



AI-assisted superresolution cosmological simulations

Yin Li^{a,b,1}, Yueying Ni^{c,d,1}, Rupert A. C. Croft^{c,d}, Tiziana Di Matteo^{c,d}, Simeon Bird^e, and Yu Feng^{f,9}

^aCenter for Computational Astrophysics, Flatiron Institute, Simons Foundation, New York, NY 10010; ^bCenter for Computational Mathematics, Flatiron Institute, Simons Foundation, New York, NY 10010; ^cMcWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213; ^dNSF AI Planning Institute for Physics of the Future, Carnegie Mellon University, Pittsburgh, PA 15213; ^eDepartment of Physics and Astronomy, University of California Riverside, Riverside, CA 92521; ^fBerkeley Center for Cosmological Physics, University of California Berkeley, Berkeley, CA 94720; and ⁹Department of Physics, University of California Berkeley, Berkeley, CA 94720

Edited by Neta A. Bahcall, Princeton University, Princeton, NJ, and approved March 11, 2021 (received for review October 26, 2020)

Cosmological simulations of galaxy formation are limited by finite computational resources. We draw from the ongoing rapid advances in artificial intelligence (AI; specifically deep learning) to address this problem. Neural networks have been developed to learn from high-resolution (HR) image data and then make accurate superresolution (SR) versions of different low-resolution (LR) images. We apply such techniques to LR cosmological N-body simulations, generating SR versions. Specifically, we are able to enhance the simulation resolution by generating 512 times more particles and predicting their displacements from the initial positions. Therefore, our results can be viewed as simulation realizations themselves, rather than projections, e.g., to their density fields. Furthermore, the generation process is stochastic, enabling us to sample the small-scale modes conditioning on the large-scale environment. Our model learns from only 16 pairs of small-volume LR-HR simulations and is then able to generate SR simulations that successfully reproduce the HR matter power spectrum to percent level up to $16 h^{-1} \text{Mpc}$ and the HR halo mass function to within 10% down to $10^{11} M_{\odot}$. We successfully deploy the model in a box 1,000 times larger than the training simulation box, showing that high-resolution mock surveys can be generated rapidly. We conclude that AI assistance has the potential to revolutionize modeling of small-scale galaxy-formation physics in large cosmological volumes.

cosmology | deep learning | simulation | super resolution

As telescopes and satellites become more powerful, observational data on galaxies, quasars, and the matter in intergalactic space becomes more detailed and covers a greater range of epochs and environments in the Universe. Our cosmological simulations (see, e.g., ref. 1) must also become more detailed and more wide-ranging in order to make predictions and test the effects of different physical processes and different dark-matter candidates. Even with supercomputers, we are forced to decide whether to maximize either resolution or volume, or else compromise on both. These limitations can be overcome through the development of methods that leverage techniques from the artificial intelligence (AI) revolution (see, e.g., ref. 2) and make superresolution (SR) simulations possible. In the present work, we begin to explore this possibility, combining knowledge and existing superscalable codes for petascale-plus cosmological simulations (3) with machine learning (ML) techniques to effectively create representative volumes of the Universe that incorporate information from higher-resolution models of galaxy formation. Our first attempts, presented here, involve simulations with dark matter and gravity only, and extensions to full hydrodynamics will follow. This hybrid approach, which will imply offloading simulations to neural networks (NNs) and other ML algorithms, has the promise to enable the prediction of quasar, supermassive black hole, and galaxy properties in a way that is statistically identical to full hydrodynamic models, but with a significant speed-up.

Adding details to images below the resolution scale (SR image enhancement) has become possible with the latest advances in deep learning (DL; ML with NN; ref. 4), including generative

adversarial networks (GANs; ref. 5). The technique has applications in many fields, from microscopy to law enforcement (6). It has been used for observational astronomical images by (7), to recover galaxy features from below the resolution scale in degraded Hubble Space Telescope images. Besides SR image enhancement, DL has started to find applications in cosmological simulations. For example, refs. 8 and 9 showed how NNs can predict the nonlinear formation of structures given simple linear theory predictions. NN models have also been trained to predict galaxies (10, 11) and 21-cm emission from neutral hydrogen (12) from simulations that only contain dark matter. GANs have been used in ref. 13 to generate image slices of cosmological models and to generate dark-matter halos from density fields (14). ML techniques other than DL find many applications, too. For example, Kamdar et al. (15) have applied extremely randomized trees to dark-matter simulations to predict hydrodynamic galaxy properties.

Generating mocks for future sky surveys requires large volumes and high accuracy, a task that quickly becomes computationally prohibitive. To alleviate the cost, recently, Dai and Seljak (16) developed a Lagrangian-based parametric ML model to predict various hydrodynamical outputs from the dark-matter density field. In other work, Dai et al. (17, 18) sharpened the particle distribution using a potential gradient descent method starting from low-resolution (LR) simulations. Note, however, that these approaches did not aim to enhance the spatial or mass resolution of a simulation.

Significance

Cosmological simulations are indispensable for understanding our Universe, from the creation of the cosmic web to the formation of galaxies and their central black holes. This vast dynamic range incurs large computational costs, demanding sacrifice of either resolution or size and often both. We build a deep neural network to enhance low-resolution dark-matter simulations, generating superresolution realizations that agree remarkably well with authentic high-resolution counterparts on their statistical properties and are orders-of-magnitude faster. It readily applies to larger volumes and generalizes to rare objects not present in the training data. Our study shows that deep learning and cosmological simulations can be a powerful combination to model the structure formation of our Universe over its full dynamic range.

Author contributions: Y.L., Y.N., R.A.C.C., T.D.M., S.B., and Y.F. designed research; Y.L. and Y.N. performed research; Y.L. and Y.N. contributed new reagents/analytic tools; Y.L. and Y.N. analyzed data; and Y.L., Y.N., R.A.C.C., T.D.M., S.B., and Y.F. wrote the paper.

The authors declare no competing interest.

Published under the [PNAS license](#).

This article is a PNAS Direct Submission.

¹To whom correspondence may be addressed. Email: eelregit@gmail.com or yueyingn@andrew.cmu.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2022038118/-DCSupplemental>.

Published May 4, 2021.

