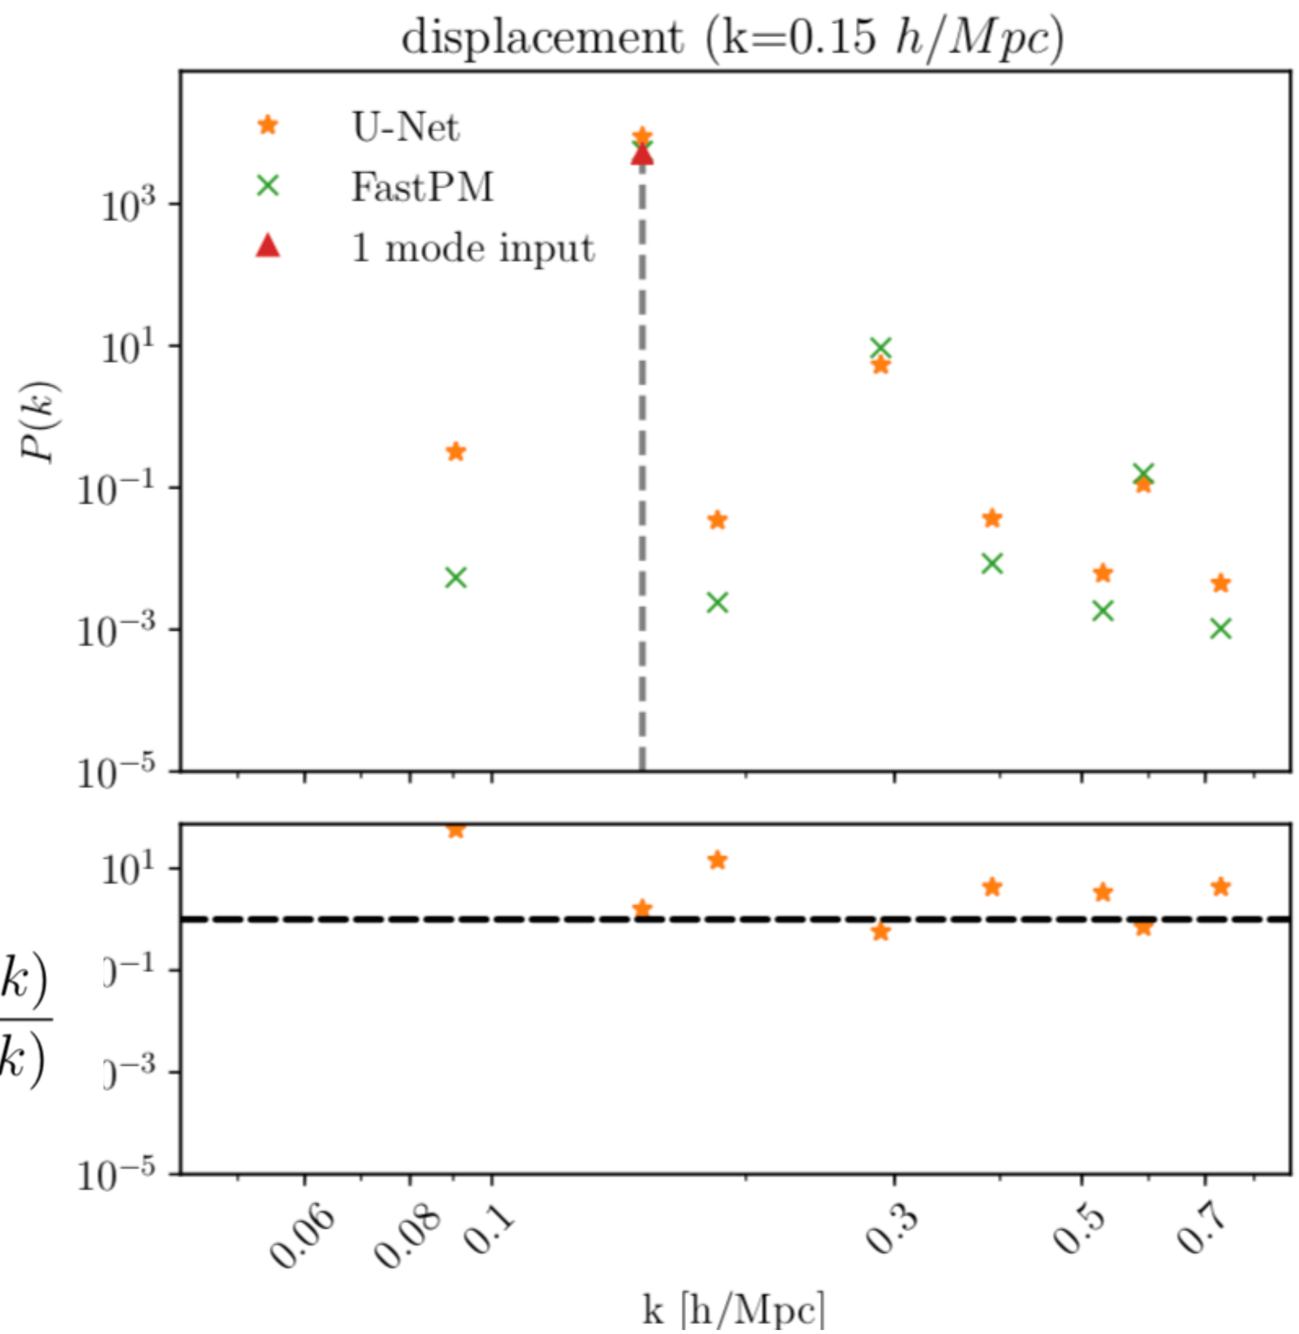

$$T(k) = \frac{P_{\mathrm{pred}}(k)}{P_{\mathrm{true}}(k)}$$

The transfer function shows that the U-Net model captures quite well at the dominate scale, which indicates the U-Net mode 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.

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I

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

Input

Prediction

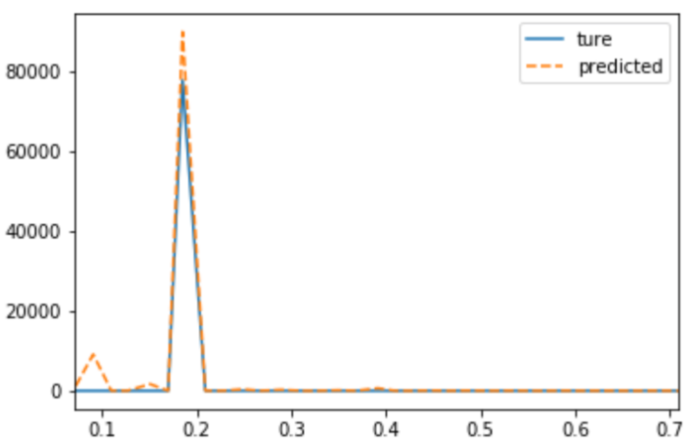

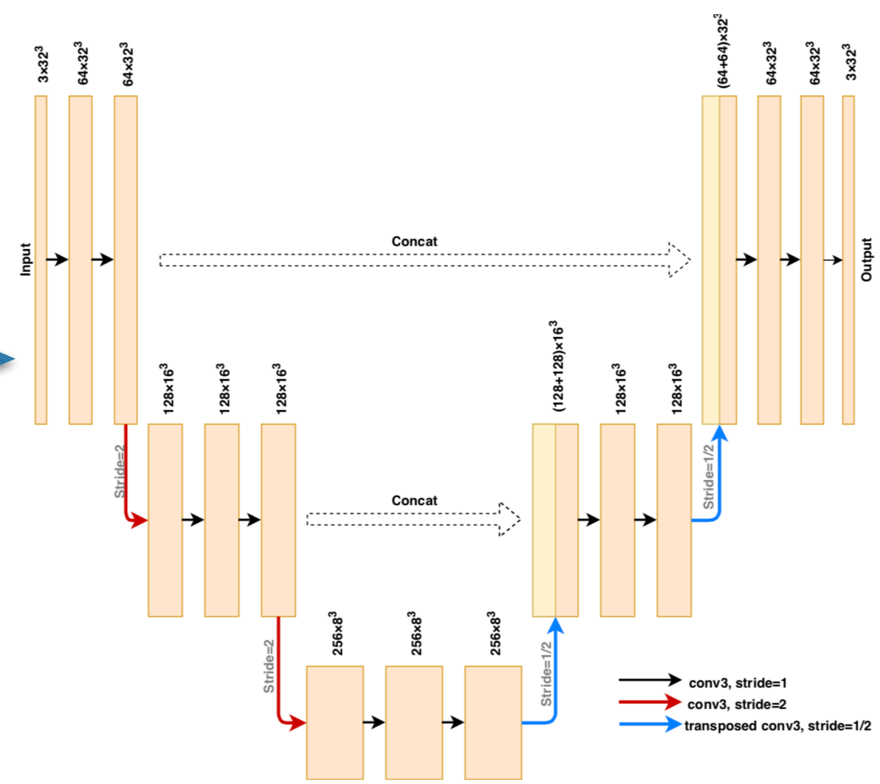

k(h/Mpc)

Inject Power
At same k,
but different phases

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I

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

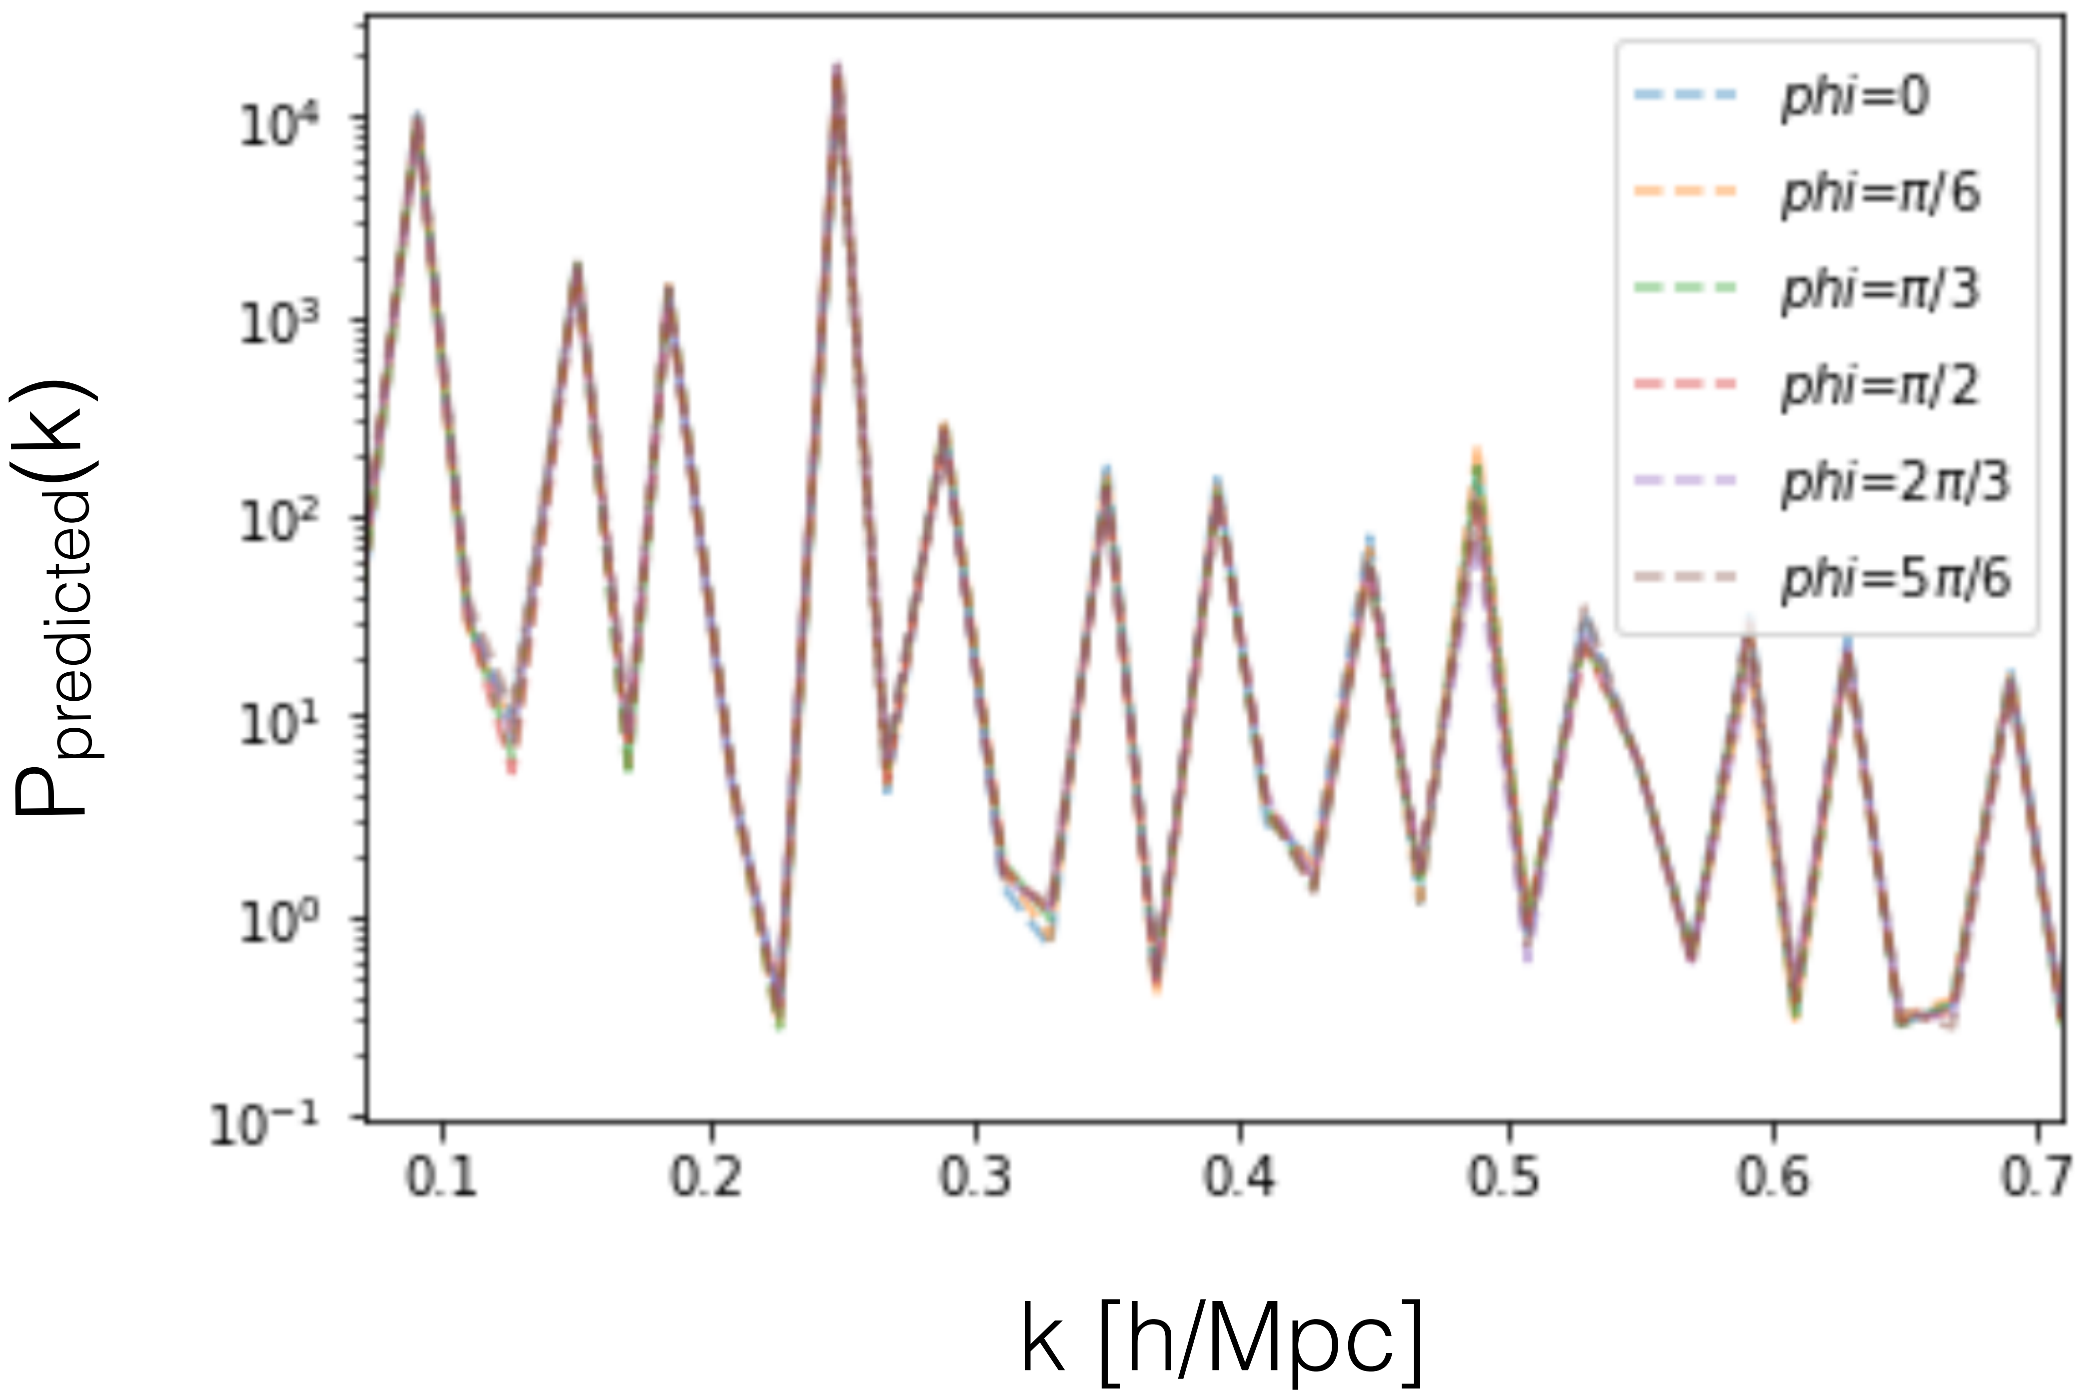

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I:

Simple analytical cases: Two modes

Fourier amplitude of the curl-free displacement field

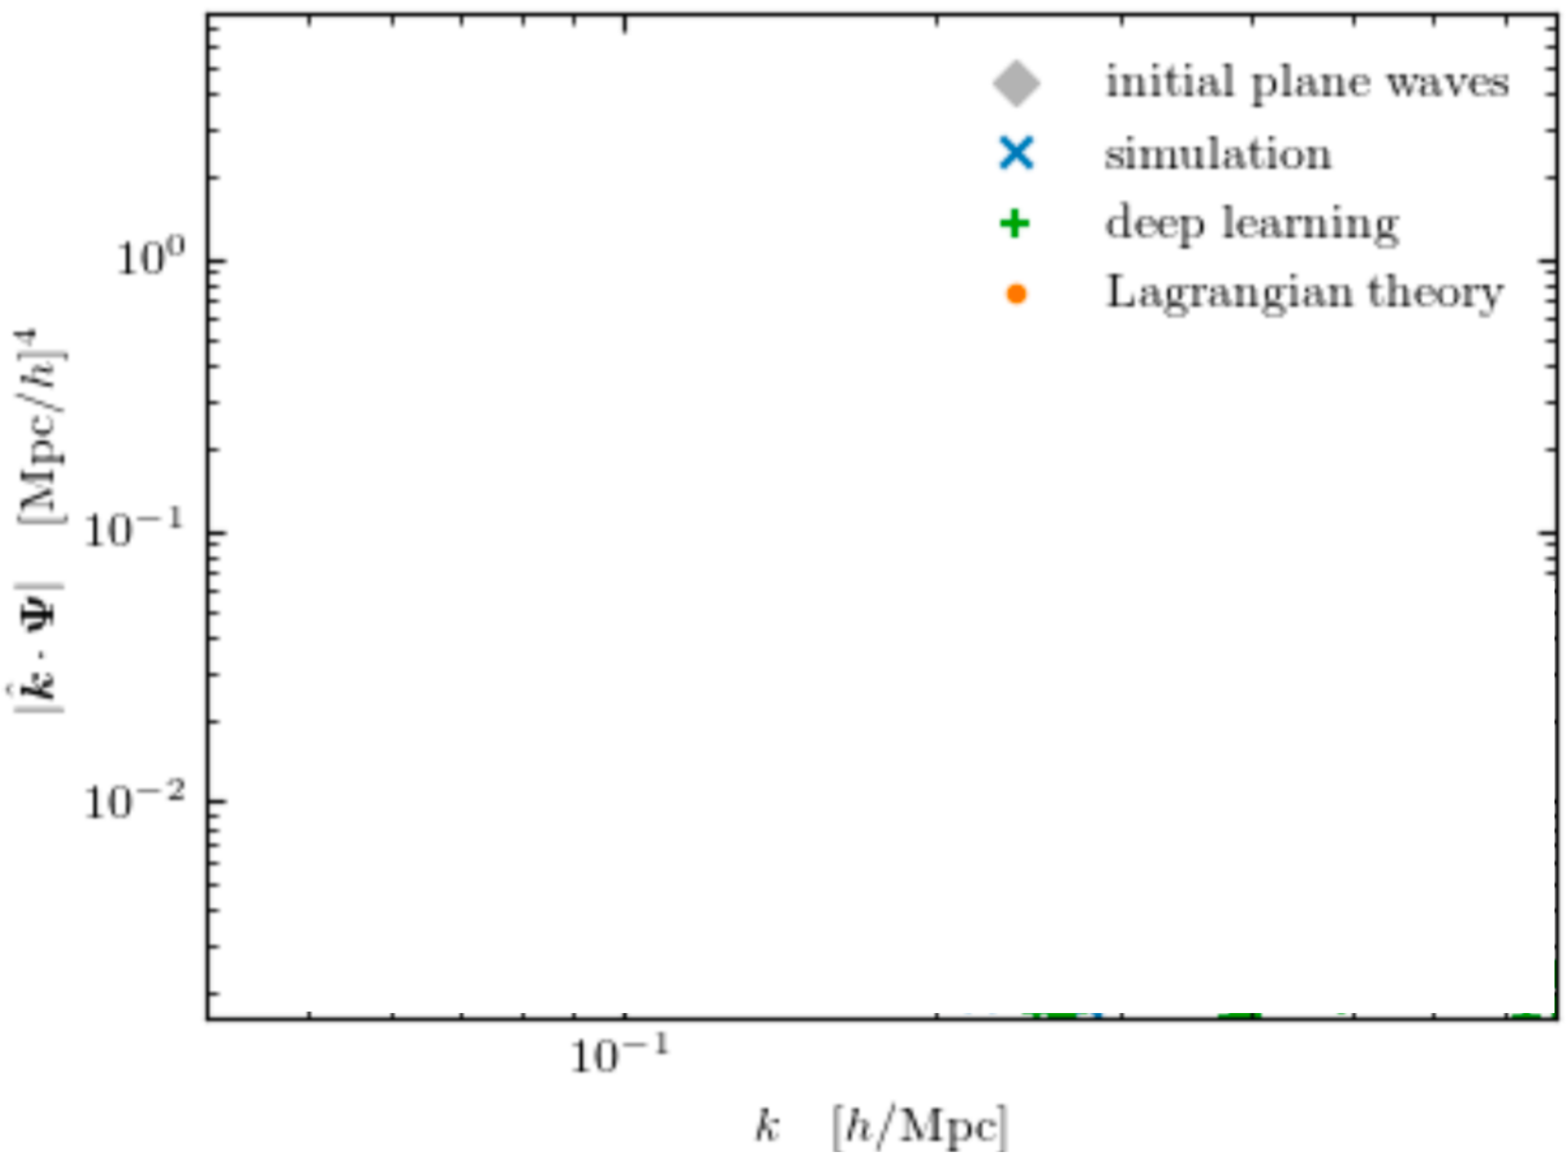

- Mode coupling of two plane waves
- Perpendicular to each other, with same amplitude and wave number as marked by grey diamond.
- Chosen to be in linear regime where perturbation theory is still valid.
- What do we expect from linear theory?
  - Plane waves stay at initial amplitude
- Generated by their interaction, new modes arise to the right of the initial modes (smaller scale).
  - The ML model is in good agreement with (Lagrangian) theory and simulations.
- According to ML, there are many more modes to the right, with similar amplitude to the simulations. But some of these can be artifacts of the simulations.

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I:

Simple analytical cases: Two modes

Fourier amplitude of the curl-free displacement field

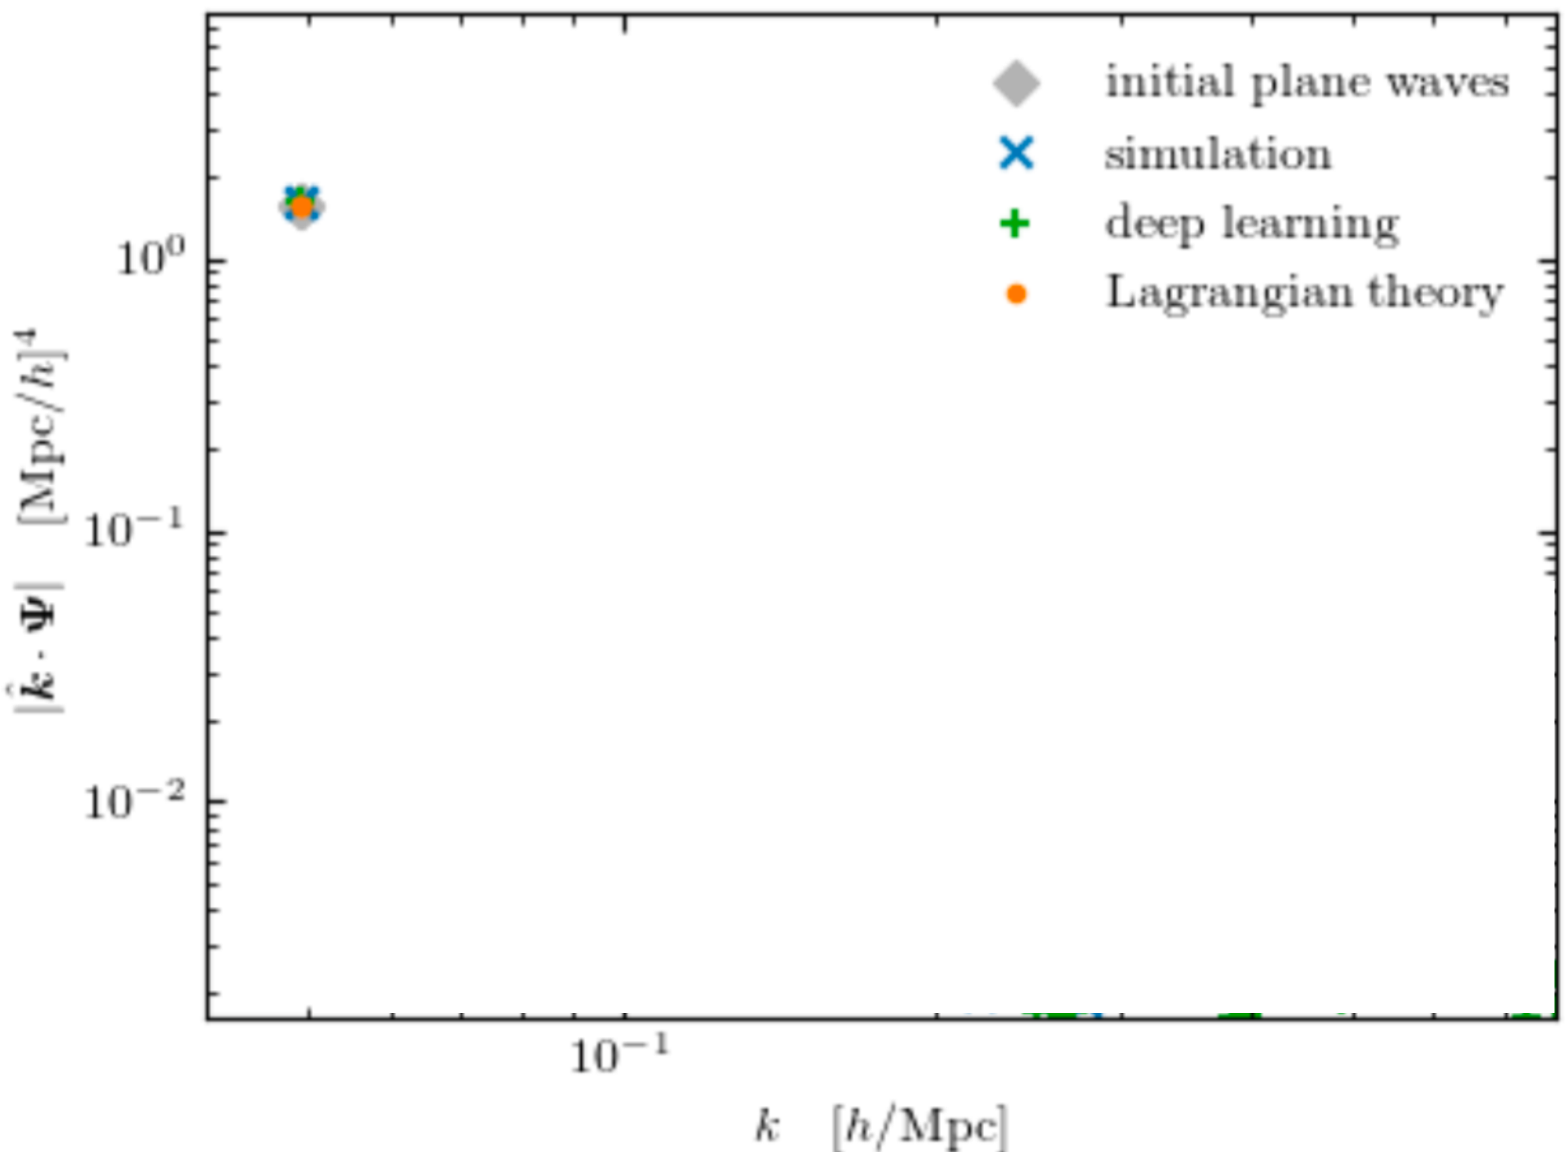

- Mode coupling of two plane waves
- Perpendicular to each other, with same amplitude and wave number as marked by grey diamond.
- Chosen to be in linear regime where perturbation theory is still valid.
- What do we expect from linear theory?
  - Plane waves stay at initial amplitude
- Generated by their interaction, new modes arise to the right of the initial modes (smaller scale).
  - The ML model is in good agreement with (Lagrangian) theory and simulations.
- According to ML, there are many more modes to the right, with similar amplitude to the simulations. But some of these can be artifacts of the simulations.

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I:

Simple analytical cases: Two modes

Fourier amplitude of the curl-free displacement field

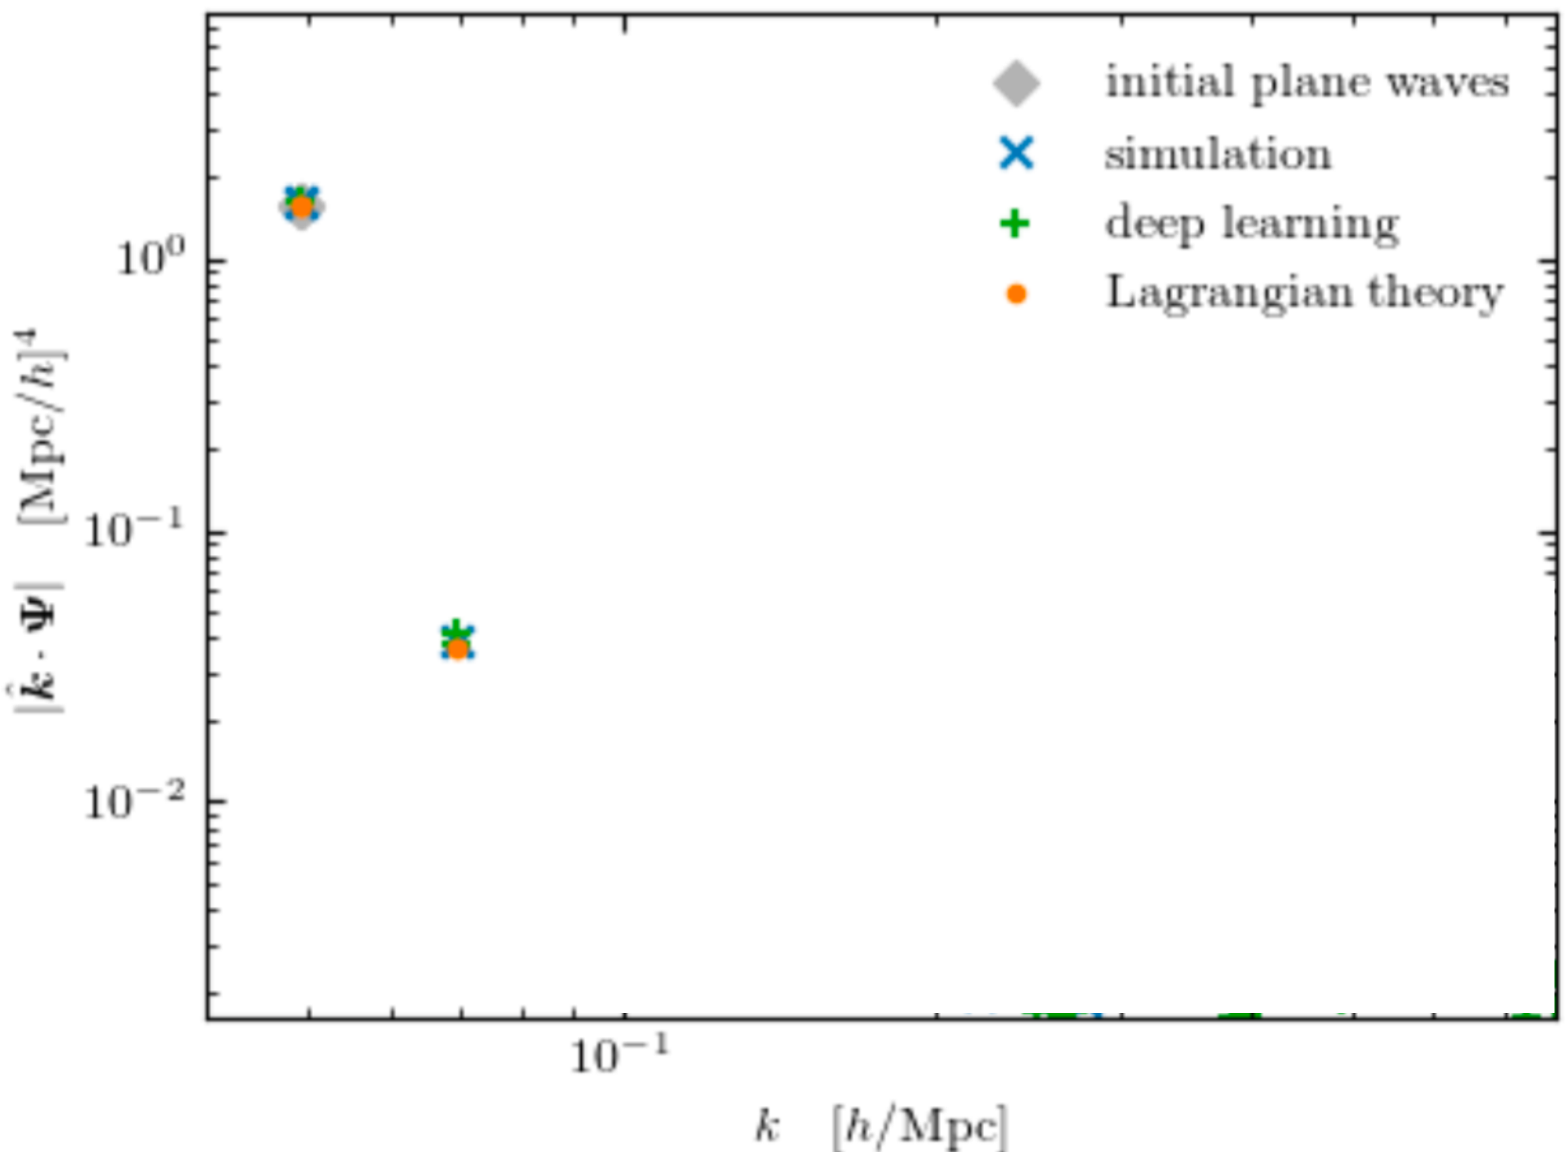

- Mode coupling of two plane waves
- Perpendicular to each other, with same amplitude and wave number as marked by grey diamond.
- Chosen to be in linear regime where perturbation theory is still valid.
- What do we expect from linear theory?
  - Plane waves stay at initial amplitude
- Generated by their interaction, new modes arise to the right of the initial modes (smaller scale).
  - The ML model is in good agreement with (Lagrangian) theory and simulations.
- According to ML, there are many more modes to the right, with similar amplitude to the simulations. But some of these can be artifacts of the simulations which were also learned?

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model I:

Simple analytical cases: Two modes

Fourier amplitude of the curl-free displacement field

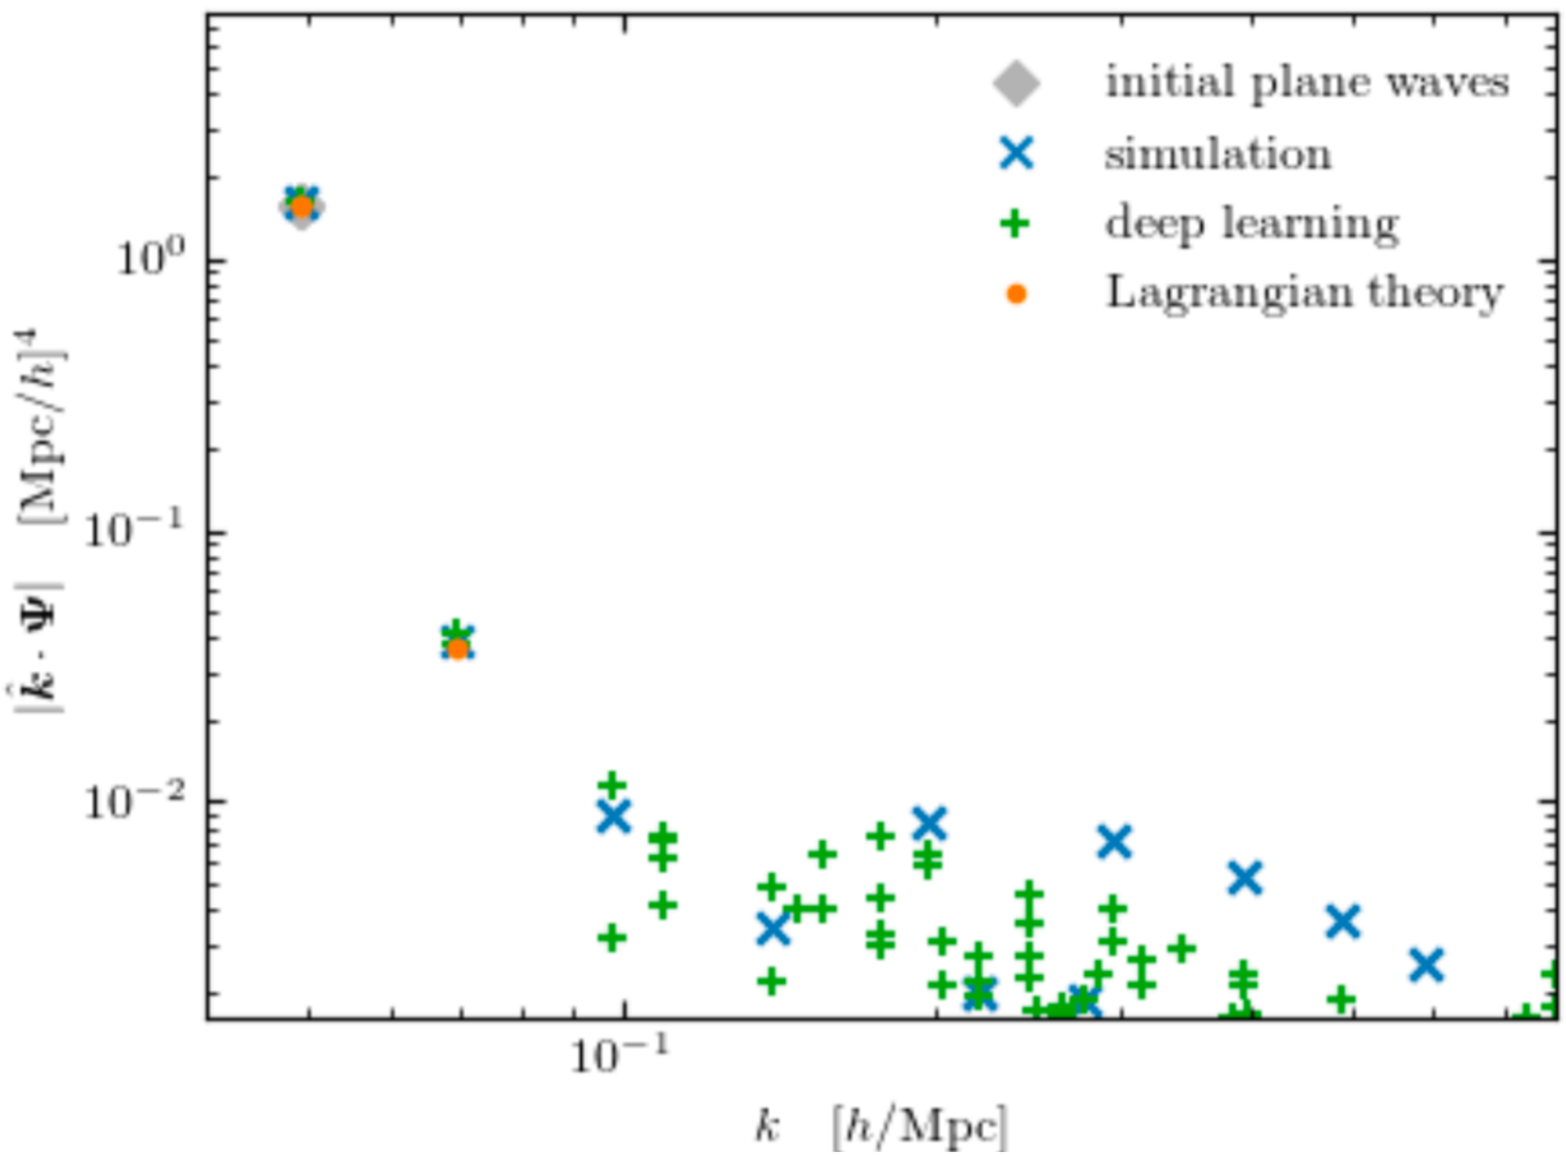

- Mode coupling of two plane waves
- Perpendicular to each other, with same amplitude and wave number as marked by grey diamond.
- Chosen to be in linear regime where perturbation theory is still valid.
- What do we expect from linear theory?
  - Plane waves stay at initial amplitude
- Generated by their interaction, new modes arise to the right of the initial modes (smaller scale).
  - The ML model is in good agreement with (Lagrangian) theory and simulations.
- According to ML, there are many more modes to the right, with similar amplitude to the simulations. But some of these can be artifacts of the simulations which were also learned?

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
  - Use simple analytical cases that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.
  - Locating invariances in the system (see Mallat 2016)
  - Locating where the information is coming from
  - … other suggestions are very welcome
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
  - Use simple analytical cases that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.
  - **Locating invariances in the system (see Mallat 2016)**
  - Locating where the information is coming from
  - … other suggestions are very welcome :)
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model II:

Locating the Invariances in the system

Input

Prediction

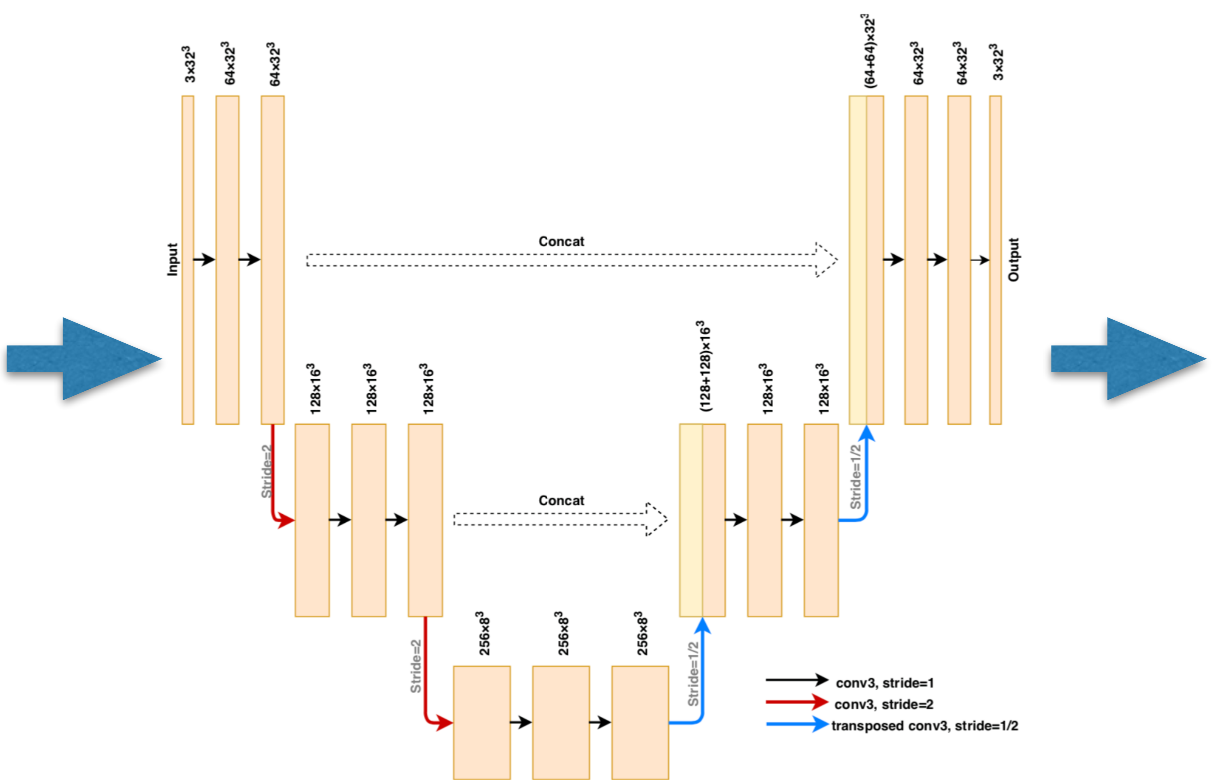

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model II

Is Rotational Invariance learnt by the model?

Input

Prediction

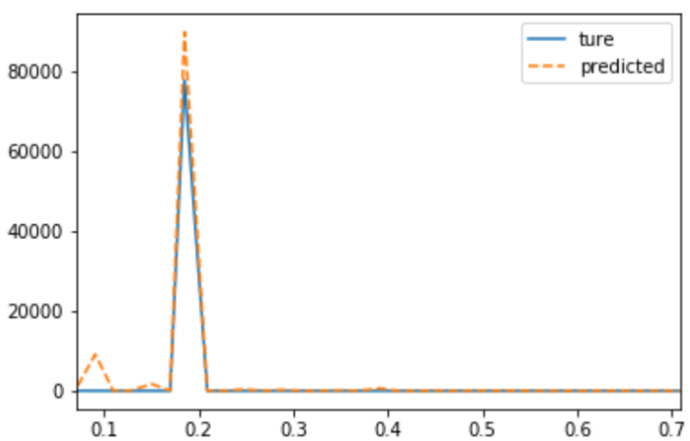

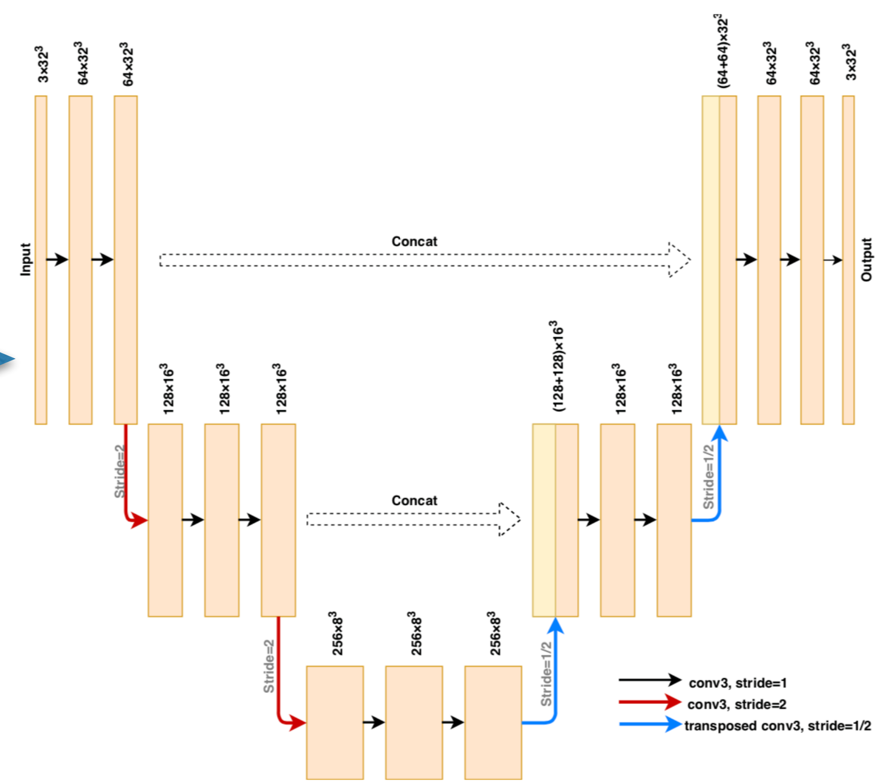

k(h/Mpc)

Inject Power

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

$= [0,0.25,0]$

$= [0,0,0.25]$

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model II

Is Rotational Invariance learnt by the model?