On the DL side, recently, Ramanah et al. (19) explored using the SR technique to map density fields of LR cosmological simulations to that of the high-resolution (HR) ones. While the goal is similar, our work has the following three key differences. First, instead of focusing on the dark-matter density field, we aim to enhance the number of particles and predict their displacements, from which the density fields can be inferred. This approach allows us to preserve the particle nature of the N-body simulations and therefore to interpret the SR outputs as simulations themselves. Second, we test our technique at a higher SR ratio. Compared to ref. 19, which increased the number of Eulerian voxels by 8 times, we increase the number of particles and thus the mass resolution by a factor of 512. Finally, to facilitate future applications of SR on hydrodynamic simulations in representative volumes, we test our method at much smaller scales and in large simulations whose volume is much bigger than that of the training data.

Results

In this work, we employ DL, the class of ML algorithms involving artificial NNs, to progressively transform an LR input simulation via multiple layers of neurons into an output that statistically reproduces the HR target. We build and train a DL model based on physical considerations and use the displacements of particles in the N-body simulations as inputs and outputs. This allows us to interpret the results as simulation realizations themselves. To help train the generative network, we employ a GAN approach to simultaneously train a discriminative network that contests with the generator in a game by evaluating its outputs. We validate our model by comparing the generated SR simulations to the authentic HR simulations, both visually and quantitatively, using summary statistics, including the power spectrum and the halo-mass function.

Visual Comparison. First, let us visually compare the generated SR results to the authentic HR simulations. We obtain the SR particle positions by moving them from their original grid locations using the predicted displacements. In the three columns of Fig. 1, respectively, we show the LR, HR, and SR dark-matter density fields at $z=0$. We visualize all dark-matter particles in blue and highlight the Friends-of-Friends (FOF) halos (see below) in orange.

As shown in Fig. 1, our SR technique is able to recover sharp features and small-scale structures starting from inputs of extremely poor resolution. It forms halos where the input LR cannot resolve them; e.g., in the figure, the LR can only form the most massive handful of halos in the box, whereas SR is able to improve the halo mass range by orders of magnitude. The SR outputs trace the large-scale environment set by the LR inputs, while appearing statistically almost indistinguishable from the HR targets on the small scales.

Our method also allows us to map the same LR input to multiple outputs. We achieve this by adding noise between layers in the generator, as explained in *Materials and Methods*. In Fig. 2, we show one HR and three random SR projections, of the same $(20 h^{-1} \text{Mpc})^3$ regions and the same LR field, at $z=2$. As expected, above and around the input-resolution scales, the HR and SR fields are consistent with each other, while being apparently different, but visually similar, on smaller scales. To demonstrate their remarkable similarities, we invite the readers to guess which panel is HR before checking the answer in the footnote of Fig. 2.

The visual tests have demonstrated that our SR technique is able to not only generate small-scale structures that statistically resemble the target, but also sample them conditioning on their larger-scale environment. These capabilities make our method extremely promising for a variety of other SR applications.

Power Spectrum. The matter power spectrum, which describes the amplitude of density fluctuations as a function of scale, is perhaps the most commonly used summary statistic in cosmology. It describes in Fourier space the two-point correlation, which completely specifies the statistical properties of Gaussian random fields. And, to the best of our knowledge, the cosmological density field is well represented after the Big Bang by a Gaussian random field and remains so in the linear growth regime. Therefore, it is standard to compare the SR and HR simulations on their matter power spectra $P(k)$.

For a periodic simulation box, the matter power spectrum can be measured by a discrete Fourier transform

$$P(k_i) = \frac{1}{N_i} \sum_{k_i < k \leq k_{i+1}} \frac{|\delta(\mathbf{k})|^2}{V}, \quad [1]$$

where V is the simulation volume and $\delta(\mathbf{k})$ is the discrete Fourier transform of the overdensity field $\delta(\mathbf{x}) \equiv \rho(\mathbf{x})/\bar{\rho} - 1$, with $\rho(\mathbf{x})$ and $\bar{\rho}$ being the matter density field and its mean value, respectively. Due to statistical isotropy, the power spectrum is only a function of the wavenumber $k = |\mathbf{k}|$ and independent of the direction of \mathbf{k} . The above estimator has already exploited this by averaging spherically within each k -bin. N_i is the number of modes falling in the i -th bin $(k_i, k_{i+1}]$.

We compute the density field by assigning the particle mass to a 512^3 mesh using the CIC (Cloud-in-Cell) scheme, for all LR, HR, and SR simulations. A common practice in power-spectrum estimation is deconvolution of the resampling window (here, CIC) after the discrete Fourier transform (20). However, this amplifies the noises and leads to artifacts in the LR results. Instead, using a large fast Fourier transform (FFT) grid for both resolutions, we avoid deconvolving the resampling window and can compare power spectra from different resolutions on an equal footing.

In Fig. 3, we compare the dimensionless power spectra $\Delta^2(k) \equiv k^3 P(k)/2\pi^2$, at $z=4, 2$, and 0. The vertical dashed lines mark the Nyquist wavenumber $k_{\text{Nyq}} = \pi N_{\text{mesh}}/L_{\text{box}}$ with $N_{\text{mesh}} = 512$ and $L_{\text{box}} = 100 h^{-1} \text{Mpc}$. The monotonically increasing $\Delta^2(k)$ is a good proxy for the variance of matter density as a function of scale, and thus a useful indicator that divides the linear and nonlinear (toward increasing k) regimes, by k_{NL} , where $\Delta^2(k_{\text{NL}}) = 1$.

Due to limited mass resolution, LR power quickly vanishes on small scales at $z=4$. This deficit persists and is partly compensated by formation of poorly resolved halos at $z=2$ and $z=0$. On the other hand, the SR power spectra successfully match the HR results to percent level to $k_{\text{max}} \approx 16 h^{-1} \text{Mpc}$ at all redshifts, a dramatic improvement over the LR predictions. This is remarkable, considering the fact that the SR model fares equally well from the linear to the deeply nonlinear regimes. While the model can learn and compare on even smaller scales using a larger FFT grid (and a larger grid for density field as input to the discriminator, as described in *Materials and Methods*), the choice of $k_{\text{max}} \approx 16 h^{-1} \text{Mpc}$ is justified by the fact that the model has reached an accuracy level comparable to that of the N-body simulations.

Halo Mass Function. Beyond the two-point correlation test—i.e., the power spectrum—we compare higher-order statistics of the fields that quantify the non-Gaussianity arising in the nonlinear regime. As the most nonlinear dark-matter structures, halos host most of the luminous extragalactic observables in the sky. Therefore, halo abundance is a natural choice when considering non-Gaussian statistics. We find the dark-matter halos in the simulations using an FOF halo finder (21) with linking-length parameter $b=0.2$. The FOF algorithm links all pairs of particles within b times their mean separation and collects each group of




Fig. 1. Two-dimensional projections of the LR, HR, and SR dark-matter density fields at $z=0$. The blue background shows the smoothed density field of all of the dark-matter particles. The particles in FOF groups are highlighted in orange to help visually identify the halos. The top images show slabs from the full box of $100 h^{-1}\text{Mpc}$ side length and $20 h^{-1}\text{Mpc}$ thickness. The middle images zoom into the orange boxes (A) from the top image, each of size $(20 h^{-1}\text{Mpc})^3$. The bottom two rows show the four zoom-in boxes (B–E), which are $(10 h^{-1}\text{Mpc})^3$ in size, to reveal even finer details. The first two columns show the LR and HR simulations, which have the same initial conditions, but a factor of $512\times$ different mass resolution. The rightmost column shows one of the SR realizations generated by our trained model. All density projections are smoothed by a Gaussian filter on a scale of $5 h^{-1}\text{kpc}$, using `gaepsi2` (<https://github.com/rainwoodman/gaepsi2>).

connected particles into one halo. We only keep halos with at least 32 particles.

In Fig. 4, we compare the halo mass functions with a $(100 h^{-1}\text{Mpc})^3$ test set simulation, at $z=4, 2$, and 0. The mass function is defined as $\phi \equiv d n / d \log_{10} M_h$, where n is the comoving number density of halos above threshold mass M_h . Poisson errors on the halo abundance are shown. These are conservative bounds because there are only small-scale differences between the HR and SR realizations (Fig. 2) due to their conditioning

on the same LR field. Due to their large particle mass and low force resolution, the LR simulations can only resolve the most massive halos above $10^{13} M_\odot$, forming no halos at $z=4$. The LR halo abundances are also underpredicted near the mass cut, with deviation significantly larger than the Poisson error estimates.

Using our GAN model, the SR field generated from the LR input has the same mass resolution as the HR field, resolving halos over the whole mass range down to $10^{11} M_\odot$. The generated SR fields predict halo populations accurately

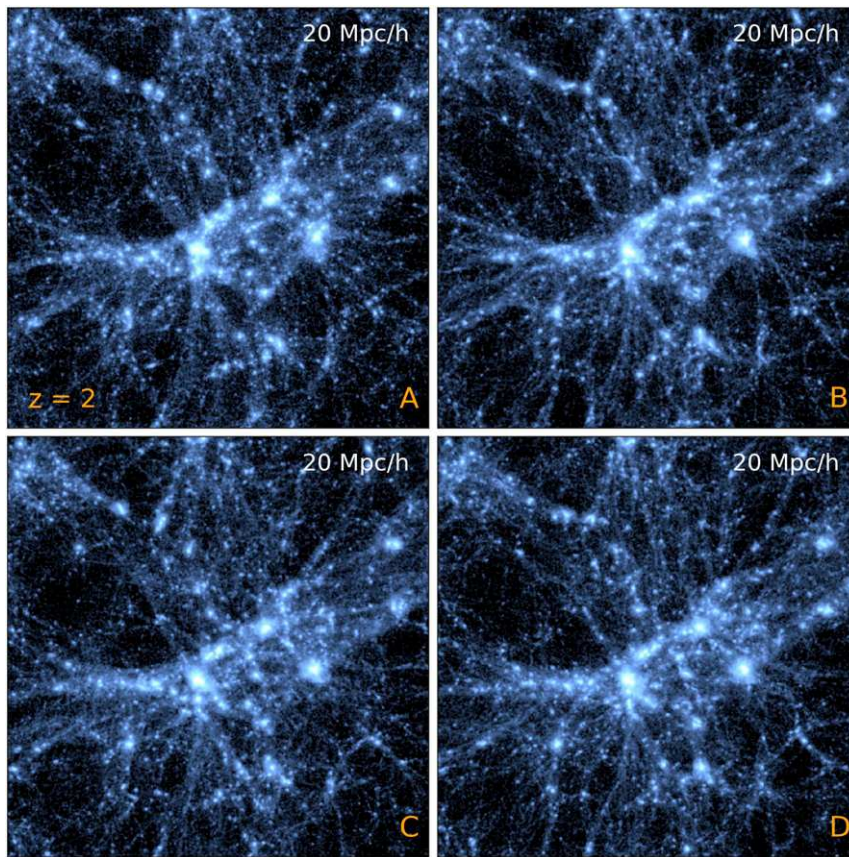


Fig. 2. Using our GAN-based algorithm, we can generate different SR realizations from the same LR inputs. We show one HR and three random SR projections, of the same $(20 h^{-1} \text{Mpc})^3$ regions and the same LR field, at $z = 2$. To demonstrate their remarkable similarities, we intentionally omit the HR and SR labels in panels A–D and invite the readers to guess which panel is HR before checking the answer in the footnote.* The HR and three SR realizations have the same large-scale structures, but are all different in their small-scale features. Yet, they appear statistically indistinguishable. The color scheme and smoothing method are the same as in Fig. 1.

*The answer to Fig. 2 is A.

compared with the HR results, to 10% level throughout the mass range.

Application to a Large Volume

As described in *Materials and Methods*, we train our GAN model on fields cropped from $(100 h^{-1} \text{Mpc})^3$ volume simulations. In all previous results, our trained model has proven to work well on a new SR realization of the same volume. Moreover, translational symmetry has been carefully preserved by our cropping and padding schemes. This implies that our method can be applied to any cosmological volumes. The additional long-wavelength modes in the large volume lead to rare massive structures that the model has not encountered in the training set. We apply our generative model to a $(1 h^{-1} \text{Gpc})^3$ LR simulation with 640^3 particles, a volume 1,000 times larger than the training-set simulations, and obtain an SR realization with $5,120^3$ particles. The generating process takes only about 16 h on one graphics-processing unit (GPU), including time for input/output (I/O). This consumes significantly less computing resources than partial differential equation-based N-body solvers.

Fig. 5 illustrates the generated $(1 h^{-1} \text{Gpc})^3$ SR result at $z = 2$, teeming with finer details throughout the whole volume compared to the same LR field. We zoom in to inspect a same halo in the SR and LR fields. It has a mass of $2 \times 10^{14} M_{\odot}$, larger than the most massive halo in the $(100 h^{-1} \text{Mpc})^3$ training sets. Remarkably, even though our GAN model has not been trained upon such massive halos, it performs well and generates them

with reasonable morphology, apparent substructure, and correct population density (see below).