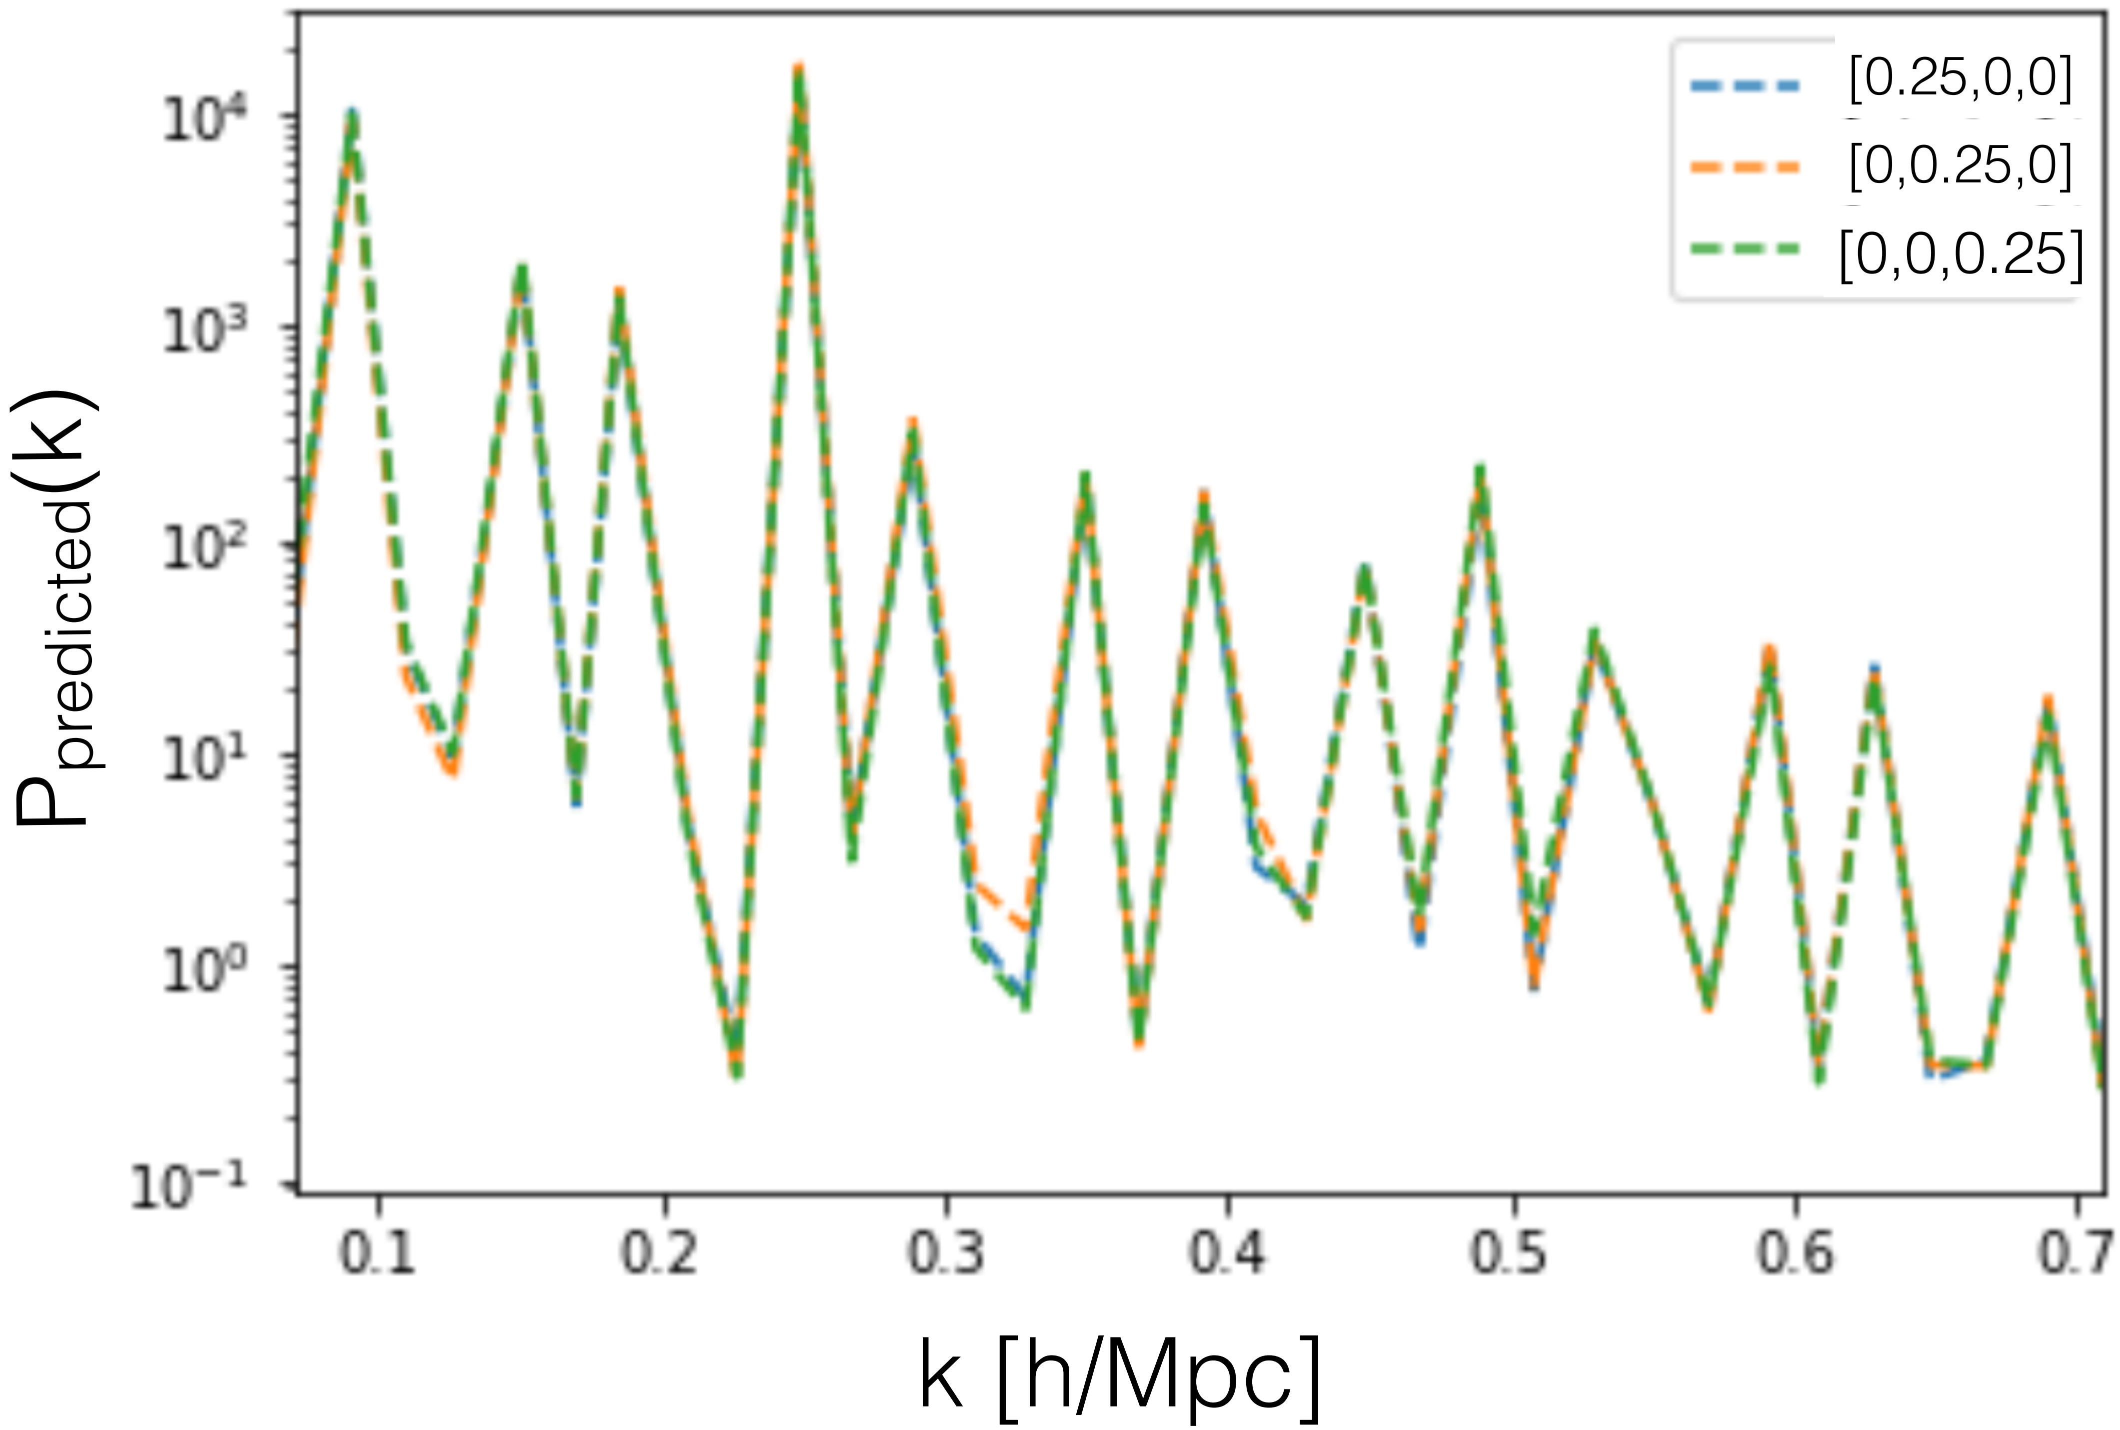

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
  - Use simple analytical cases that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.
  - Locating invariances in the system
  - **Locating where the information is coming from**
  - … other suggestions are very welcome :)
- More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?

Interrogating the machine learning model III

Locating where the information is coming from

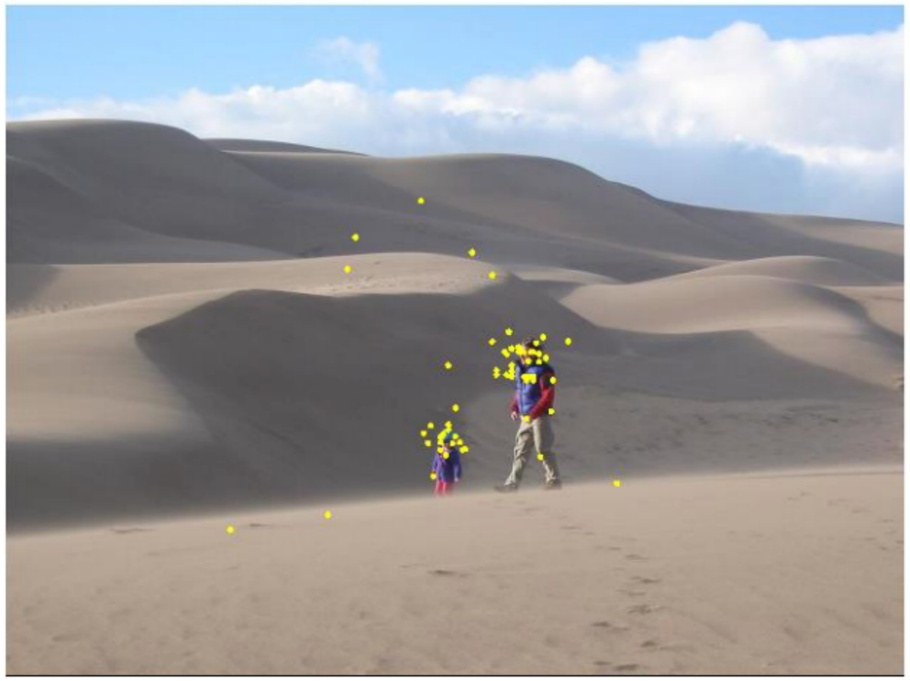

input image

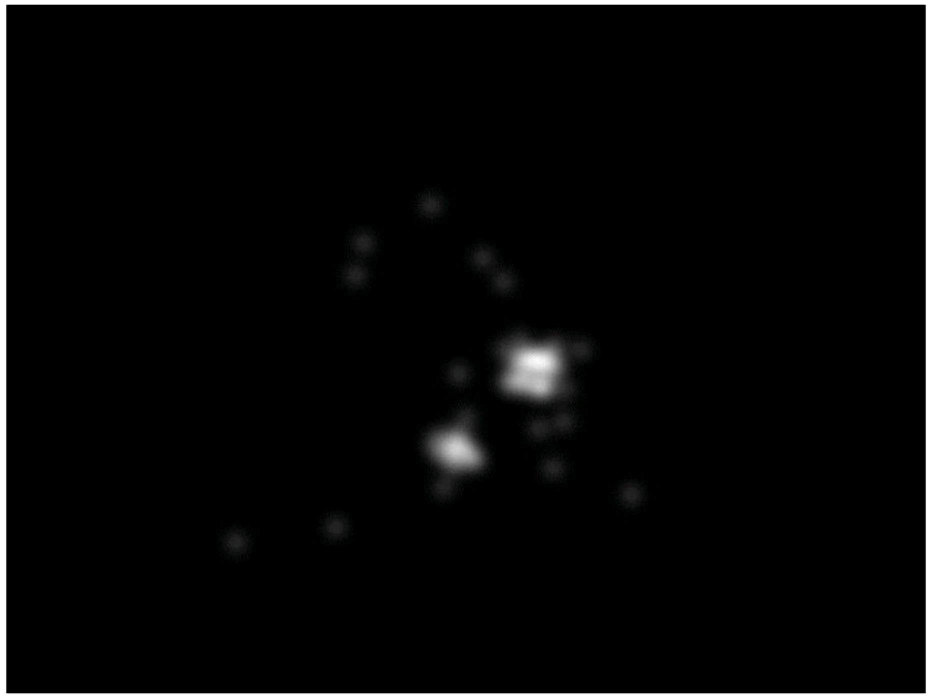

saliency map

- Introducing Saliency map
- Plots the derivative of the output with respect to the input

$$S_c(I) \approx w^T I + b,$$

$$w = \left.\frac{\partial S_c}{\partial I}\right|_{I_0}.$$

Karen Simonyan, Andrea Veldaldi & Andrew Zisserman 2013

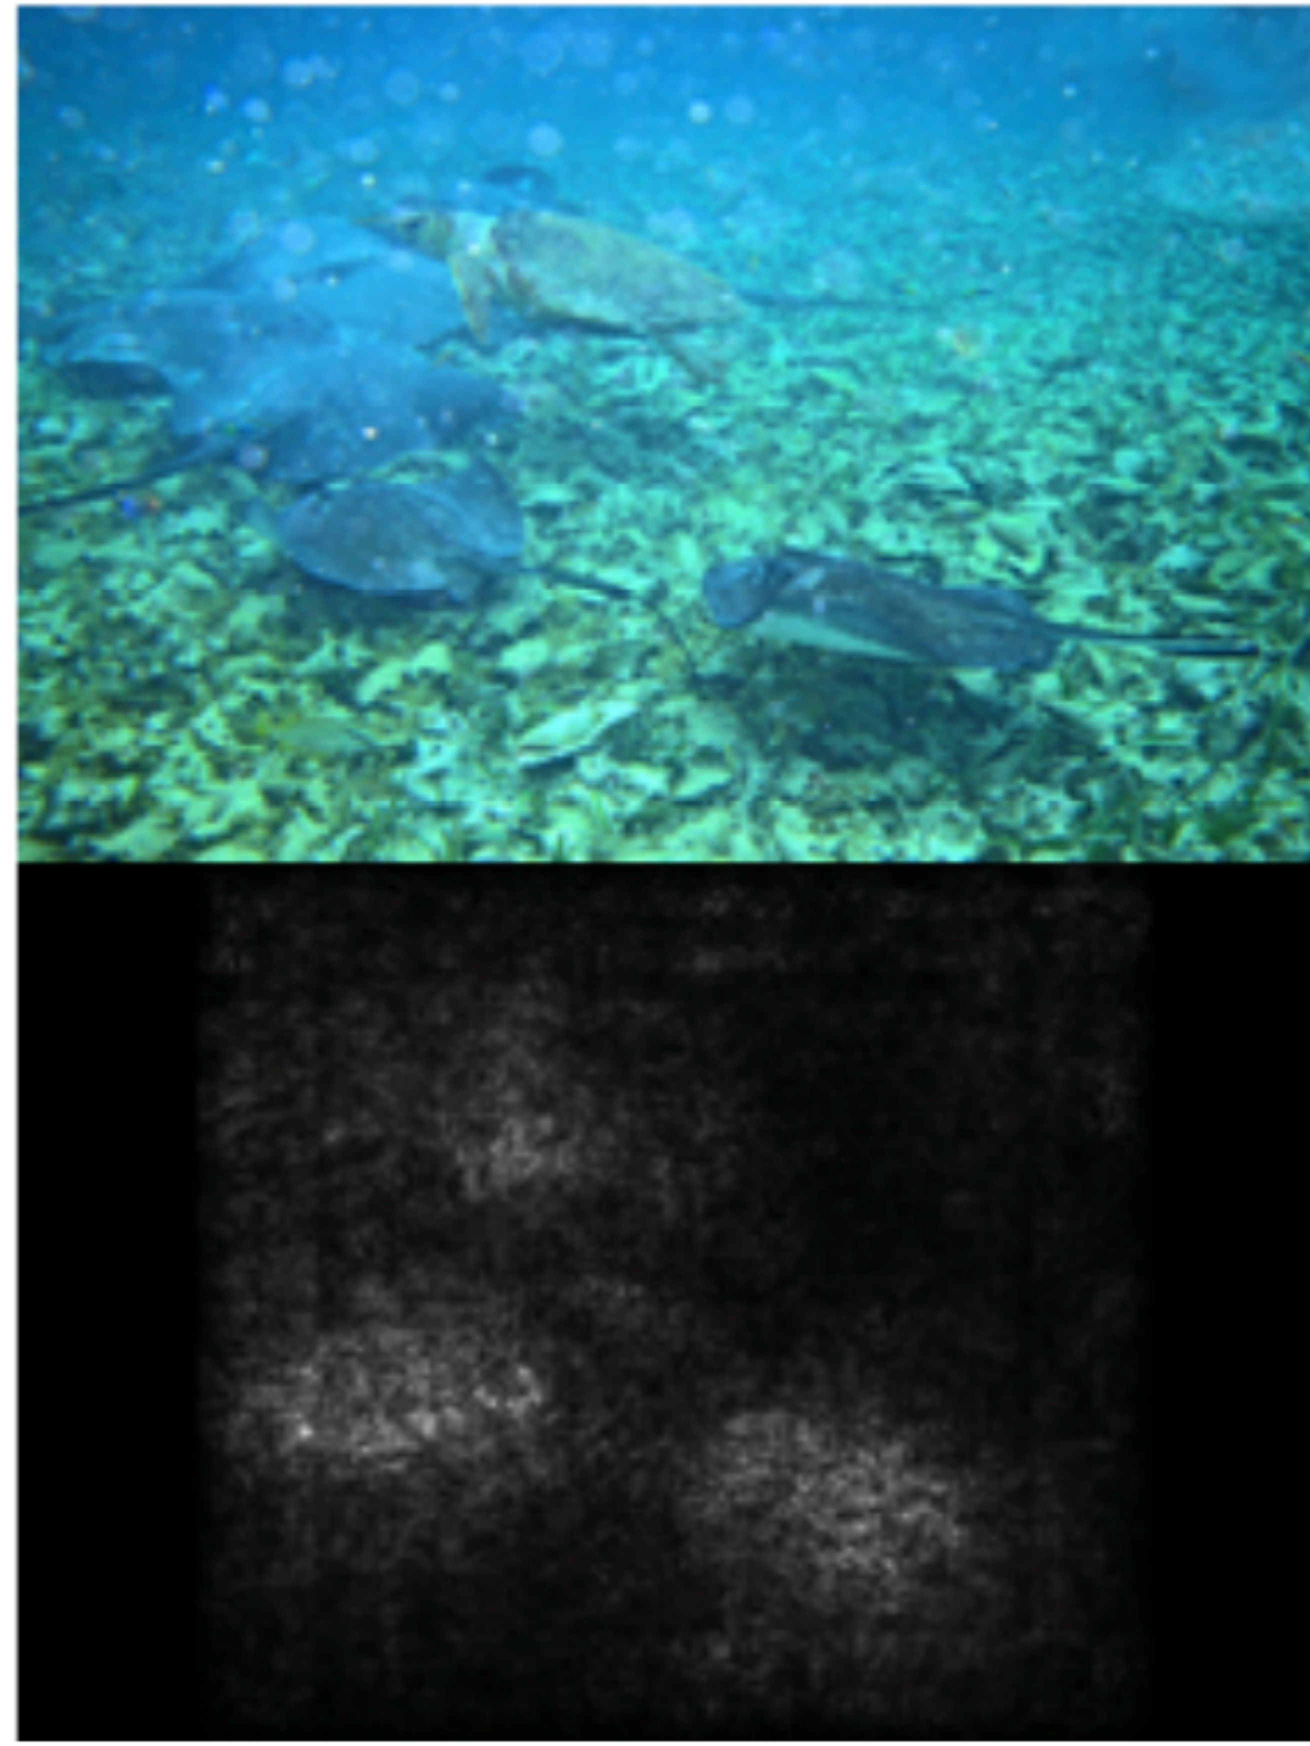

Shirley Ho

Karen Simonyan, Andrea Veldaldi & Andrew Zisserman 2013

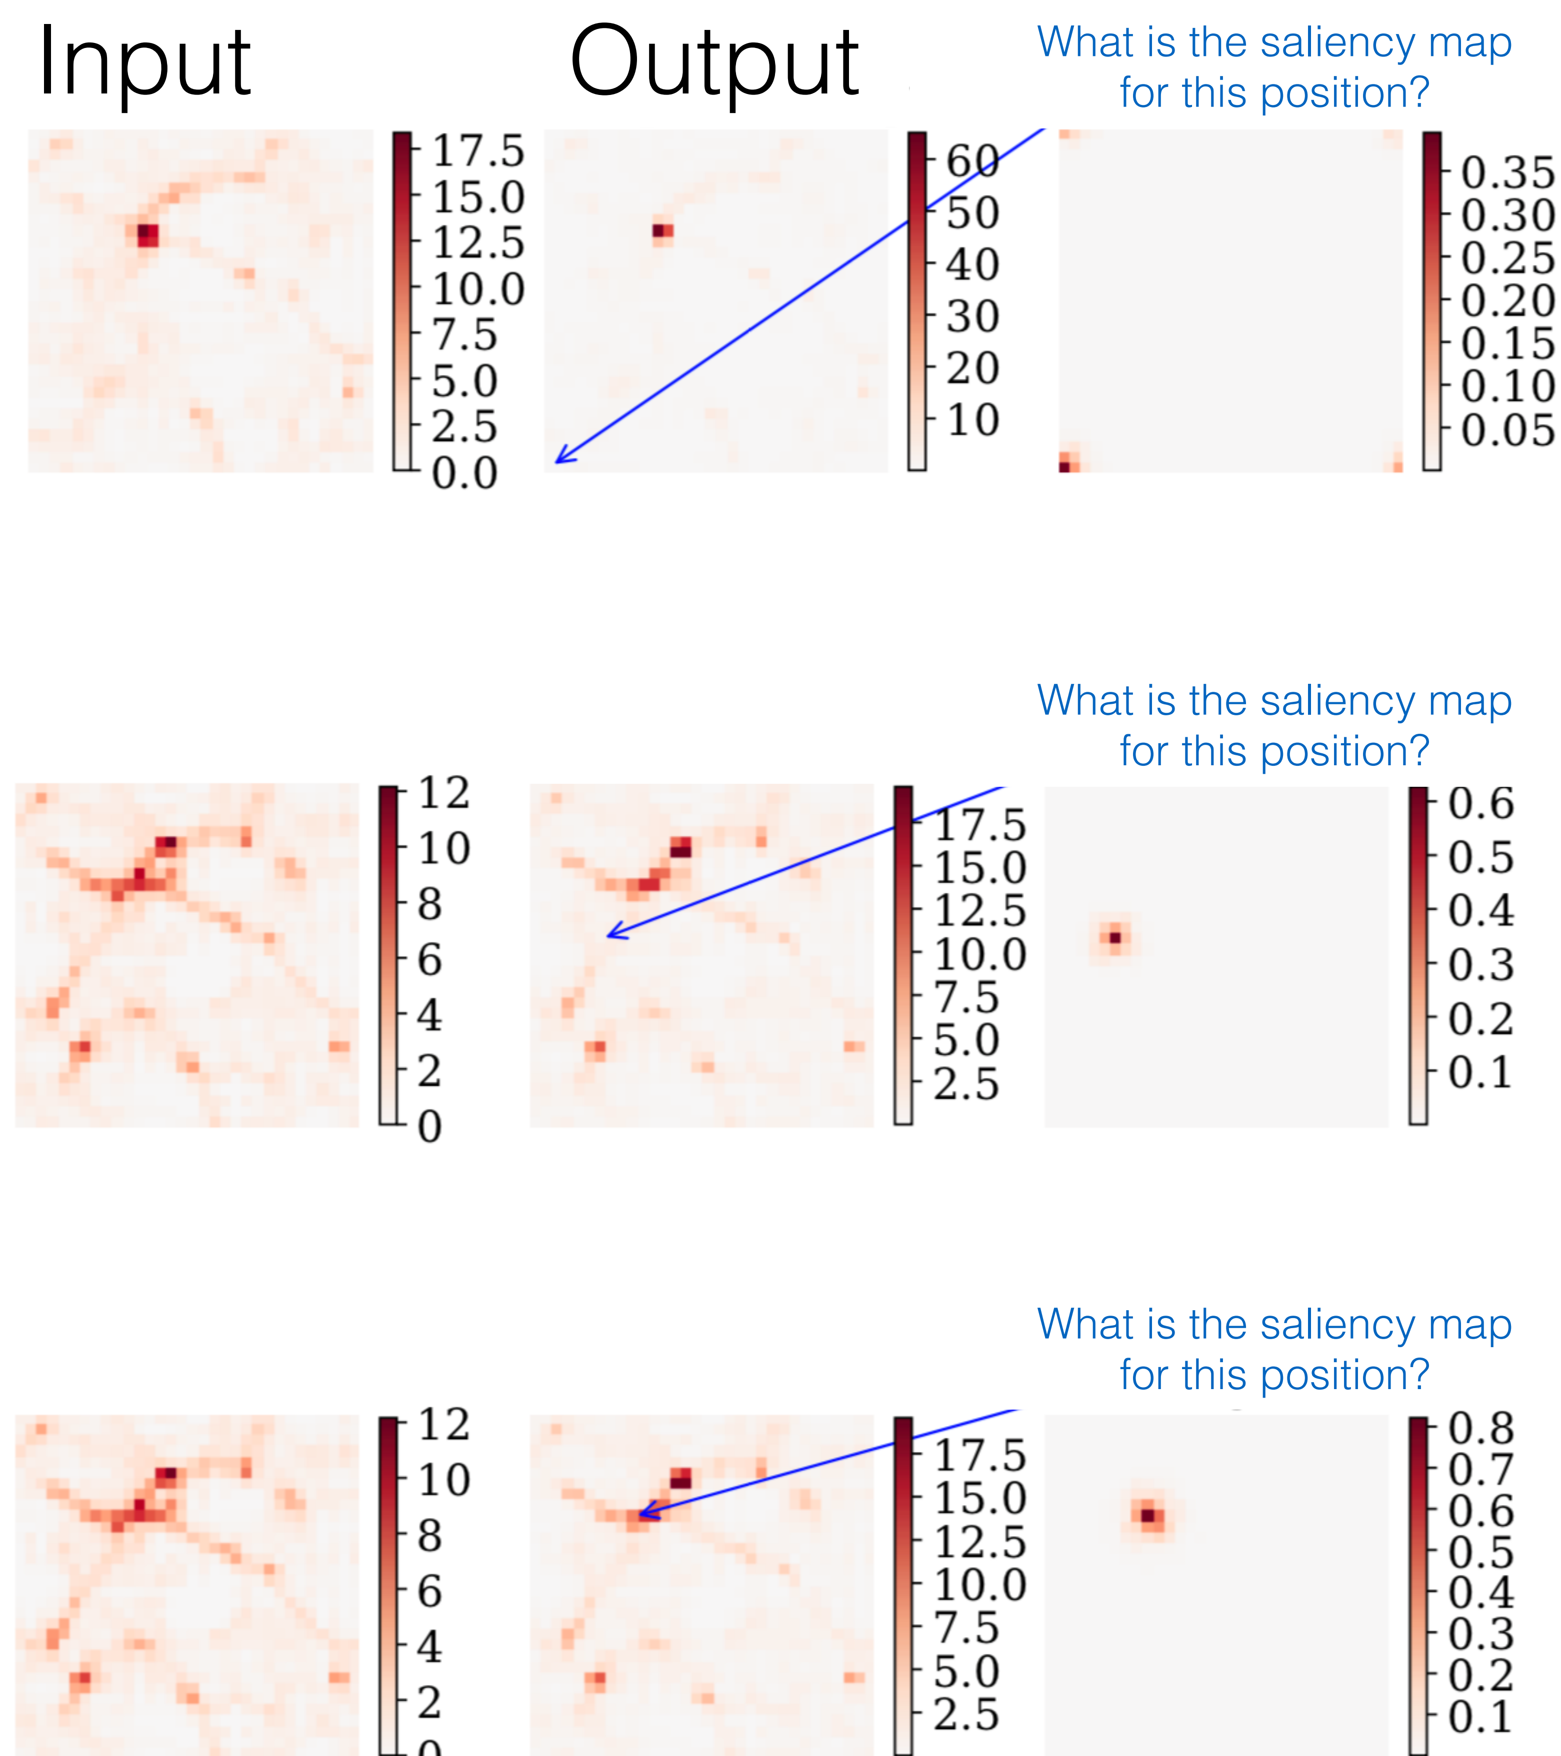

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbakhsh, Barnabas Poczos

Interrogating the machine learning model

- What can we check to understand what the network has learned that seems to understand?
    - Use simple analytical cases that are not in the training-set explicitly and see whether these cases agree with our physics as we know it.
    - Locating invariances in the system
    - Locating where the information is coming from
    - … other suggestions are very welcome :)
- **More provocatively: Can machine learning interpolate and generalize from data just like human and find new physical laws?**

Interrogating the machine learning model IV

Can the model extrapolate instead of just interpolate?

Input

Prediction

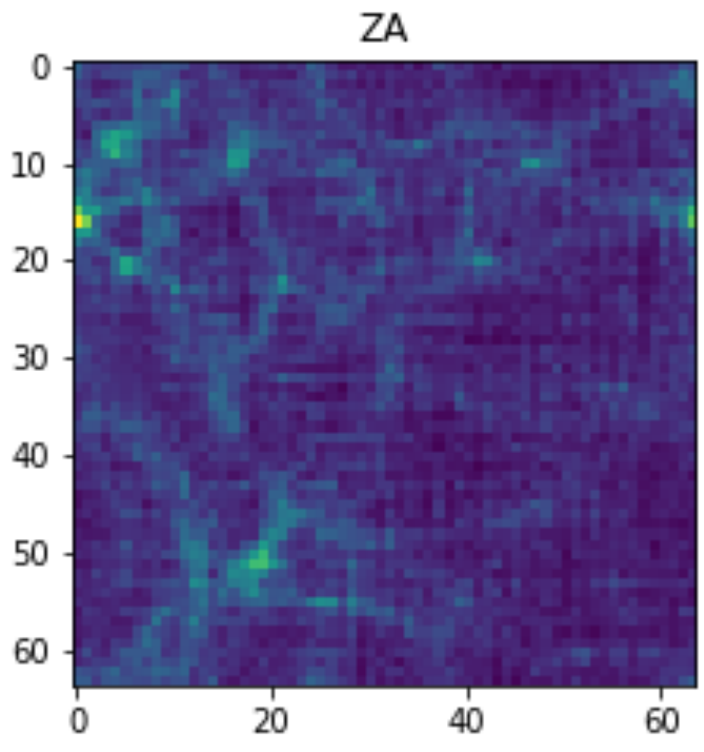

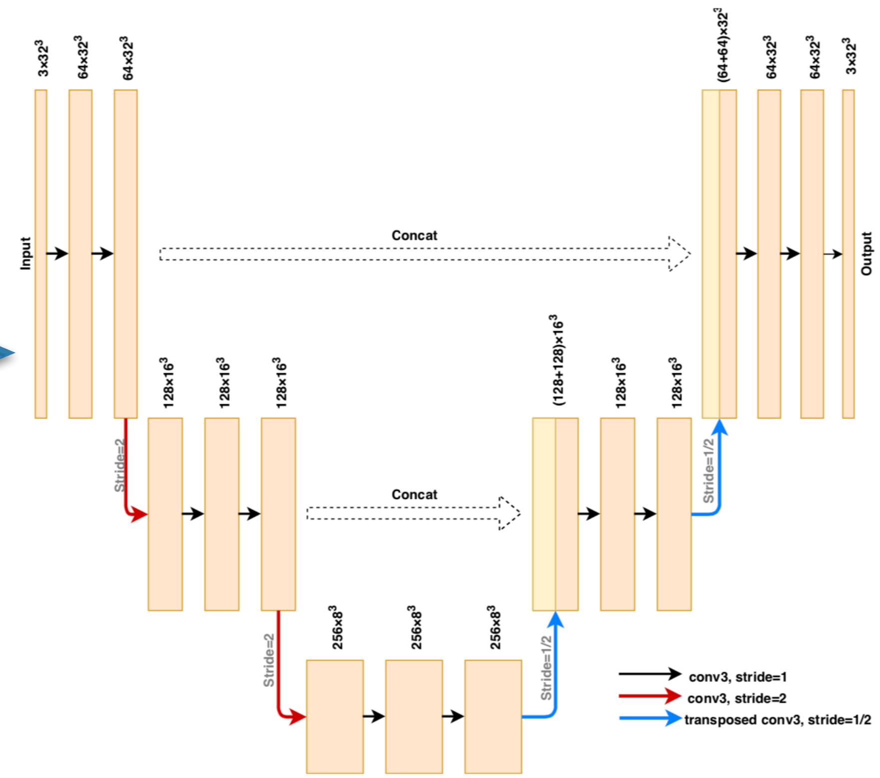

ZA maps of Different cosmology

Dark matter density parameter = [0.1 - 0.5]

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model IV

Can the model extrapolate instead of just interpolate?

Input

UNET
Slight variant to Residual NN

Prediction

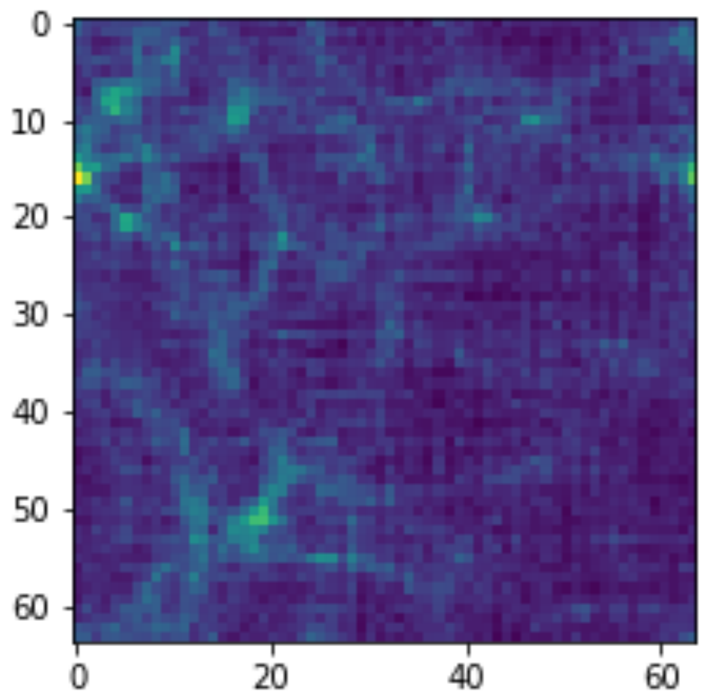

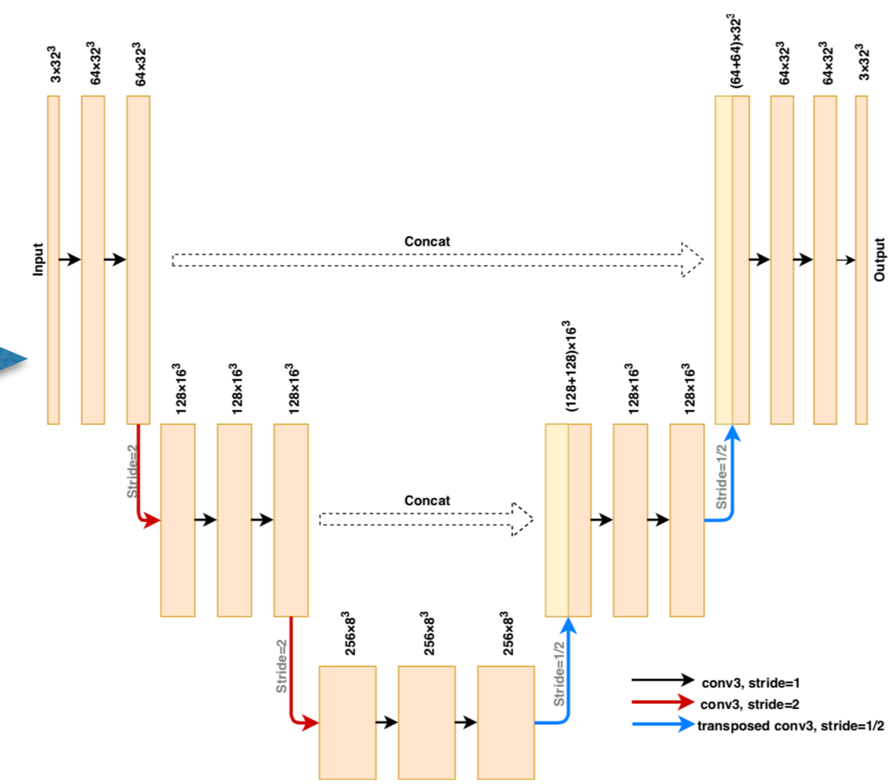

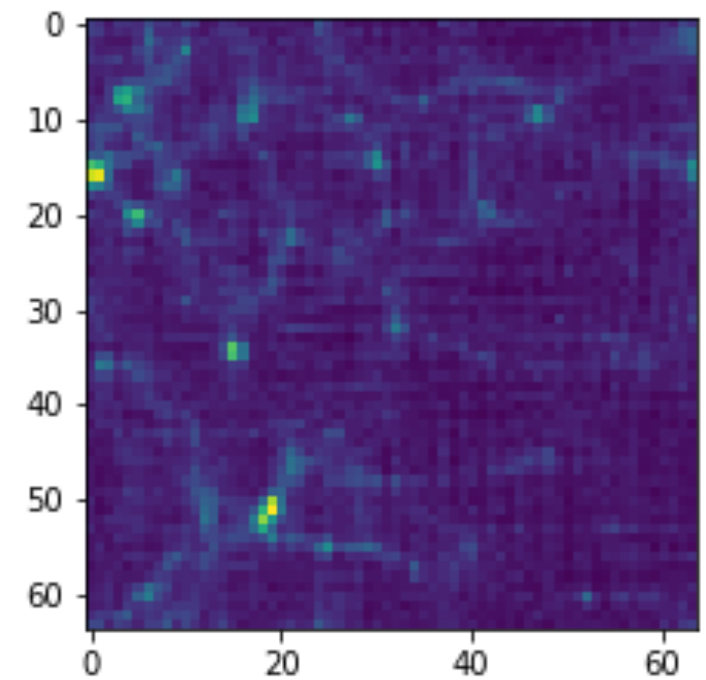

ZA maps of Different cosmology

Dark matter density parameter = [0.1 - 0.5]

Can the learned model "extrapolate" and predict simulations that do not have the same Cosmological parameters?

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Interrogating the machine learning model IV

Can the model extrapolate instead of just interpolate?

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.

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

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

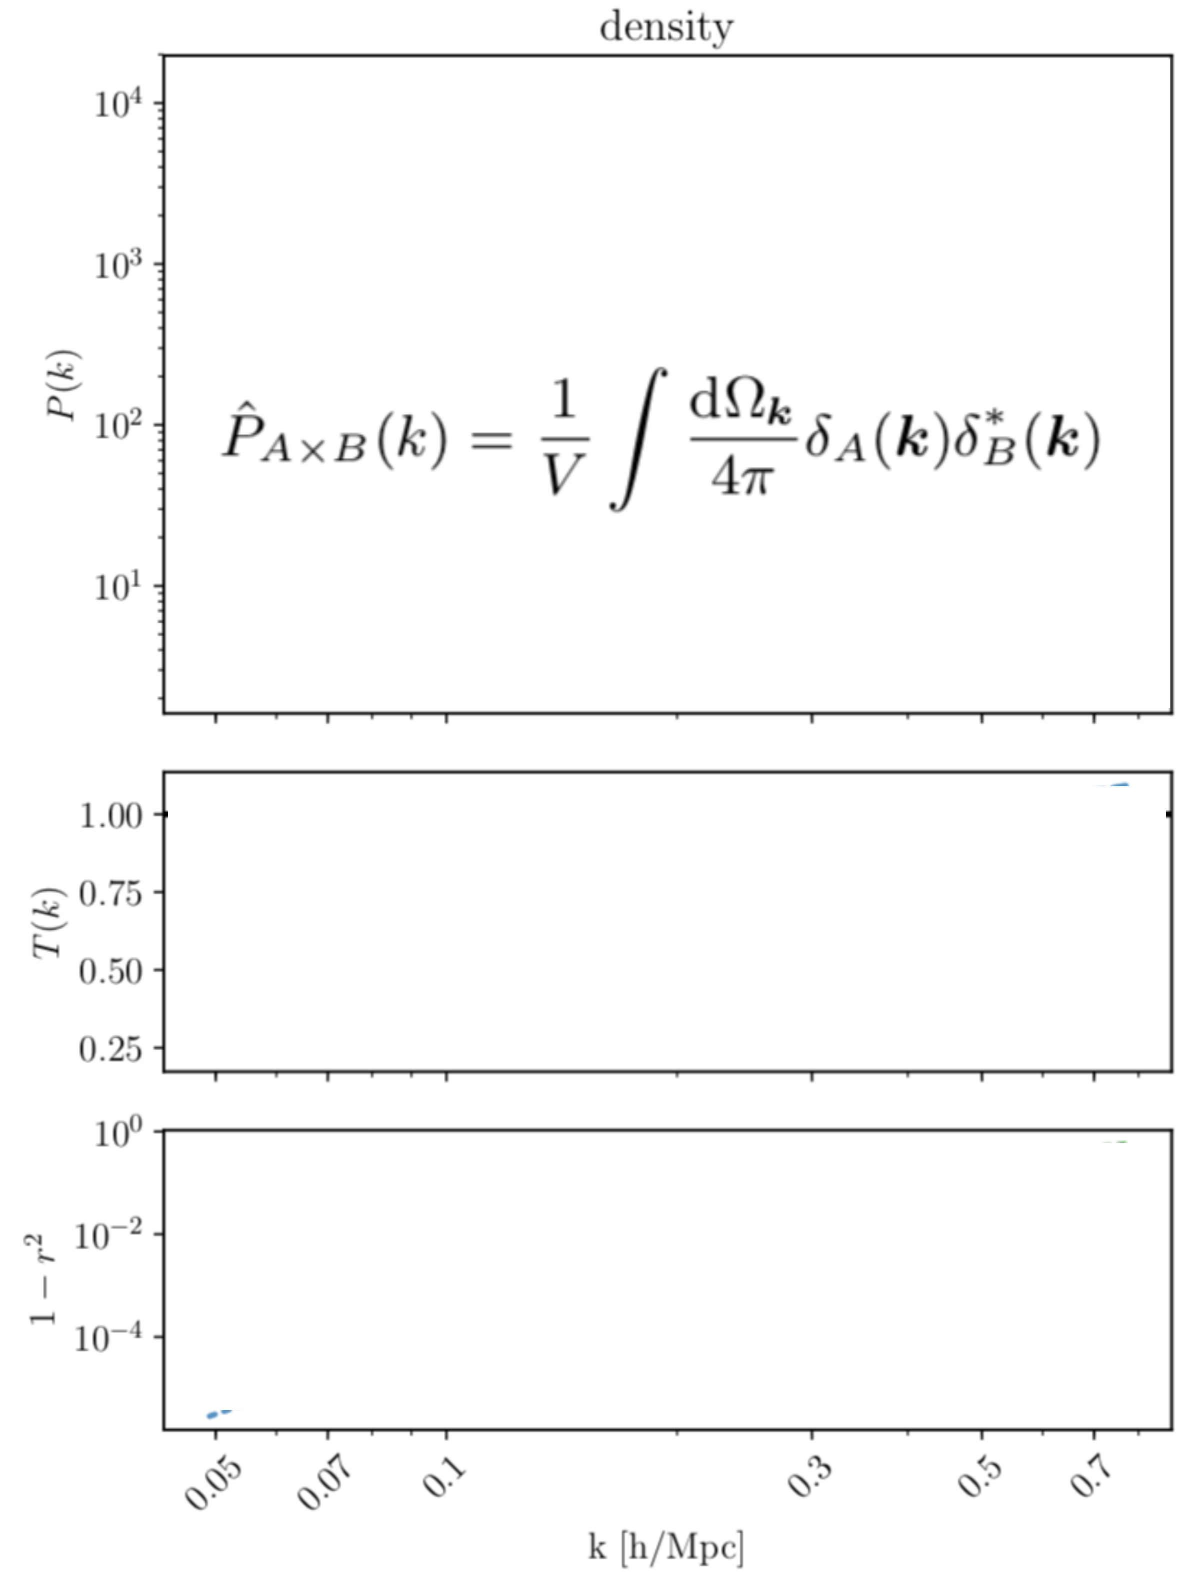

(b)Two point analysis for density field

Solid Line: Simulation /Truth
Long Dashed line: Prediction using ML
Short Dashed Line: Analytical approximation (2LPT)

Shirley Ho

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

Interrogating the machine learning model IV

Can the model extrapolate instead of just interpolate?

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.