Quantitatively, we identify FOF halos in both the LR and SR large volumes and compare their resultant halo mass functions in Fig. 6. Again, the LR halos can only be well resolved above $10^{13} M_{\odot}$, while the SR halo abundance agrees remarkably well with both the $(1 h^{-1} \text{Gpc})^3$ LR and the previous $(100 h^{-1} \text{Mpc})^3$ HR mass functions. The SR prediction even has the right abundance for the massive halos above $10^{14} M_{\odot}$, which are too rare to be present in the training volume. Our model generalizes and extrapolates surprisingly well to predict structure formations on scales over which it has not been directly trained.

This exercise demonstrates the potential of our method to tackle the fundamental challenge of constructing large-volume cosmological simulations at otherwise unfeasibly high resolution, with a dramatically lower footprint in computing resources. It opens up an avenue for constructing large-volume cosmological (eventually hydrodynamical) simulations and creating large ensembles of associated mocks at scales comparable to the current and future surveys, for maximal scientific return.

Discussion

We have shown how a GAN SR architecture inspired by StyleGAN2 can be used to enhance cosmological simulations so that they reproduce the appearance and statistics of much higher-resolution models. Our approach has similarities to the




Fig. 3. Dimensionless matter power spectrum $\Delta^2(k)$ comparison on LR (purple), HR (blue), and SR (black) test realizations, at $z = 4, 2,$ and 0 . The vertical dashed lines mark the Nyquist wavenumber. LR power quickly vanishes on small scales at $z = 4$. This deficit persists and is partly compensated by formation of poorly resolved halos at $z = 2$ and $z = 0$. The SR result is a dramatic improvement, with the SR power spectra matching the HR curves remarkably well, within a few percent on all scales at all redshifts.

work of Ramanah et al. (19), who showed how SR dark-matter density fields can be generated. Among the differences with this work are our use of particle displacements as the inputs and outputs of our modeling. By generating a field of SR particle displacements, we effectively create a whole simulation, which can be analyzed (for example, by carrying out halo finding) in the same manner as a full simulation run at the higher resolution. This in contrast with methods, such as ref. 19 or 13, which generate a density field.

One can ask whether there is fundamental limit to the level of SR enhancement that our approach is able to produce. In principle, one should be able to generate simulations completely using a GAN, without an LR model for conditioning, e.g., refs. 13 and 22. Logically, such an approach would need proportionally more training data to produce satisfactory results, as it would be gen-

erating large-scale modes as well. At the opposite end, Ramanah et al. (19) have added SR details to their simulations by enhancing by a factor of two in length scale. In this paper, we have shown that our SR modeling works well to produce small-scale structures a factor of 8 in length (and 512 in mass) below the LR scale. At these scales, one can begin to find bound structures, such as subhalos.

The main advantage in the use of SR simulations over full HR runs is the potentially huge reduction in the computational resources involved. As an example, for the $(100 h^{-1} \text{Mpc})^3$ volume simulations of our test set, it takes approximately 560 core hours to run a 512^3 particle HR model to $z = 2$, while the 64^3 -particle LR run takes only about 0.6 core hours, a factor of $\sim 1,000$ faster. We can achieve even more dramatic speed up when applying the SR model to a larger volume. For a




Fig. 4. FOF halo mass function comparison among LR (purple), HR (blue), and SR (black) test simulations, at $z = 4, 2,$ and 0 . The LR simulations can only resolve the most massive halos ($M_h \gtrsim 10^{13} M_\odot$). Our generated SR fields have the same mass resolution as the HR fields, resolving halos all of the way down to $10^{11} M_\odot$. Their halo populations closely match the HR results, at the 10% level throughout the mass range.




Fig. 5. Illustration of the $(1 h^{-1} \text{Gpc})^3$ volume SR with 5120^3 particles generated from the input 640^3 LR field using our GAN model. Colors are the same as in Fig. 1. The large panel shows a slice through the full box, $1 h^{-1} \text{Gpc}$ in length and $20 h^{-1} \text{Mpc}$ in thickness. The top left inset in the white box shows the same field from the LR simulation. In the two red inset panels, we show a $(20 h^{-1} \text{Mpc})^3$ zoom-in region around a massive halo in both the SR and LR field. That central halo has a mass of $2 \times 10^{14} M_{\odot}$, larger than the most massive halo in the $(100 h^{-1} \text{Mpc})^3$ volume training set. The generating process only takes about 16 h on one GPU.

$(1 h^{-1} \text{Gpc})^3$ cosmological volume as shown above, only 500 core hours are needed to run a 640^3 -particle LR simulation to $z = 2$. To run a $5,120^3$ HR counterpart with an N-body code would be daunting and require dedicated supercomputing resources. Our trained GAN model, on the other hand, only takes 16 h with 1 GPU (including the I/O time) to generate a $5,120^3$ SR field for the $(1 h^{-1} \text{Gpc})^3$ volume, a tiny fraction of the cost of the HR counterpart. An additional advantage, made apparent by our processing the simulations in distinct chunks, is that the SR enhancement can be applied where it is needed—e.g., in the vicinity of a specific galaxy cluster or supermassive black hole. The data storage required can therefore also be much less than for a full HR run.

The stochastic nature of the StyleGAN implementation is also an interesting feature that can be exploited. As we have seen in Fig. 2, it is possible to sample multiple “realizations” of the small-scale clustering, conditioning on the large-scale modes, by varying the input noise component. This opens up the possibility to improve the statistical inference of cosmology from the small-scale clustering of galaxies, by jointly sampling the small-scale modes with the cosmological parameters.

To train and test our model, we have used simulations of the same cosmology, and this simplifies the learning task. It is expected that the LR to HR mapping should depend weakly on cosmology, and introducing such dependence should make the model adaptable and more accurate for a range of parameters spanned by the training set. Furthermore, to fully capture the state of a simulation, our model can be used to learn particle

velocities as well. We leave those cosmology dependence and velocity improvements for future work.

A central question for GANs (and DL in general) is how to determine if enough training data are being used. Different diagnostics can be employed, and the dependence of accuracy on training set size can be evaluated. We have seen that our SR simulations are at least able to reproduce the power spectrum and halo mass function well with a very limited set of training data. A related issue is the fact that in astrophysics, we often deal with rare objects (e.g., quasars or galaxy clusters), and there may not even be a single example in the training data. In the future, one should examine how best to assemble a training dataset that sufficiently covers all cases that will be studied. As an example, separate universe or constrained simulations—e.g., refs. 23–25—could be used to enhance the number of high-mass objects in the training data.

So far, though, we have demonstrated that our method is capable of generating SR simulations 1,000 times larger than the training sets. Our large SR simulation does successfully reproduce the halo mass function and power spectrum over several orders of magnitude. It also leads to the formation of visually reasonable large and rare galaxy clusters. Future work will further evaluate the quality of generated rare objects using other measures.

Particularly in the way we have generated them (as full particle-position datasets), the SR simulations could be used for many purposes that would have needed HR simulations. These include mock galaxy catalogs for large-scale structure




Fig. 6. Halo mass function comparison of the $(1 h^{-1}\text{Gpc})^3$ LR (blue), the $(1 h^{-1}\text{Gpc})^3$ SR (black and generated from the former), and the $(100 h^{-1}\text{Mpc})^3$ HR (orange) simulations at $z = 2$. Limited by mass resolution, the LR run only resolves the most massive halos with $M_h \gtrsim 10^{13} M_\odot$. Also, limited by volume, the HR run lacks halos above $M_h \gtrsim 10^{14} M_\odot$. Our SR realization is able to match both the resolved LR and HR abundances over the whole mass range. We emphasize that our model generalizes surprisingly well to successfully predict the mass function of the massive halos that it has not been trained on, for over three orders of magnitude in abundance, demonstrating its extrapolating power in modeling structure formations.

studies—e.g., ref. 26. Our current SR models are dark matter only, but could be used for mocks in the same manner as HR models with resolved halos. As we discuss below, extensions of this work will mean that it will be possible to include hydrodynamic and star-formation effects, through training using HR models, which have these physics incorporated. SR mock catalogs could then be made that are more complex—for example, including reionization (27). This type of mock making would have some similarities with the “painting” of galaxies onto dark-matter simulations by Agarwal et al. (28), except that, unlike that paper, our method is conditioned on the entire density distribution rather than a few halo properties and would also add SR structure to the model.

Straightforward extensions of this work can be imagined, with different levels of sophistication. Particle data can be used to generate SR three-dimensional (3D) gas, star, and dark-matter particle distributions. Because the prediction will be of SR particles, these would effectively function as full simulations, as we have mentioned previously. In order to achieve this, training will have to be carried out on full physics models—e.g., refs. 3 and 29. It should also be possible to train on HR models run with different hydrodynamics algorithms [e.g., ENZO (30)] than are used in the LR models (as long as the initial conditions are identical). In such SR calculations, we will still be adding SR details to simulations after the LR models have been run, as in this paper.

Another extension is to replace the LR scheme, currently a full N-body solver, with faster alternatives, such as another DL model—e.g., D3M (8)—or a pseudo N-body solver—e.g., FastPM (31) or FlowPM (32). The computational cost of an SR realization can be further reduced. We will also be able to compute the full sensitivity matrix of the LR-SR model from

back-propagation, which, in turn, allows us to use the method in Hamilton Monte Carlo or similar sampling methods (e.g., refs. 33 and 34) for a joint study of the initial condition and the cosmology parameters.

Beyond this, an improvement will be to run AI algorithms “on the fly,” alongside the hydrodynamic simulation codes. This will allow coupling and feedback between the physical processes at superresolved scales and the LR scales. For the most accurate modeling, this will be important—for example, star-formation-driven winds or metal pollution from supernovae can spread far from a galaxy. This means that the SR structure on galaxy scales should ideally be able to affect the properties of the LR simulation. To achieve this, it may also become necessary to carry out training on the fly as HR and LR simulations run together. We leave development of these methods to future work.

The ongoing revolution in AI has already changed many fields of science. Applications of these techniques are becoming widespread in cosmology also. SR enhancement, which we have explored here, has many attractive features: Conventionally generated large-scale Fourier modes are fully linked to rapidly generated small-scale structures. This will allow consistent full-volume simulations of the Universe to be produced, covering an unprecedented dynamic range.

Materials and Methods

A Survey of Models. A variety of DL models have been applied to tackle SR tasks, ranging from early convolutional NNs (CNNs) plus simple loss-function-based supervised approaches [e.g., SRCNN (35)] to more recent unsupervised methods such as GANs [e.g., SRGAN (36)].

To generate SR outputs, both types of models use fully convolutional networks, which apply multiple convolution layers with learnable kernels of finite size. The difference between the two lies in their loss functions. By applying a simple loss function—e.g., the ℓ_2 or ℓ_1 norm—on the difference between the HR output and target, the first approach is a fully supervised learning task, and therefore easy to train. However, the drawback is that those simple loss functions typically lead to blurry output images, due to the fact that the target almost always contains information not present in the input—i.e., high-frequency features in Fourier space.

As a remedy, more complex loss functions have been used to match the output to the target on high-level features (37). Known as perceptual loss or content loss, it feeds the output and target separately to another NN (pretrained on some other image dataset) and compares their feature maps at intermediate layers, still by a simple loss function. Because those intermediate results contain higher-level features compared to the raw output or target, this method is able to generate less blurry images. However, such models are still deterministic and generate only one output for each fixed input, whereas, in principle, each input could map to infinitely many outputs due to the variability in the high-frequency modes.

The output image quality can be further improved by unsupervised learning methods, such as GAN (36). A GAN (5) is a class of DL system in which two NNs contest with each other in a game: The generative network generates candidates, while the discriminative network evaluates them. Given a training set, this technique learns to generate new data with the same characteristics as the training set. GANs can be used to generate entirely new data from initially random inputs, or, in the case of SR, HR images from LR ones. When the discriminator receives, in addition to SR or HR, the LR data as input, it is able to identify true and generated data conditioning on its large-scale knowledge. This technique is known as the conditional GAN (cGAN) (38).

To solve the one-to-many problem in the input-to-target mapping, a model needs more than the LR input to generate nondeterministic output. Different ways of introducing random factors have been attempted. For example, the image-to-image translation model pix2pix (39) uses test-time dropout to add uncertainty to its output. More successfully, the image-generation models StyleGAN and StyleGAN2 (40, 41) have achieved state-of-the-art results on generating human faces. It achieves a nondeterministic mapping by simply adding noise after every convolution layer. In our architecture, we use the same noise mechanism as in StyleGAN2 to add stochasticity to the output.

In designing our NN model, we are inspired mostly by StyleGAN2. Even though this model was originally designed for image-generation tasks, it does so by successively upsampling some initial (constant) LR input,

a process extremely similar to our SR task. We have also been influenced by the SR model SRGAN and the image-to-image translation model pix2pix.