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

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

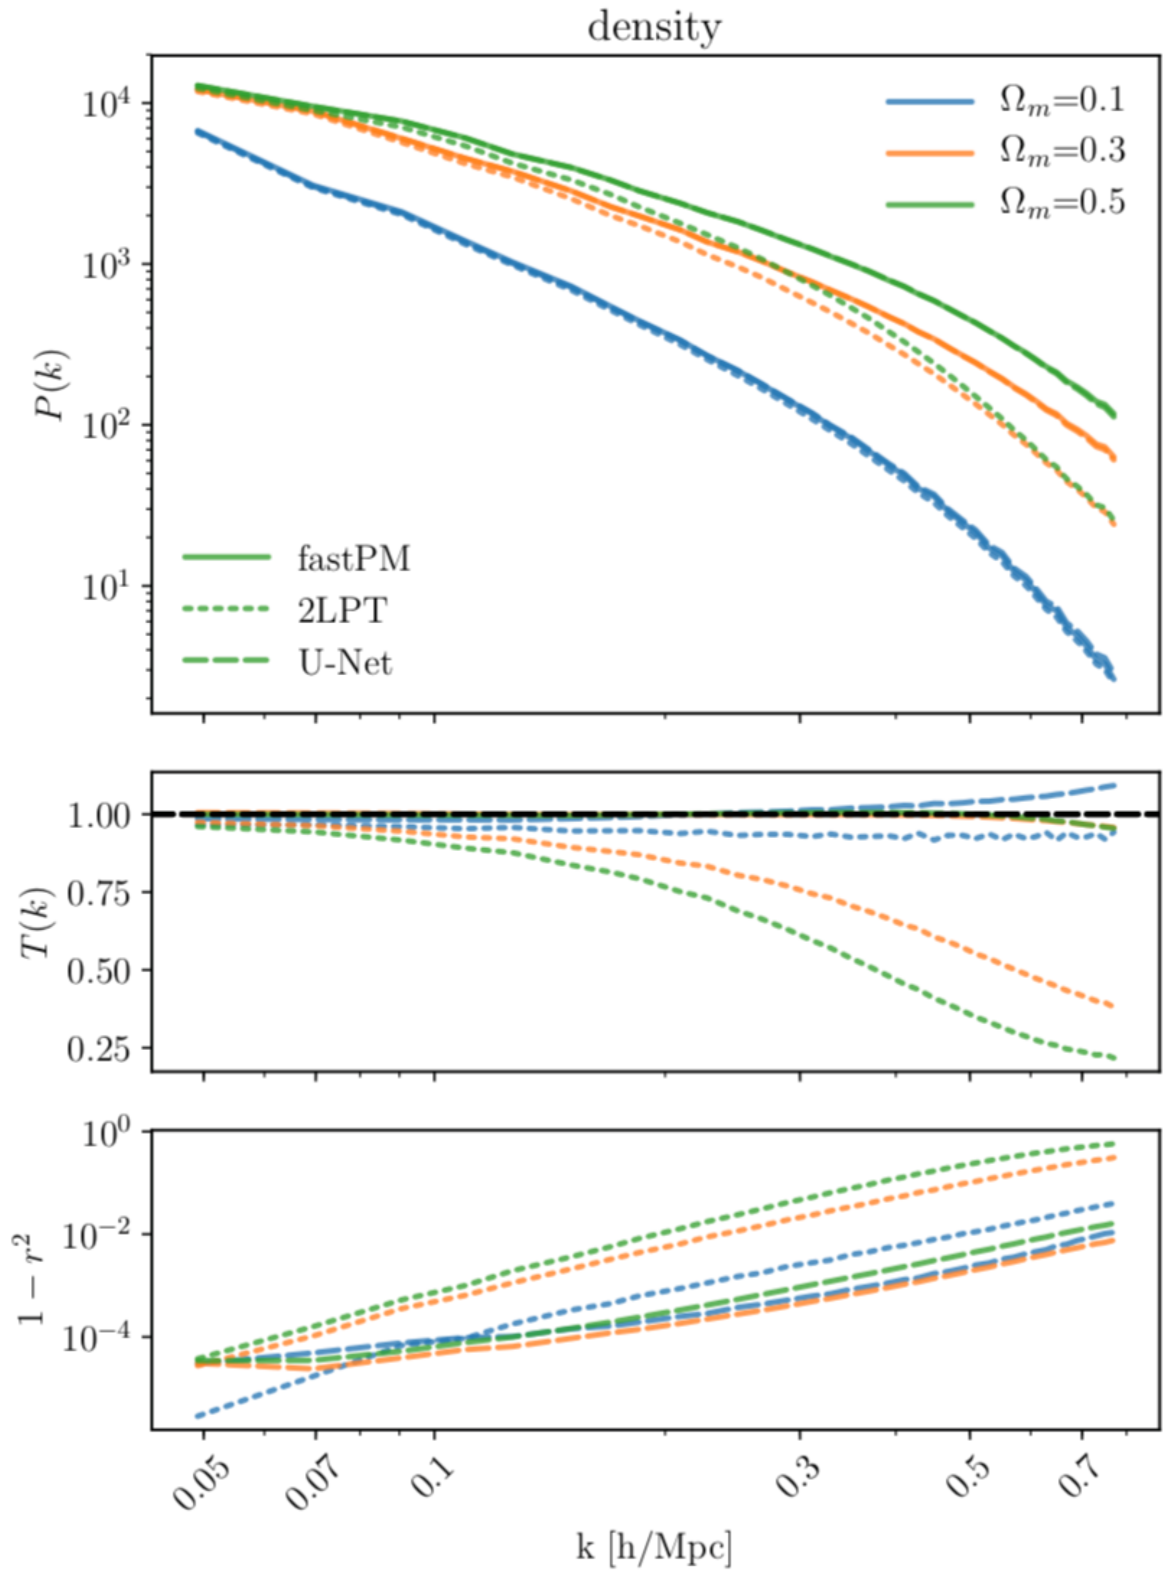

(b)Two point analysis for density field

Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos 2018, arxiv:1811.06533

What the heck happened?

- We didn’t need to overlap the training set with test set.
- We did not explicitly use any transfer learning or meta-learning
- Maybe there is overlap in information somewhere between these universes?
- Maybe the Universe is fairly simple, so that the generalization and extrapolation by the network is ‘easy’?
- Maybe I will finally get famous ?

My possible climb to fame?

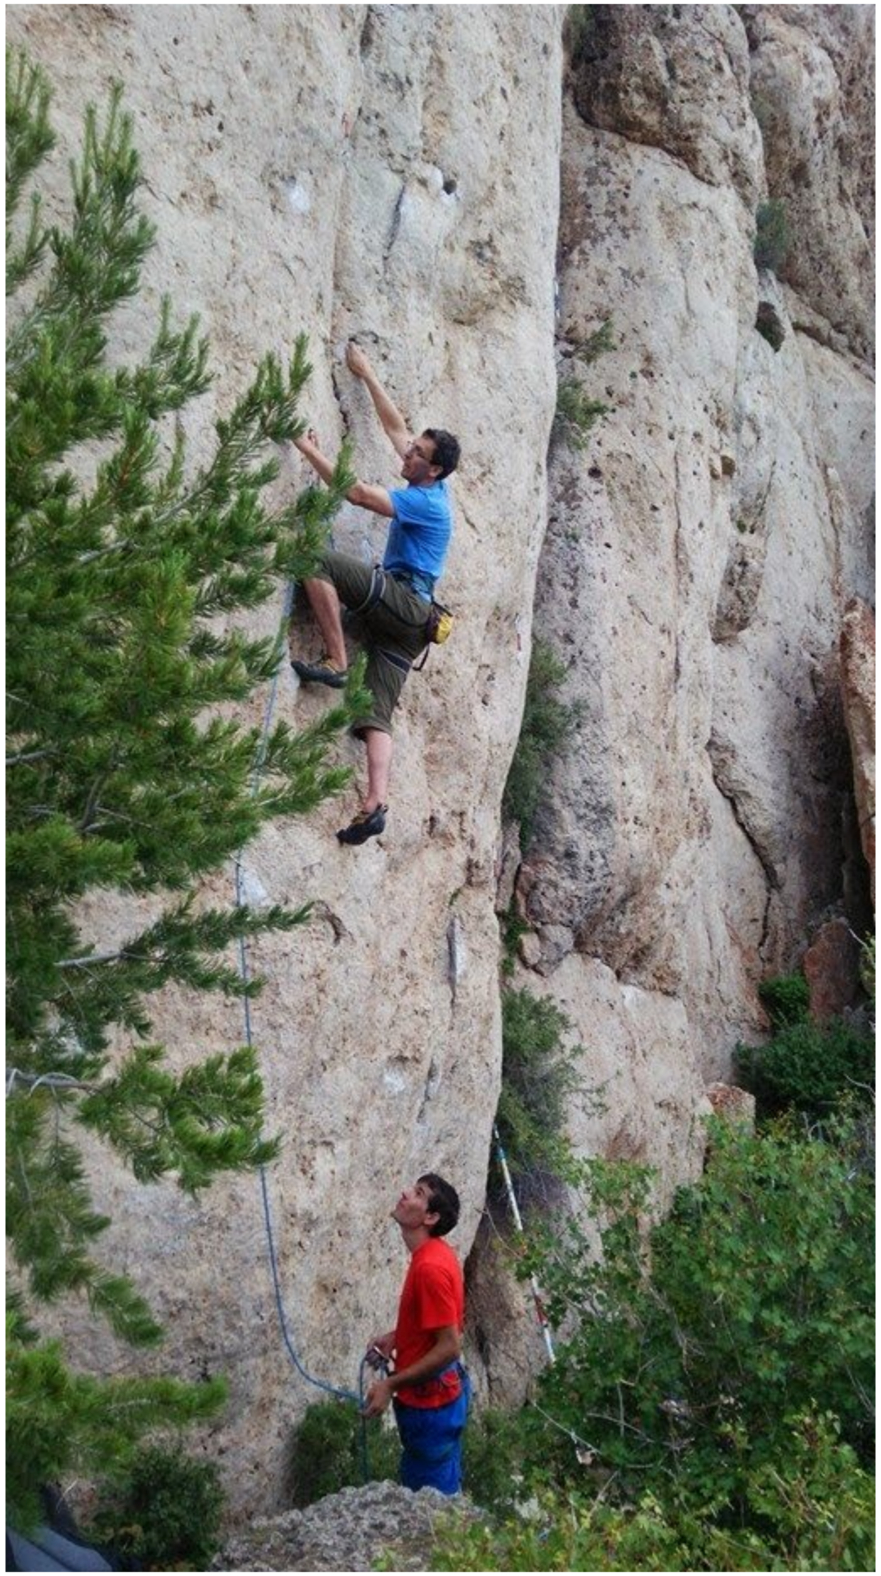

My possible climb to fame?

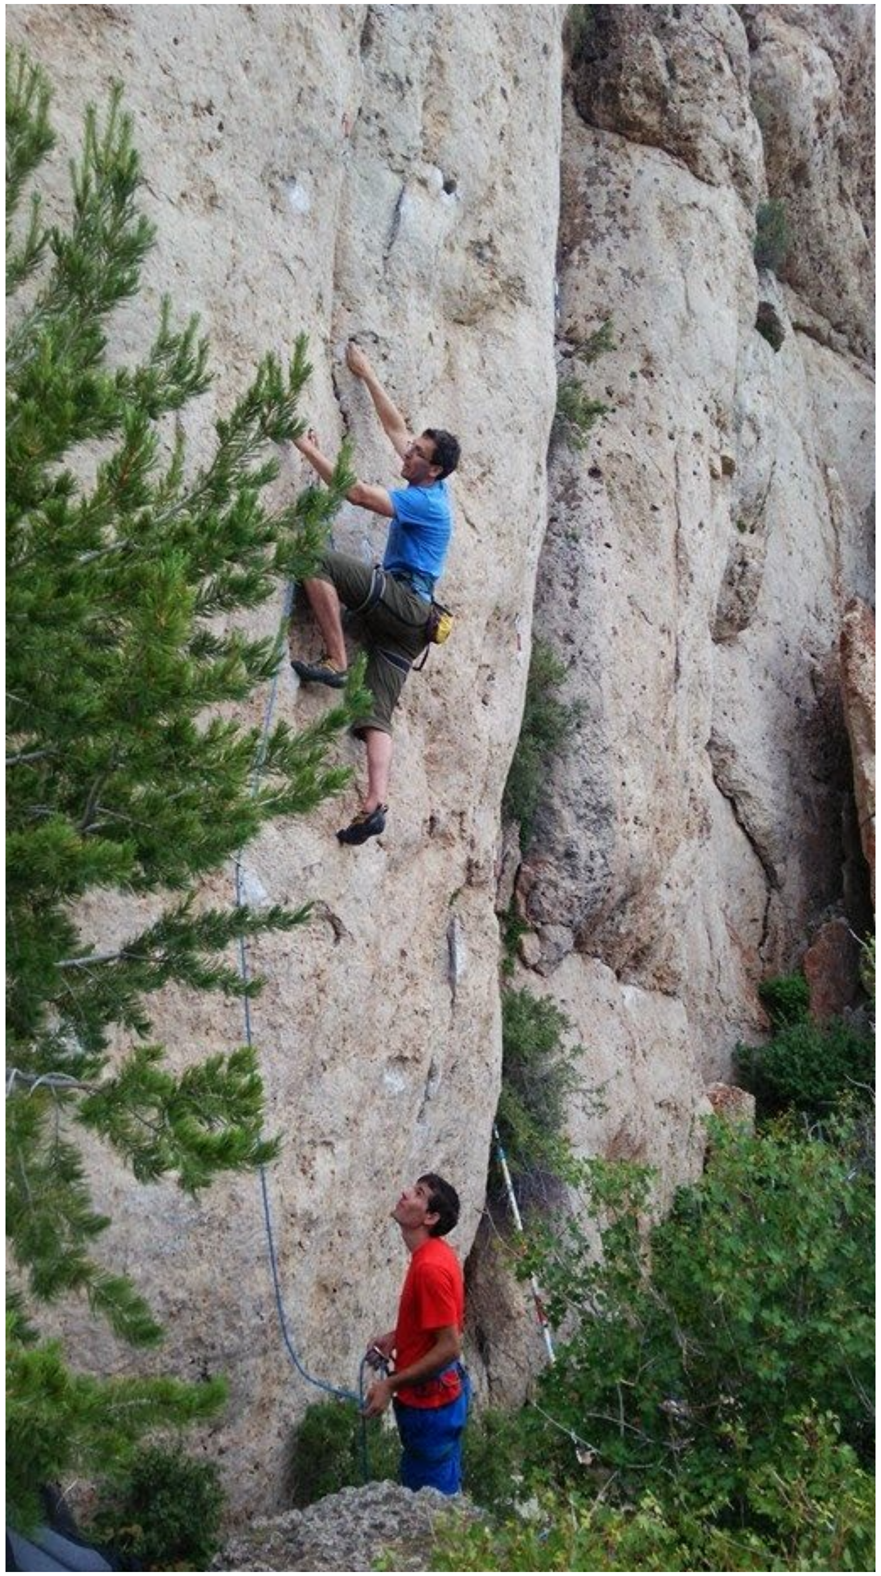

- Understanding Machine Learning?
- Compressing the learned model into physical laws?
- Discover new laws of nature?

Conclusions

- There is immense hype, and probably immense potential for Machine Learning in everything field today, ranging from playing Go, image recognition to health-care.
- We may be able to use machine learning to help advance physics and astrophysics
    - In cosmology, we used deep neural networks to predict cosmological parameters with some successes
    - We can start to use machine learning to be an approximate simulator in not only small number systems, but also relatively complex systems like our Universe.
- But we need to understand what is happening under the hood to fully employ machine learning.
- Furthermore, physical datasets can also provide an interesting playground for understand machine learning as we have a much better understanding of the natural world than the random pictures taken off facebook.
- We have more questions than answers. But that's why it is exciting !

It seems like physics are being learned by the model…

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 ?

Possible reason ?

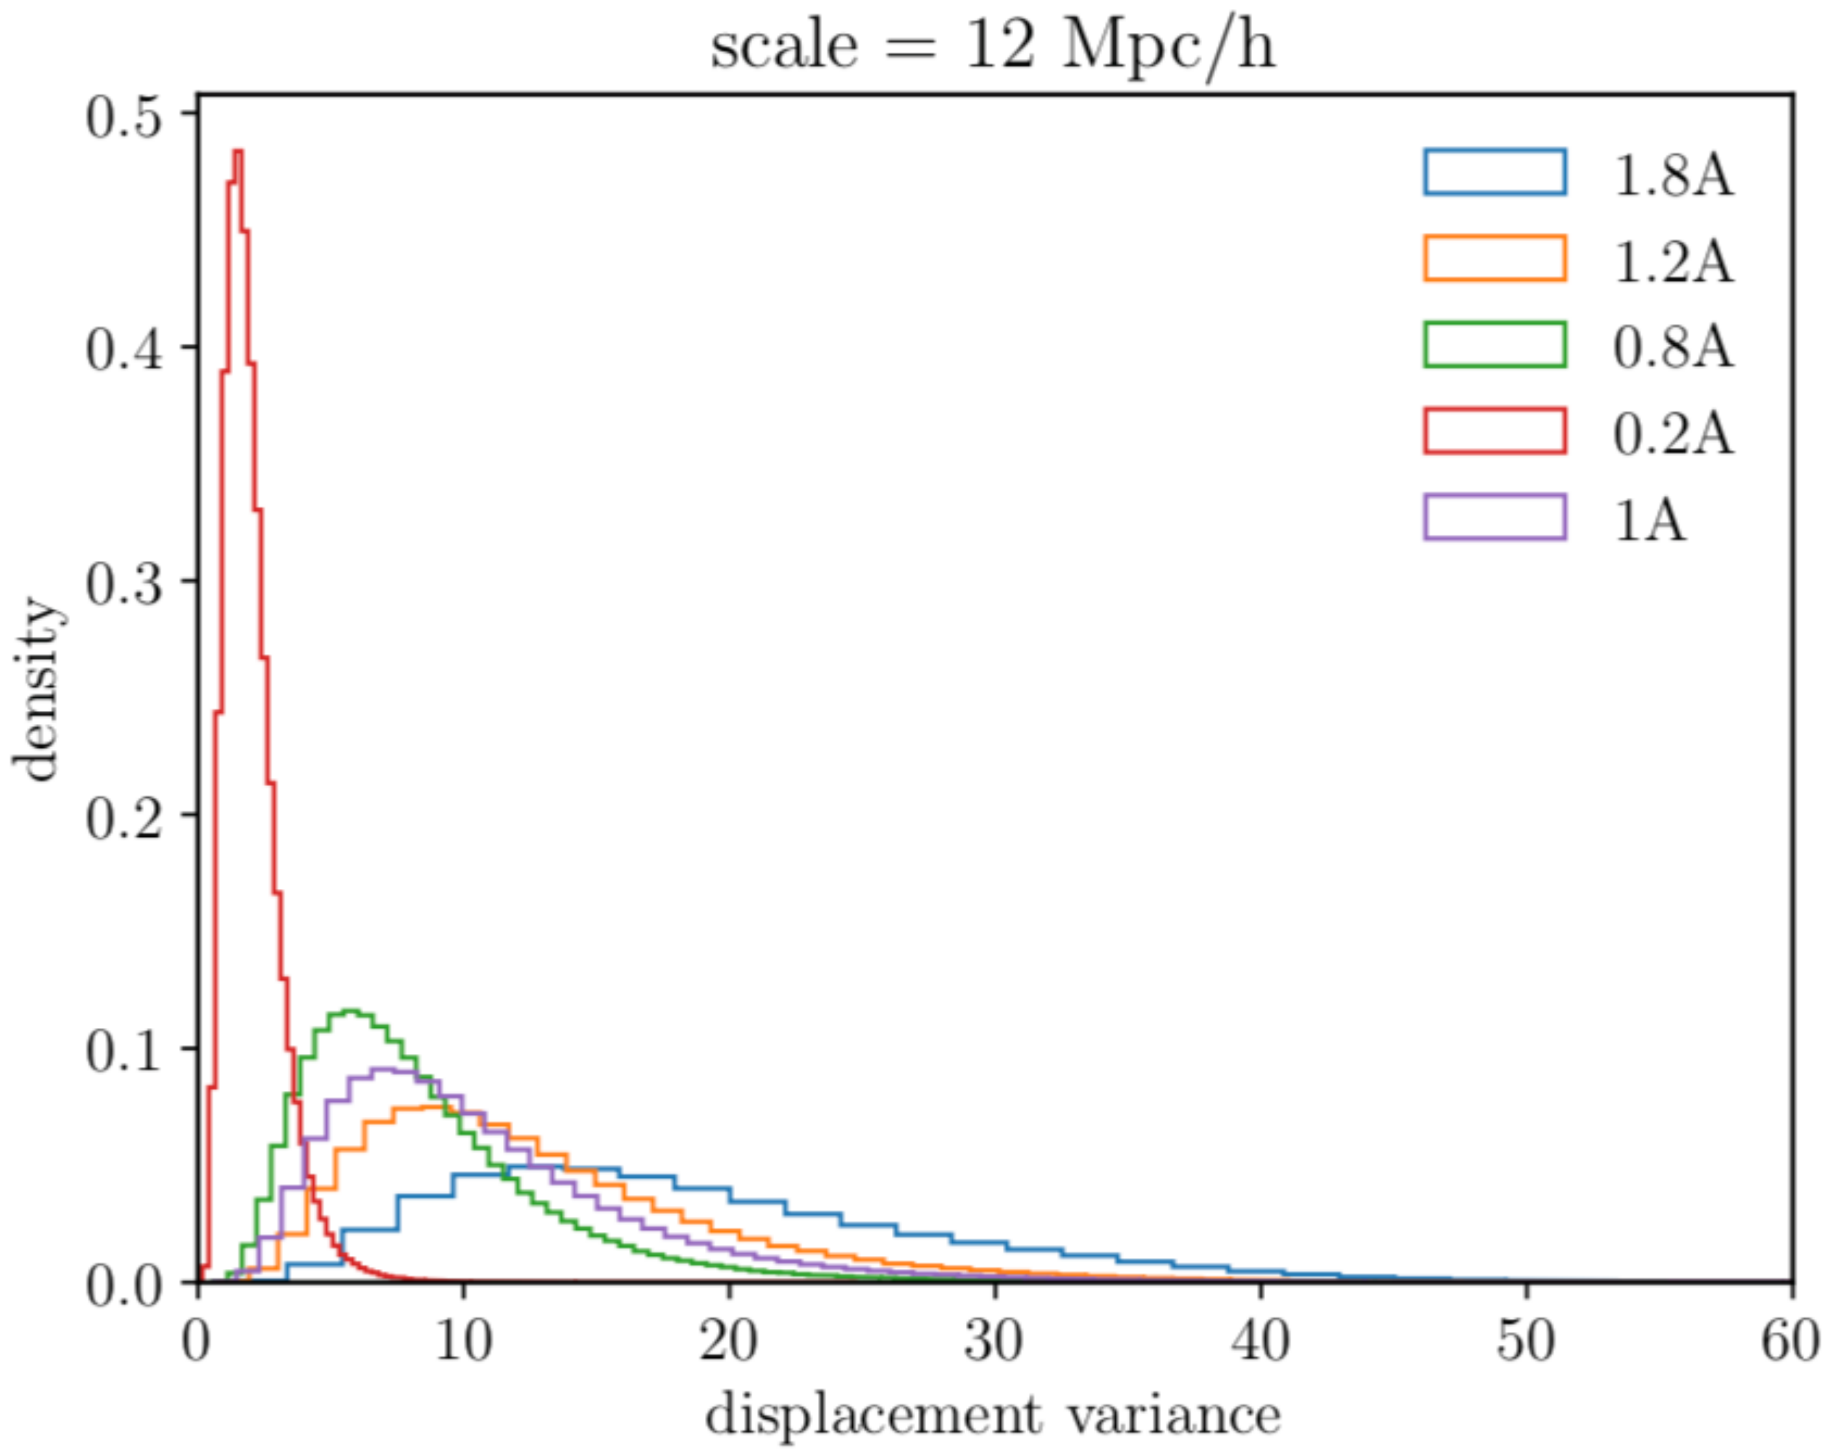

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

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

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

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

Dashed Line -> (2LPT) Theoretical predictions

Preliminary results from: Siyu He, Yin Li, Yu Feng, S.H., Siamak Ravanbaksh, Barnabas Poczos

Power-spectrum of Density field

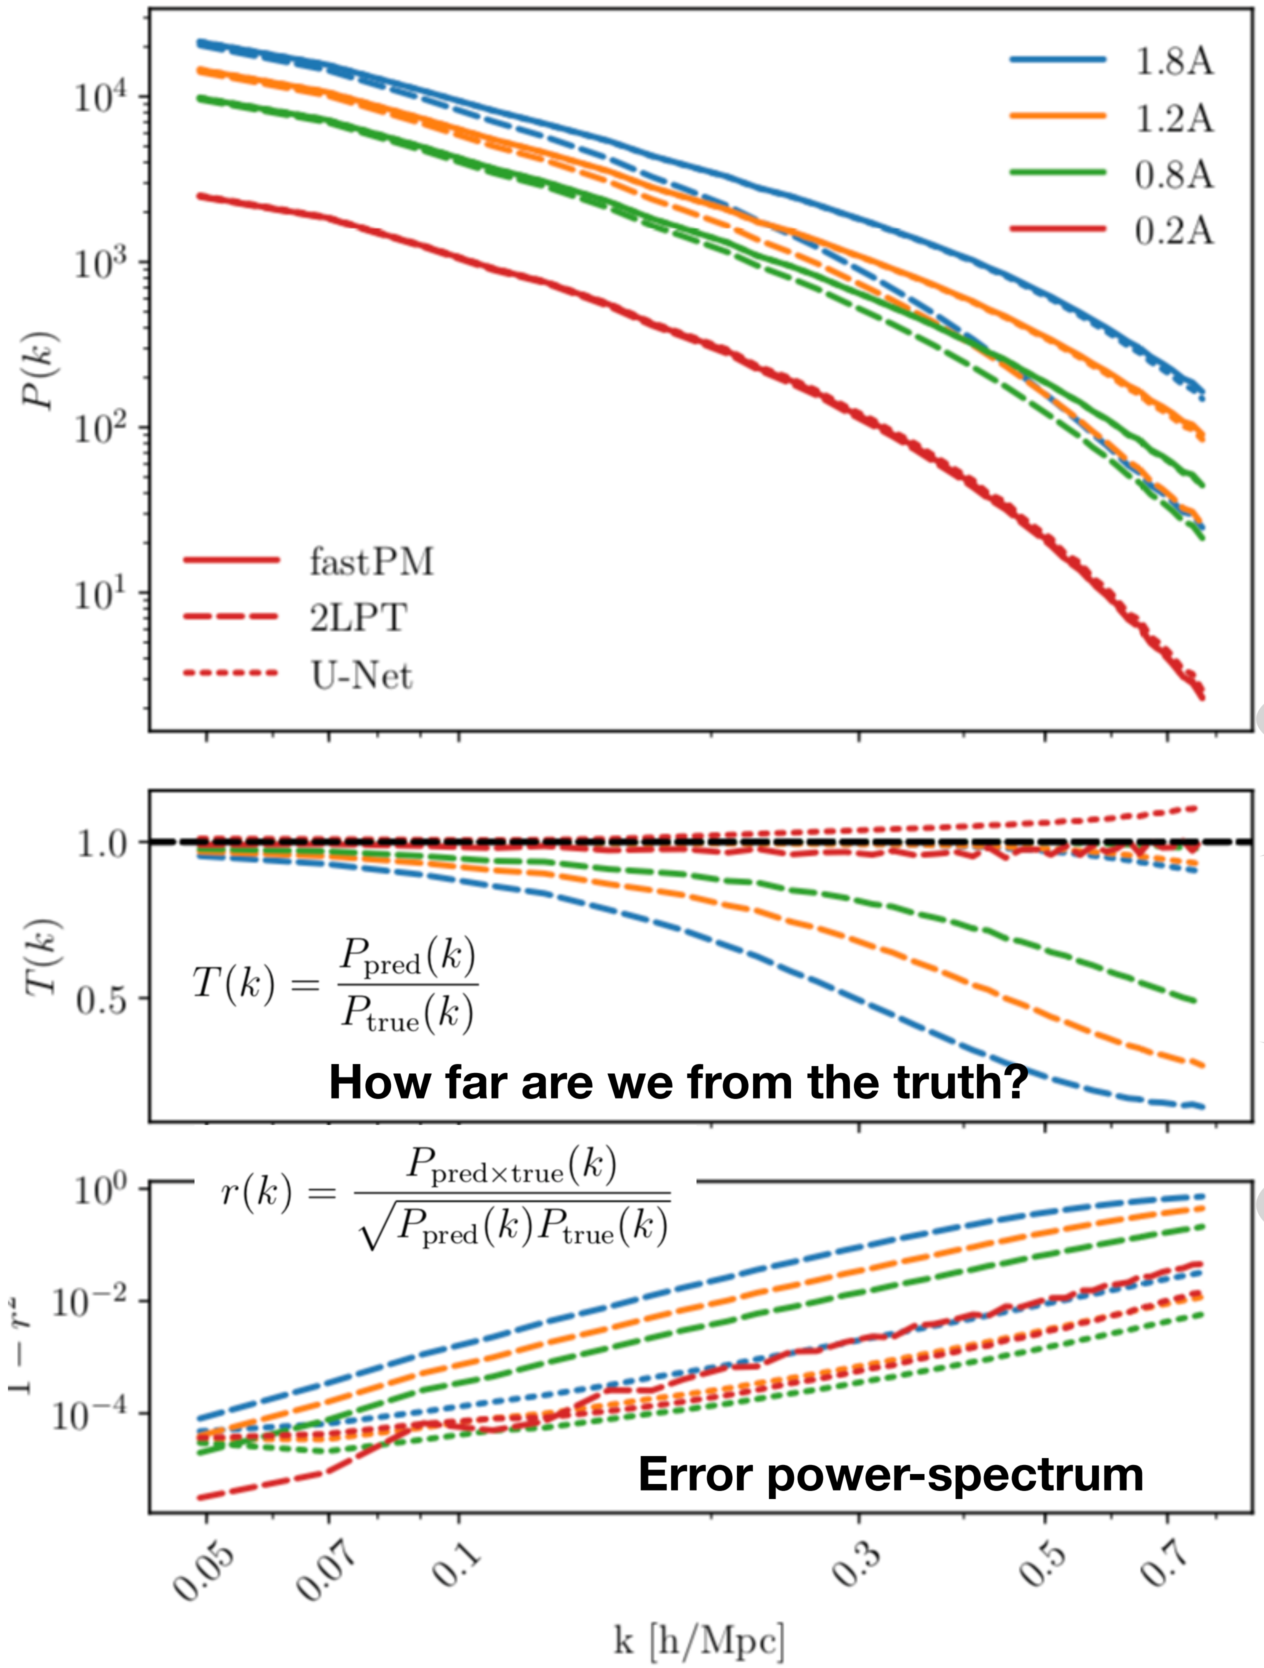

(a) Results from the density field

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

Dashed Line -> (2LPT) Theoretical predictions

Preliminary results from: Siyu He, Yin Li, Yu Feng, S.H., Siamak Ravanbaksh, Barnabas Poczos

Power-spectrum of Density field

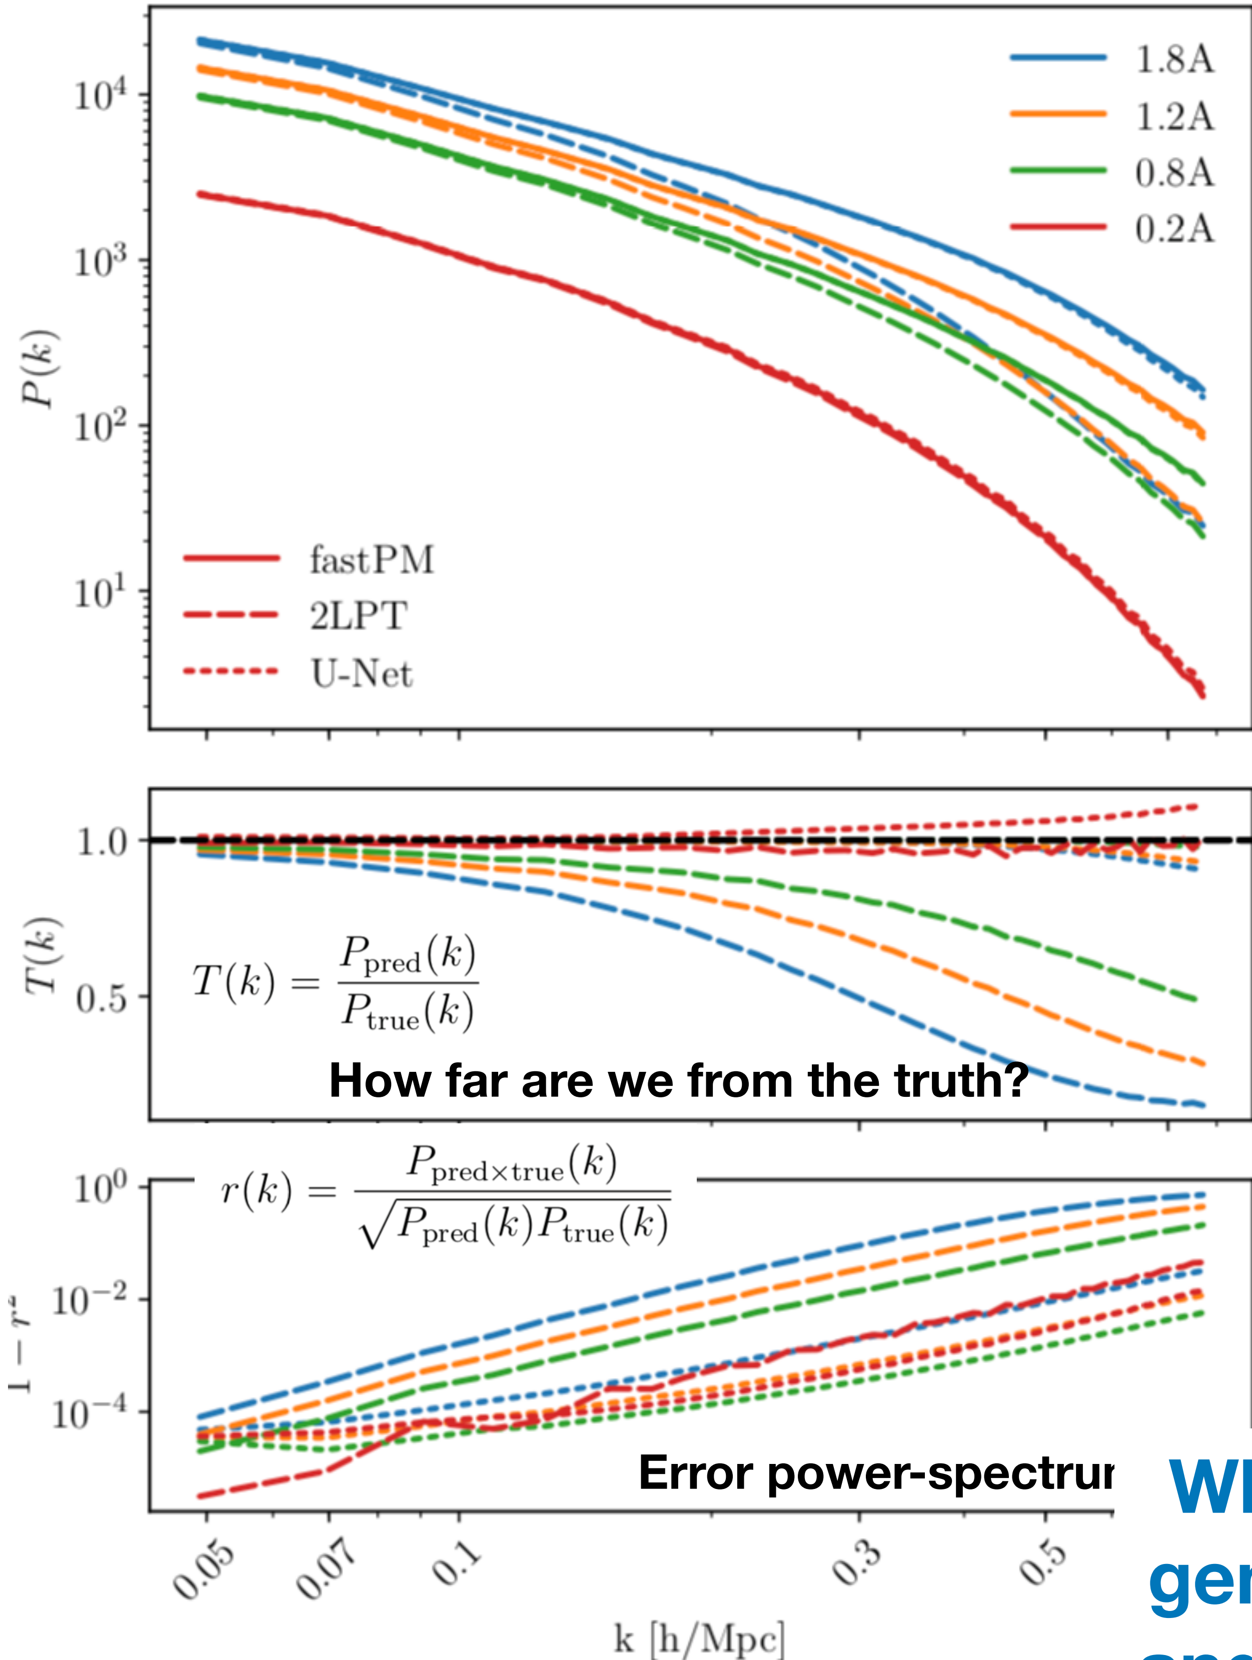

(a) Results from the density field

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

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.

Preliminary results from: Siyu He, Yin Li, Yu Feng, S.H., Siamak Ravanbaksh, Barnabas Poczos

Improving on Machine Learning algorithms:
Introducing Deep Residual Neural Net

ImageNet Large Scale Visual Recognition Competition

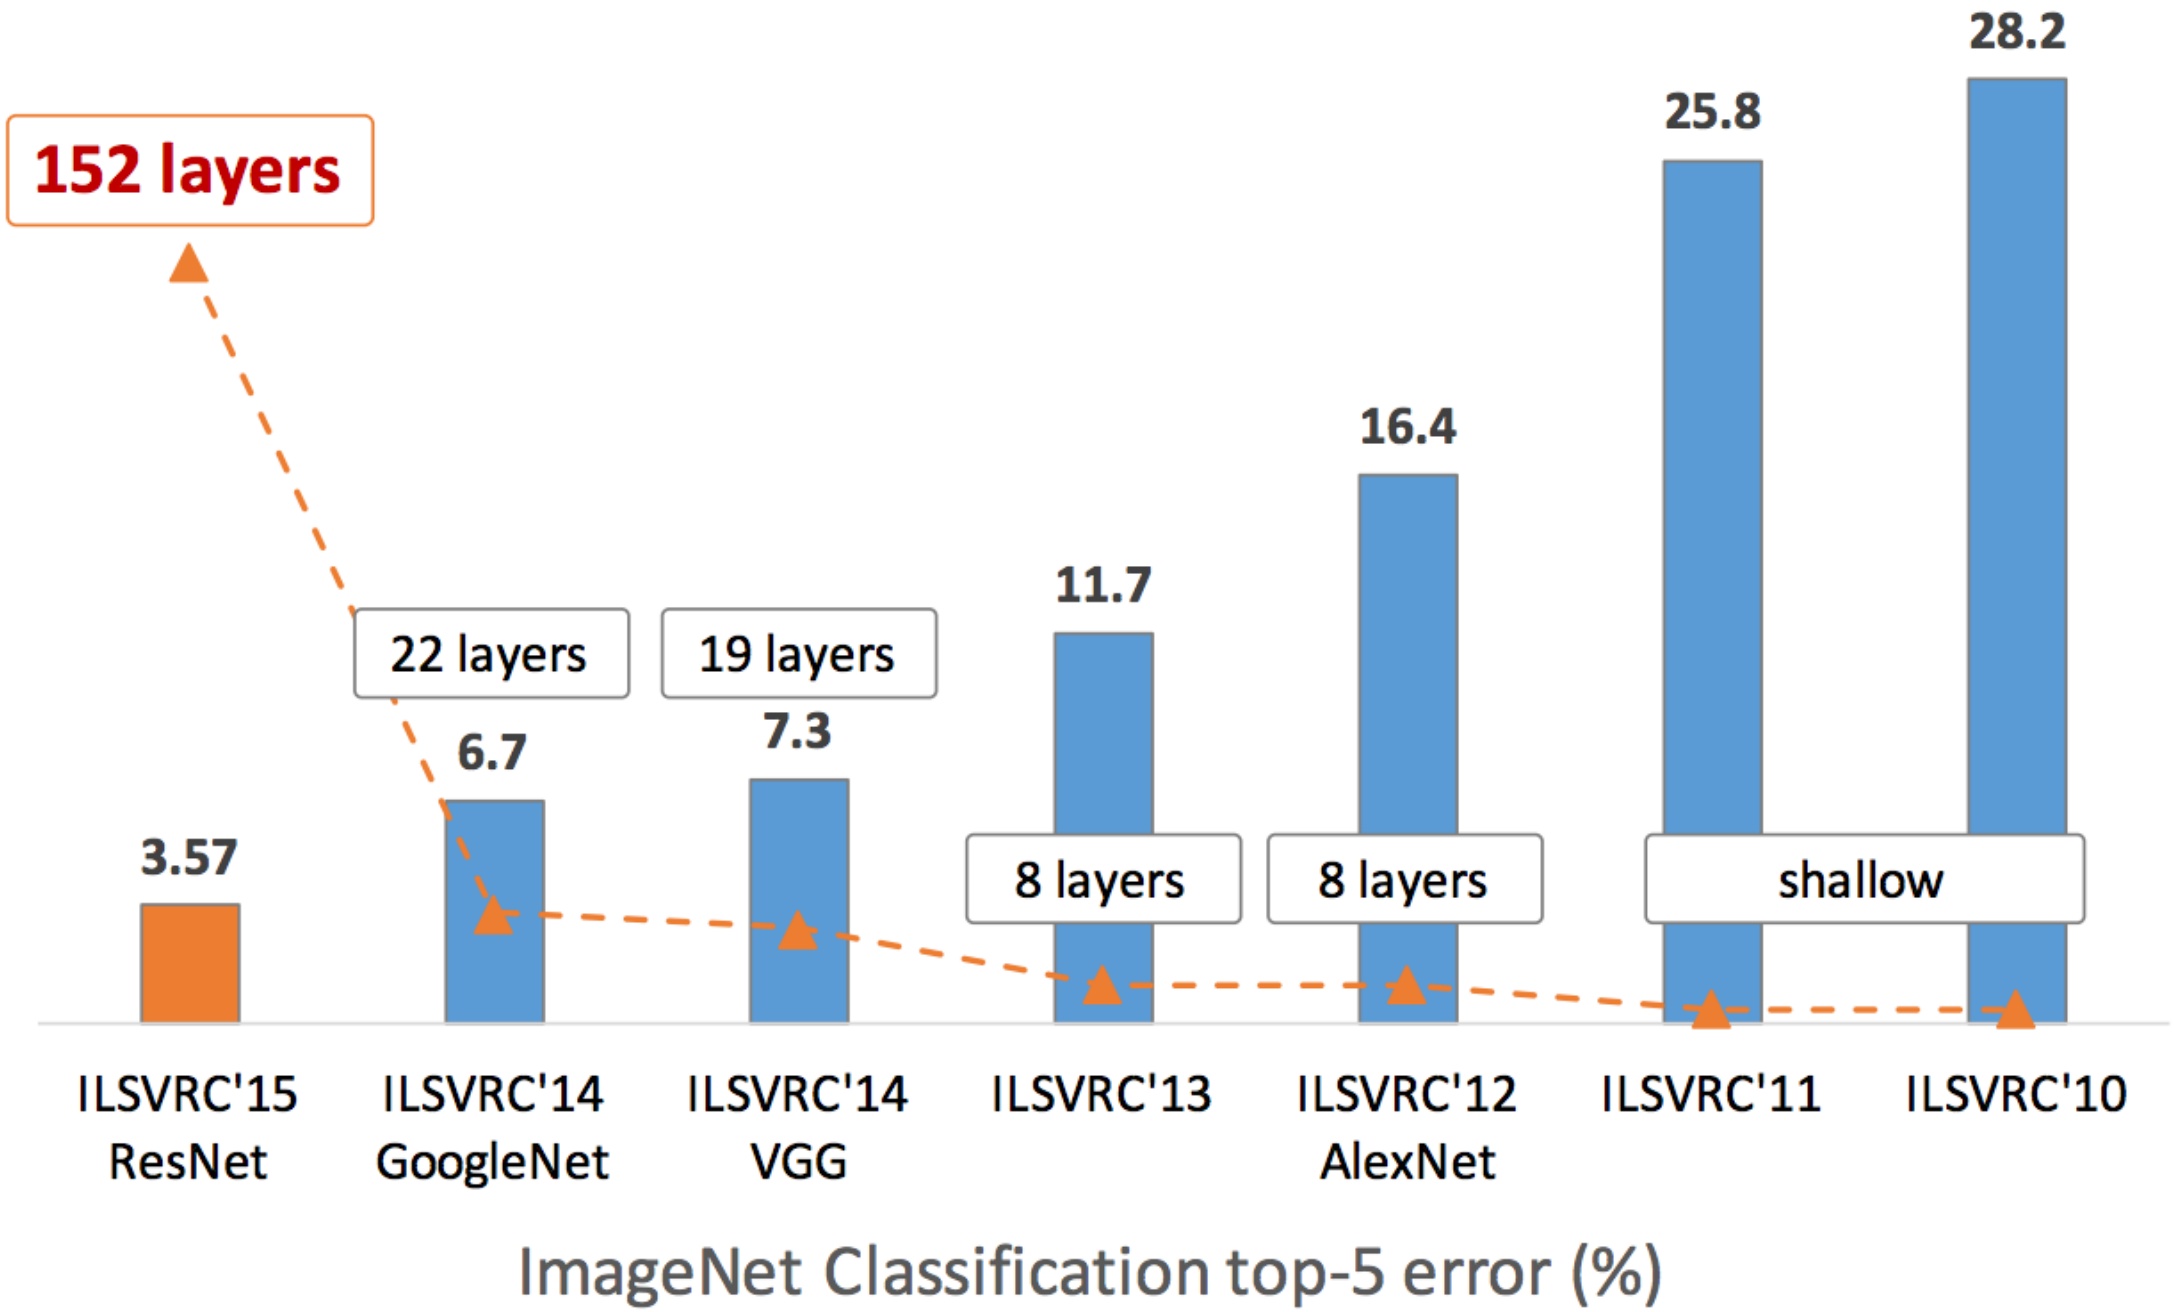

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". CVPR 2016.