Physical Considerations. We perform the SR simulation task in the Lagrangian description. Most often, the particles are originally located on a uniform grid, so we can structure a displacement field from the N-body particles as a 3D image with three channels. Each channel corresponds to one component of the displacement vector, and the value in each voxel provides the displacement of the particle originally from that voxel.

The DL model takes the LR particle displacements as the input and outputs a possible realization of their HR counterparts. Therefore, the outcome can be viewed as a higher-resolution simulation with more particles and higher mass resolution. Thus, we name them the “SR simulations.”

Limited by the size of GPU memory and the fact that 3D data consume more memory than lower-dimensional tasks, we cannot feed the whole simulations into the GPU during training and testing. To overcome this, we crop each simulation into smaller spatial chunks. While doing so, it is desirable to preserve the translational symmetry that arises naturally from the fully convolutional networks. We explain the procedure below in more detail.

In addition to translational symmetry, the equations of motion of an N-body system are also invariant under rotation. Under the periodic boundary conditions imposed in numerical cosmological simulations, the full 3D rotation group $SO(3)$ reduces to a finite group of 48 elements, known as the octahedral group O_h . Each group element can be decomposed into a permutation of the simulation box axes and reflections along those axes. During training, we feed the NNs with input and output pairs that are randomly transformed by these group elements. Such data augmentation greatly enlarges the training dataset and better enforces the said symmetry in the trained model.

With many more particles in the HR simulations, they can contain small-scale information that is not present in their LR counterparts. To generate SR simulations that statistically match the HR targets, a generator NN requires extra stochasticity in addition to that present in the LR input. Also, this additional stochasticity needs to be transformed in a way such that it results in the right correlations on different scales. To this end, we add white noise to the intermediate feature maps, with learnable amplitudes, throughout all stages of the NN. Noise injected at early stages is later upsampled and can then introduce correlations across multiple voxels.

Network Architecture and Loss Function. The design of our NNs is mostly inspired by StyleGAN2 (41). The generator is illustrated in Fig. 7. The ladder-like structure of the generator upsamples the resolution by a factor of two at each of its rungs. See the figure legends for more architectural details.

With the LR simulations l as inputs, our generator G produces SR outputs $G(l)$ to mimic HR target simulations h . To train the generator, one simple option is to minimize the voxel-by-voxel mean squared error (related to the ℓ_2 norm) between $G(l)$ and h

$$L_2 = E_{l,h} [\|G(l) - h\|_2], \quad [2]$$

where $E_{l,h}$ is the expected value while sampling all LR and HR simulation pairs. With the L_2 loss, the generator model is easier and faster to train, and we can produce output tightly correlated with the HR target on large scales. However, below the input-resolution scale, L_2 leads to blurry results that lack high-frequency features, as explained above when introducing different DL models.

This problem can be addressed by replacing the simple loss function with a discriminator network, which provides higher-level feedback to the generator on its performance. As shown in Fig. 8, our discriminator utilizes residual connections (42) that perform well on classification tasks. See the figure legends for more architectural details. Given either a generated SR or an original HR sample as input, the discriminator D learns to assign a score for it being real. In the original GAN, this score is a probability to be compared to the known values (zero or one) using the binary cross entropy.

Wasserstein GAN [WGAN (43)] introduced a different score, which can be used to evaluate the Wasserstein distance, the distance between the distributions of real and fake images by optimal transport. The discriminator (usually referred to as the critic in WGAN) is subject to constraint of it being Lipschitz-continuous with a Lipschitz constant of one—i.e., $|D(x_1) - D(x_2)| \leq |x_1 - x_2|$. Compared to the vanilla GAN, WGAN is empirically superior, for it requires less tuning for training stability and has a loss function that converges as the generated image quality improves. However, it requires more computation per batch to maintain the Lipschitz constraint. This is most often achieved by adding a gradient-penalty regularization term to the WGAN loss function [WGAN-gp (44)]. We train our networks using the WGAN-gp method and only penalize the critic gradient every 16 batches for training efficiency.

Because the HR and LR images must correlate on large scales, and high-frequency features should depend on the low-frequency context, we can

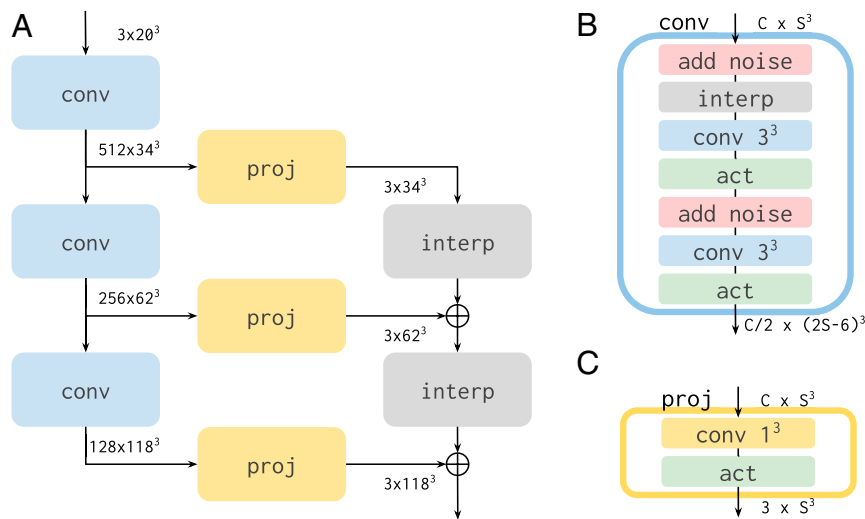


Fig. 7. Generator network architecture, inspired by StyleGAN2. The whole network structure is shown in A, with components enlarged in B and C. The colored plates are different operations, connected by lines or arrows from the input to the output. The sizes (channel number \times spatial size) of the input, intermediate, and output tensors are next to the arrows. The generator takes the shape of a ladder, where each rung upsamples the data by $2\times$. The left rail consists of consecutive convolution blocks (“conv” in blue plates) operating in the high-dimensional latent space and is projected (“proj” in yellow plates) at every step to the low-dimensional output space on the right rail. The projected results are then upsampled by linear interpolation (“interp” in gray plates), before being summed into the output. A key ingredient is the addition of noise (on red plates), which adds stochasticity absent from the input at each level of resolution. The added noises are then transformed into high-frequency features by the subsequent convolutions and activations. The kernel sizes of the convolutions are labeled in their plates (that distinguish them from the conv block). Note that with a kernel size one, “conv 1^3 ” is simply an affine transformation along the channel dimensions, and thus a convolution only in the technical sense. All activation functions (“act” in green plates) are Leaky ReLU with slope 0.2 for negative values. All conv blocks have the same structure as shown in B, except the first one, which starts with an additional 1^3 convolution and an activation.