Kaimin He (Facebook Research) et al. 2016

Improving on Machine Learning algorithms: Introducing Deep Residual Neural Net

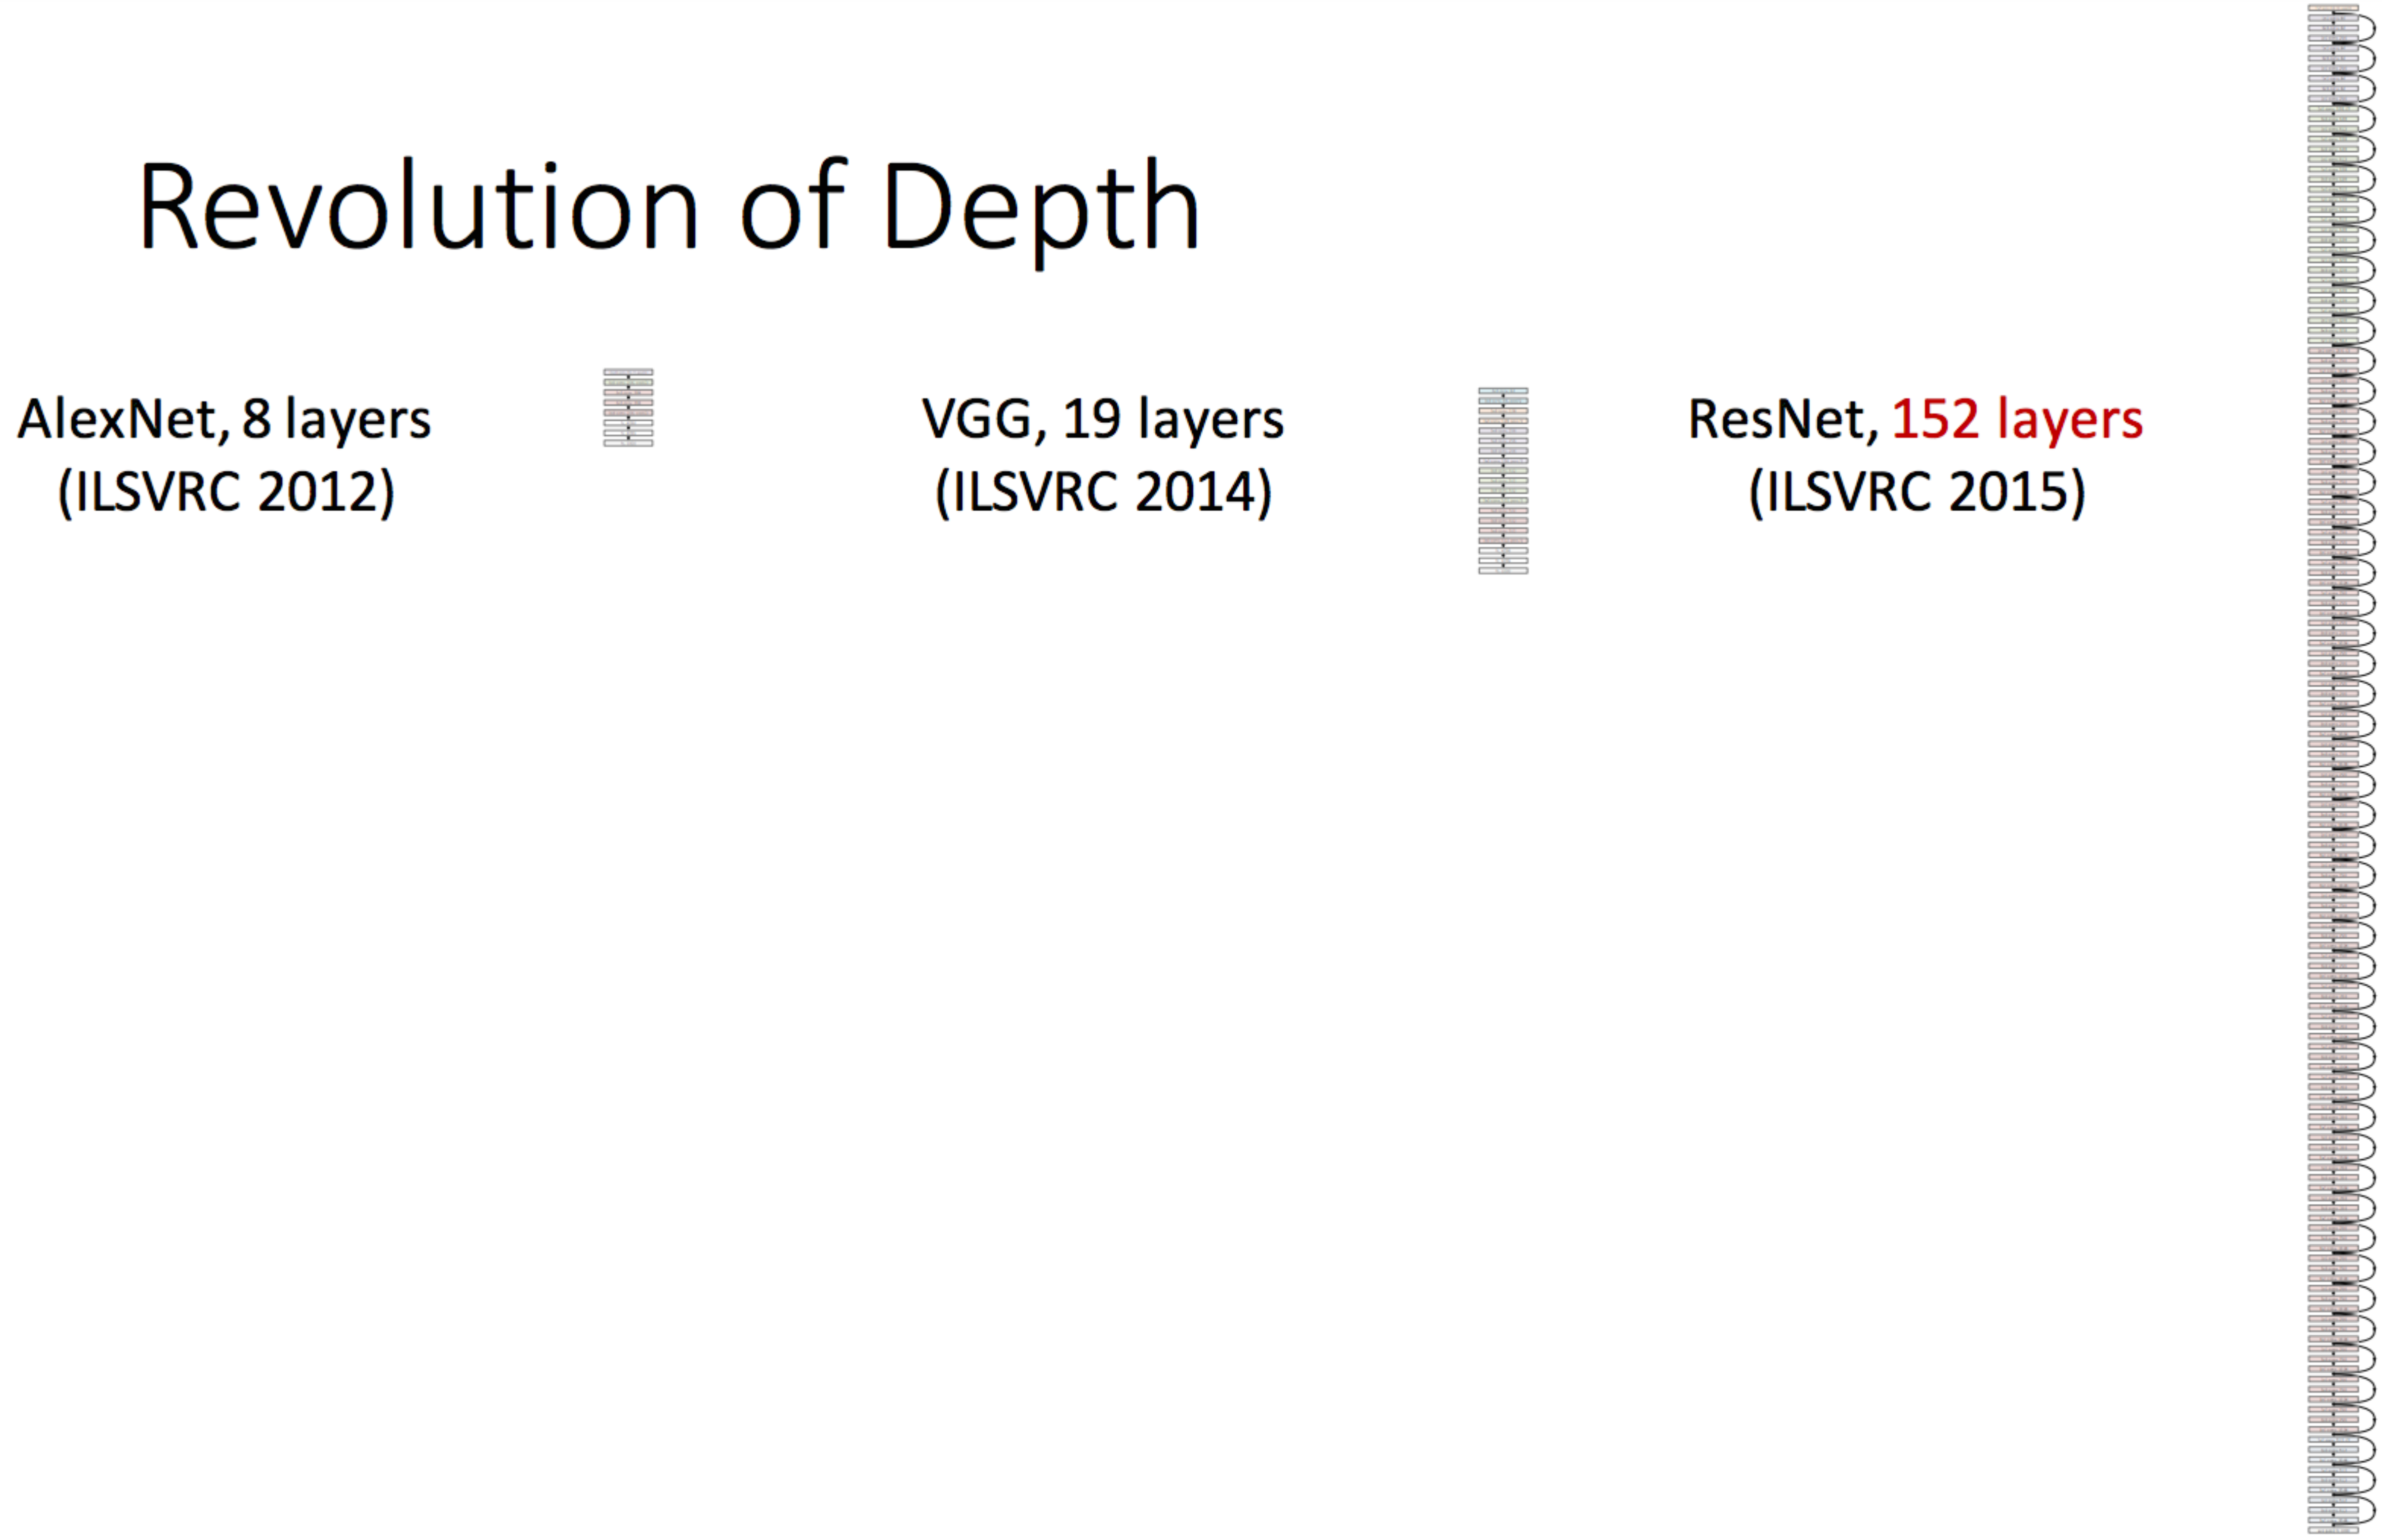

Kaimin He (Facebook Research) et al. 2016

Improving on Machine Learning algorithms: Introducing Deep Residual Neural Net

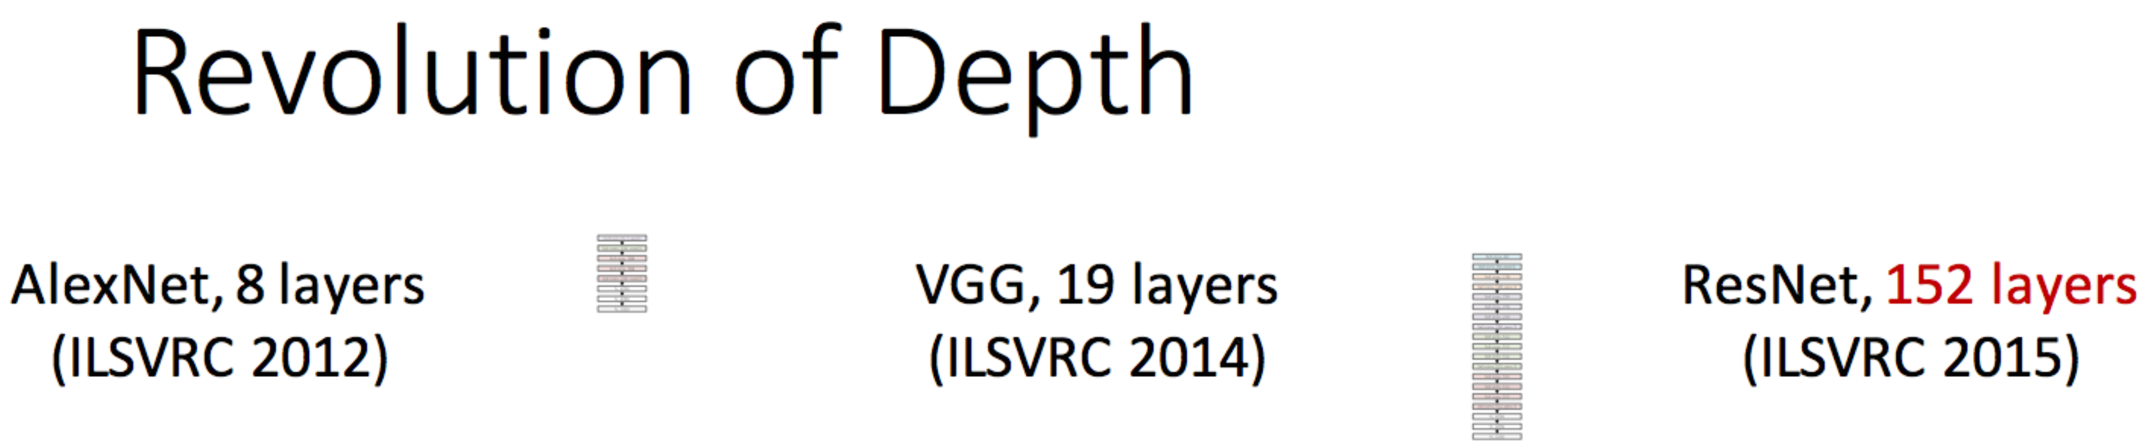

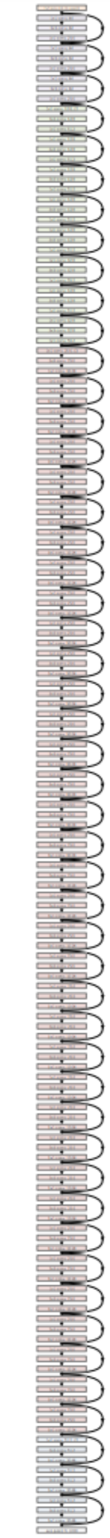

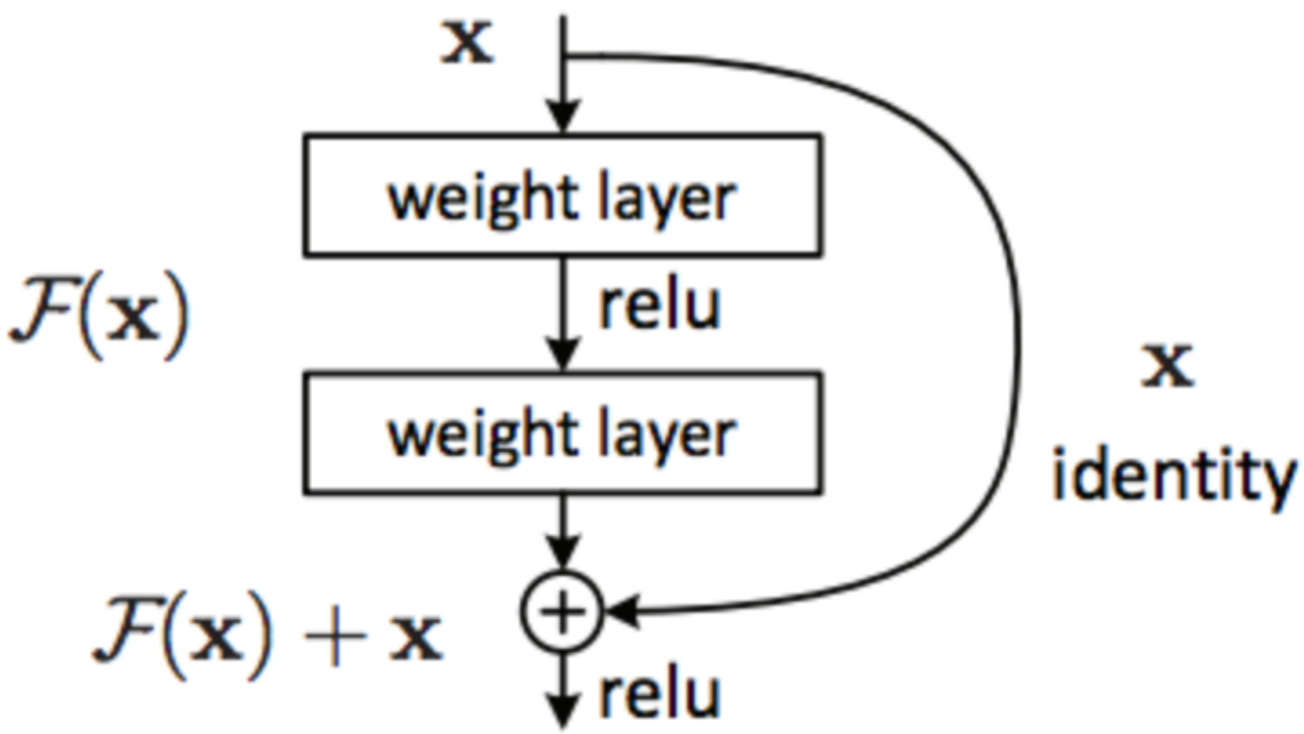

Try to fit for F(x) instead,
desired mapping:

$$H(x) = F(x) + x$$

Kaimin He (Facebook Research) et al. 2016

Improving on Machine Learning algorithms: Introducing Deep Residual Neural Net

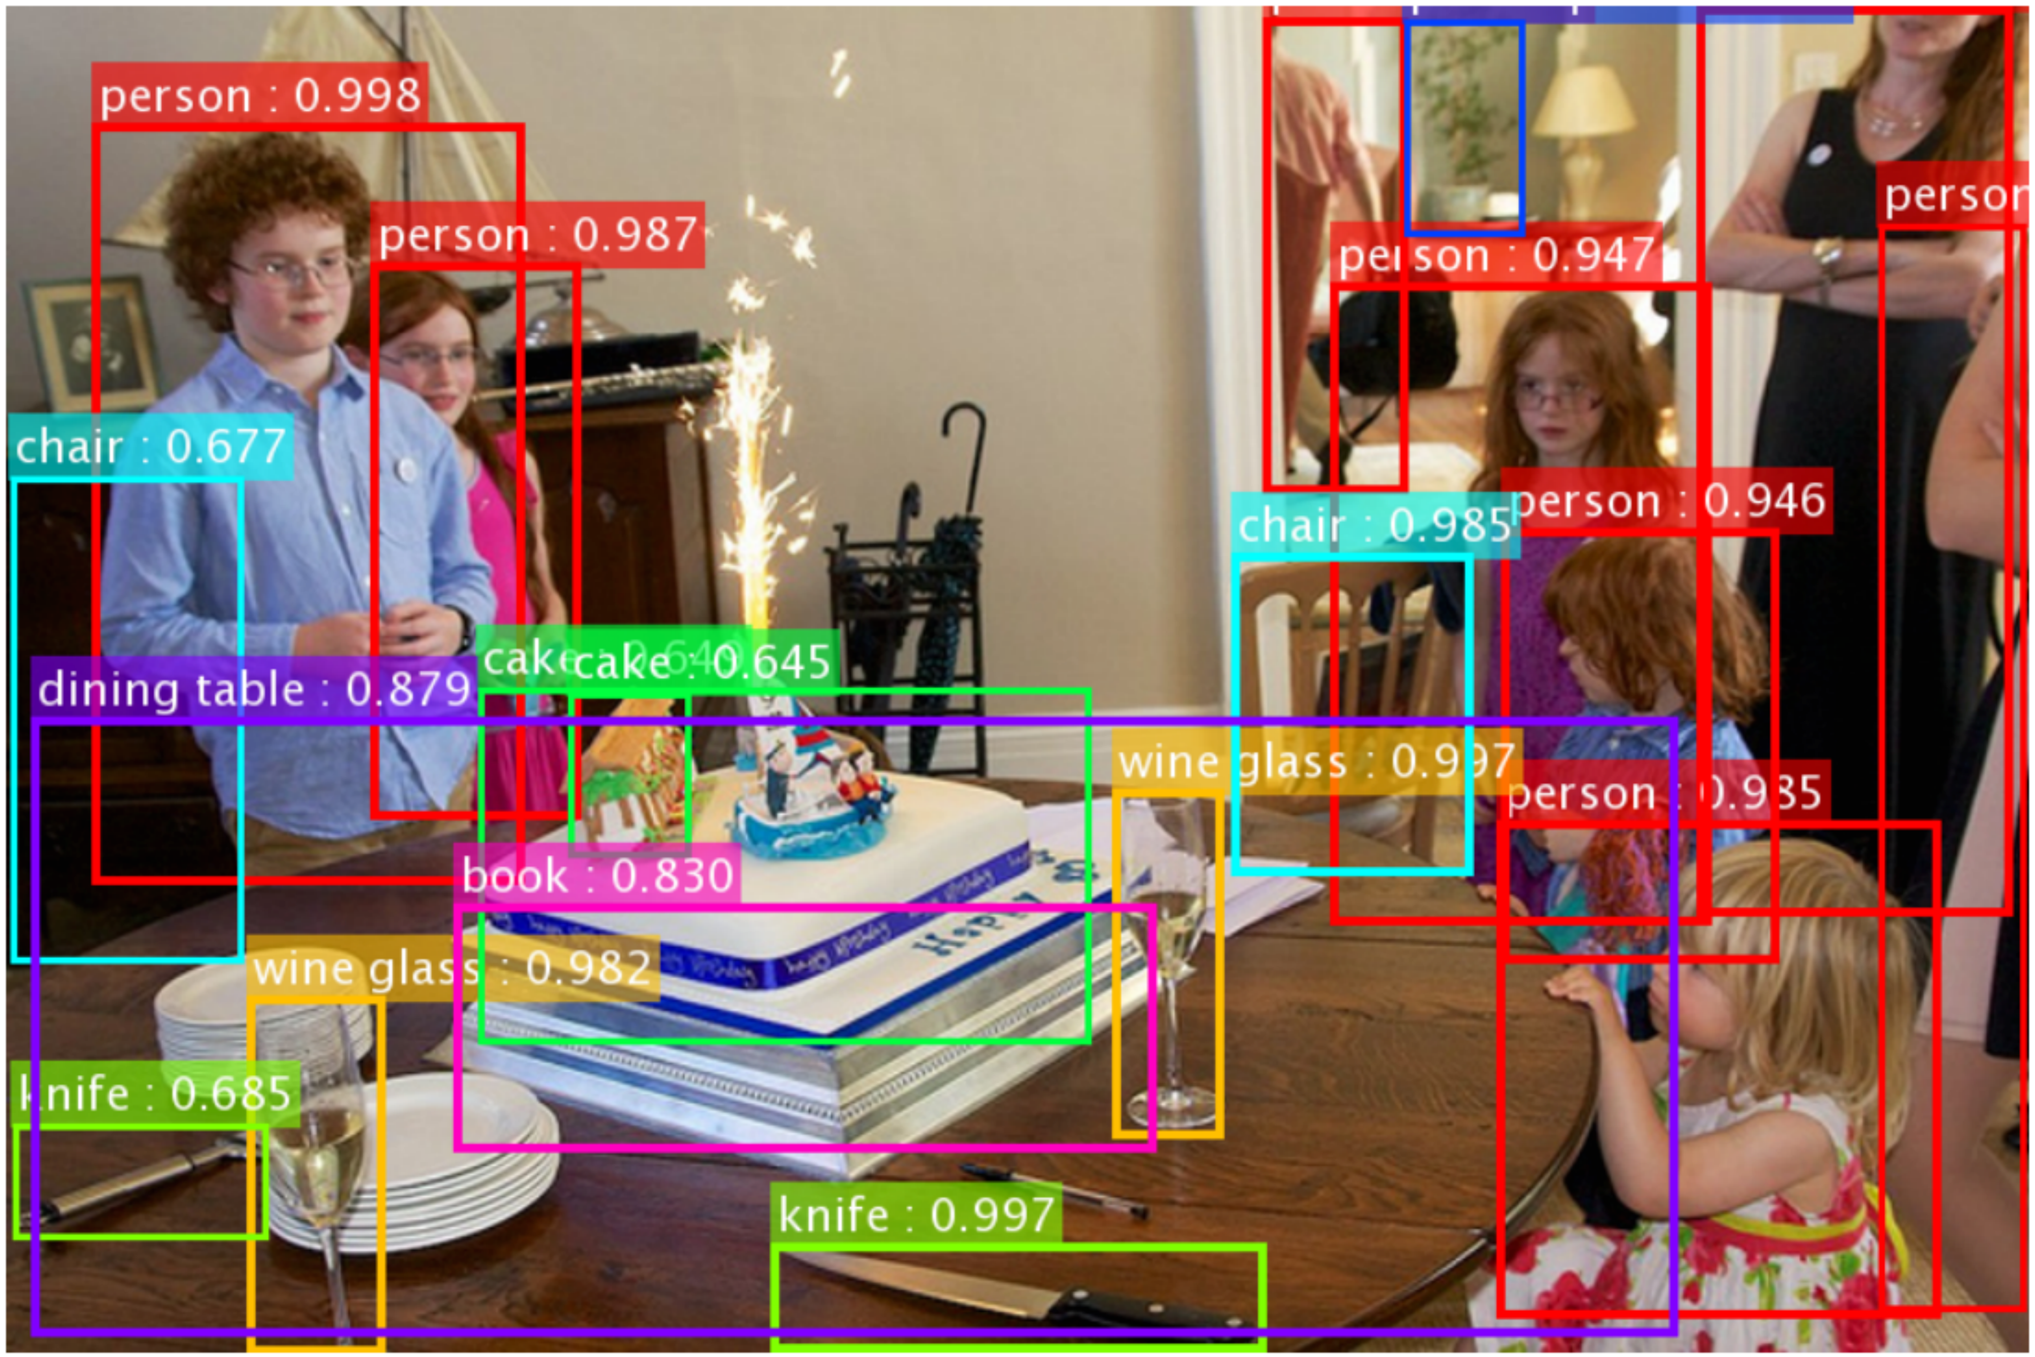

ResNet's object detection result on Common Object in Context

Kaimin He (Facebook Research, now Microsoft Research Asia) et al. 2016

Where do we go from here?

- Better Prediction possible? Improving the algorithms.
- **Can we interpret the model learnt in Machine Learning?**

Where is this extra information coming from ?

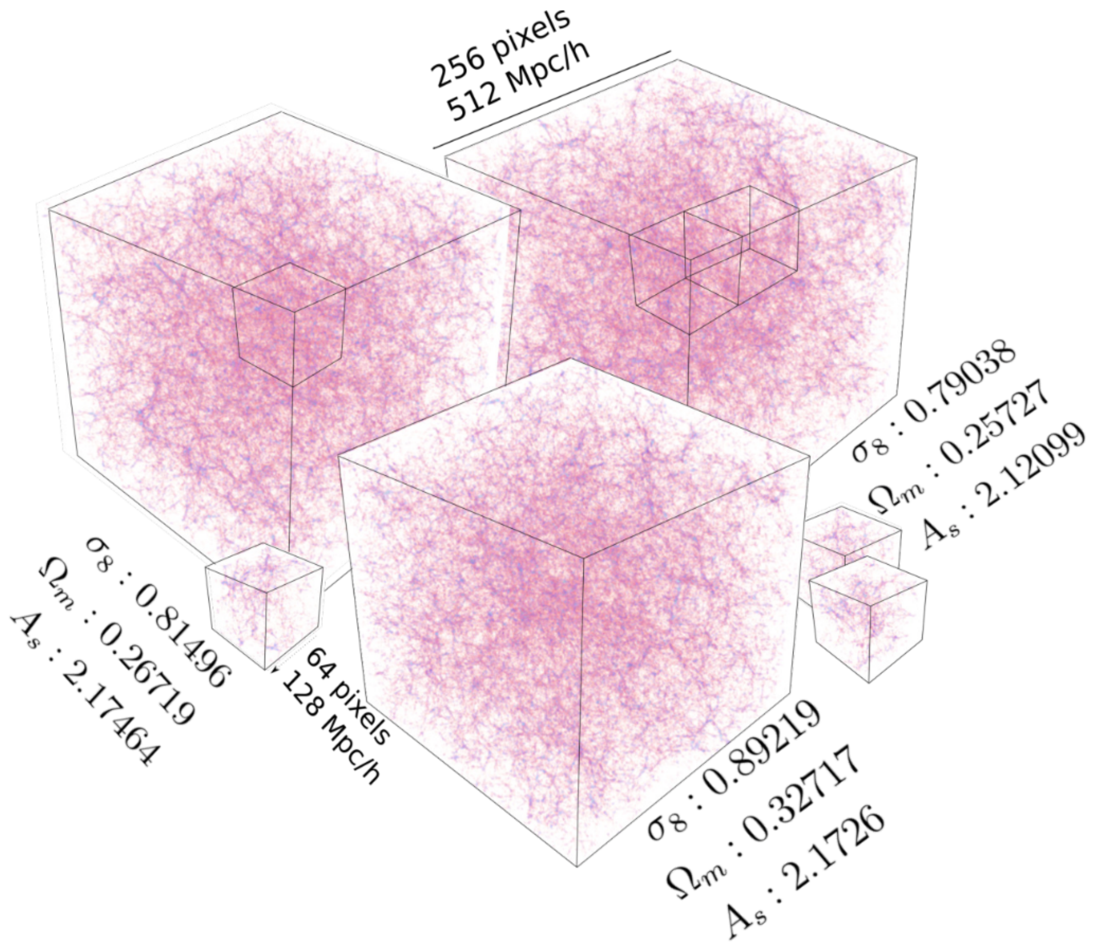

Figure 1. Dark matter distribution in three cubes produced using different sets of parameters. Each cube is divided into small sub-cubes for training and prediction. Note that although cubes in this figure are produced using very different cosmological parameters in our constrained sampled set, the effect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

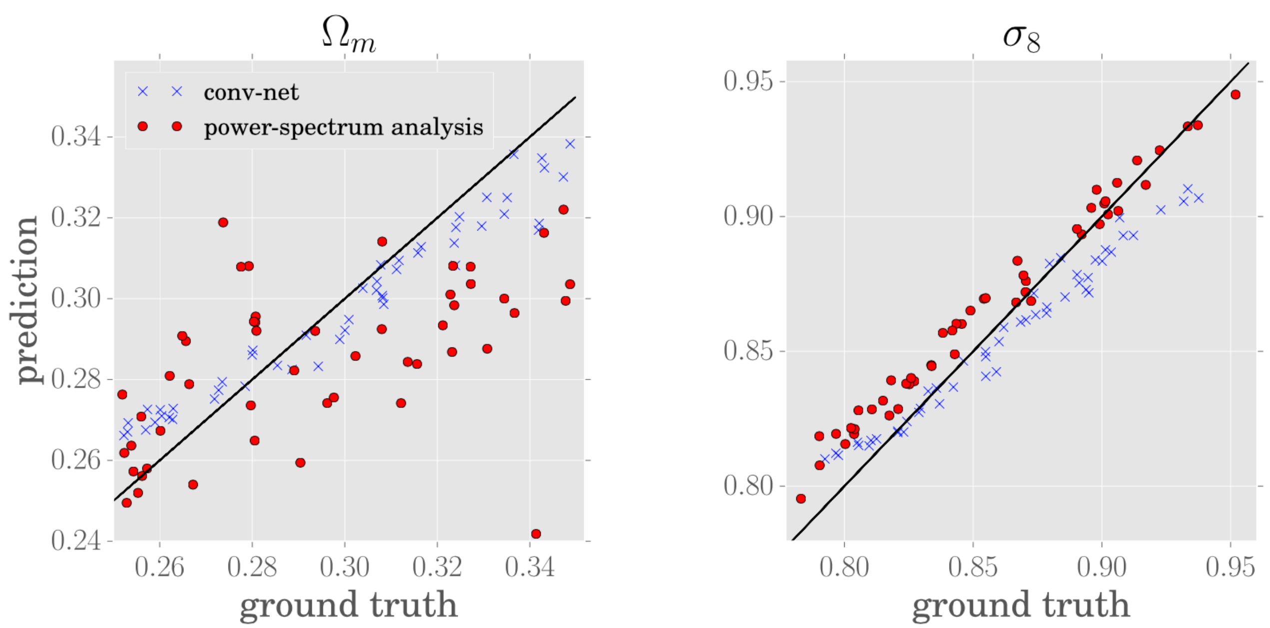

Figure 2. Prediction and ground truth of Ω_m and σ_8 using 3D conv-net and analysis of the power-spectrum on 50 test cube instances.

Analytical physics (2nd order Lagrangian Perturbation Theory) vs Computer Simulation (N-body/complex simulations)

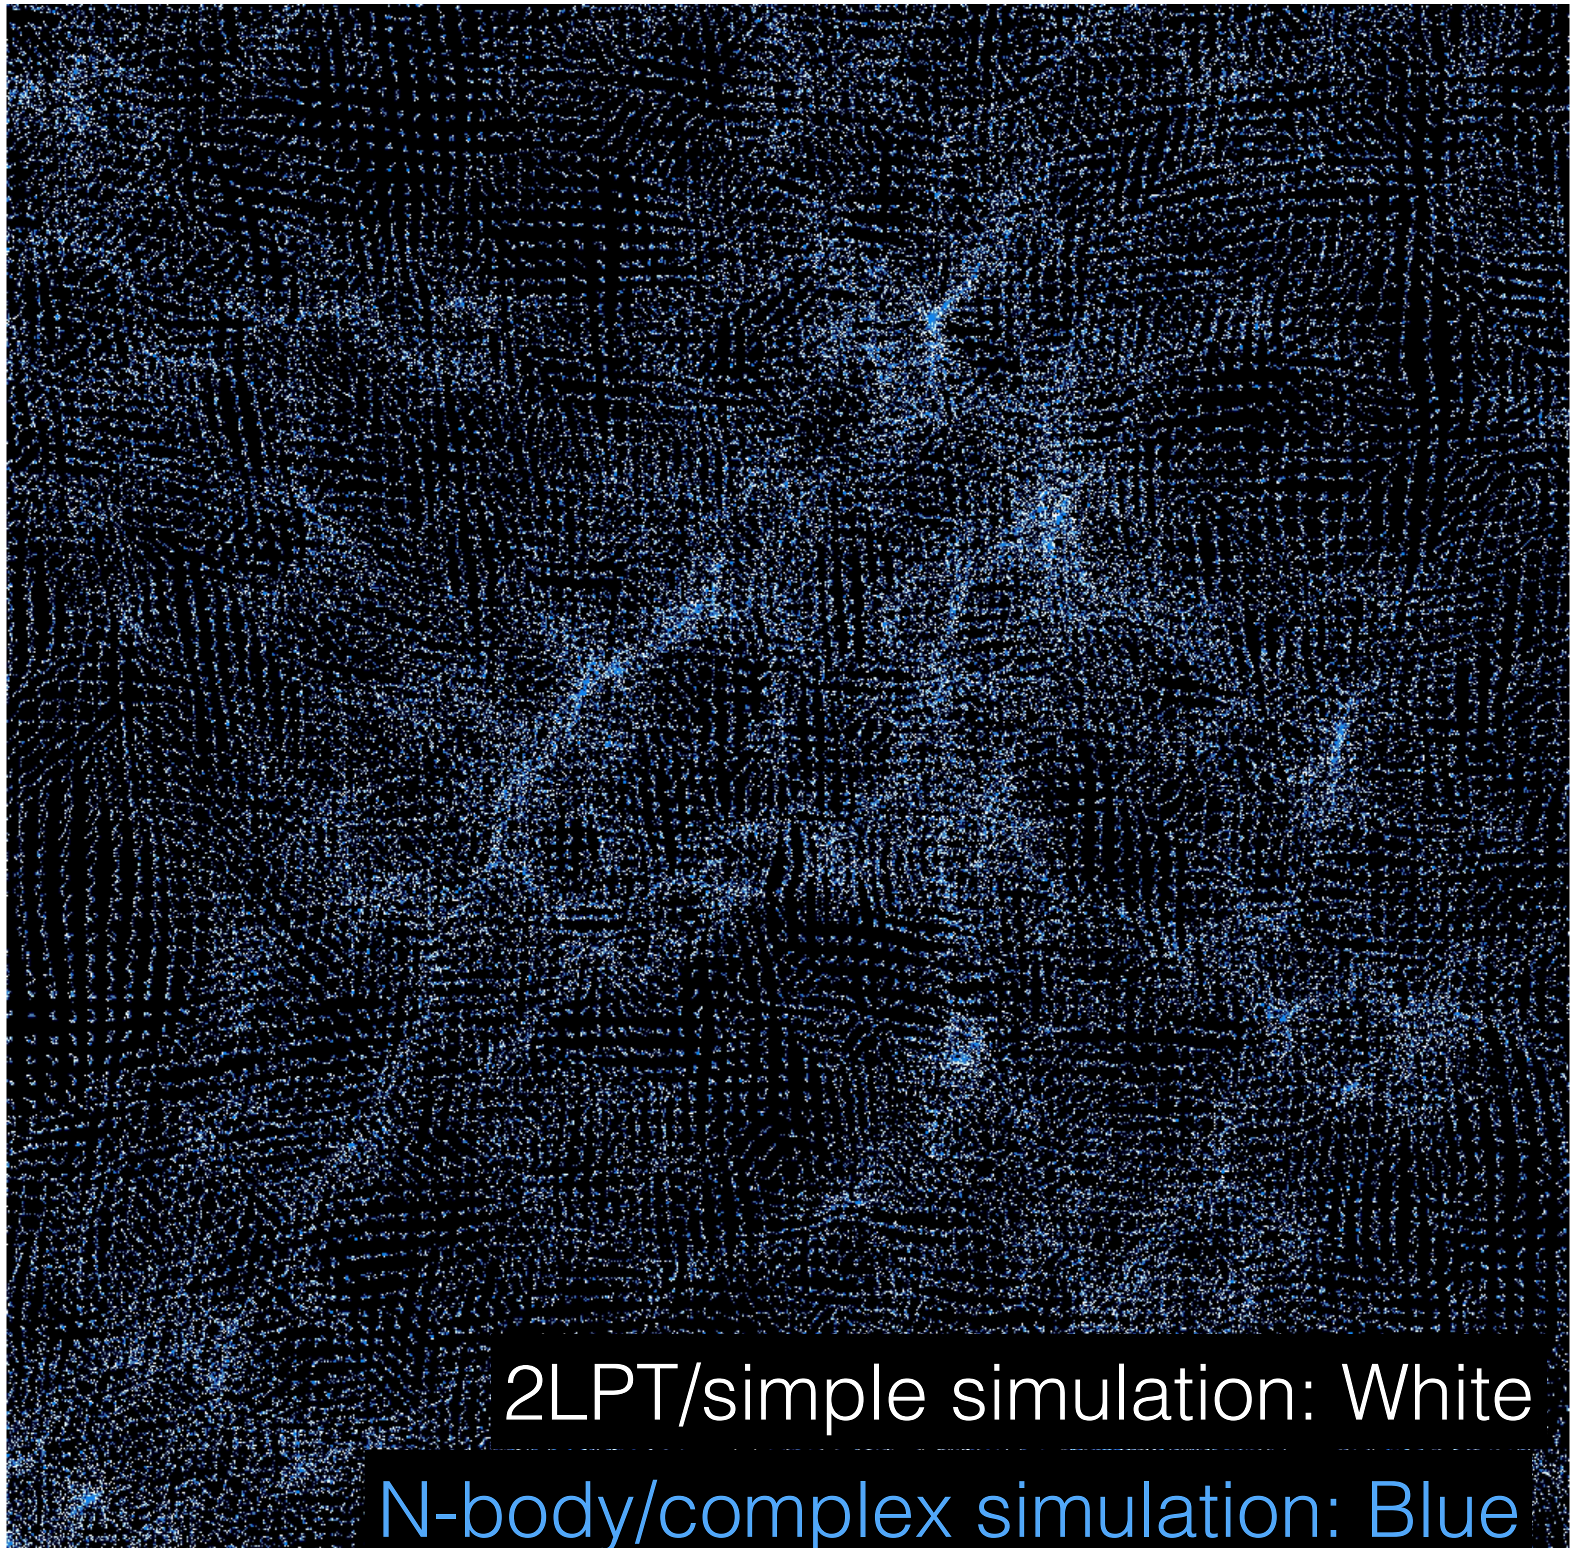

Movie Credit: Andrea Klein, Junier Oliver, Hy Trac

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

Training

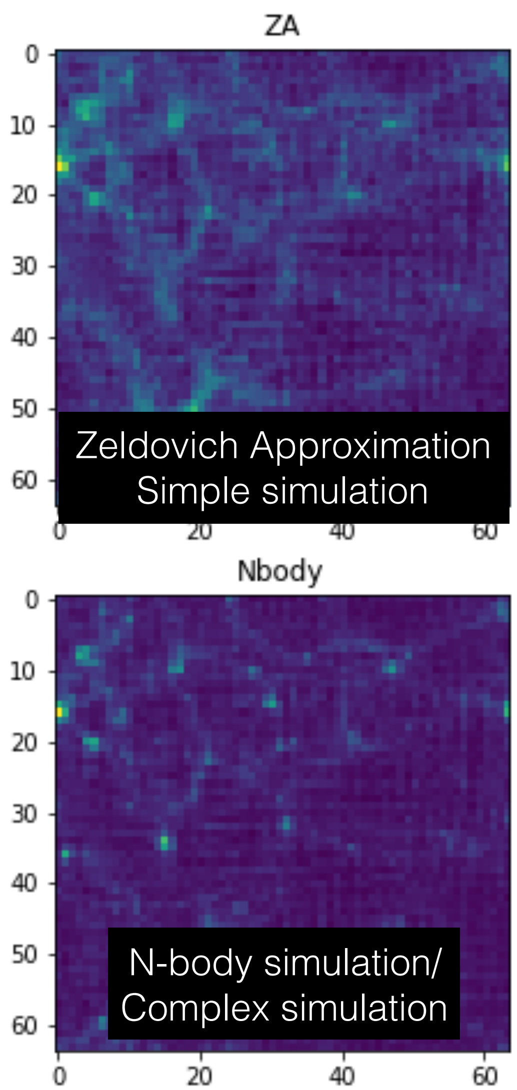

10,000 pairs of
[Simple,complex] simulations
For training

UNET
Slight variant to Residual NN

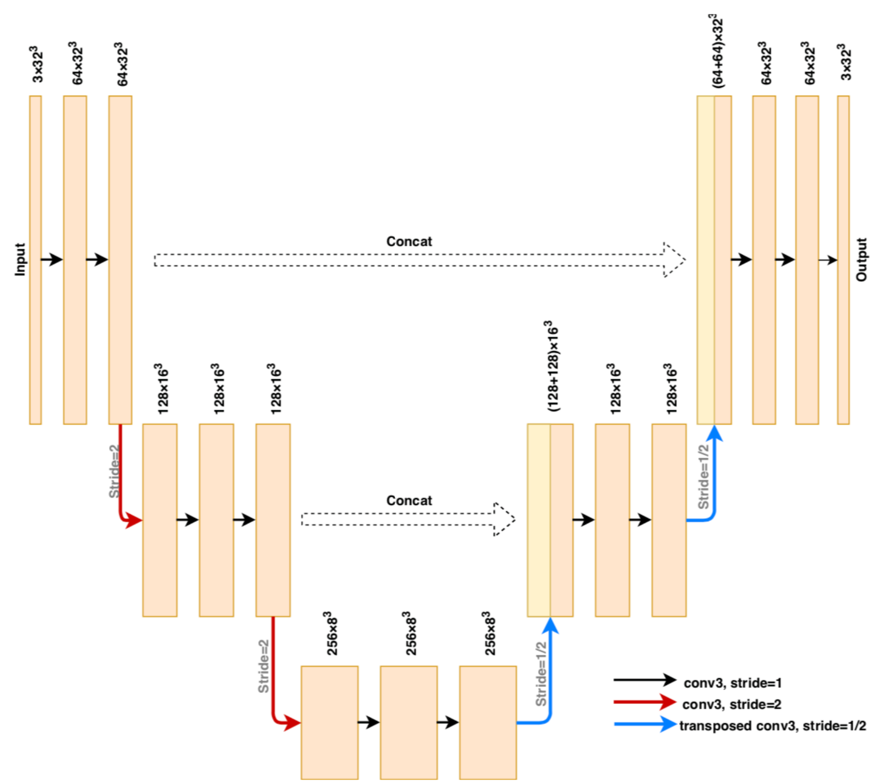

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input

Slight variant to Residual NN

Prediction

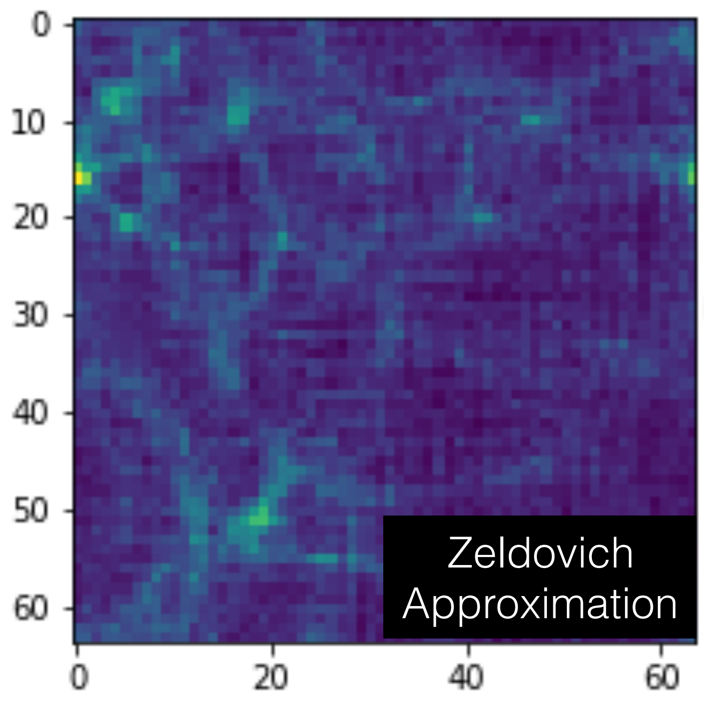

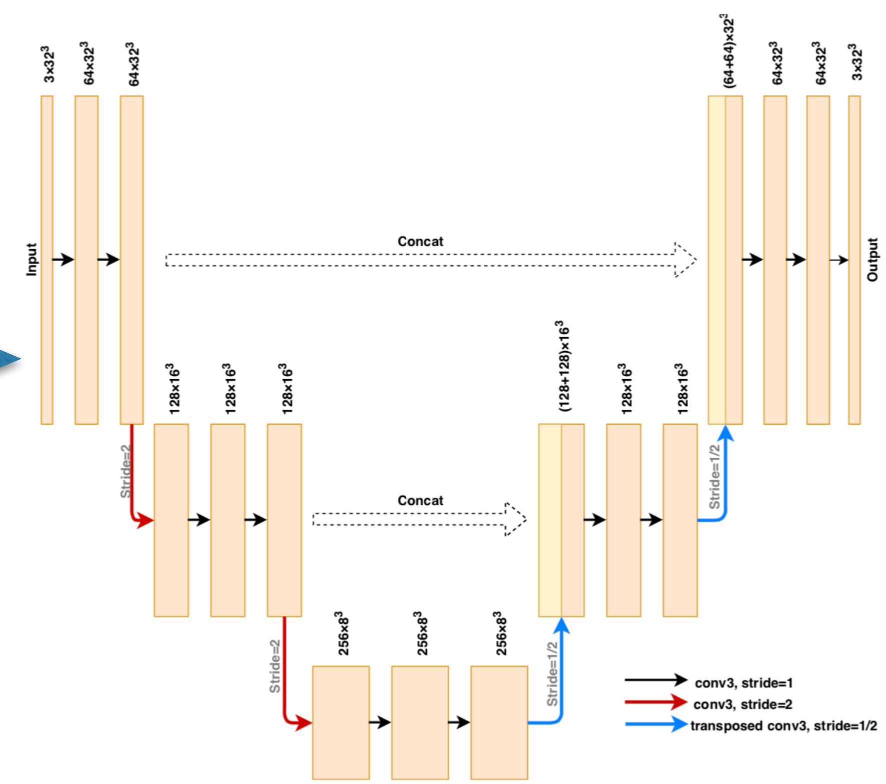

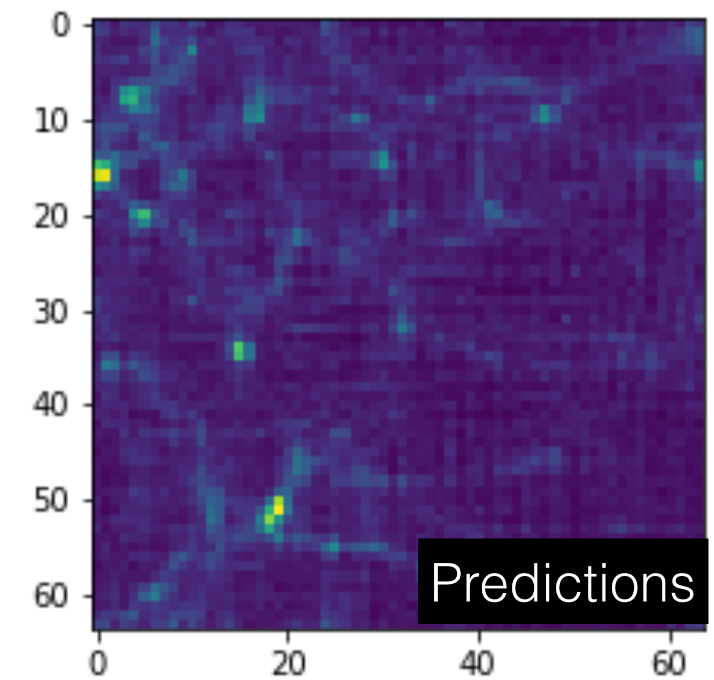

Simple simulation

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Density fields quick visual comparison

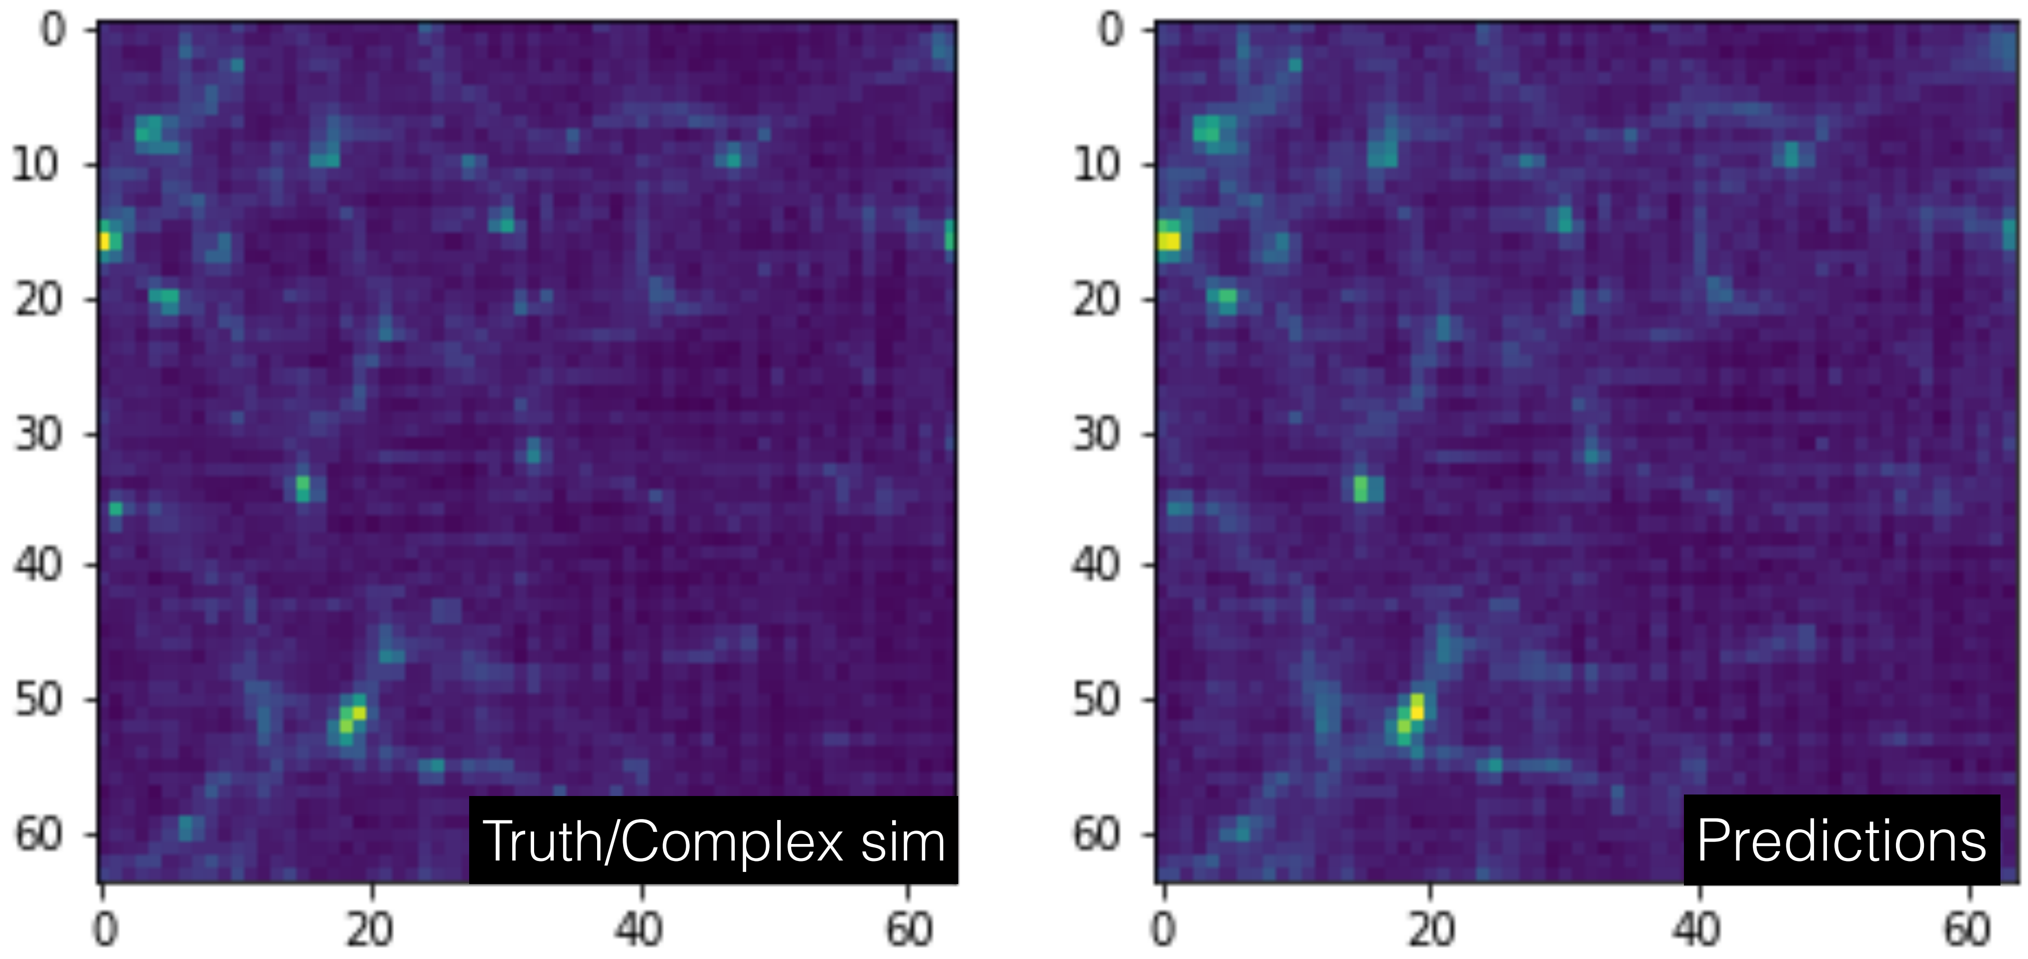

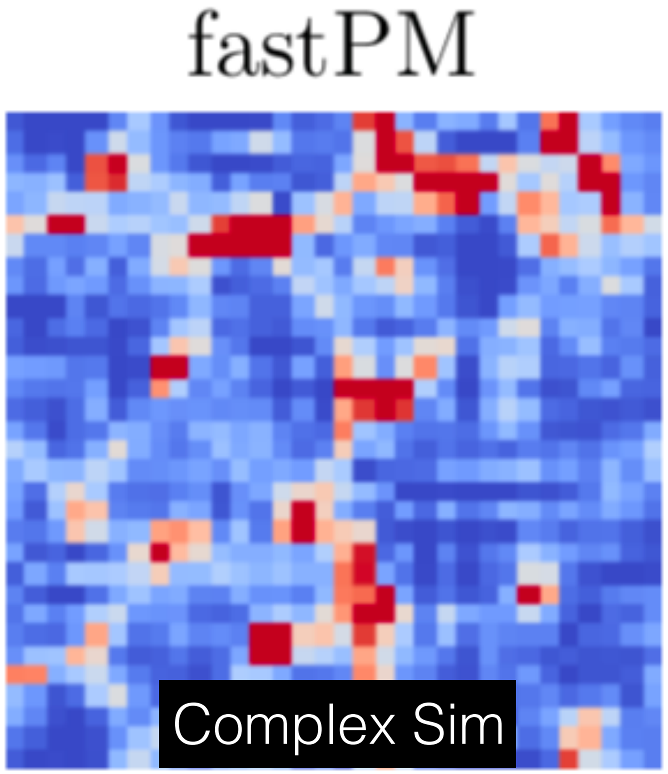

Let’s see how well the simple 2LPT would predict

Density field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

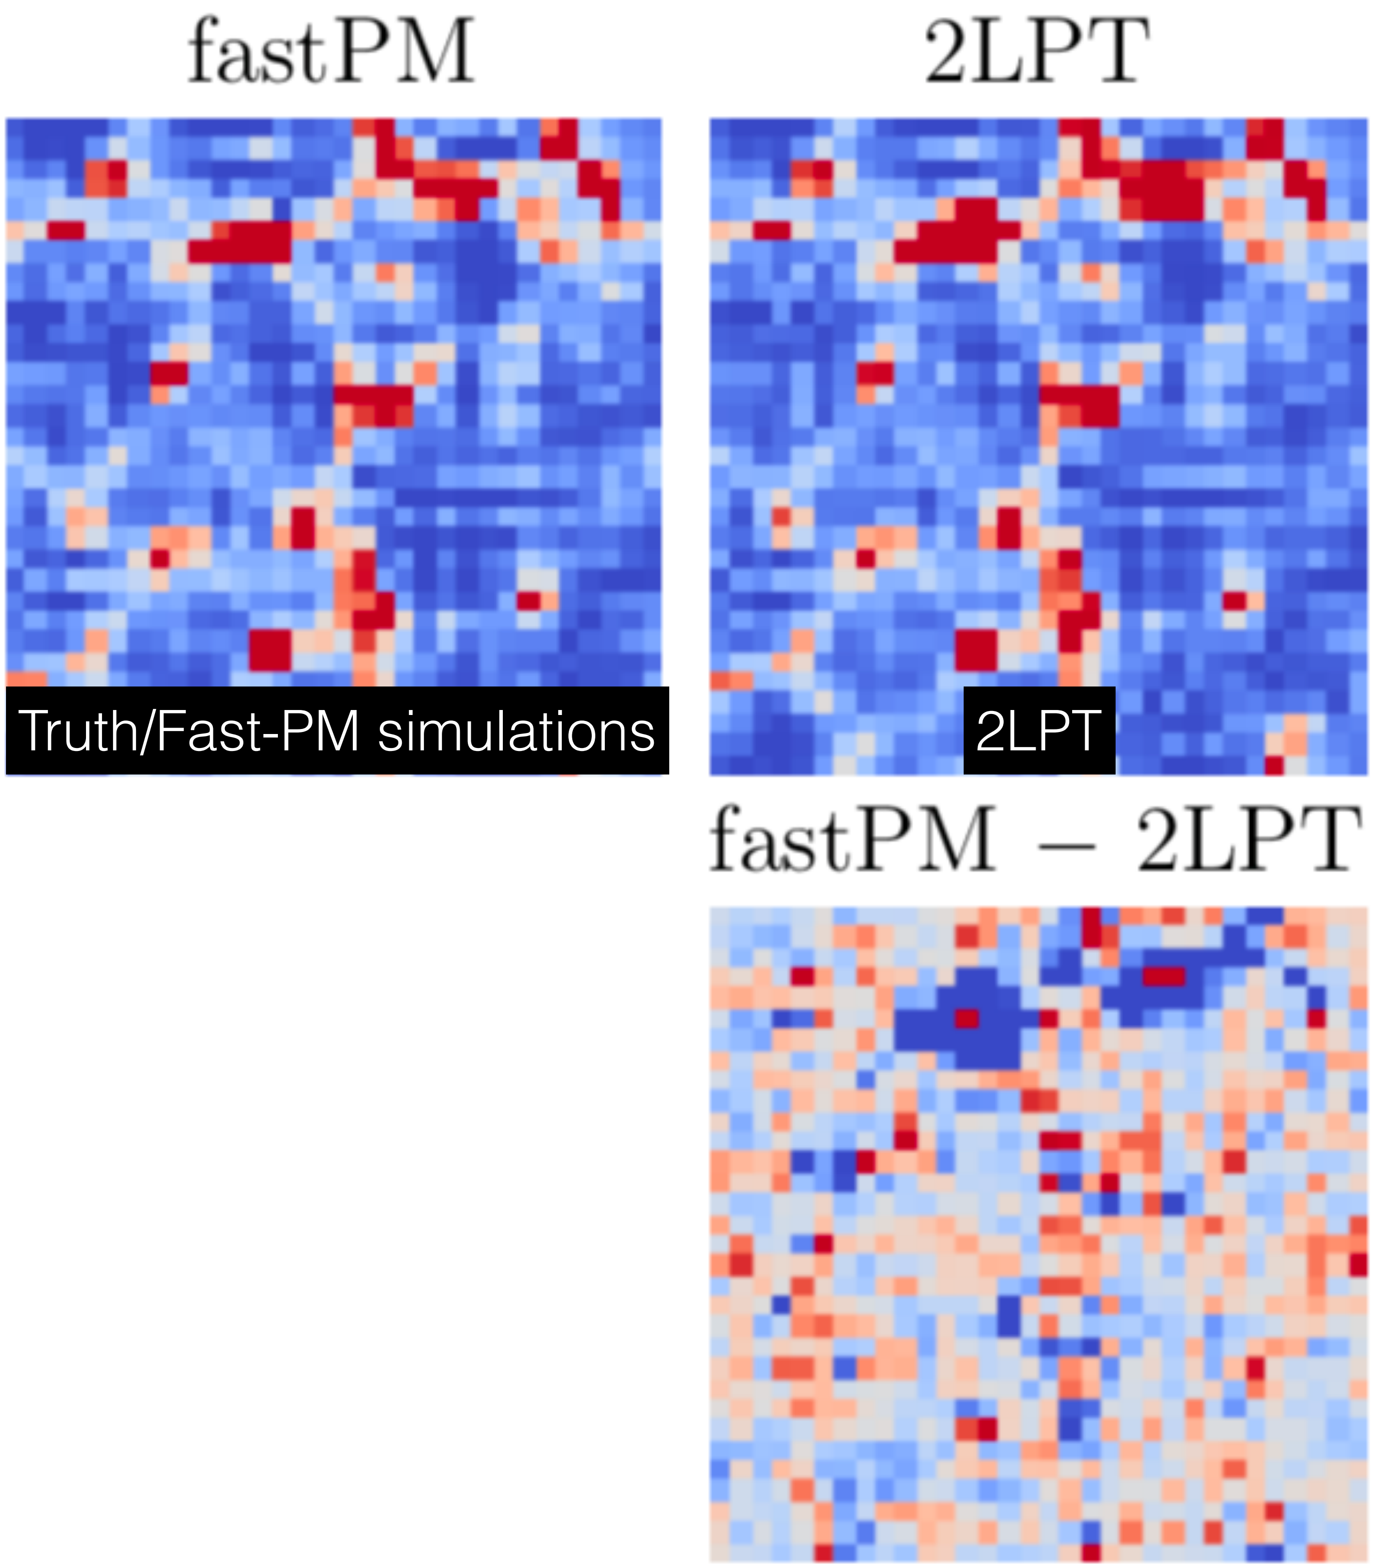

Density field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

fastPM

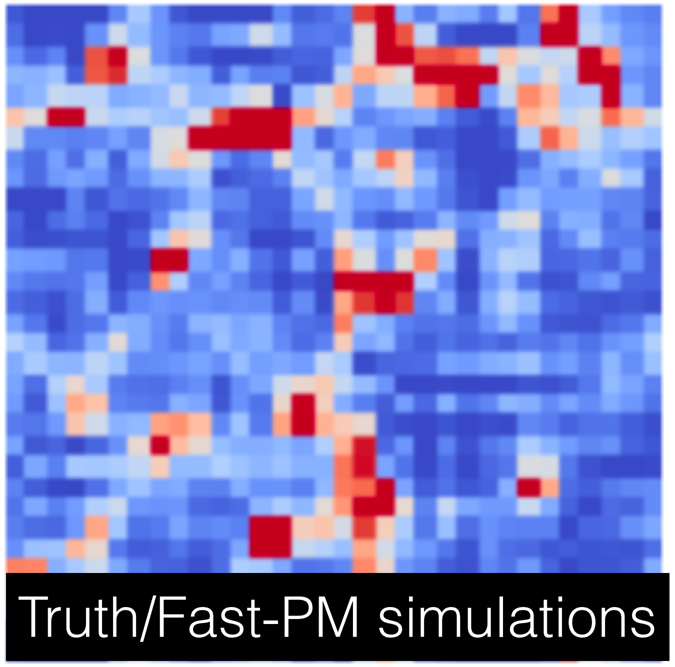

Now let’s compare the ML predictions with the truth!

Density field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

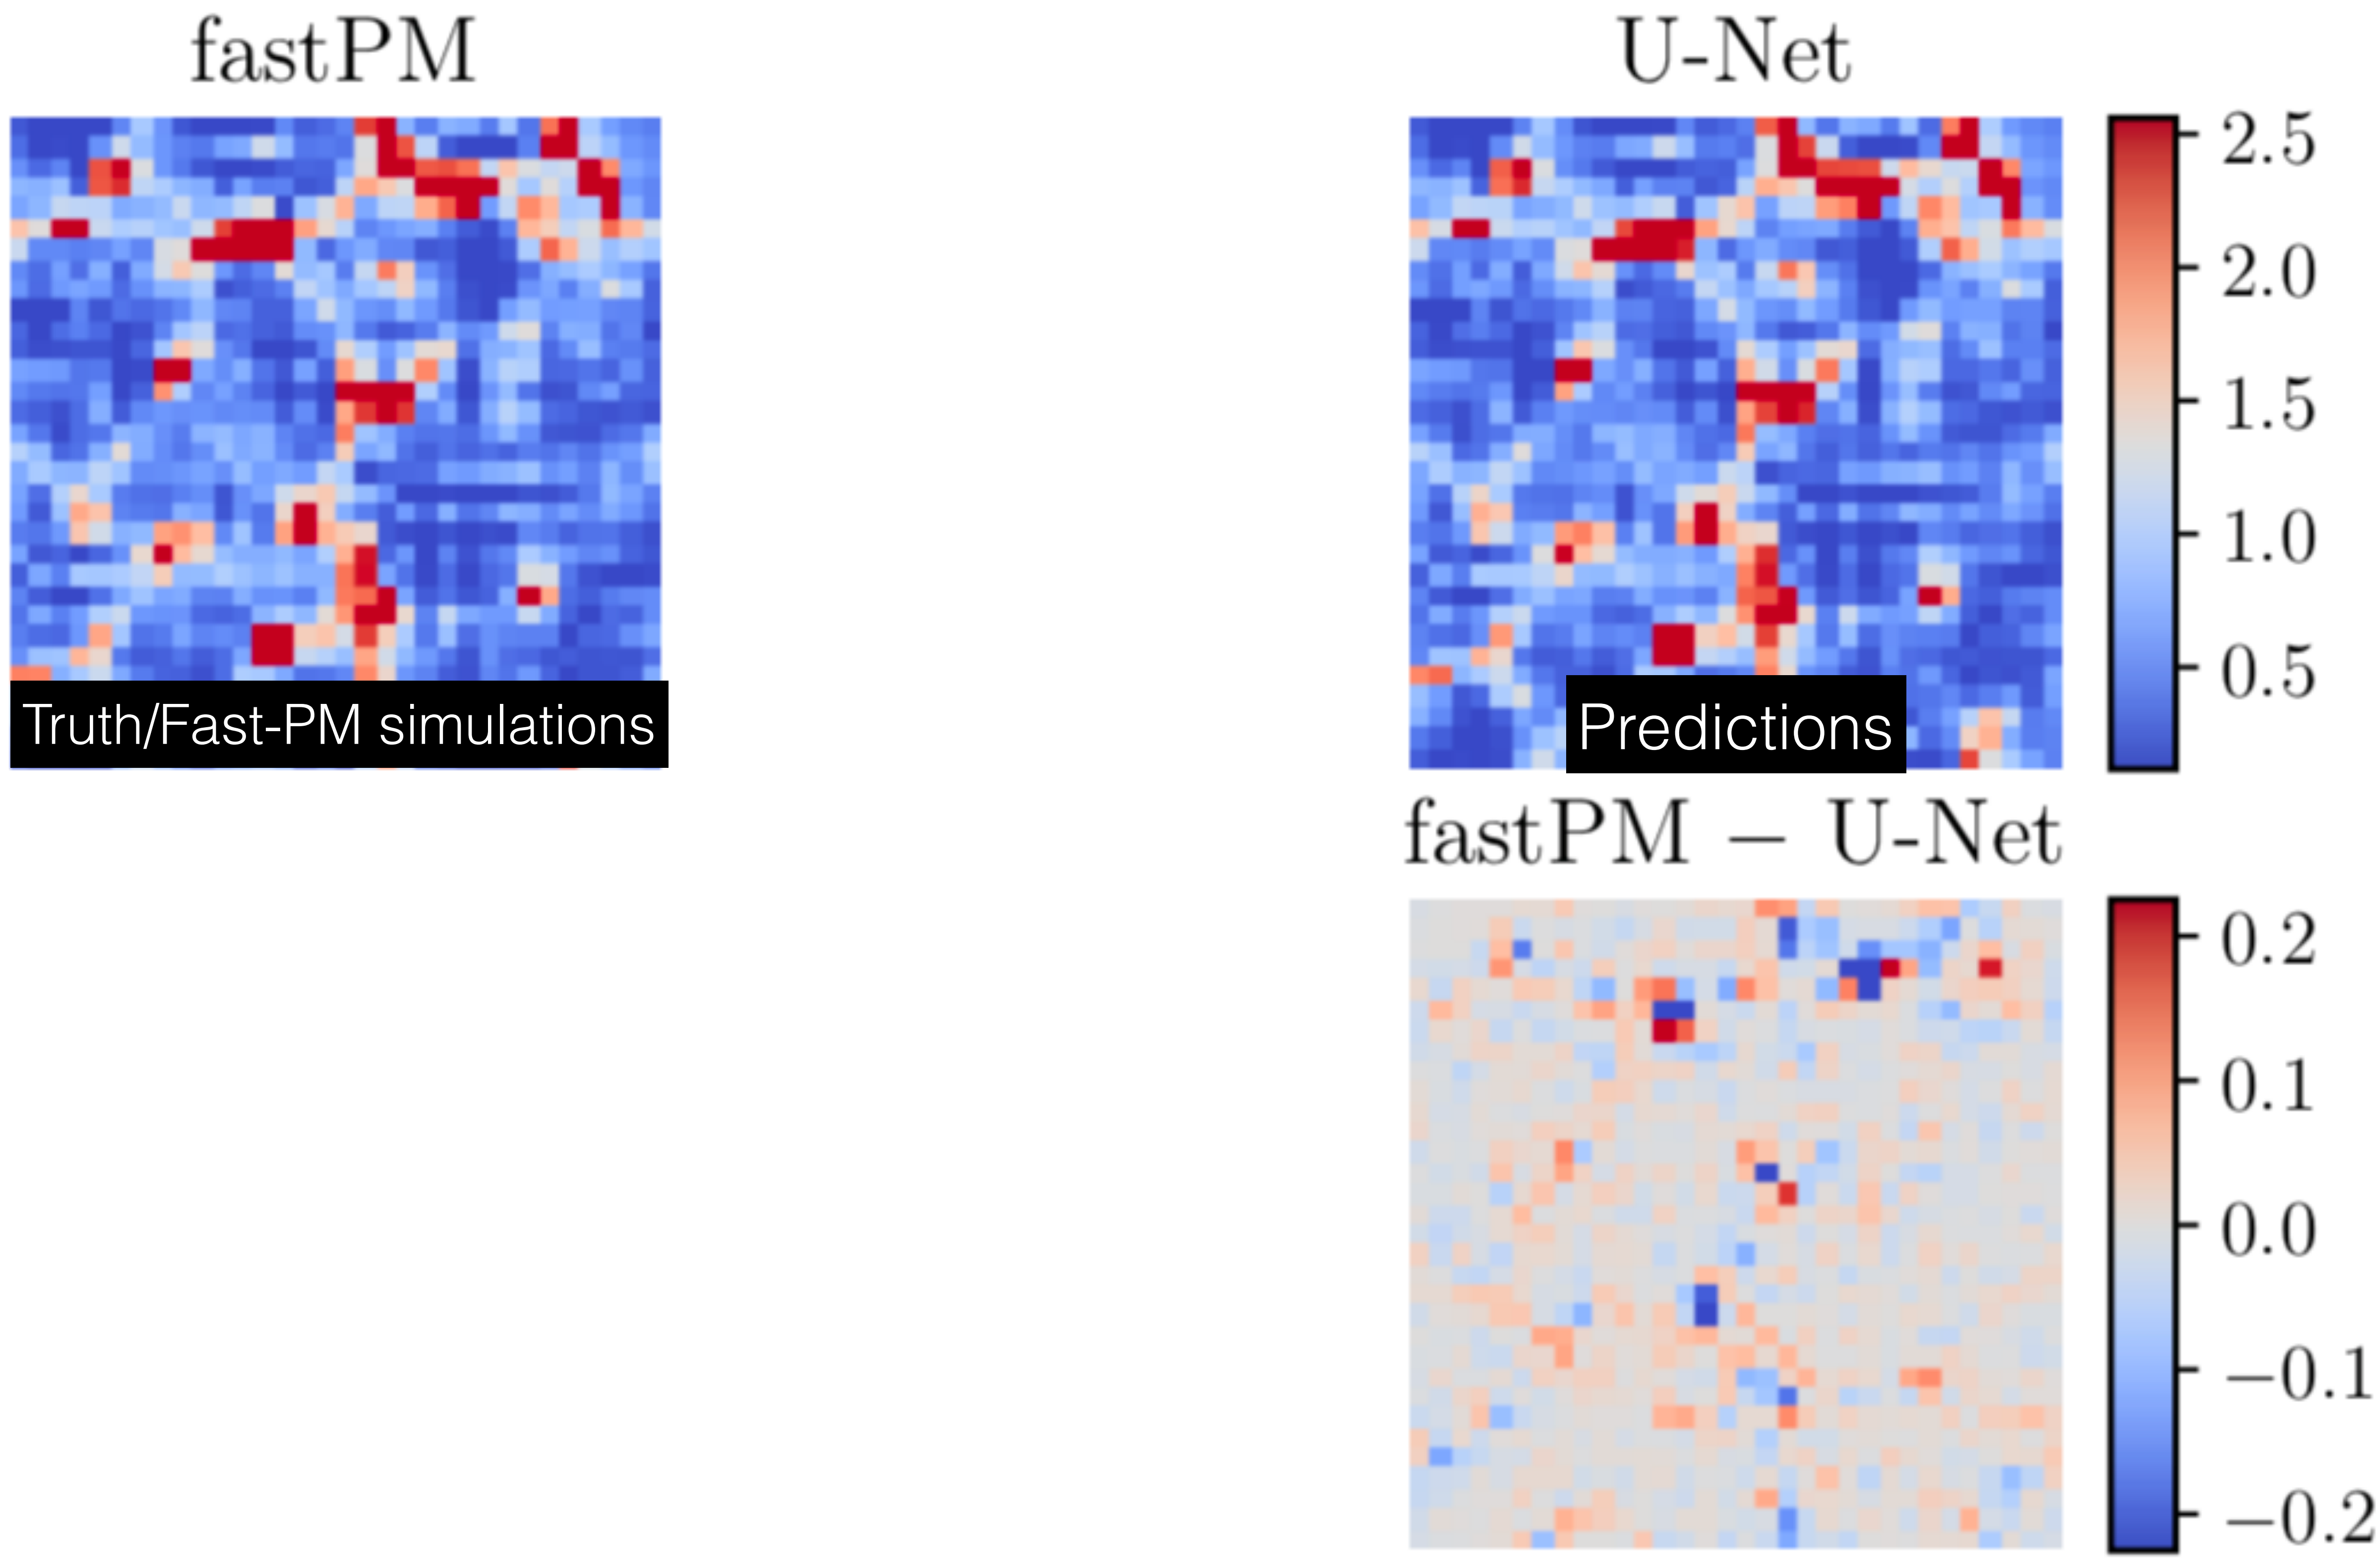

Density field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

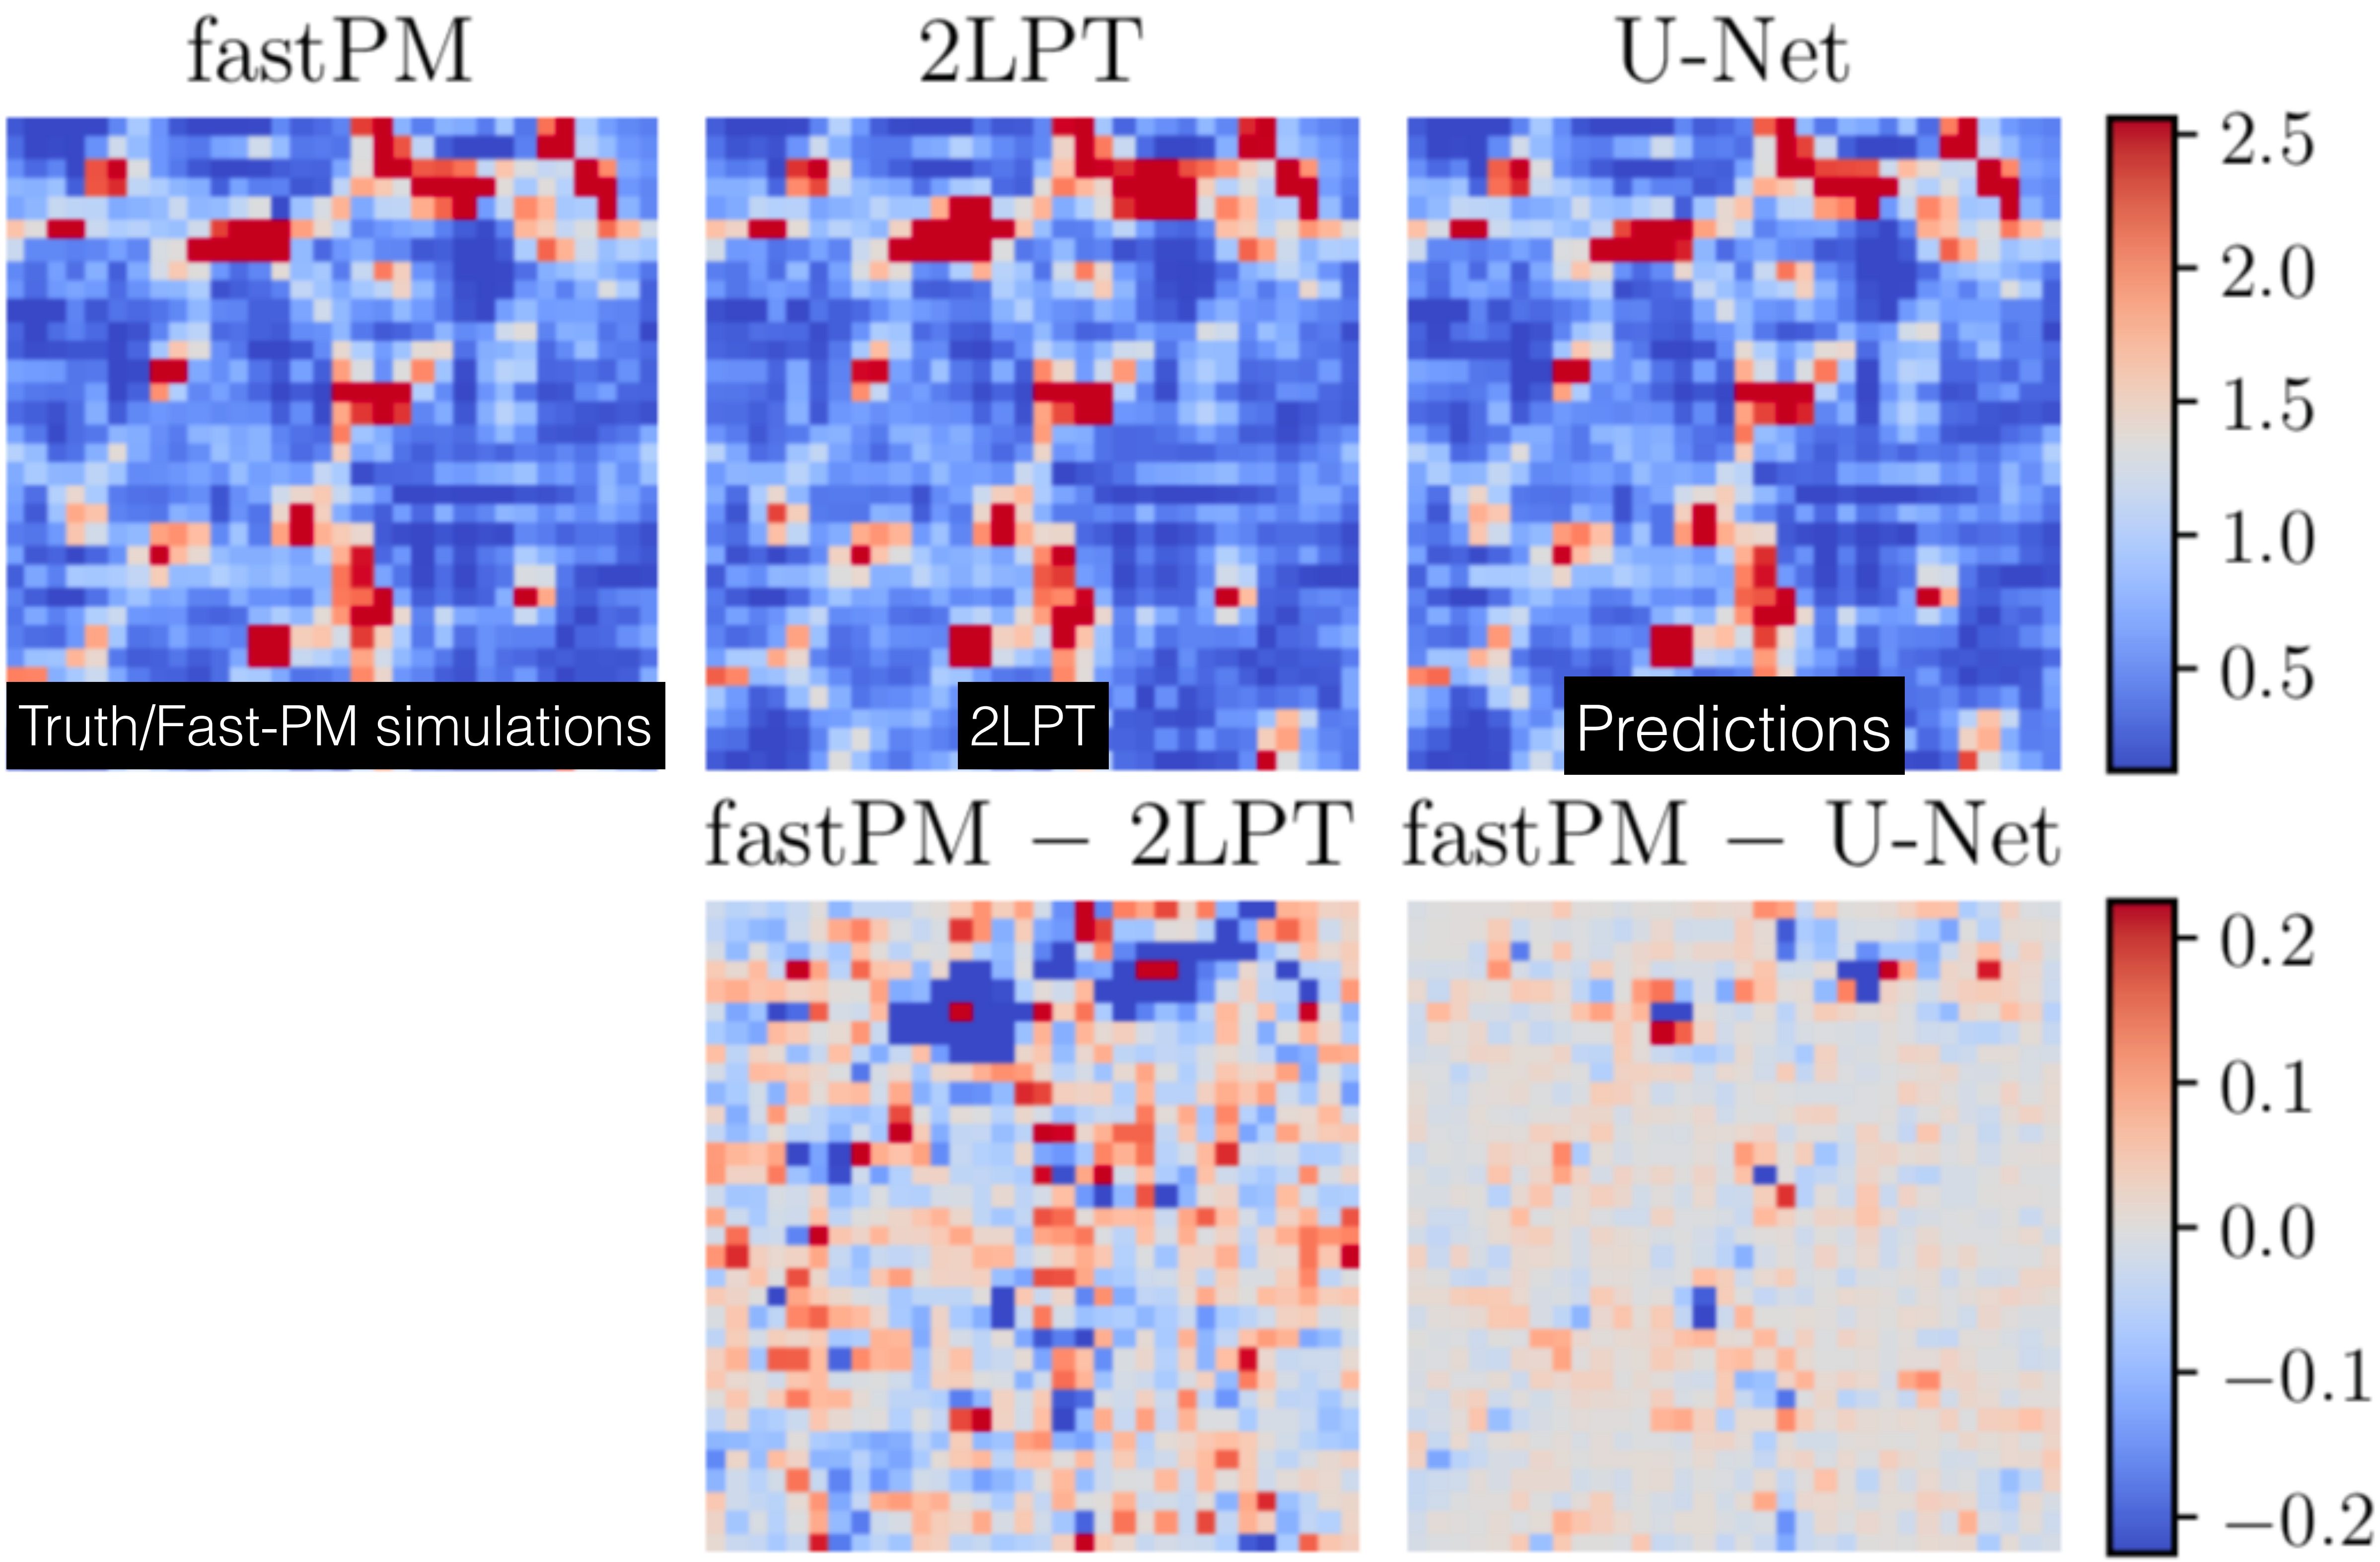

Density field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

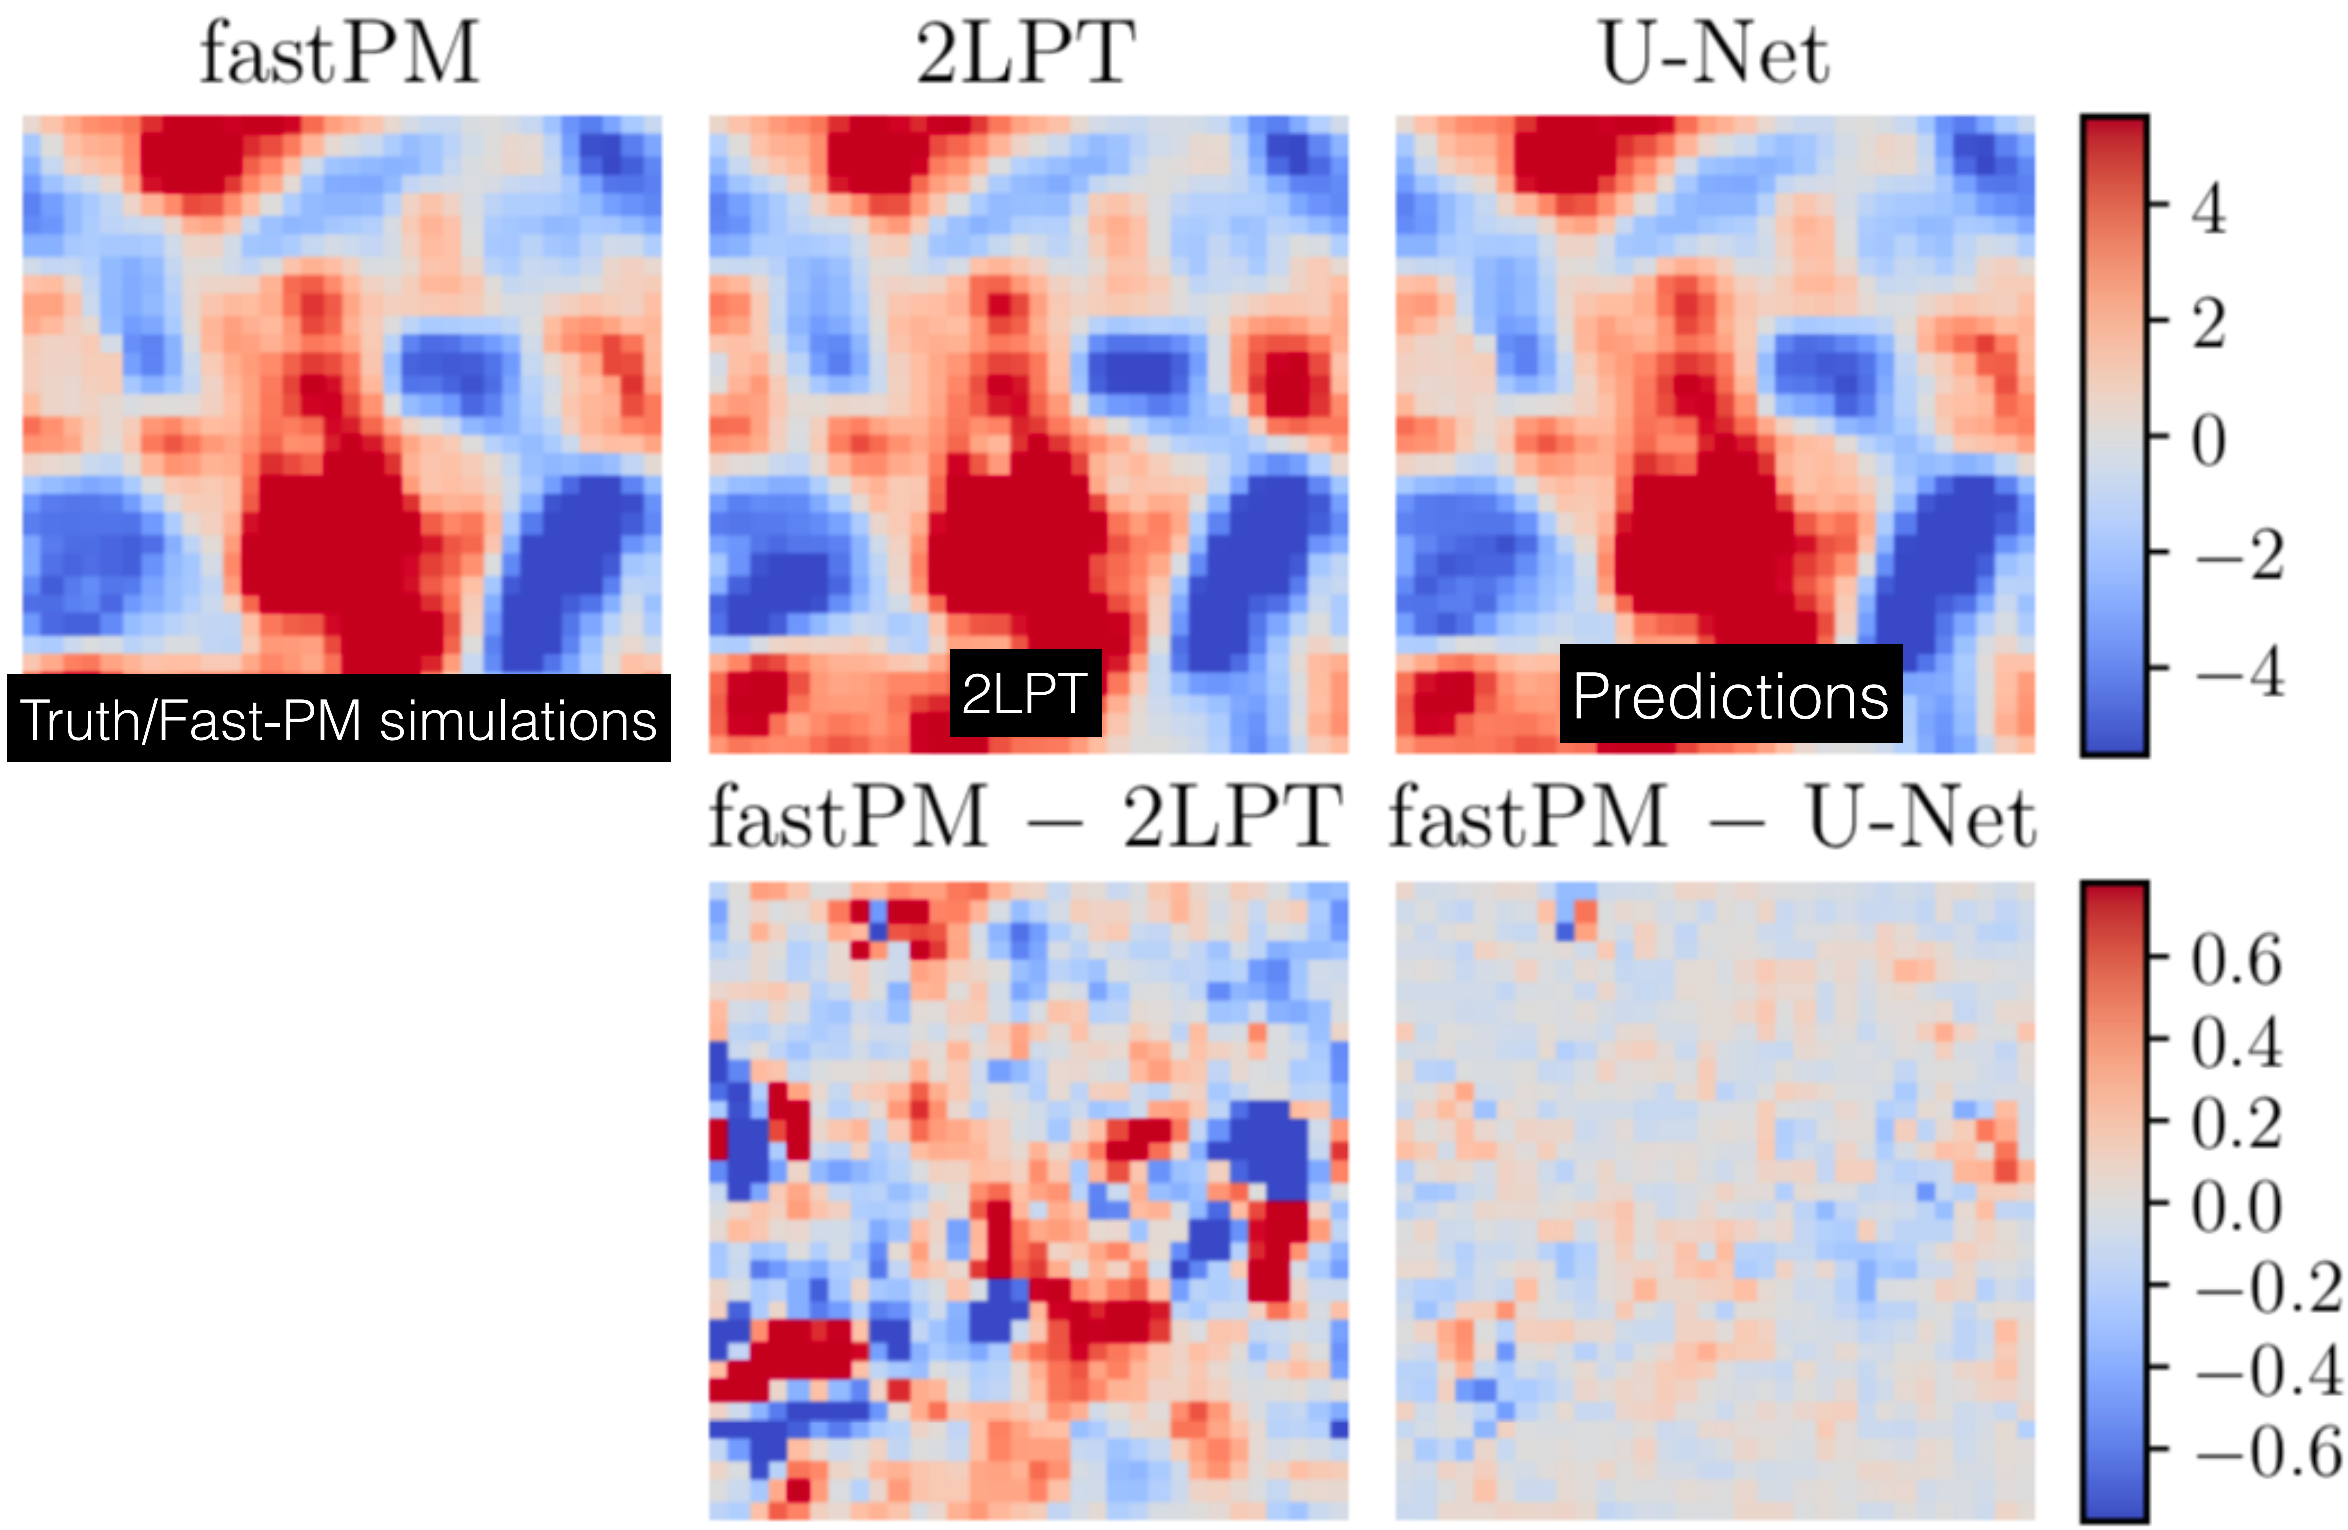

Displacement field comparisons

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

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.