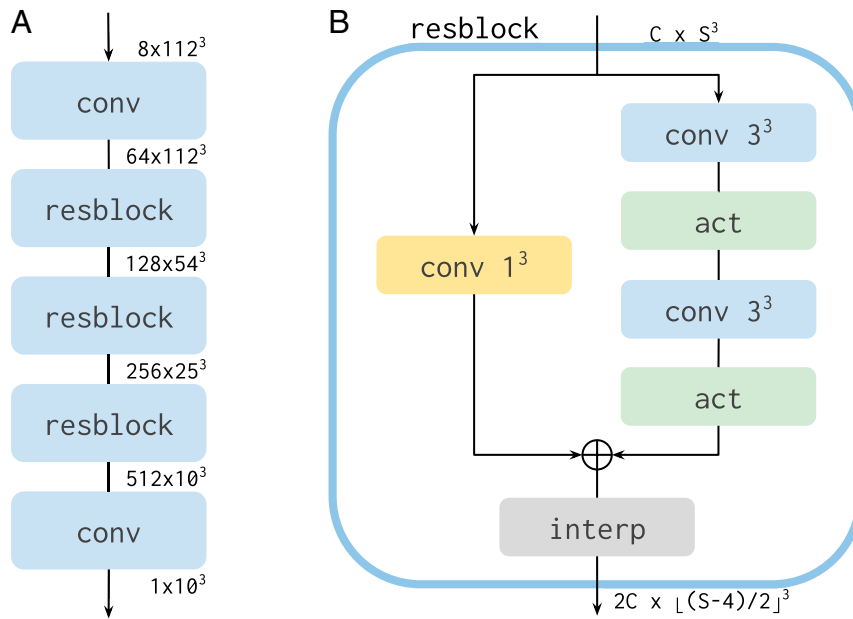


Fig. 8. Discriminator (critic) network architecture, inspired by the StyleGAN2. The whole network structure is shown in A, with the residual block enlarged in B. The residual block consists of two branches: The so-called “skip” branch is only a 1^3 convolution, and, on top of that, the other branch has real convolutions and activations to learn the “residuals.” It then downsamples by $2\times$ the sum of the two branches with linear interpolation. Aside from the residual blocks, the first conv block includes a 1^3 convolution followed by an activation, and the last conv block has a 1^3 convolution to double the channel, an activation, and a 1^3 convolution to reduce the channel to 1. See Fig. 7 for other details.

make the discriminator more powerful by giving it the LR images as additional input, making our model a cGAN. To achieve this, we concatenate the upsampled LR images (by linear interpolation) to both the SR output and HR target, respectively, as inputs to the discriminator. In addition, we also concatenate the density fields, computed with a differentiable CIC operation, from all three displacement fields. This allows the discriminator to see directly the Eulerian pictures of formed structures, thus greatly enhancing its capability. We find this addition crucial to generating visually sharp images and accurate predictions of the small-scale power spectra.

The final adversarial loss function we use is

$$L_{\text{WGAN-gp}} = E_i[D(I, G(I))] - E_{i,h}[D(I, h)] + \lambda E_{i,h} \left[\left(\|\nabla_i D(I, h)\|_2 - 1 \right)^2 \right], \quad [3]$$

where the first line gives the Wasserstein distance and the last term is the gradient penalty. i is a random sample drawn uniformly from the line segment between pairs of real (h) and fake ($G(I)$) samples. We refer the readers to ref. 44 for more details on WGAN-gp. During training, we update the discriminator to minimize $L_{\text{WGAN-gp}}$, involving all three terms, and update the generator to maximize it when only the first term takes effect.

In both the generator and the discriminator, we use convolution layers with no padding. Recall that the inputs and outputs of our networks are cropped parts of the bigger simulations, as limited by the size of GPU memory. Convolutions with zero padding or other forms of nonperiodic padding break translational invariance. Without padding in the convolution layers, the outputs of the generator are smaller than a simple factor-of-eight scaling. We compensate for this by adding extra padding to the inputs of the generator.

N-body Simulation Dataset. To train and validate our SR model, we use N-body simulations that only contain dark matter interacting via gravity. The dynamics of dark matter are evolved by using MP-GADGET (3, 45). It is an N-body and hydrodynamics cosmological simulation code optimized to run on the most massively parallel high-performance computer systems. MP-Gadget was used to run the BlueTides simulation (45) on BlueWaters, the only cosmological hydrodynamic code that has carried out full machine runs, scaling to 648,000 cores and producing more than 6 petabytes of useable data. In all simulations, the gravitational force is solved with a split Tree-PM approach, where the long-range forces are computed from a particle-mesh

method, and the short-range forces are obtained with a hierarchical octree algorithm.

We use $1/30$ of the mean spatial separation of the dark-matter particles as the gravitational softening length. The simulations have the WMAP9 cosmology with matter density $\Omega_m = 0.2814$, dark-energy density $\Omega_\Lambda = 0.7186$, baryon density $\Omega_b = 0.0464$, power-spectrum normalization $\sigma_8 = 0.82$, spectral index $n_s = 0.971$, and Hubble parameter $h = 0.697$. We train our model separately on simulation snapshots, at redshift $z = 4, 2$, and 0 , of different levels of nonlinearity. We expect that it is harder for the model to learn to form nonlinear structures than linear ones, and these redshifts allow us to test the performance degradation due to nonlinearity.

For training and testing, we run, respectively, 16 and 1 LR-HR pairs of dark-matter-only simulations with box size of $(100 h^{-1} \text{Mpc})^3$, and 64^3 and 512^3 particles for LR and HR, respectively. The mass resolution is $m_{\text{DM}} = 2.98 \times 10^{11} M_\odot/h$ for LR and $m_{\text{DM}} = 5.8 \times 10^8 M_\odot/h$ for HR. So, our SR task is to enhance the spatial resolution by $8\times$ and mass resolution by $512\times$. We also run a $(1 h^{-1} \text{Gpc})^3$ LR simulation, of otherwise the same configuration as the smaller LR runs, to test deploying our model to a larger volume.