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

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

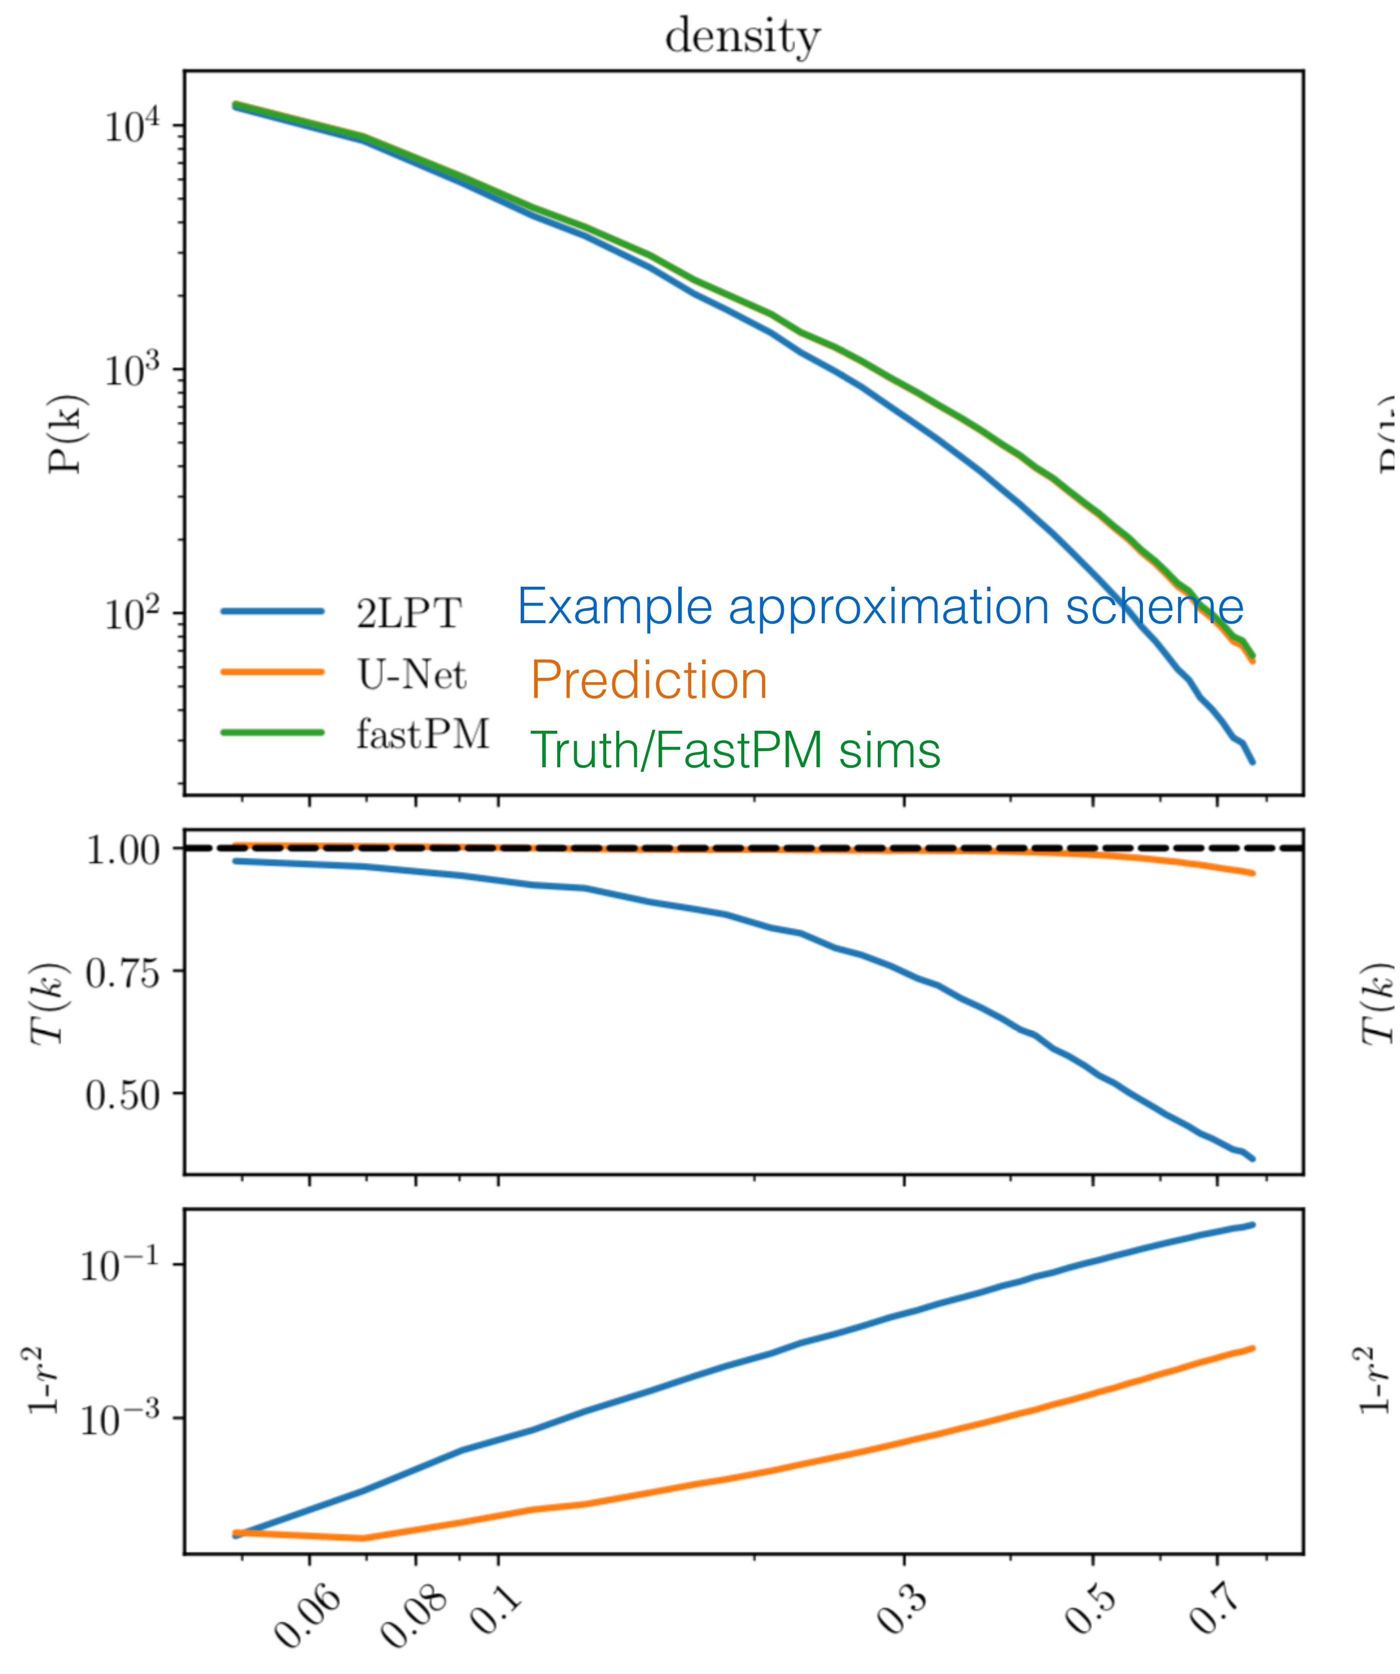

(a) Results from the density field

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

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:

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

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input

Slight variant to Residual NN

Prediction

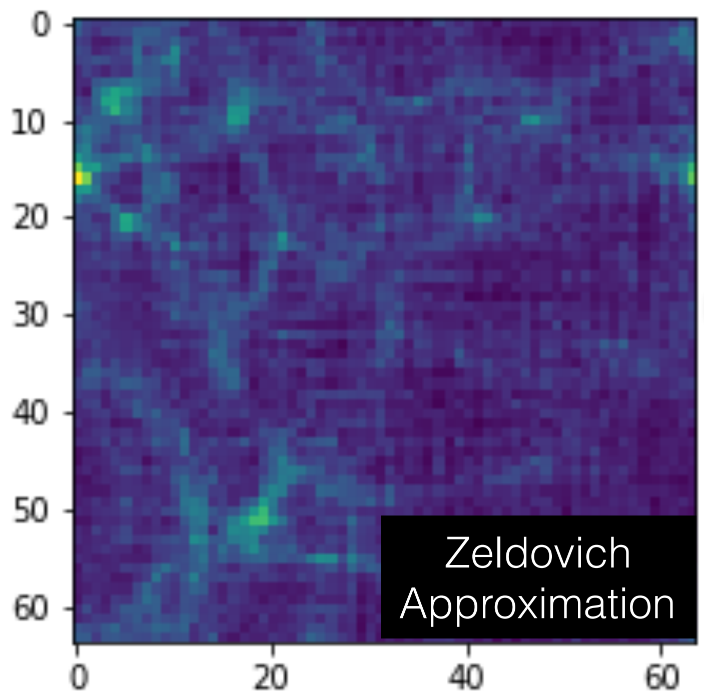

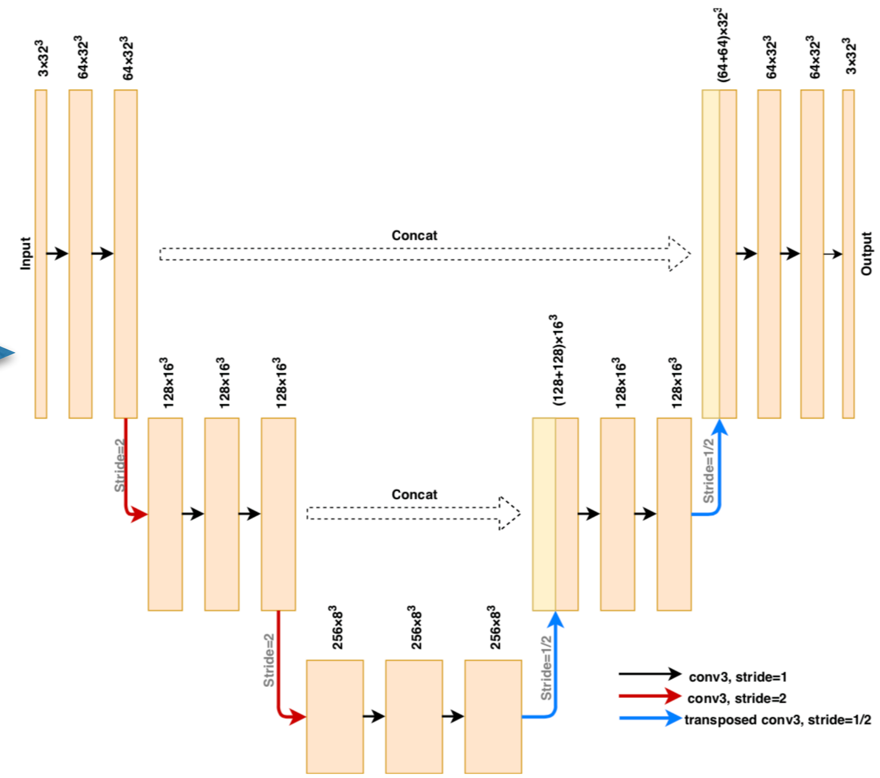

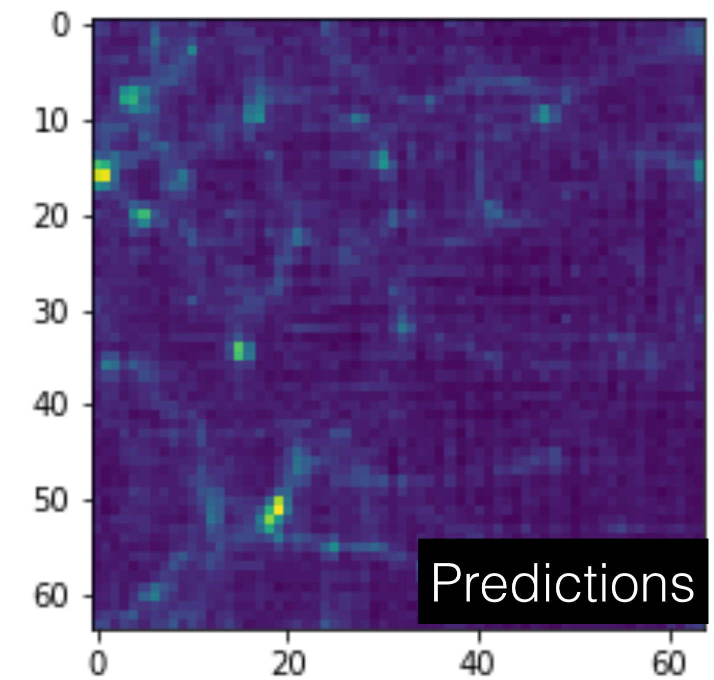

ZA boxes

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input | Slight variant to Residual NN | Prediction

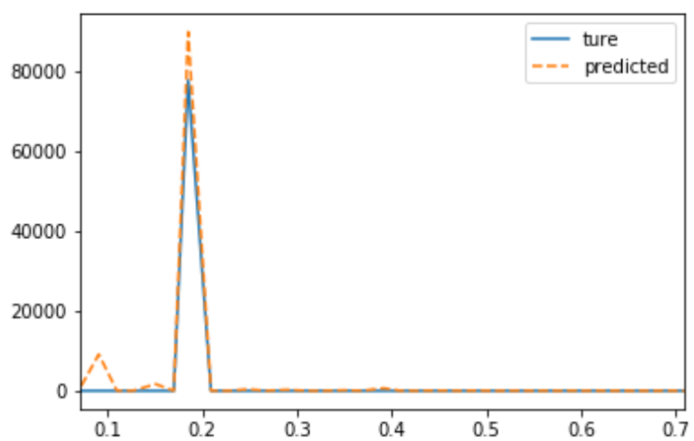

k(h/Mpc)

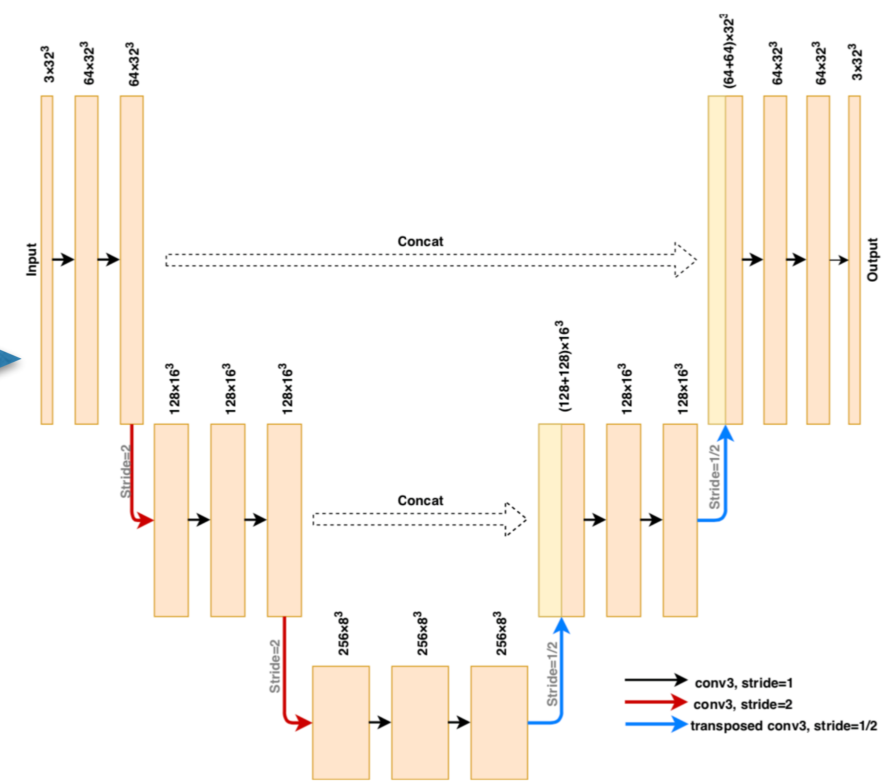

Inject Power
at one scale as input

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Input mode: Pancake…

What happen if we have power only one scale?

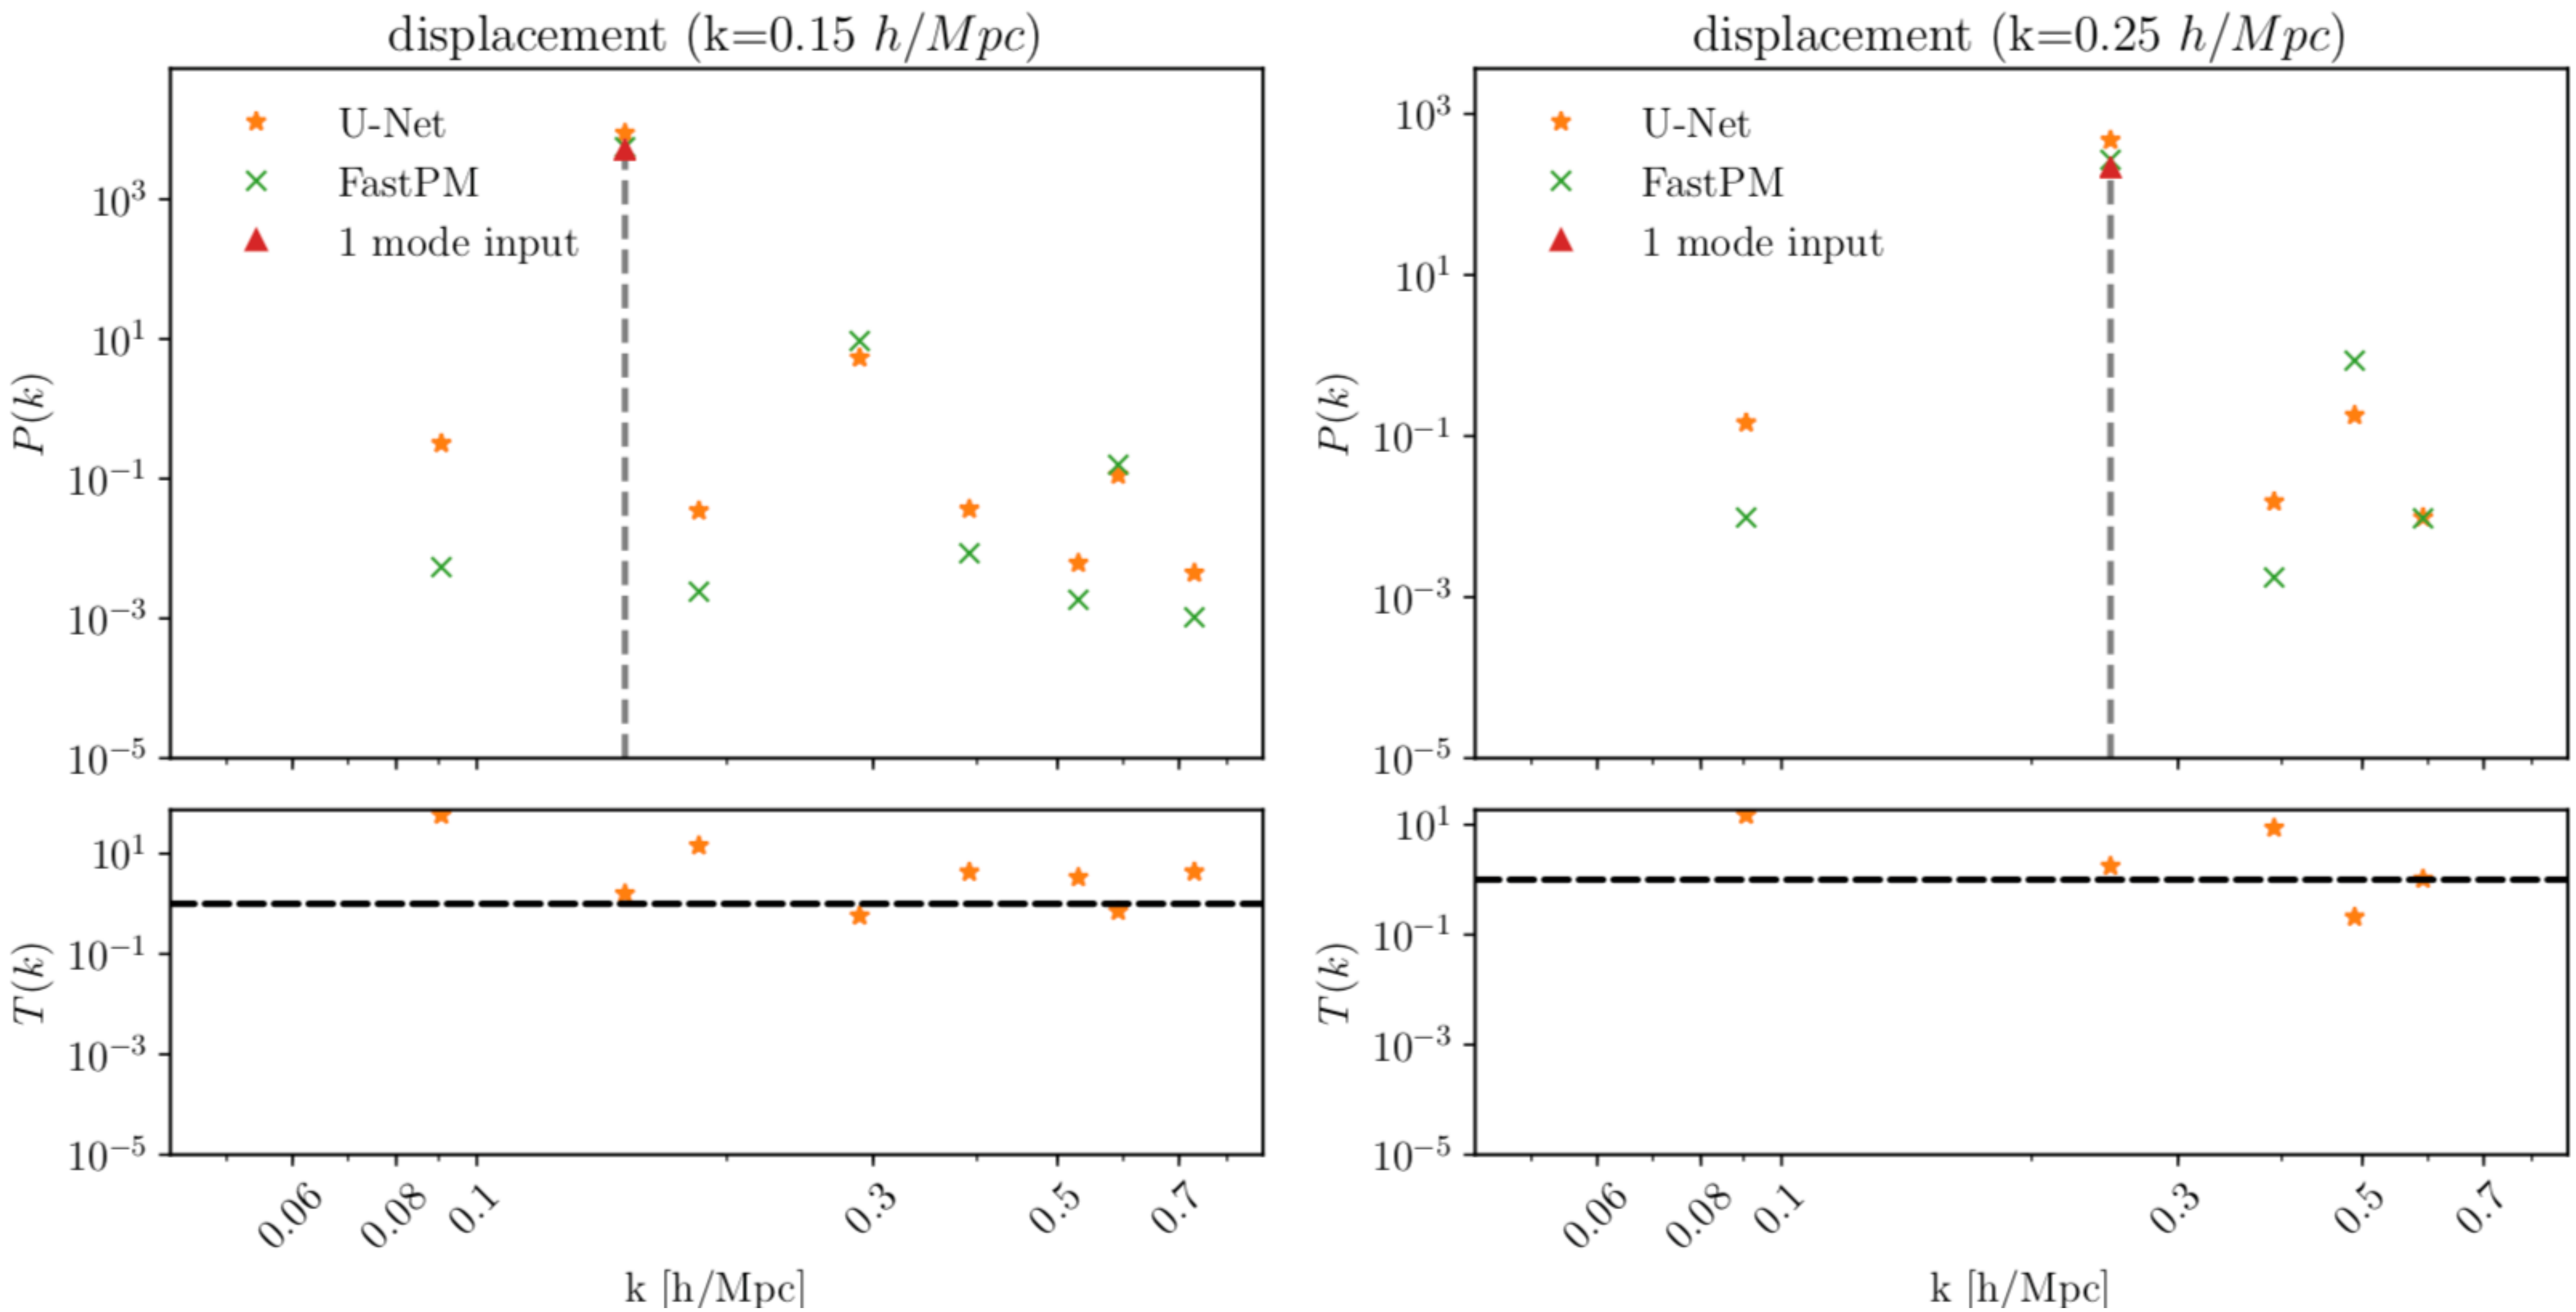

The transfer function shows that the U-Net model captures quite well at the dominate scale, which indicates the U-Net mode is able to capture scale information. The U-Net model also captures the other modes of FastPM that are two orders smaller than the dominant mode and come from the numerical artifact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

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input

Slight variant to Residual NN

Prediction

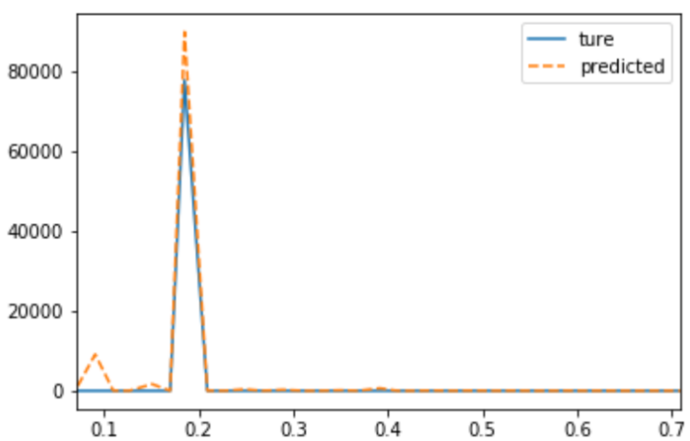

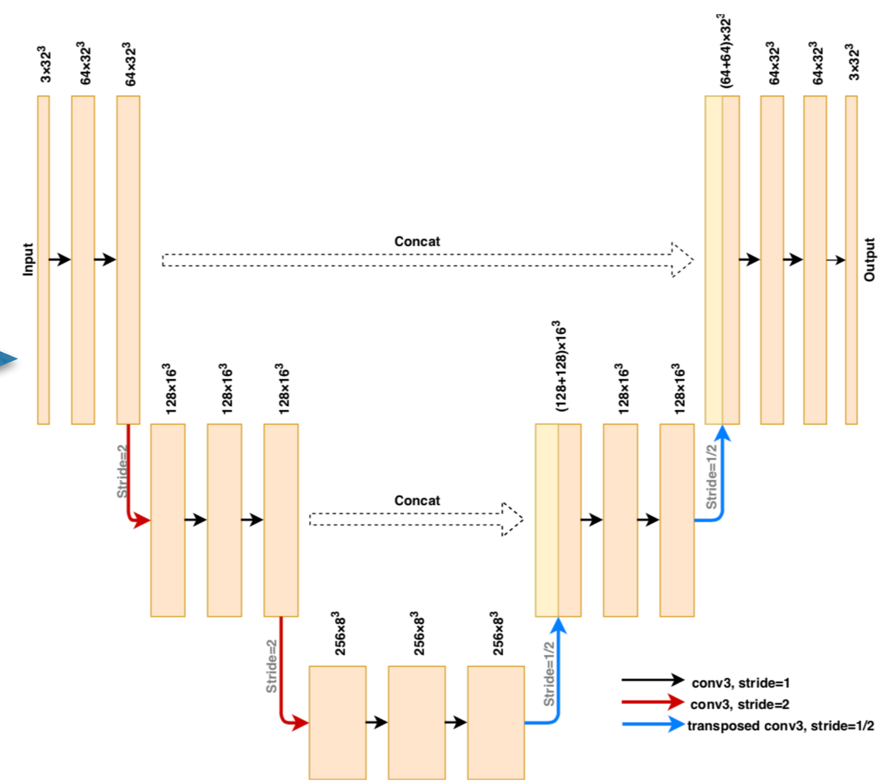

k(h/Mpc)

Inject Power

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

$= [0, 0.25, 0]$

$= [0, 0, 0.25]$

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Is Rotational Invariance learnt by the model?

Yes, *predicted power is the similar no matter which orientatio*n: rotational invariance is learnt!

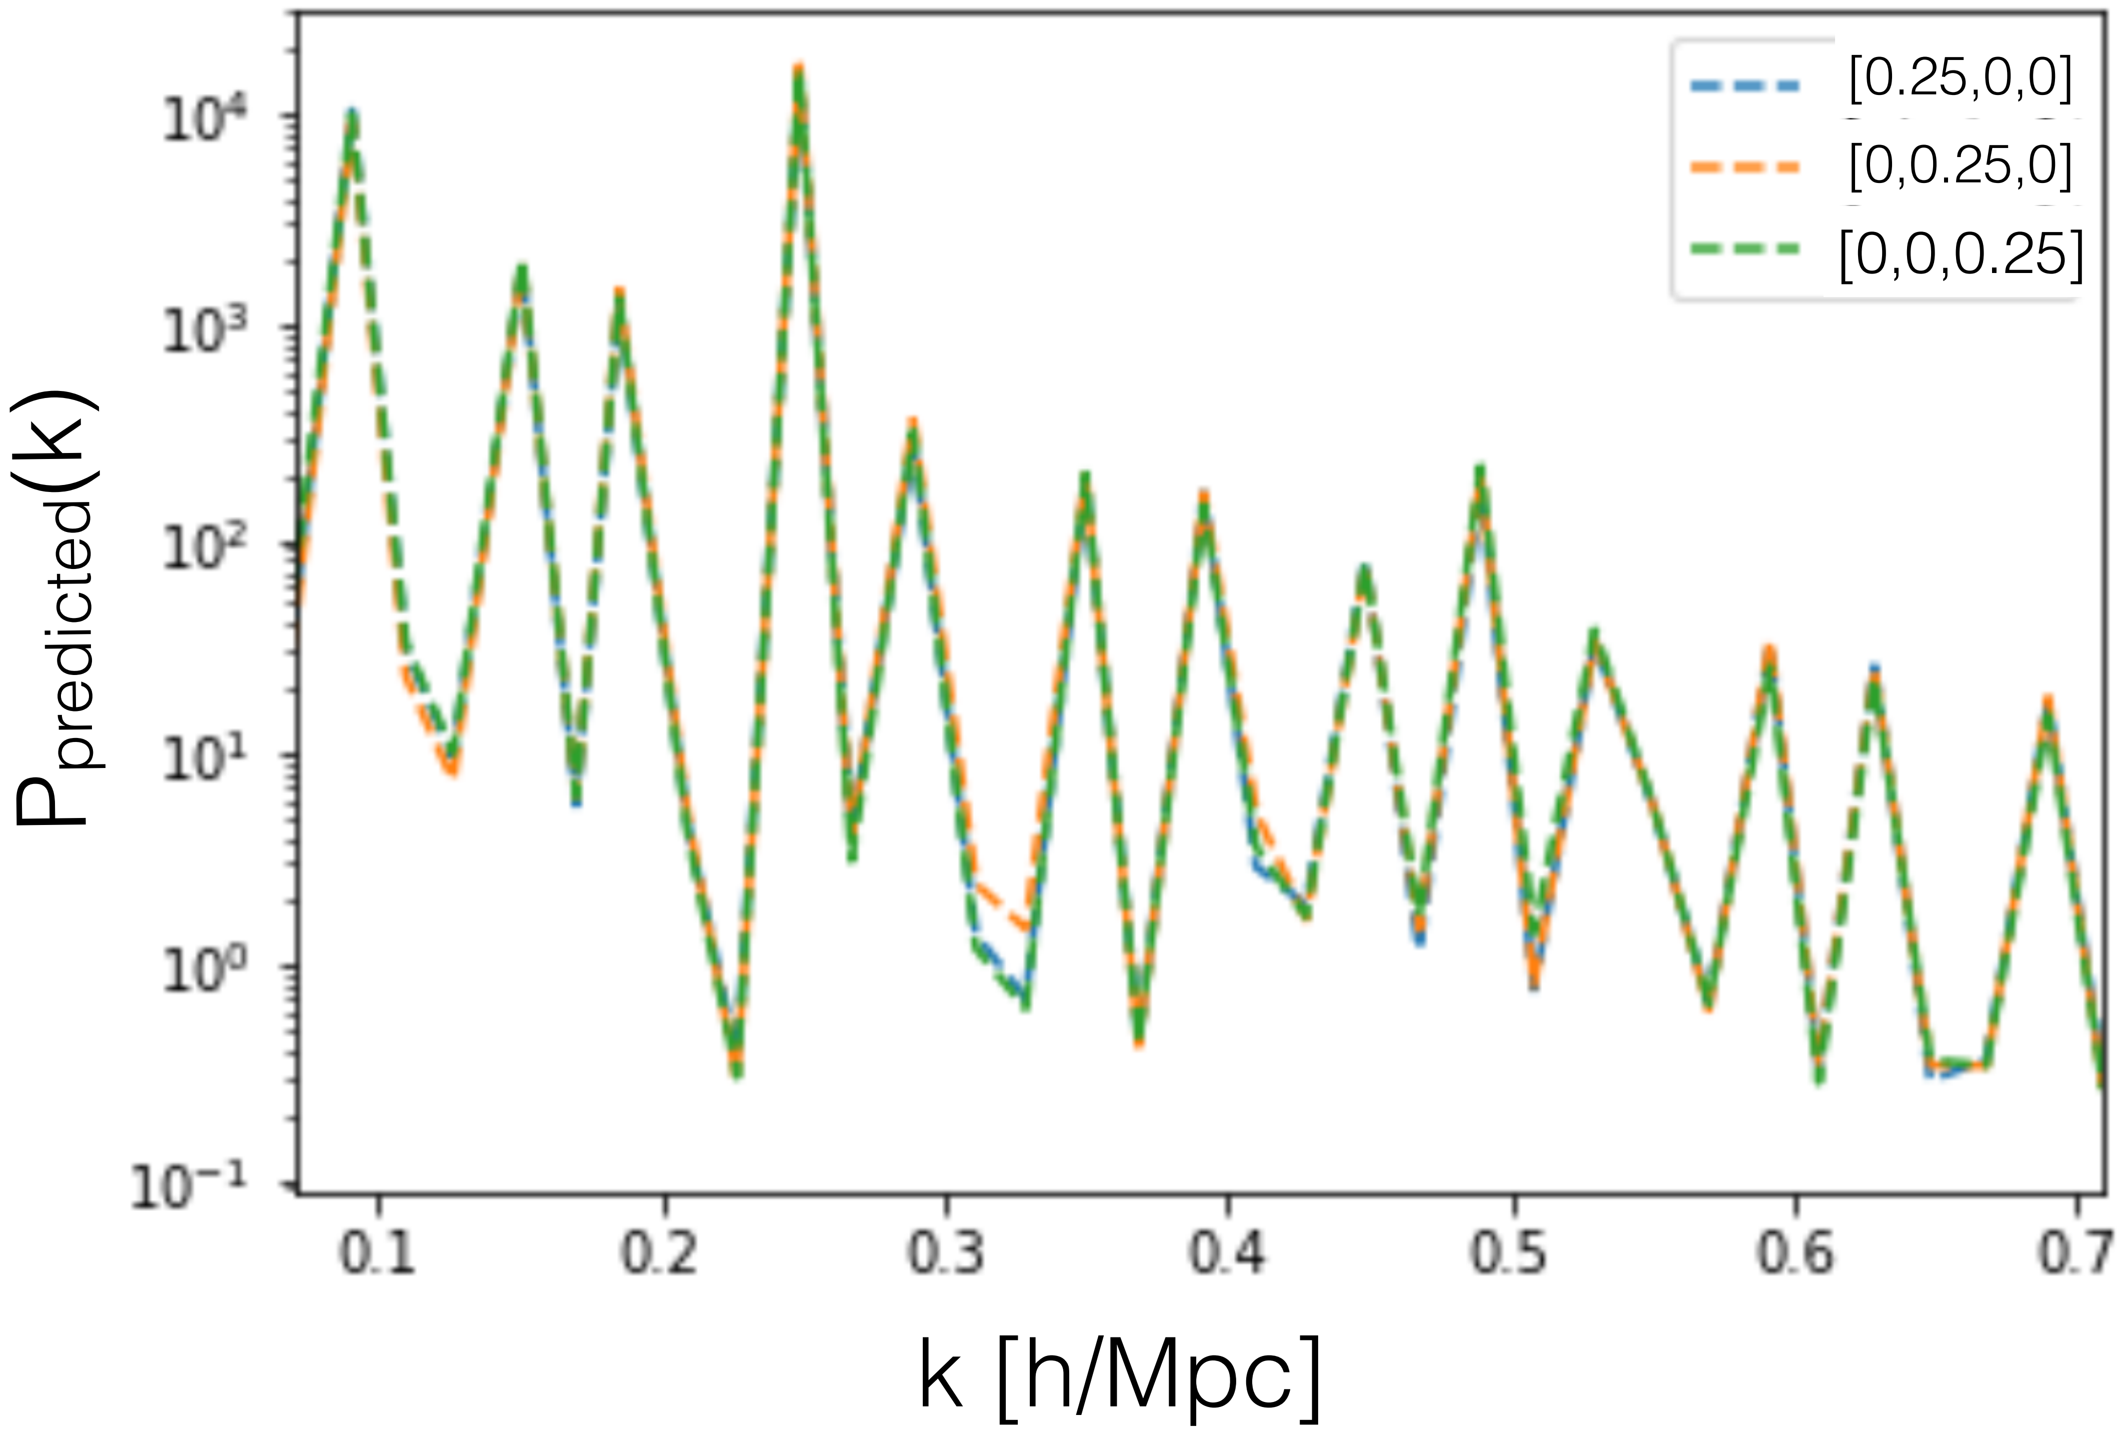

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input | Slight variant to Residual NN | Prediction

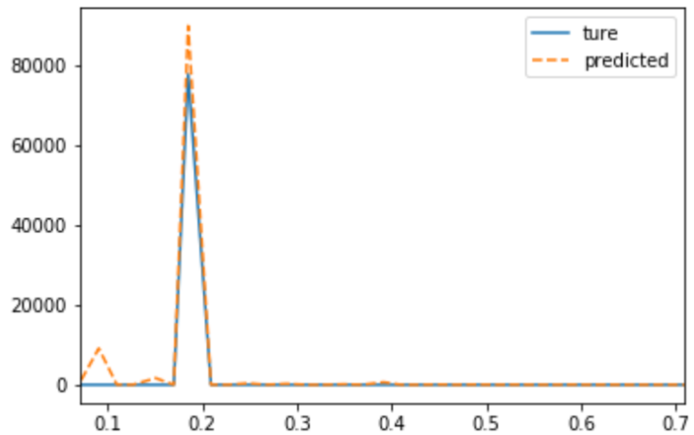

k(h/Mpc)

Inject Power
At same k,
but different phases

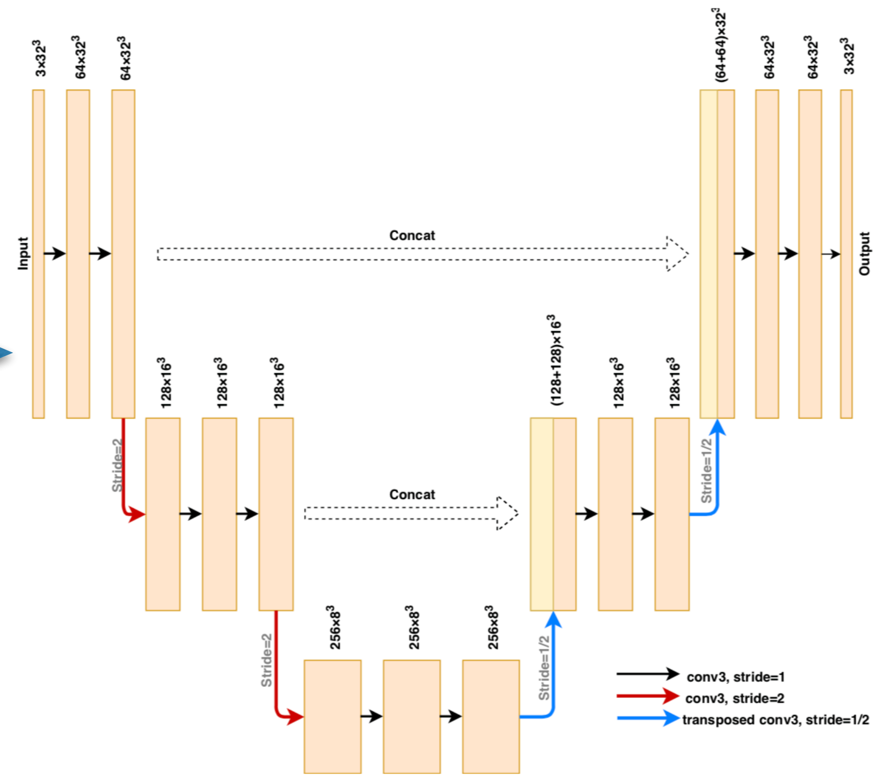

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

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

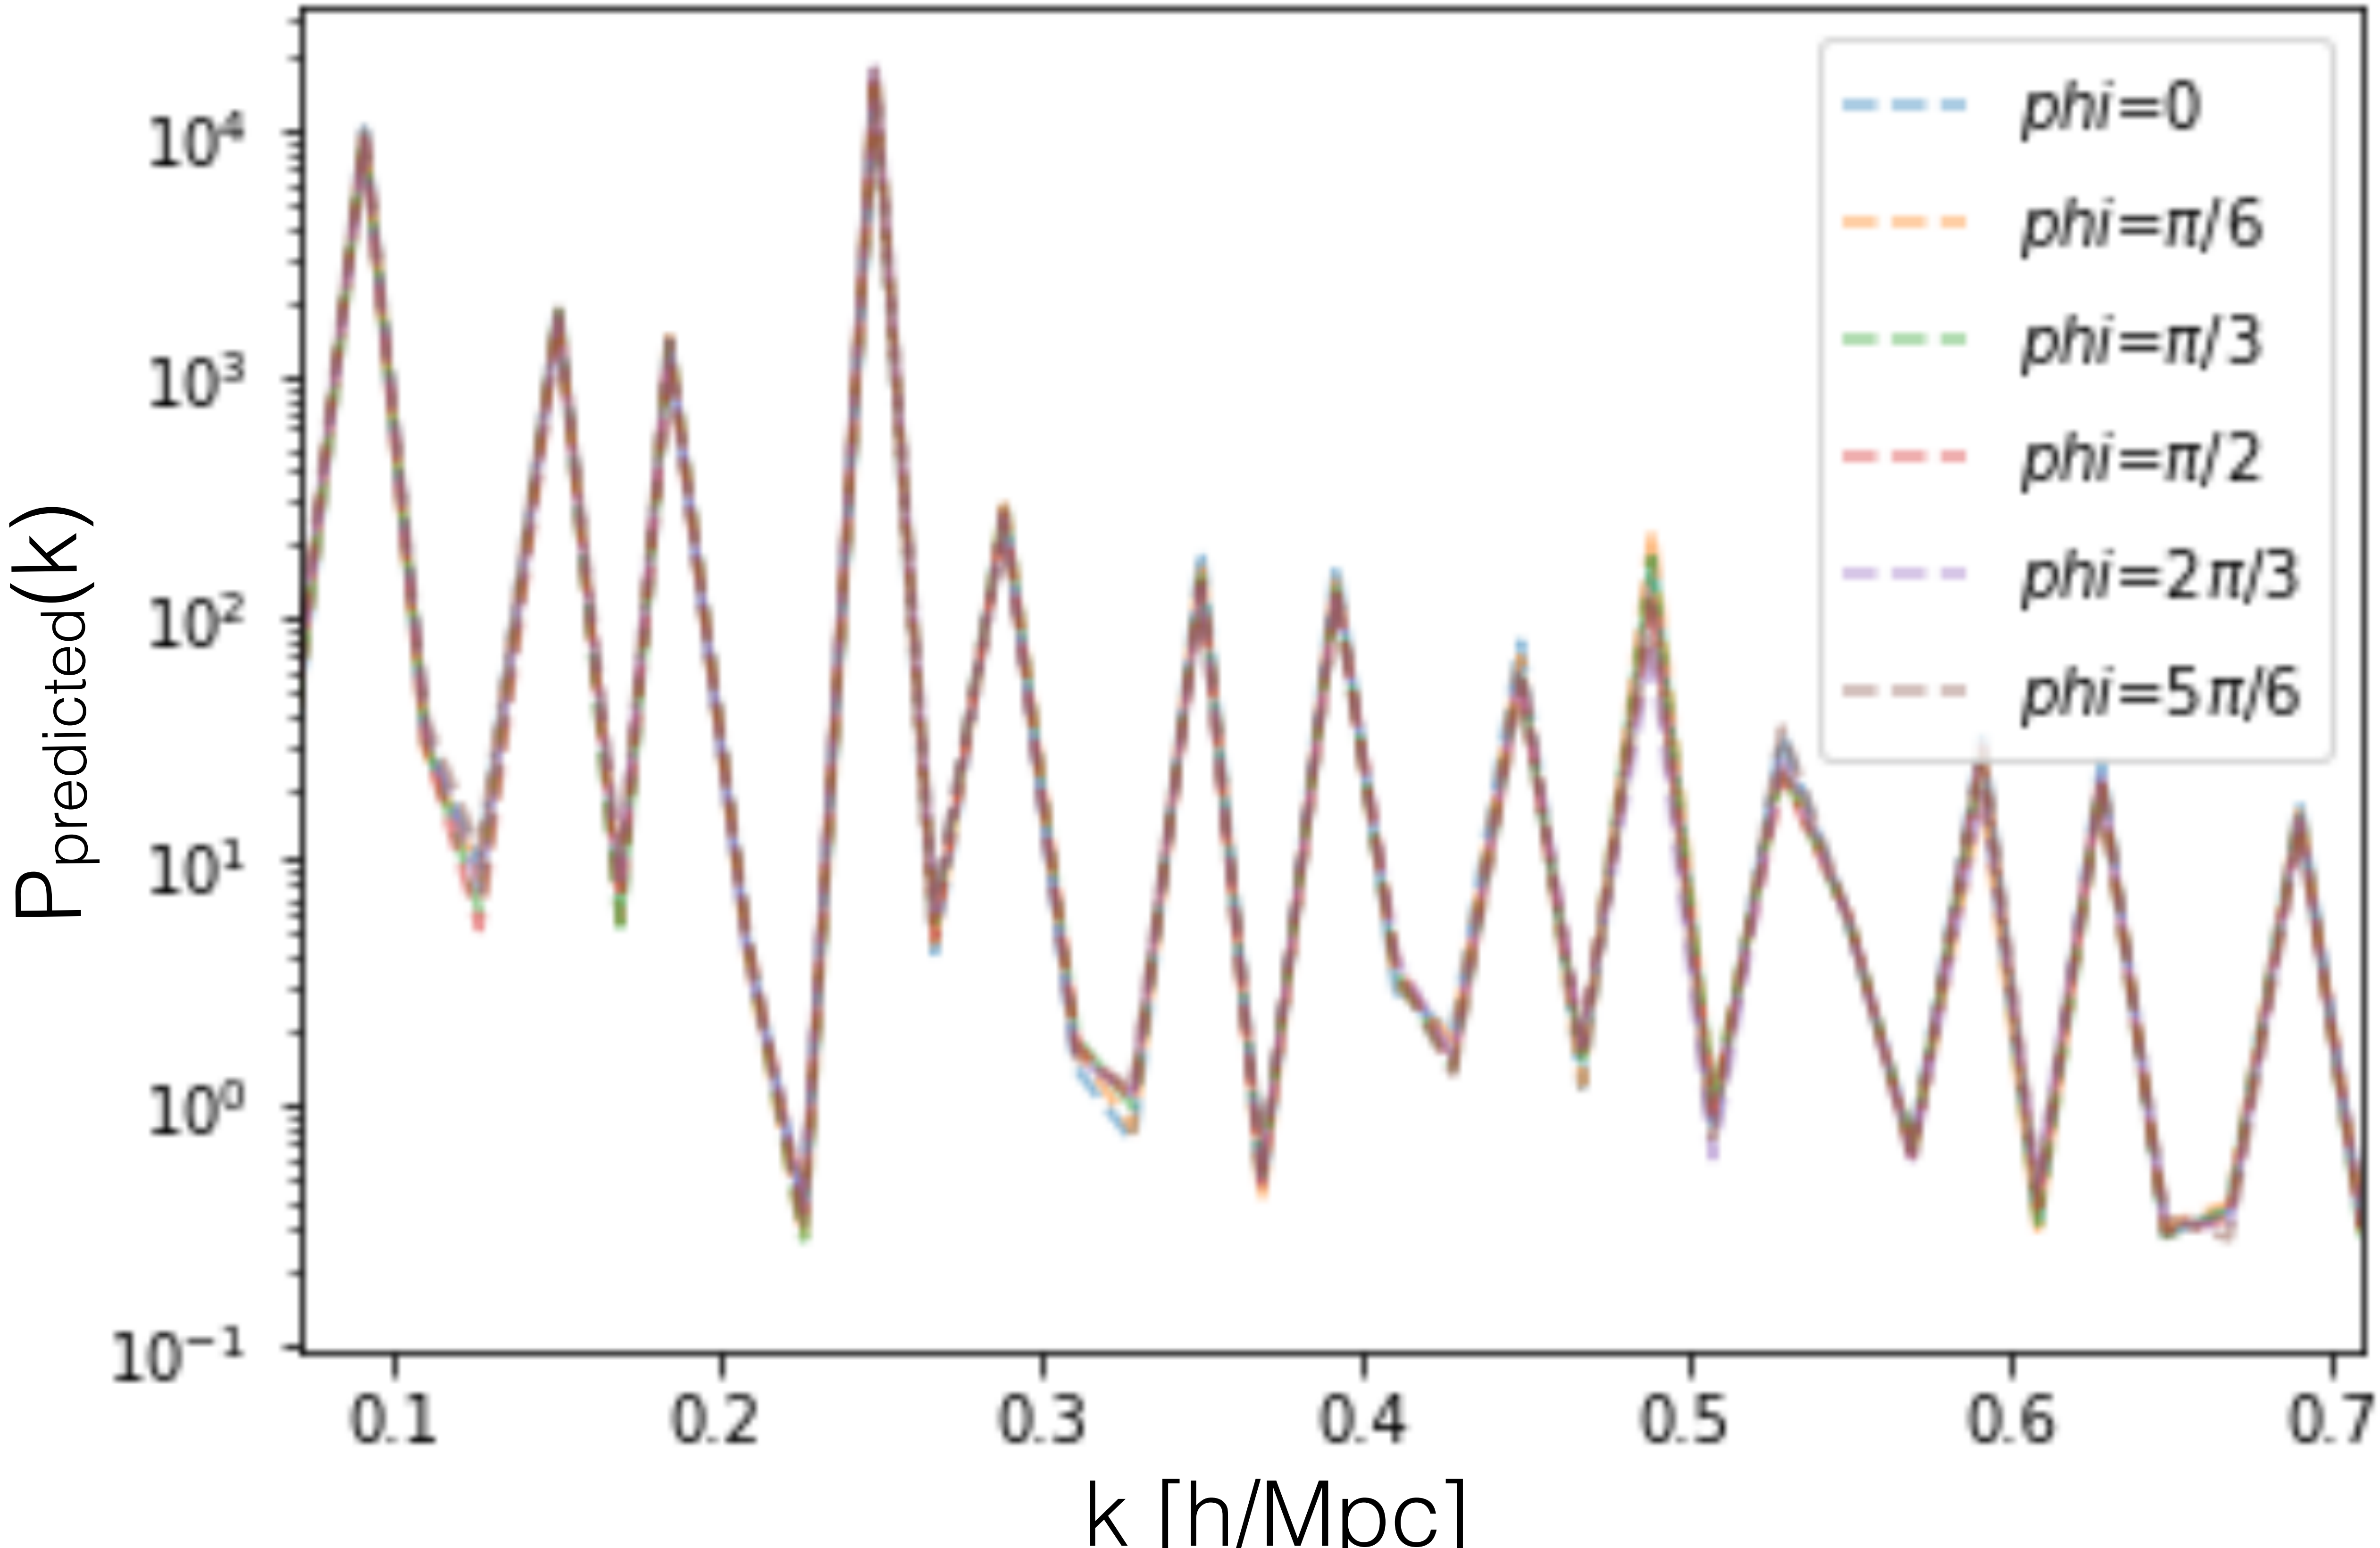

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Predicting from Zeldovich Approximation fields to Fast-PM simulated fields

UNET

Input

Slight variant to Residual NN

Prediction

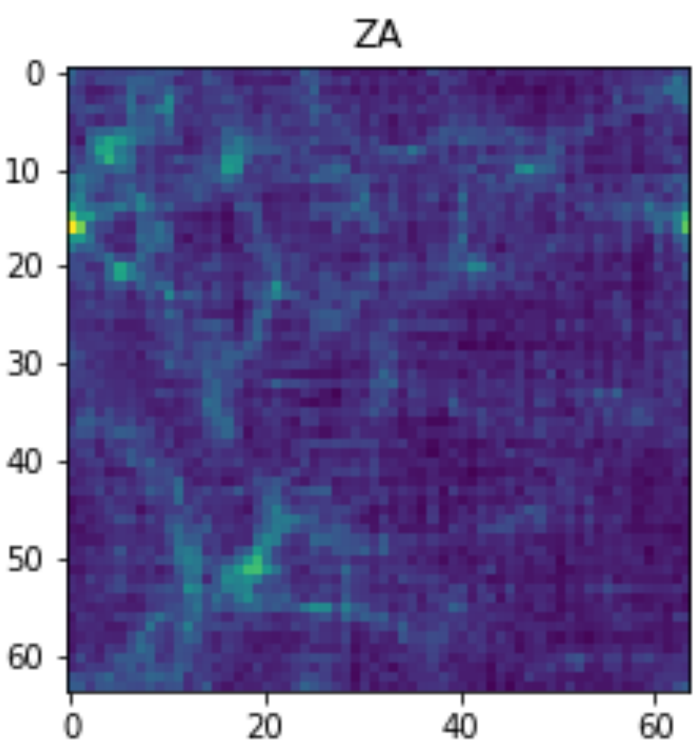

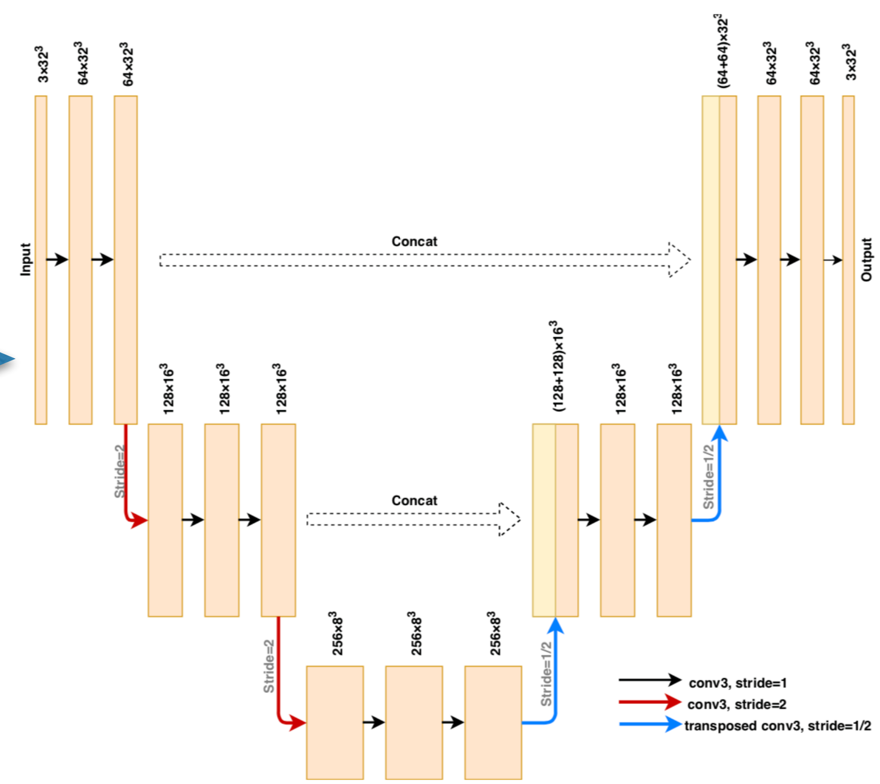

ZA maps of

Different cosmology:

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

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H**., Siamak Ravanbaksh, Barnabas Poczos

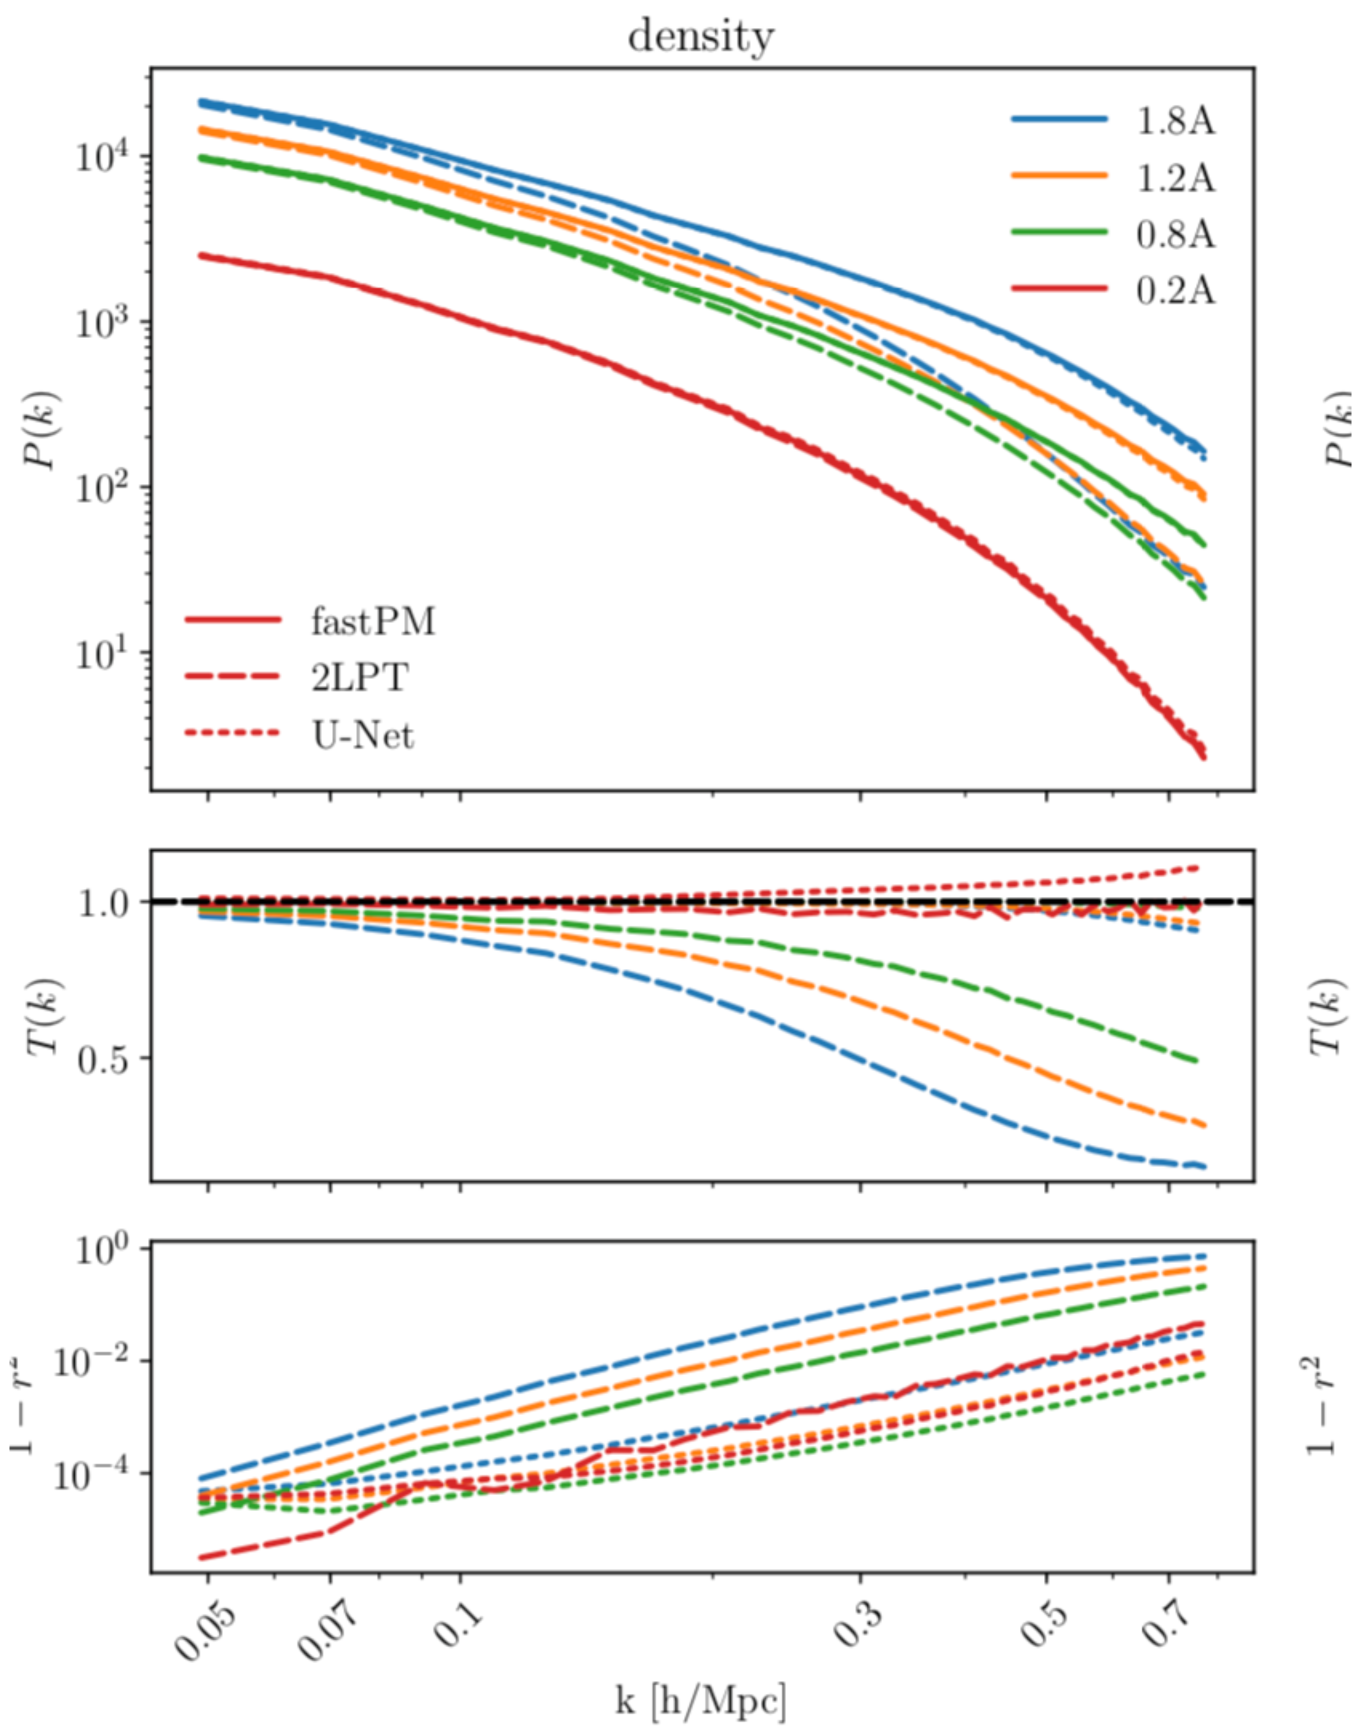

(a) Results from the density field

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

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

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

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 ?

Possible reason ?

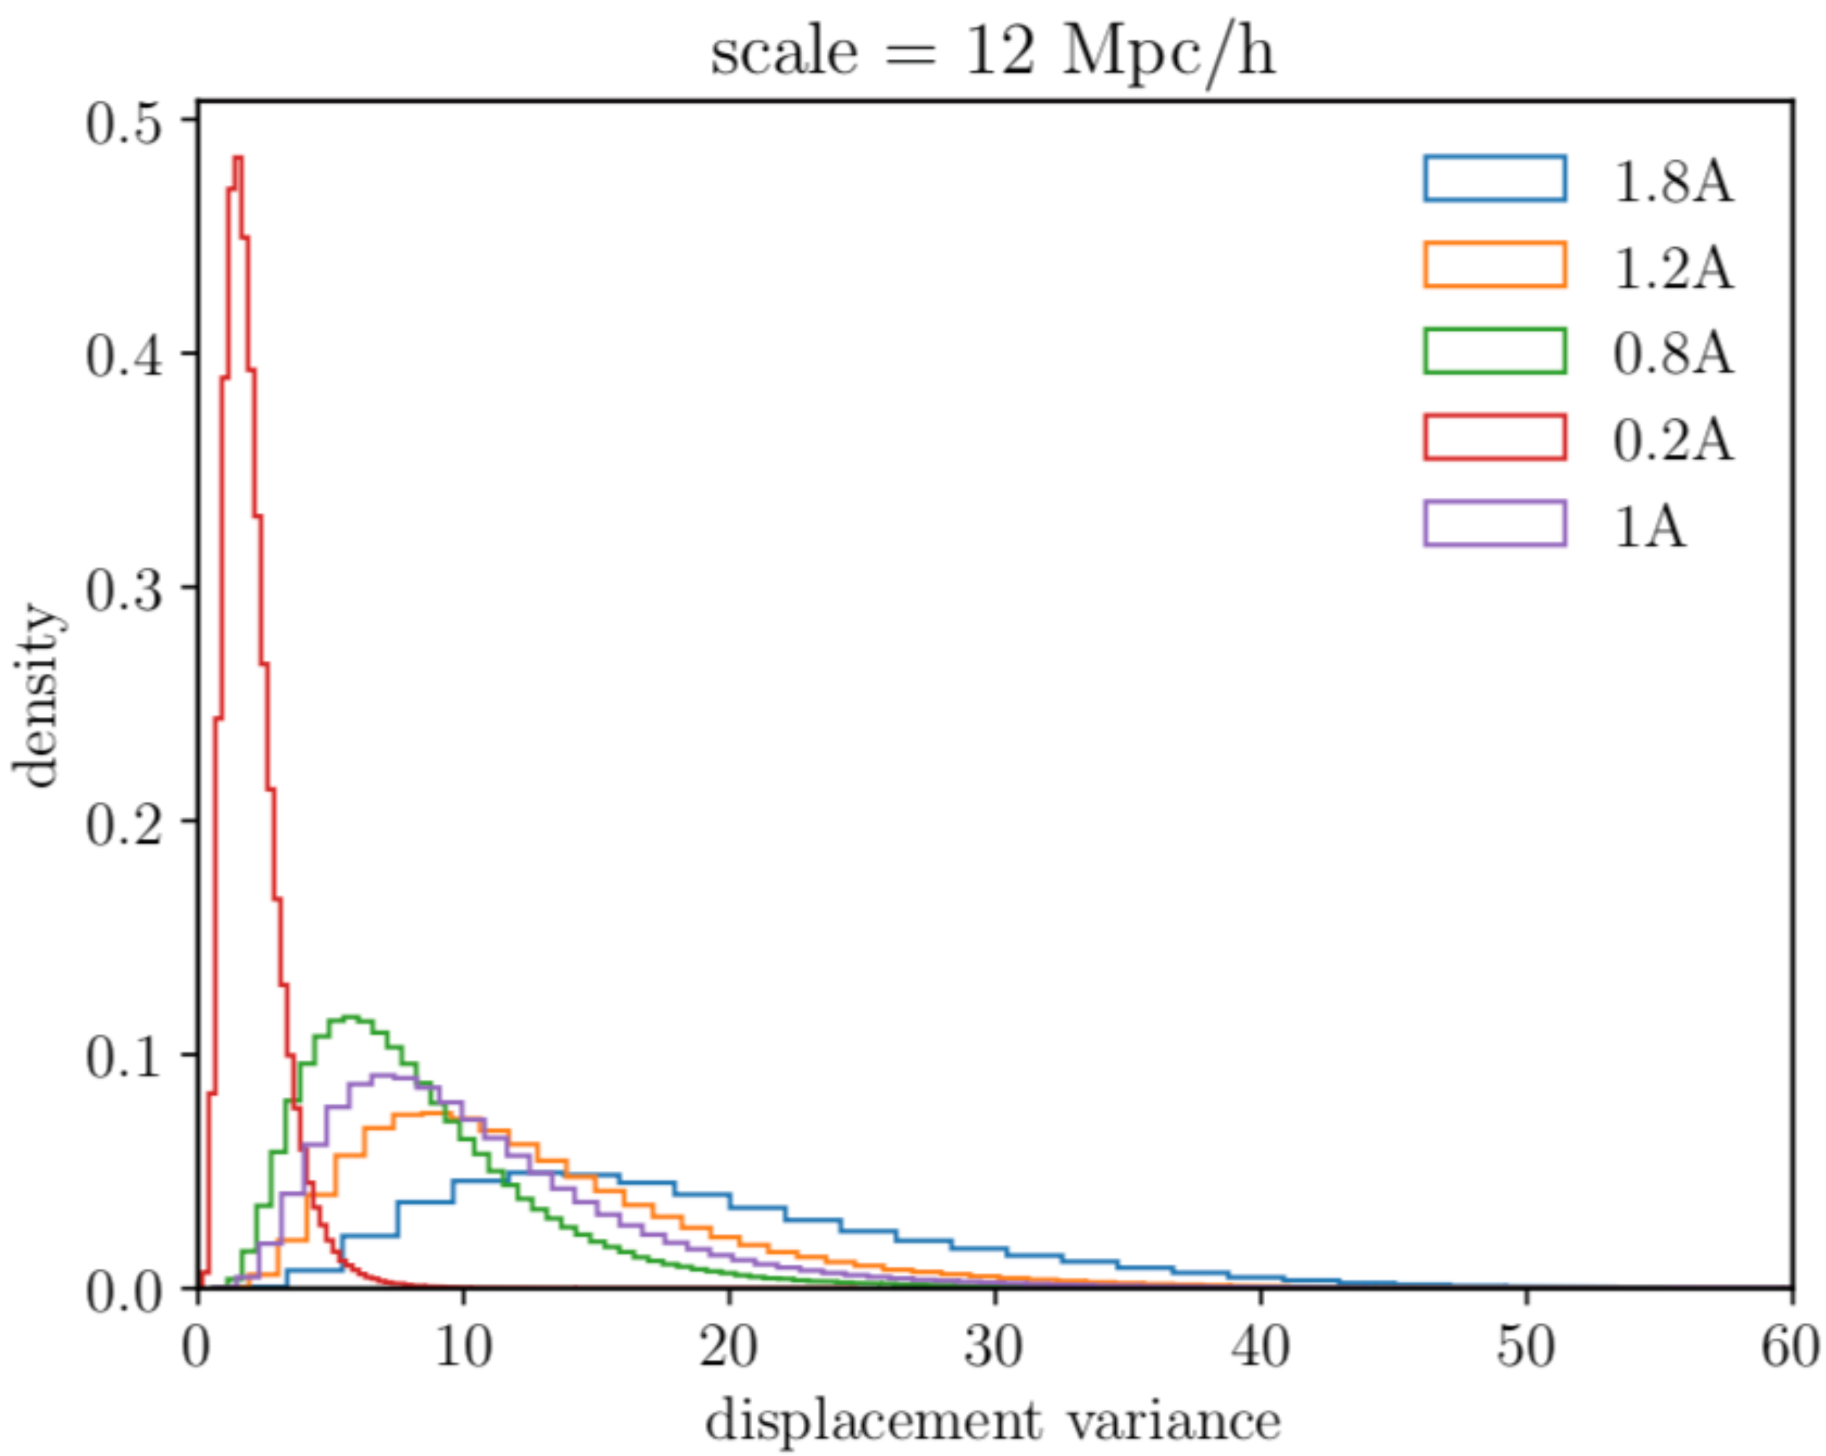

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Conclusion

- First foray into learning cosmological parameters from LSS simulations
- First foray into unraveling the blackbox called Deep Neural Net.
- Predicts N-body like simulations in minutes (*post training).
- Model learn about power at each scales, rotational invariance, phase preservation.
- Does the model generalize and learn real physical laws?
- Or does it generalize from the various “island universes” with different cosmological parameters?

Let's talk about what we expect first

- At large scales: physics are completely linear, and can be fully represented by the analytical inputs, so the large scales should be preserved at the output

- At small scales: physics are not well modeled by linear theory, so we expect that the model predict both small scale power and large scale powers.

- We have predictions from perturbation theory to higher order, but these are not complete, but we hope to use these as guidance/prior. Another interesting question: What is the best way to incorporate intuition / prior knowledge in the network?

Learning Physics from ML?

What happen if we have power only one scale?

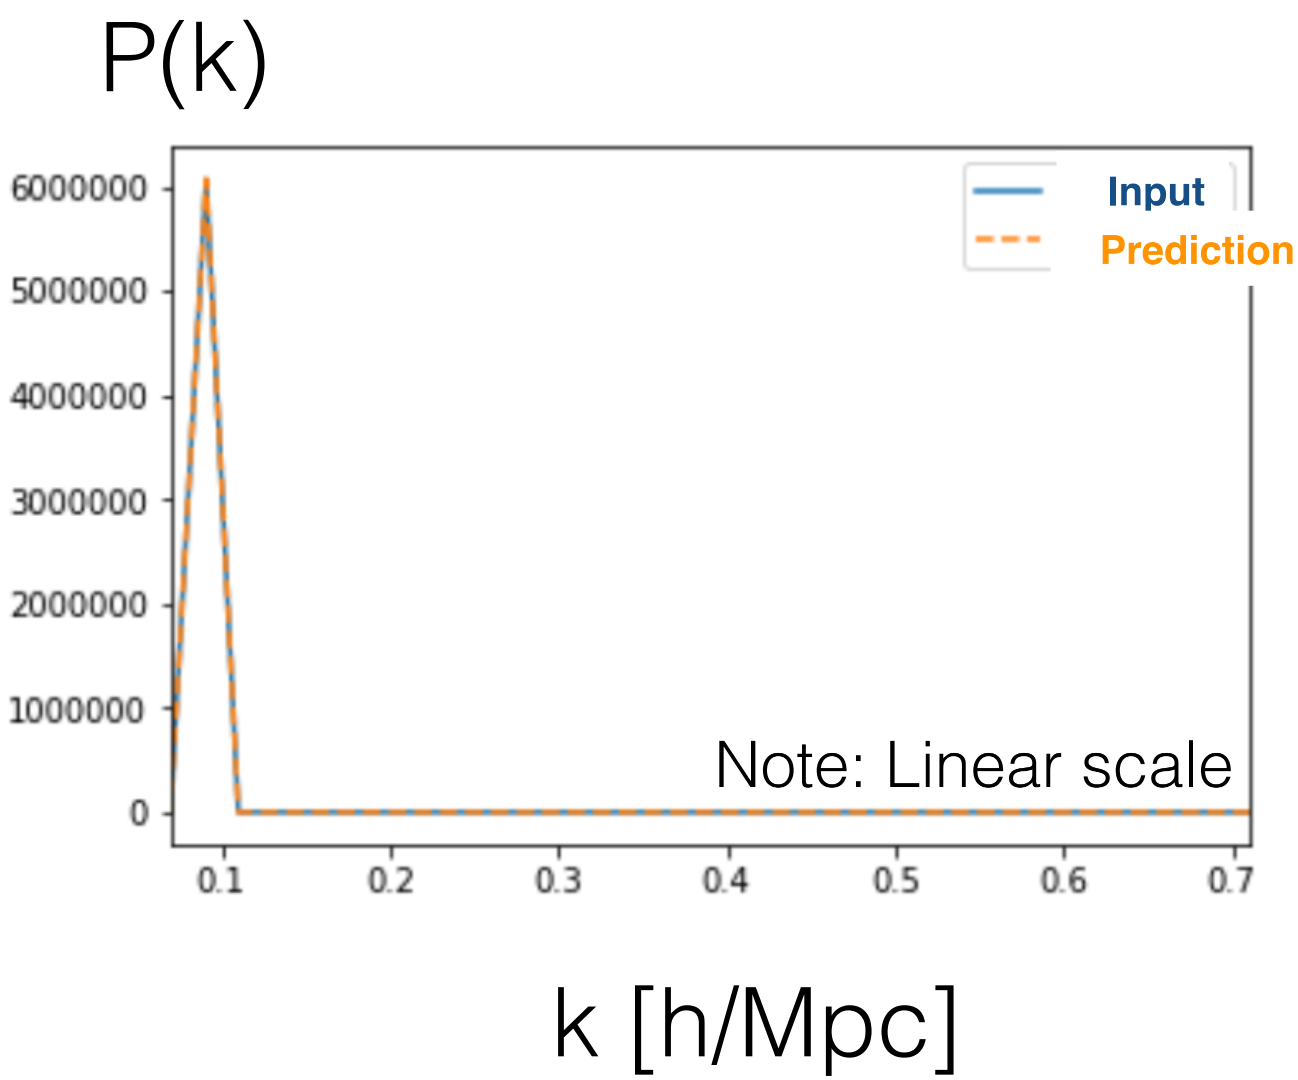

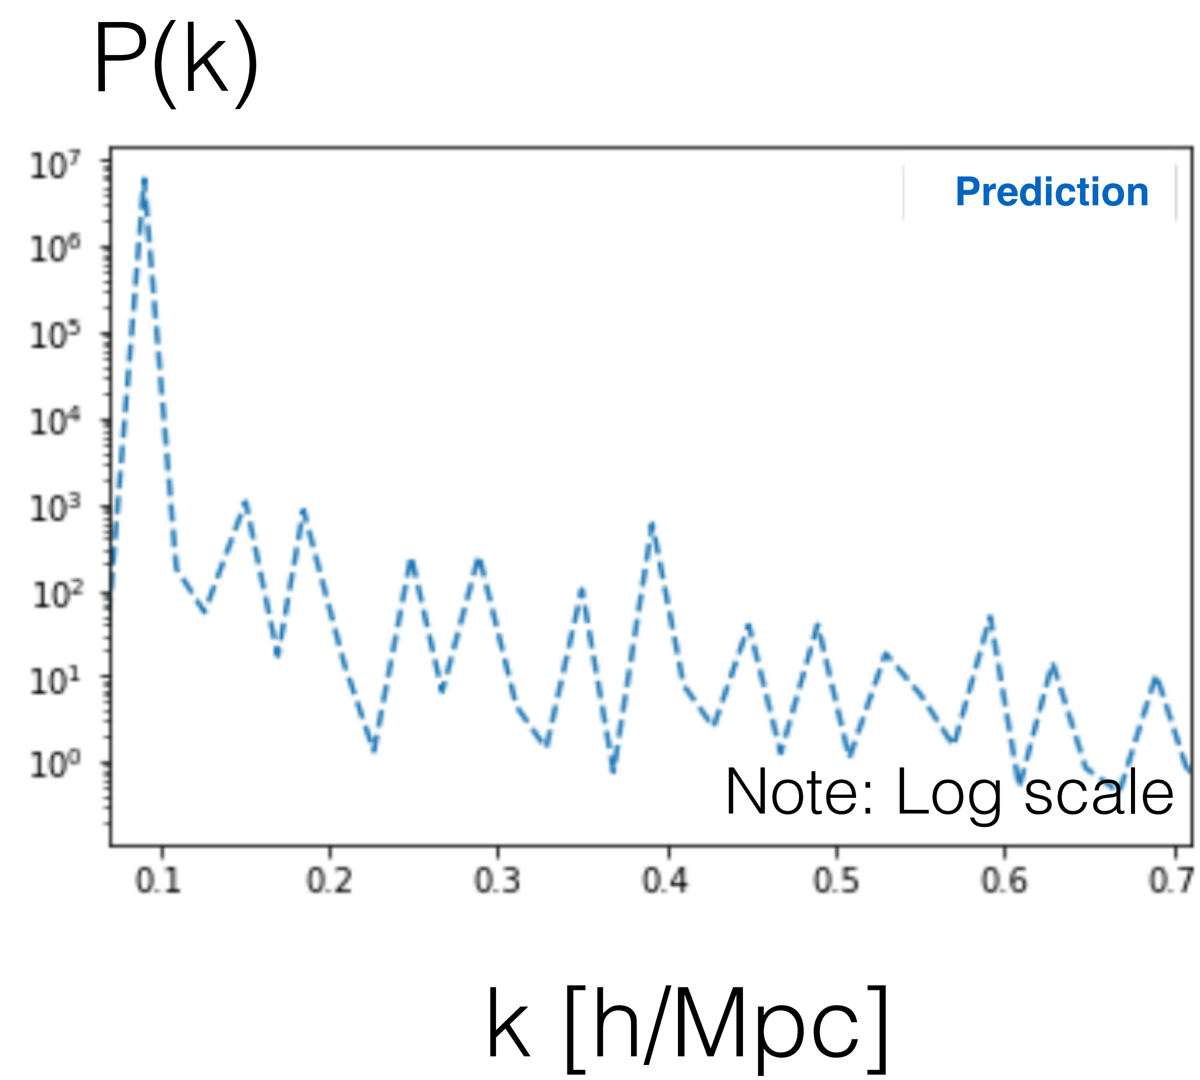

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Learning Physics from ML?

Power at one large scale gives power at multiple scales

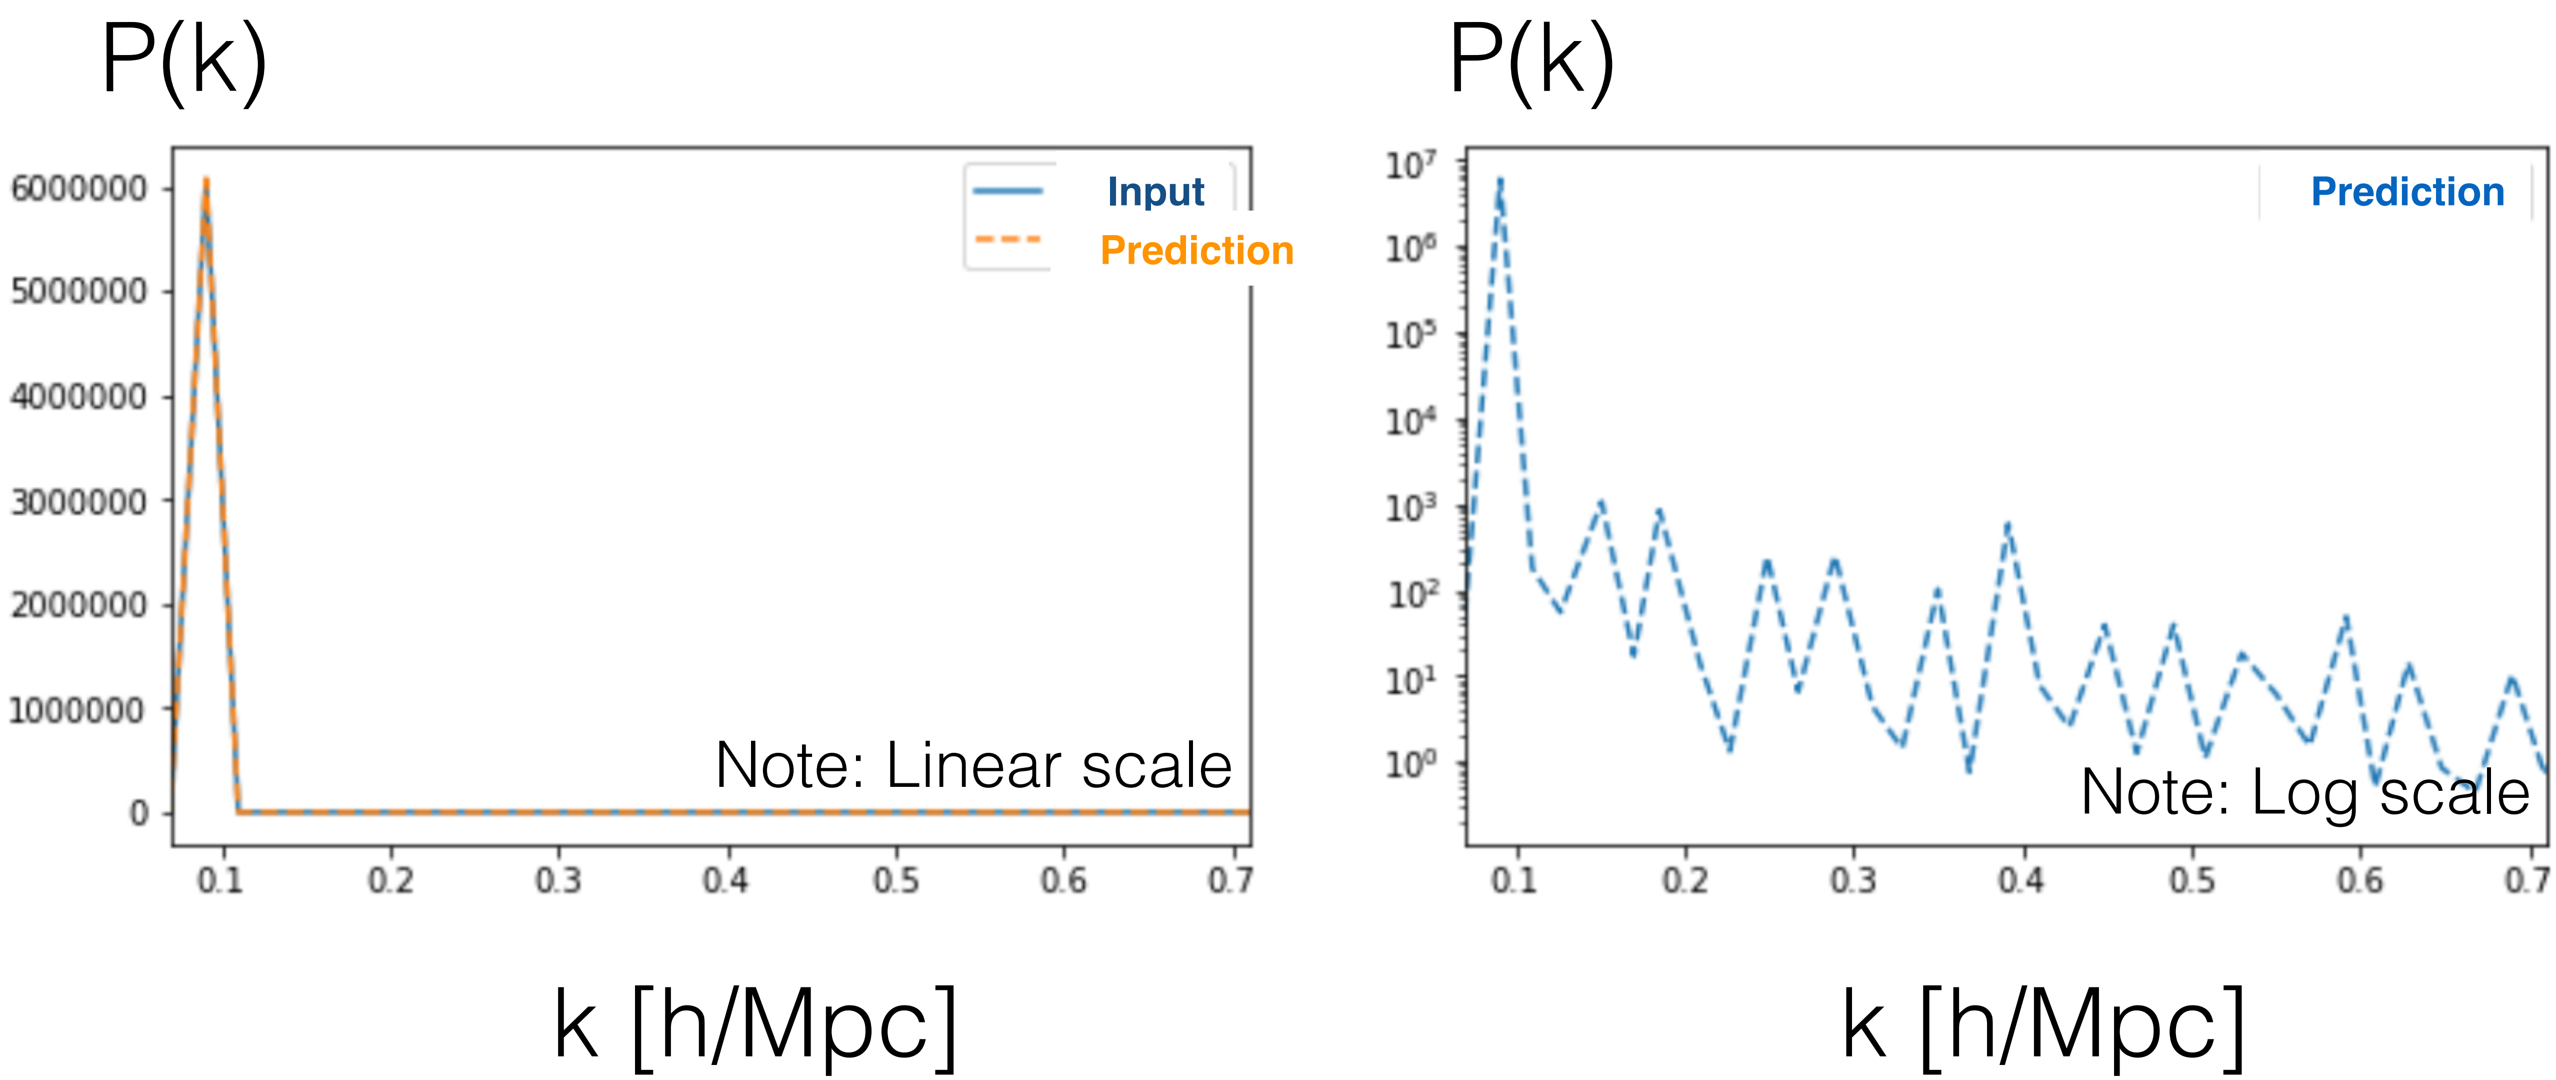

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Moving the input mode to smaller scales.

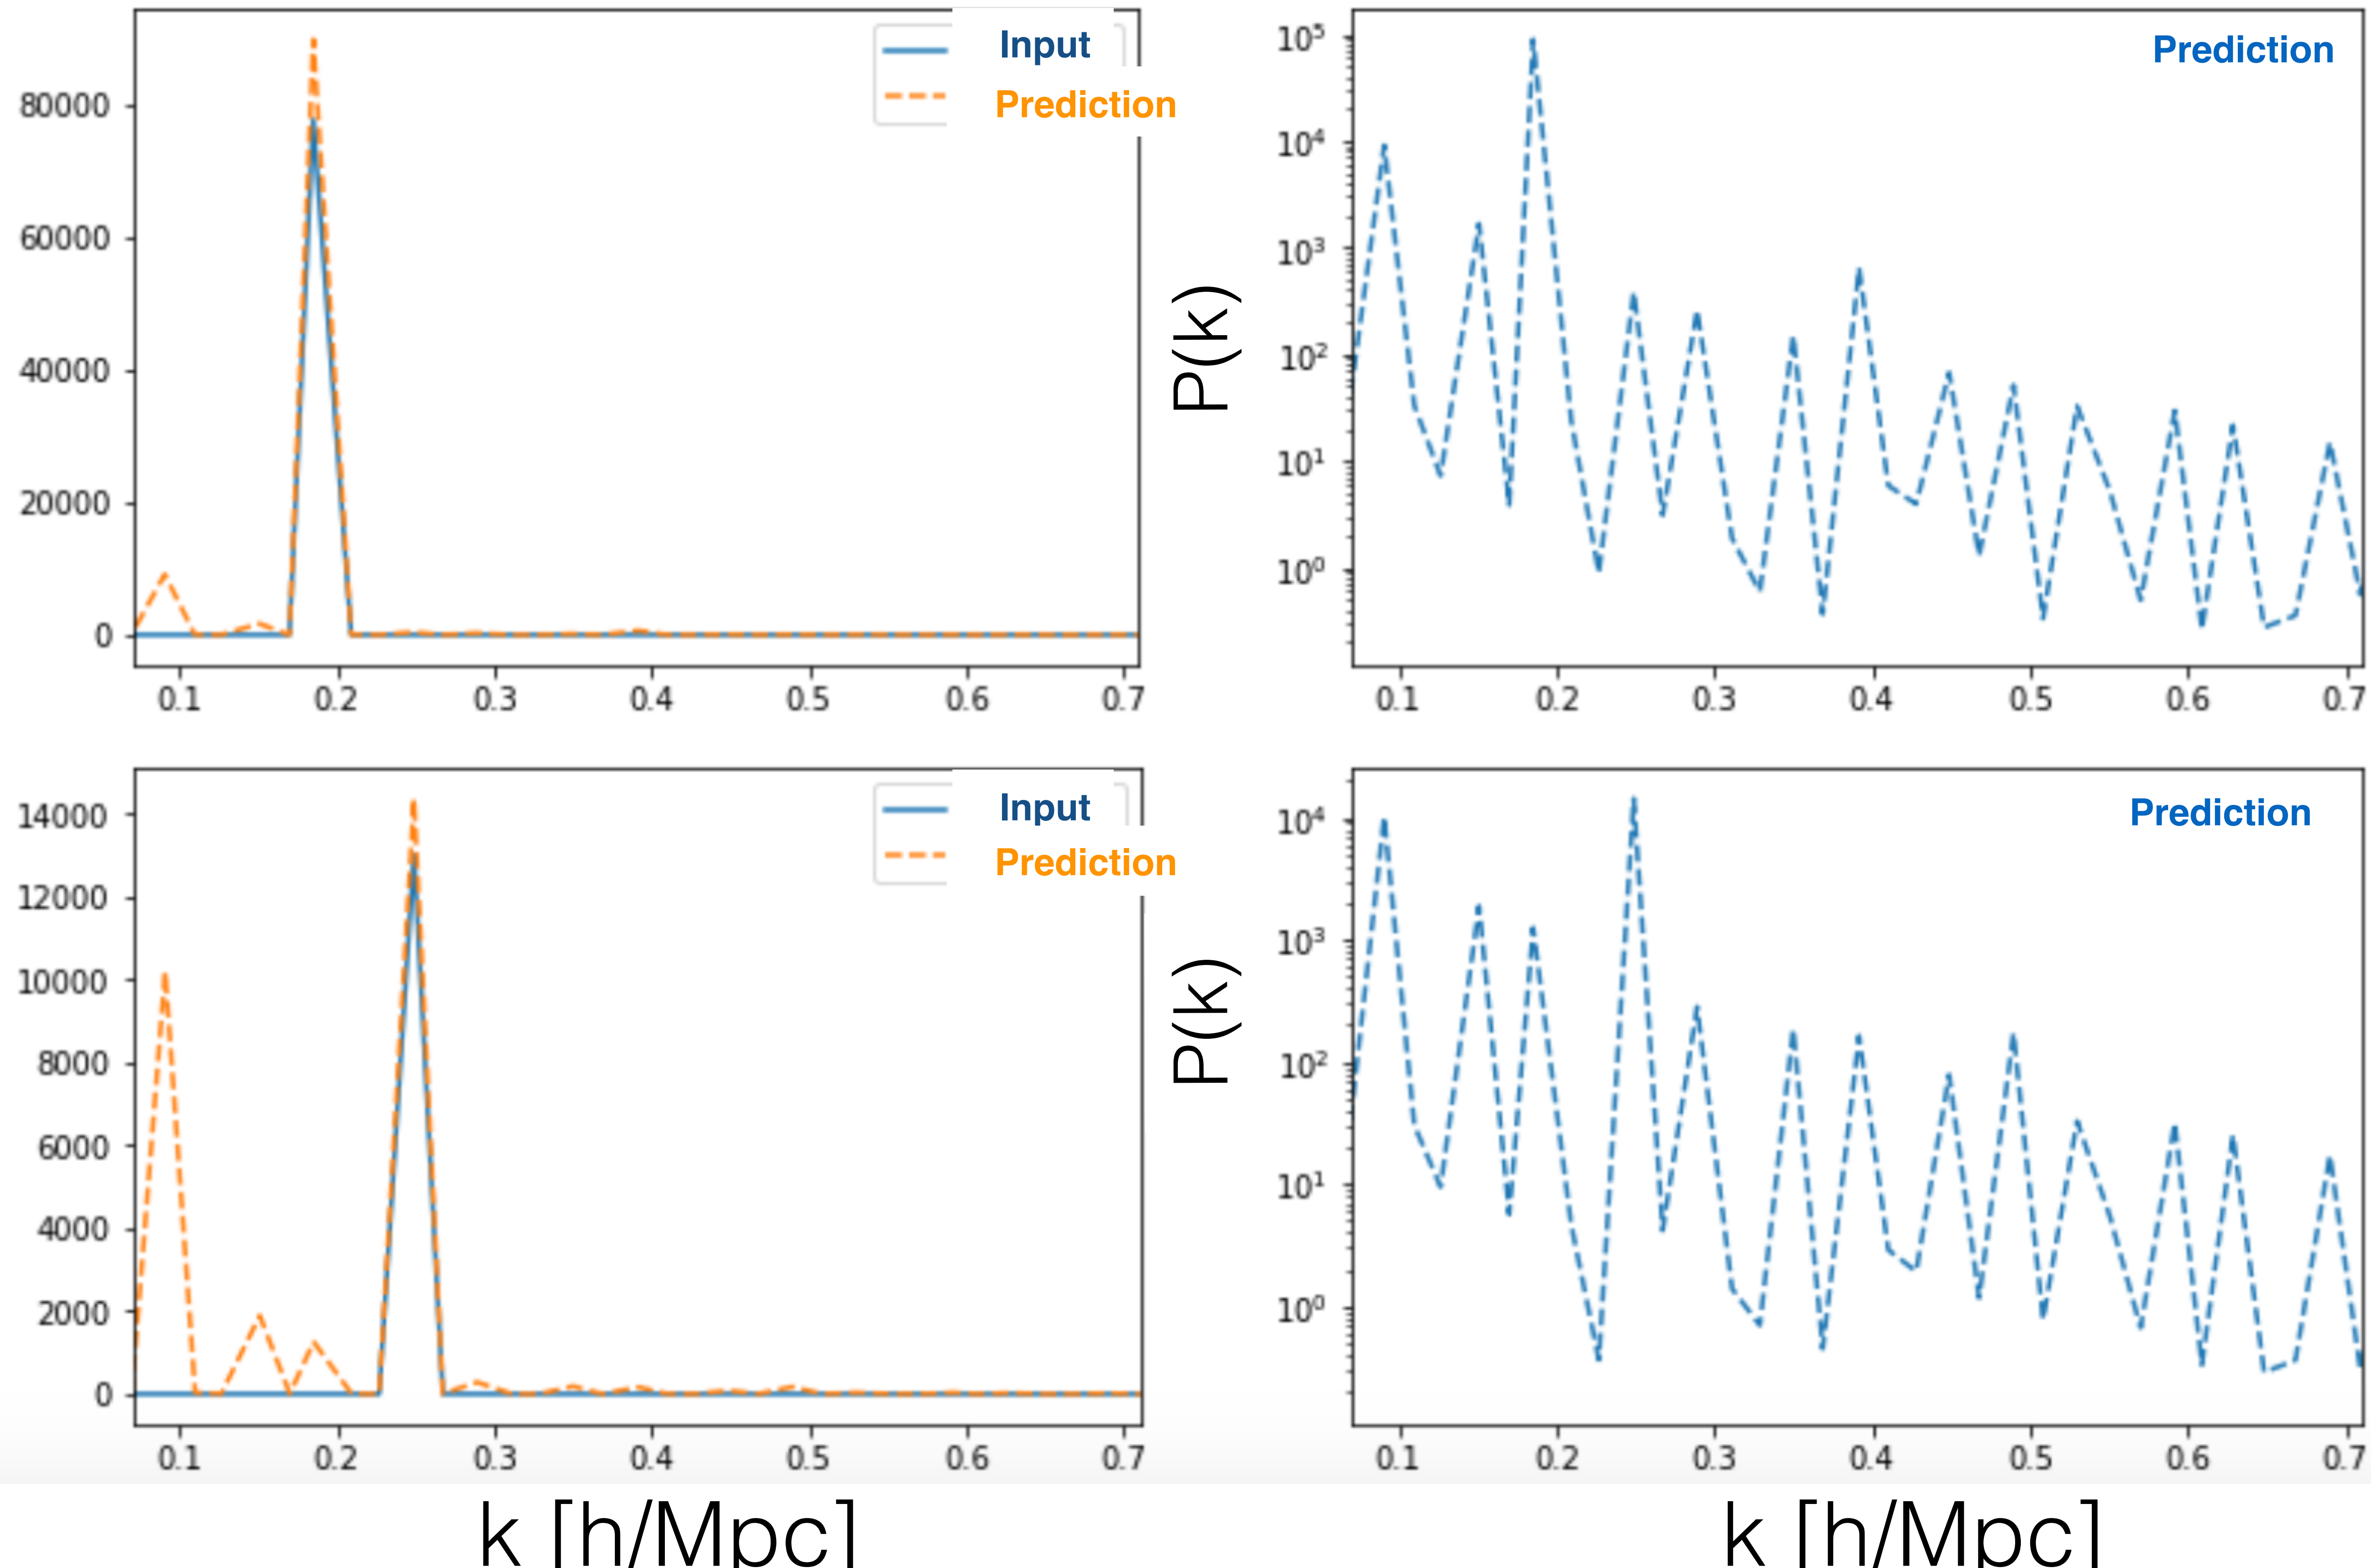

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Is Rotational Invariance learnt by the model?

Aka: If I input same power at modes at k_x, k_y, k_z independently they should give the same power

Is Rotational Invariance learnt by the model?

Yes, *predicted power is the similar no matter which orientatio*n: rotational invariance is learnt!

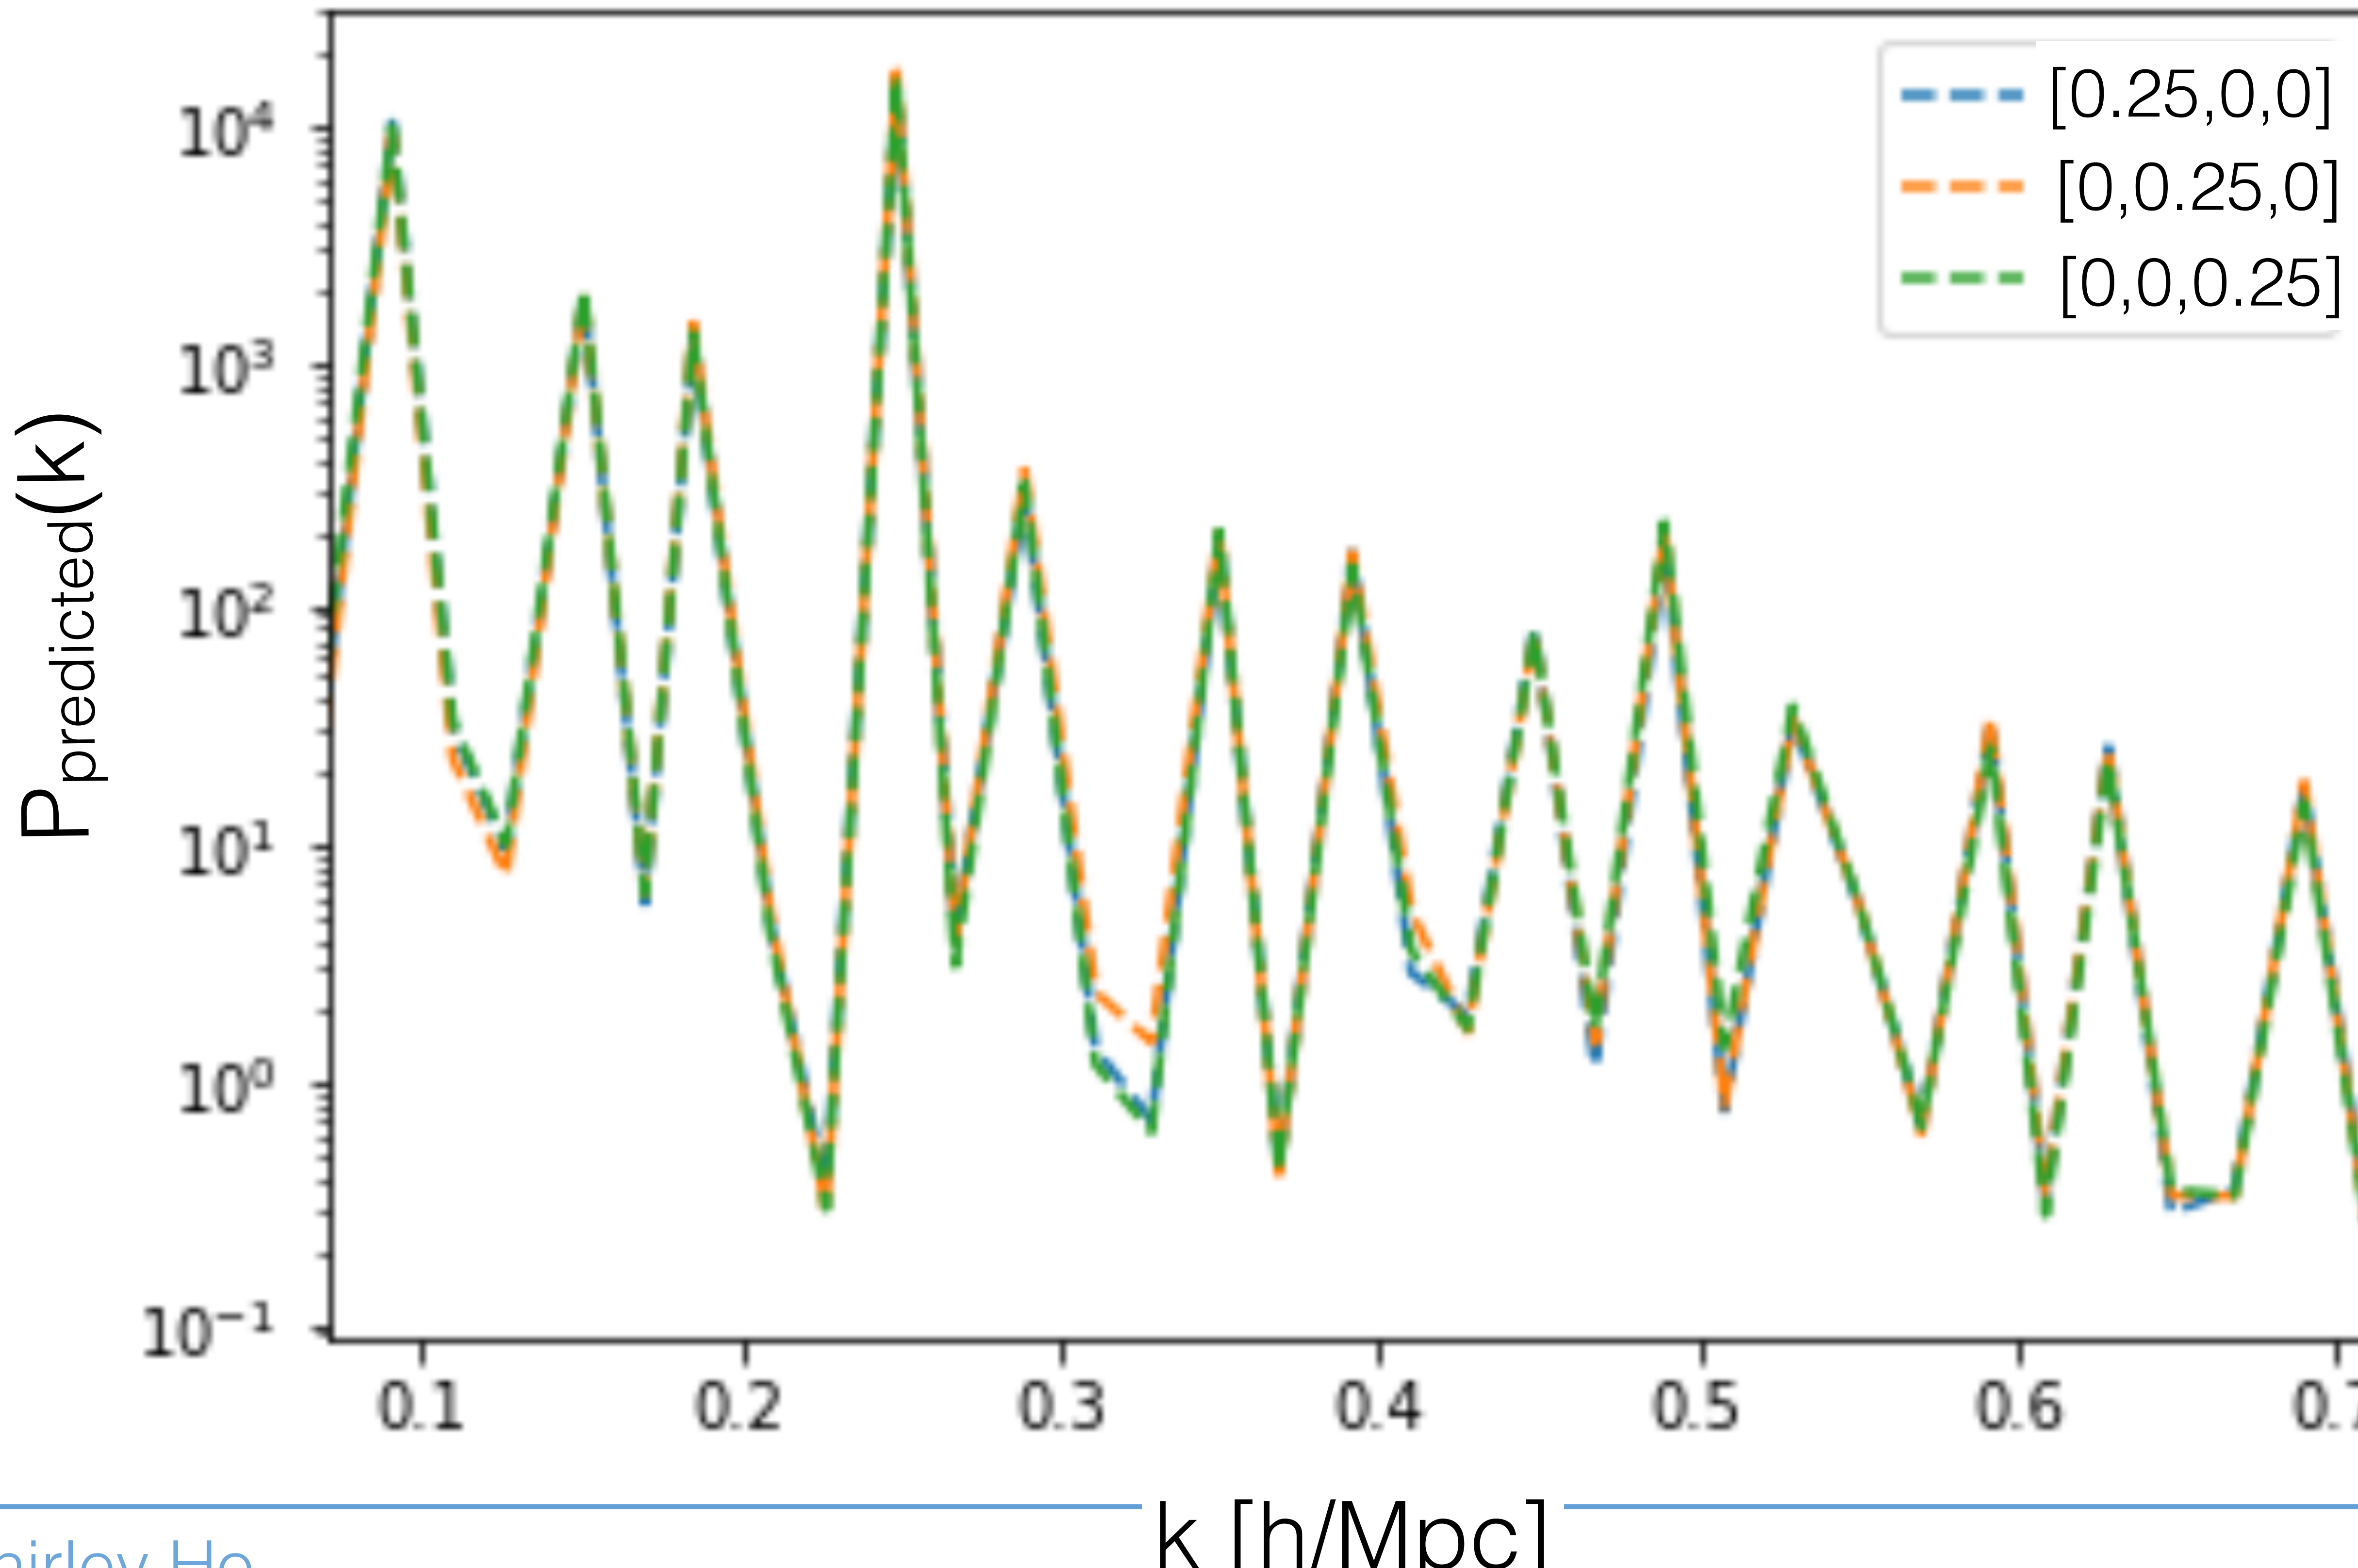

What happens if we change the phase of the input mode?

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

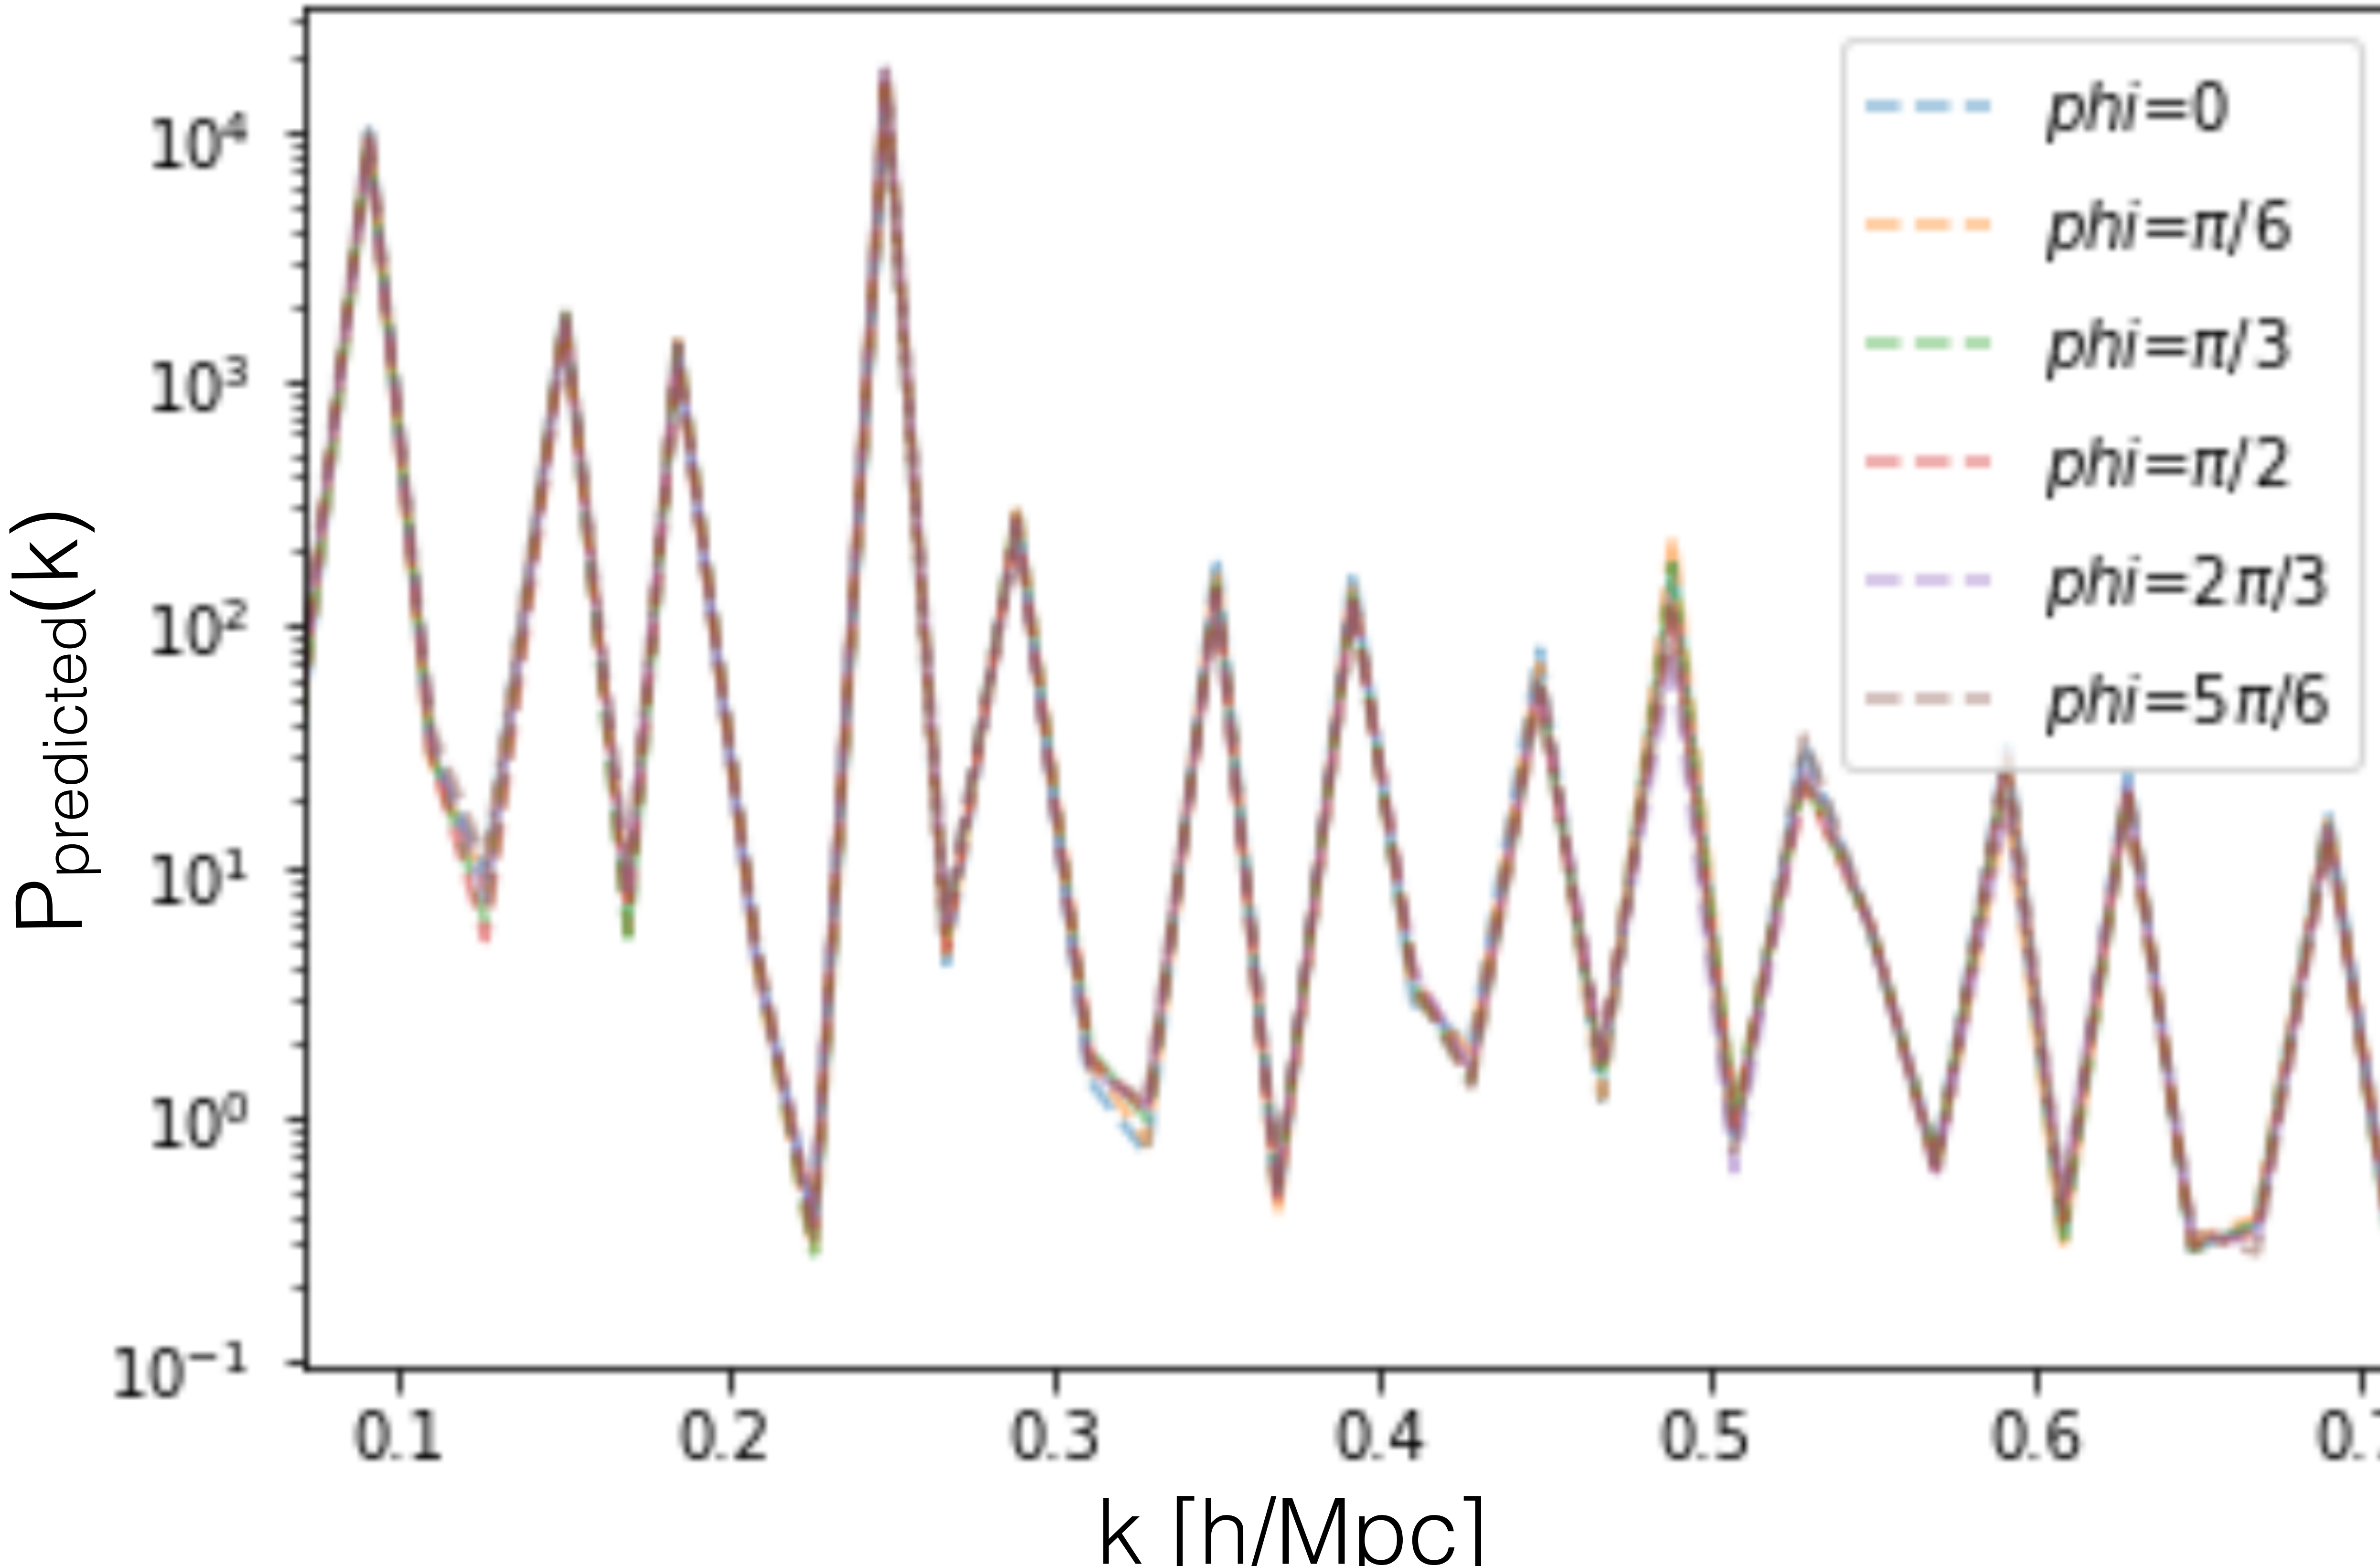

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

What did the model learn so far?

- power at one single mode gives power at many scales
- rotational invariance is learnt by the model
- power at different phases predicts the same power
- slight excess at large scales that are not expected (possible to fix with different models?)

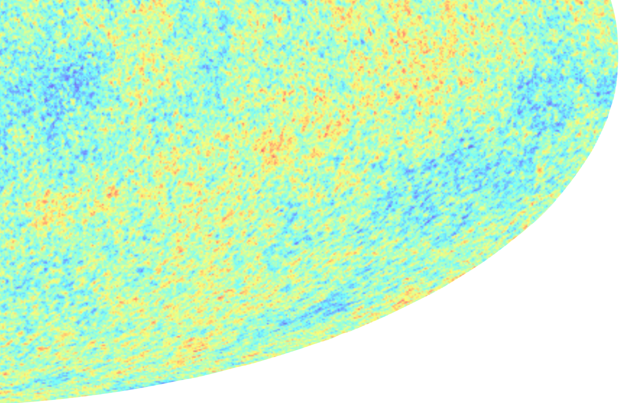

Looking forward

- Improving the models, and see if the excess power at large scale will go away
- Compare what the model has learnt to classical theory (LPT/2LPT/CLPT/EFT..)
- Discover new physics with Machine Learning!/?
- Combine LSS and CMB [with realism] and learn more about our Universe!

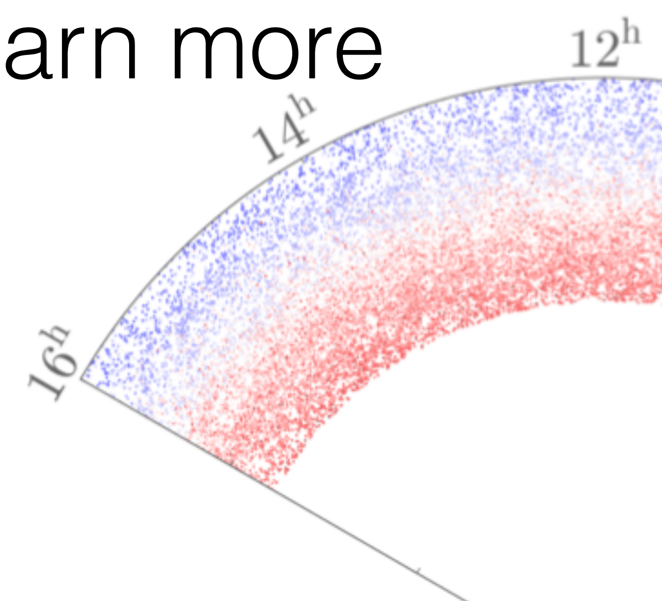

Analytical physics (Zeldovich Approximation/ Lagrangian Perturbation Theory) vs Computer Simulations (N-body)

Input

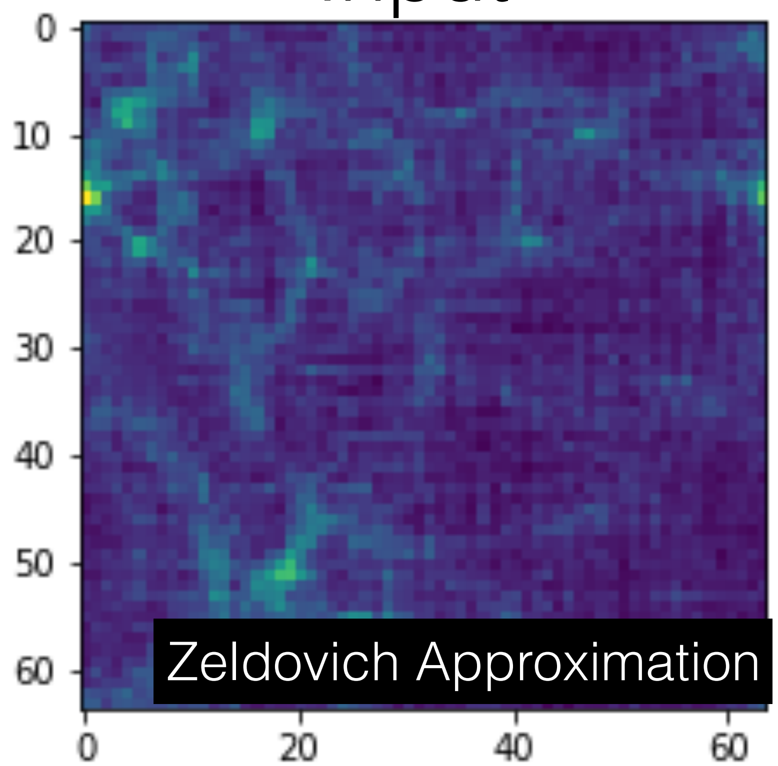

Input

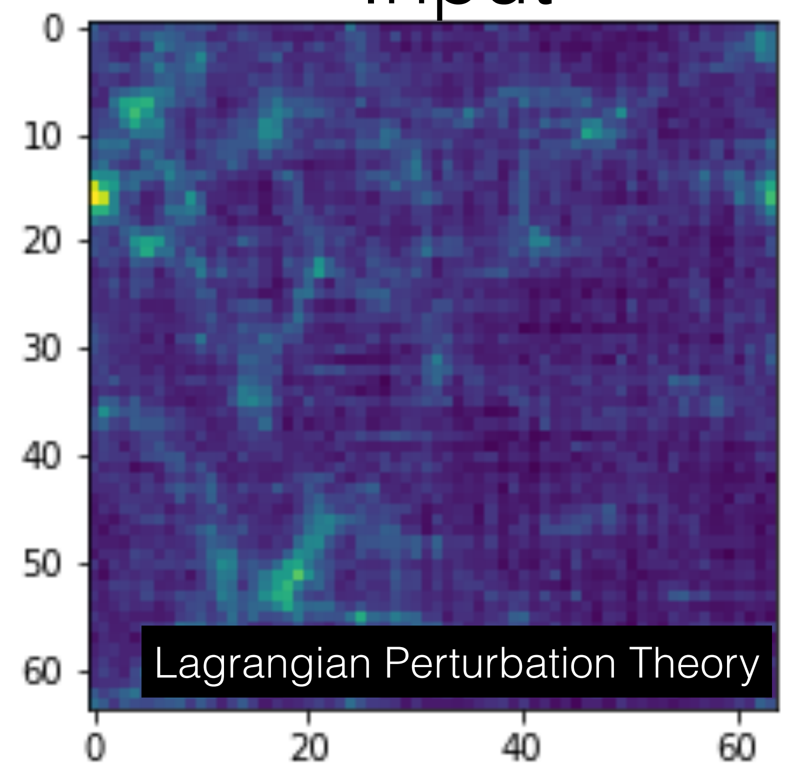

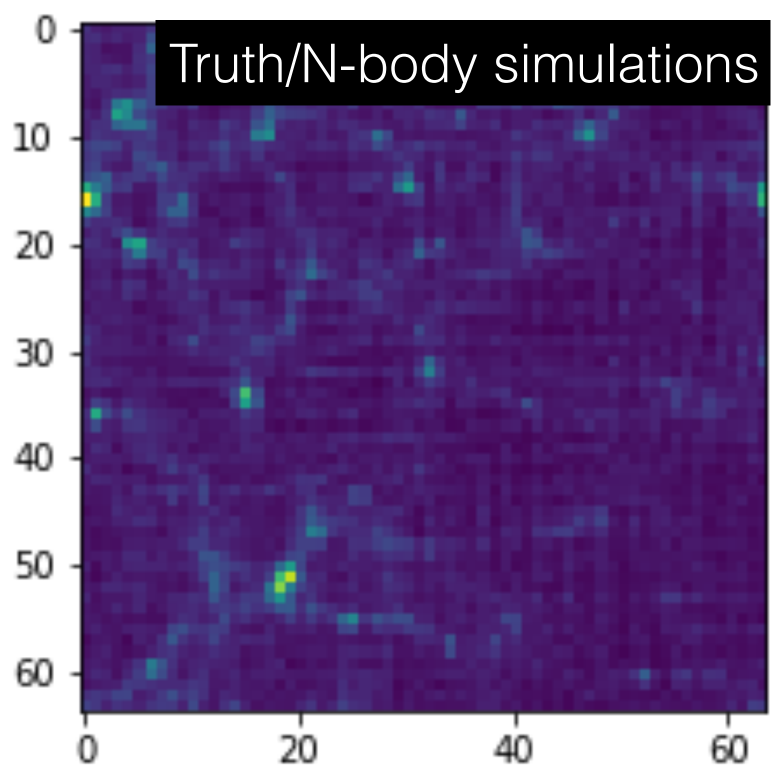

Prediction/Output

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos

Analytical physics (Zeldovich Approximation/ Lagrangian Perturbation Theory) vs Computer Simulations (N-body)

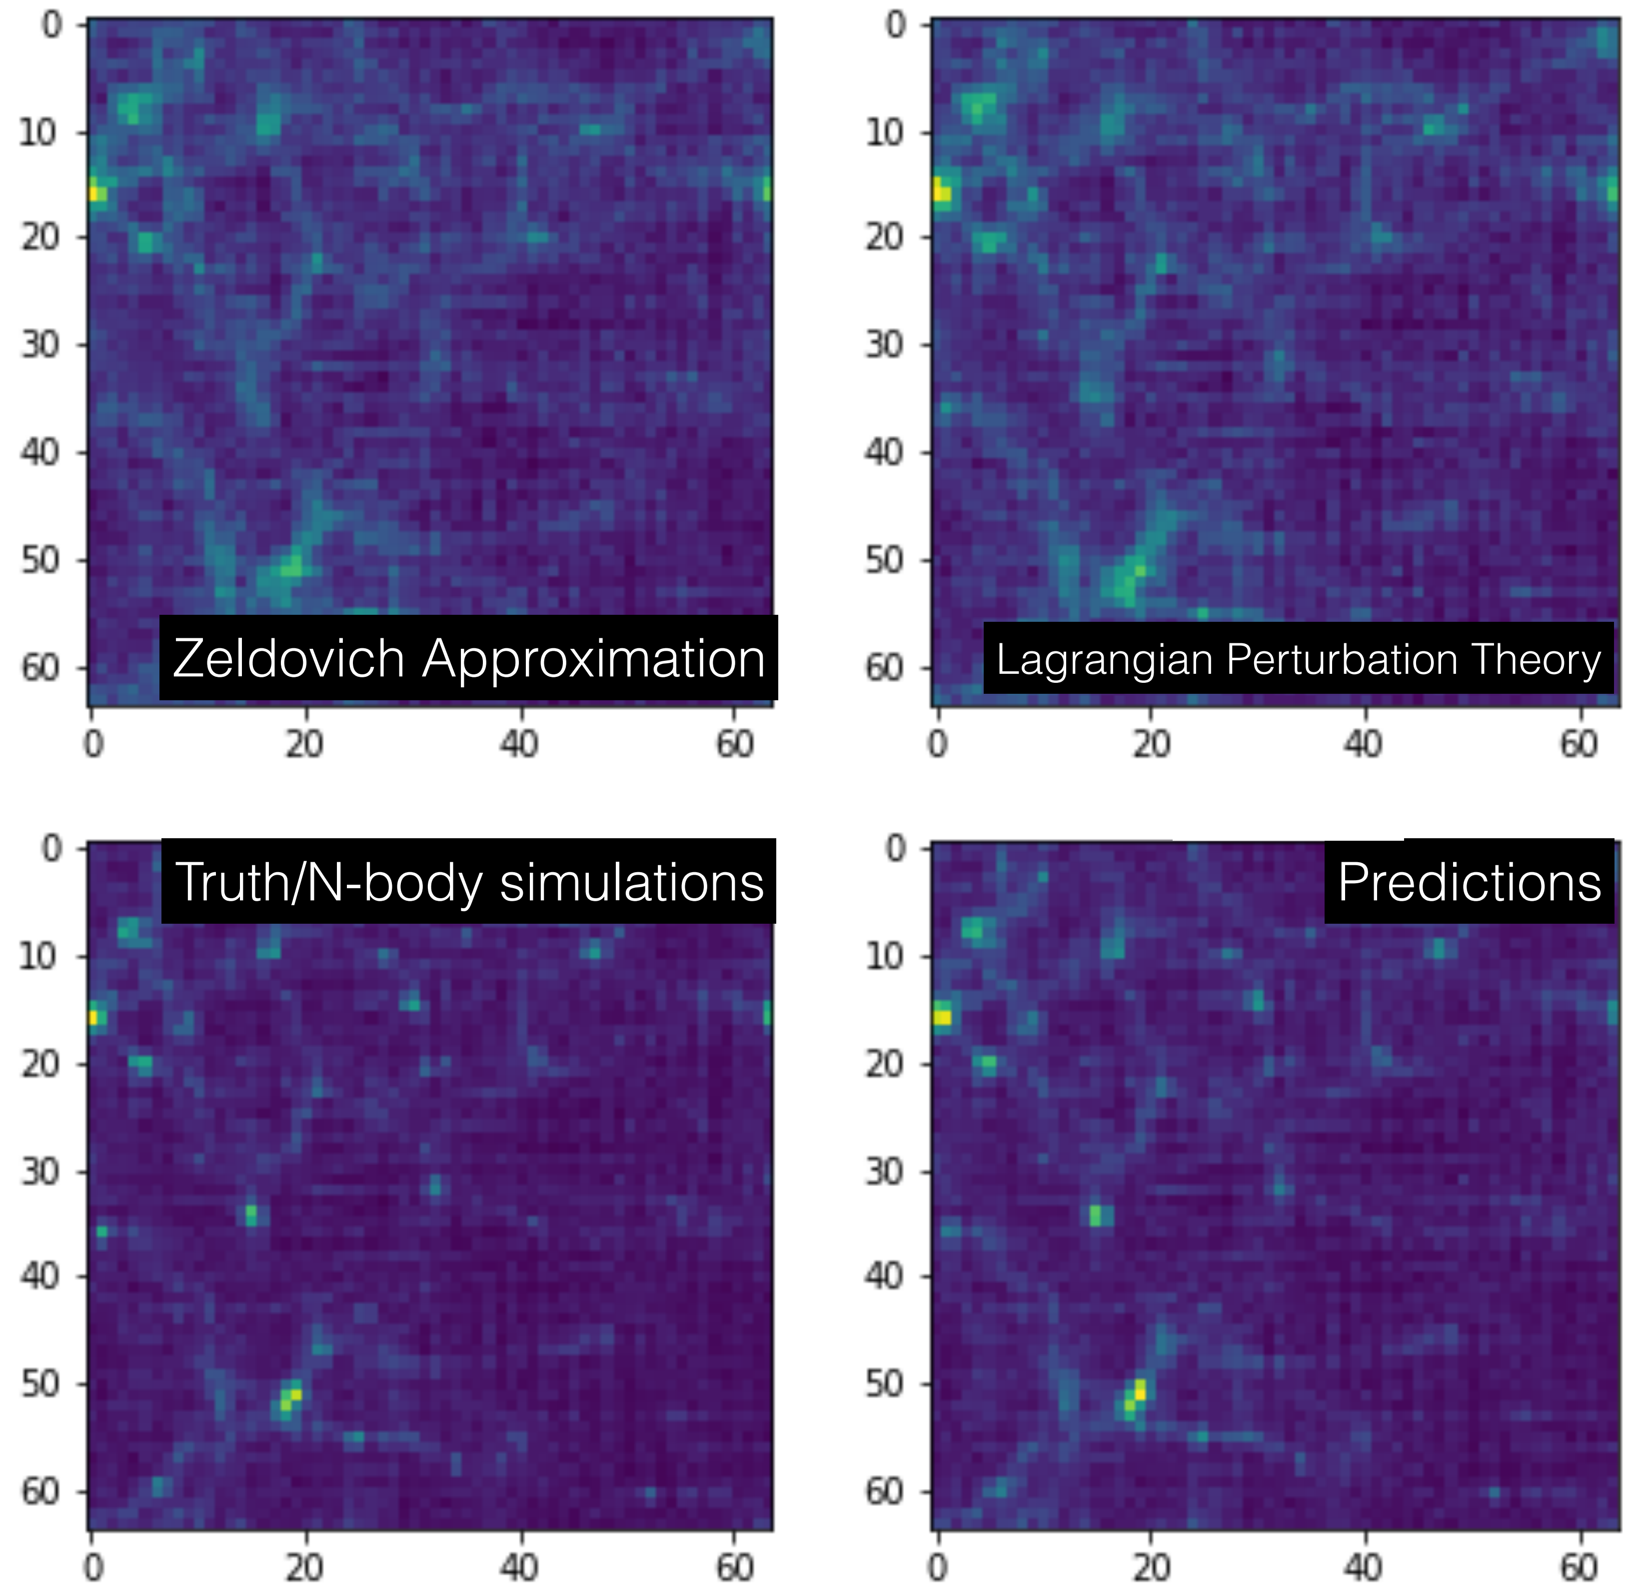

Preliminary results from: Siyu He, Yin Li, Yu Feng, **S.H.**, Siamak Ravanbaksh, Barnabas Poczos