Training and Testing. Our GAN model is trained upon the displacement field in Lagrangian space. We first preprocess the training data by converting the particle position to their displacement vector with respect to the initial grid. Then, for each mini-batch, we crop a 14^3 grid from the LR displacement field as input and pad three cells on each side to compensate for the loss of voxels during the generator convolution layers. This transforms into an SR output of size 118^3 through the generator (see sizes annotated in Fig. 7), out of which we crop the inner part to match the corresponding $(8 \times 14)^3 = 112^3$ grid from the HR target displacement field. Therefore, all of the mini-batches for training are about $(22 h^{-1} \text{Mpc})^3$ in size. We then concatenate the trilinear interpolations ($8\times$ upsampling per dimension) of the LR inputs to the SR outputs or the HR targets. The results are six channel images to be taken by the discriminator and transformed in a similar way, as shown in Fig. 8. Data augmentation is applied for each mini-batch.

We train our NN model using the Adam optimizer with learning rate 1×10^{-5} and exponential decay rates $\beta_1 = 0$ and $\beta_2 = 0.99$. We start with five epochs of supervised training with the simple loss function given in Eq. 2. This allows the generator to quickly learn to generate SR outputs that are consistent with LR inputs above the input resolution. We then proceed to the adversarial training and update the generator and the discriminator alternately, for 150 epochs.

For the test set, we use a new pair of 64^3 LR and 512^3 HR simulations with the same volume of $(100 h^{-1} \text{Mpc})^3$. The test set shares the same cosmological parameters as the training sets, but is a new realization with different initial conditions. Therefore, our comparison is free of overfitting to the training dataset. Remarkably, we are able to deploy our model to LR input of $1,000\times$ bigger volume and demonstrate the scaling ability of our approach.

The procedure of generating the SR simulation from the LR input goes as follows. We first preprocess the LR input by converting the particle position to the displacement field of shape 3×64^3 . Then, we use our trained GAN model to generate the SR displacement field with shape 3×512^3 and obtain the particle positions by moving them from their original positions on a lattice by these displacement vectors. The generation process from LR to SR is done in chunks due to the limit of GPU memory. We crop the LR input into pieces, generate their corresponding SR field, and stitch the output patches together to exactly match the shape of the target. The generated full SR field is periodically continuous through this procedure, thanks to the translational symmetry preserved by our GAN model.

We have implemented our model and training in map2map (<https://github.com/eelregit/map2map>), a general NN framework to transform field data.

Data Availability. Code and model implemented with map2map, a pytorch-based NN to emulate field data, have been deposited in GitHub (<https://github.com/eelregit/map2map>).

ACKNOWLEDGMENTS. The Flatiron Institute is supported by the Simons Foundation. We acknowledge the Texas Advanced Computing Center at The University of Texas at Austin for providing high-performance computing resources that have contributed to the research results reported within this paper. T.D.M. was supported by NSF Grants ACI-1614853 and AST-1616168 and NASA Astrophysics Theory Program (ATP) Grants 19-ATP19-0084 and 80 NSSC20K0519. T.D.M. and R.A.C.C. also were supported by NASA ATP Grants 80 NSSC18K101 and NNX17AK56G; and R.A.C.C. was supported by NSF Grant AST-1909193. S.B. was supported by NSF Grant AST-1817256. This work was also supported by the NSF AI Institute: Physics of the Future, NSF Grant PHY-2020295.

- M. Vogelsberger, F. Marinacci, P. Torrey, E. Puchwein, Cosmological simulations of galaxy formation. *Nat. Rev. Phys.* **2**, 42–66 (2020).
- S. J. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach* (Pearson Series in Artificial Intelligence, Pearson, Upper Saddle River, NJ, ed. 4, 2020).
- Y. Feng *et al.*, The formation of Milky Way-mass disk galaxies in the first 500 million years of a cold dark matter universe. *Acta Pathol. Jpn.* **808**, L17 (2015).
- I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- I. Goodfellow *et al.*, Generative adversarial nets. *Adv. Neural Information Processing Systems* **27**, 2672–2680 (2014).
- Z. Wang, J. Chen, S. C. H. Hoi, Deep learning for image super-resolution: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10.1109/tpami.2020.2982166 (2020).
- K. Schawinski, C. Zhang, H. Zhang, L. Fowler, G. K. Santhanam, Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *MNRAS* **467**, L110–L114 (2017).
- S. He *et al.*, Learning to predict the cosmological structure formation. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 13825–13832 (2019).
- P. Berger, G. Stein, A volumetric deep convolutional neural network for simulation of mock dark matter halo catalogues. *Mon. Not. Roy. Astron. Soc.* **482**, 2861–2871 (2019).
- C. Modi, Y. Feng, U. Seljak, Cosmological reconstruction from galaxy light: Neural network based light-matter connection. *JCAP* **2018**, 028 (2018).
- X. Zhang *et al.*, From dark matter to galaxies with convolutional networks. arXiv [Preprint] (2019). <https://arxiv.org/abs/1902.05965> (Accessed 2 April 2019).
- D. Wadekar, F. Villaescusa-Navarro, S. Ho, L. Perreault-Levasseur, Hinnet: Generating neutral hydrogen from dark matter with neural networks. arXiv [Preprint] (2020). <https://arxiv.org/abs/2007.10340> (Accessed 28 July 2020).
- A. C. Rodríguez *et al.*, Fast cosmic web simulations with generative adversarial networks. *Comp. Astrophys. Cosmol.* **5**, 4 (2018).
- D. K. Ramanah, T. Charnock, G. Lavaux, Painting halos from cosmic density fields of dark matter with physically motivated neural networks. *Phys. Rev. D* **100**, 043515 (2019).
- H. M. Kamdar, M. J. Turk, R. J. Brunner, Machine learning and cosmological simulations II. Hydrodynamical simulations. *Mon. Not. Roy. Astron. Soc.* **457**, 1162–1179 (2016).
- B. Dai, U. Seljak, Learning effective physical laws for generating cosmological hydrodynamics with Lagrangian Deep Learning. arXiv [Preprint] (2020). <https://arxiv.org/abs/2010.02926> (Accessed 7 October 2020).
- B. Dai, Y. Feng, U. Seljak, A gradient based method for modeling baryons and matter in halos of fast simulations. *JCAP* **2018**, 009 (2018).
- B. Dai, Y. Feng, U. Seljak, S. Singh, High mass and halo resolution from fast low resolution simulations. *JCAP* **2020**, 002 (2020).
- D. K. Ramanah, T. Charnock, F. Villaescusa-Navarro, B. D. Wandelt, Super-resolution emulator of cosmological simulations using deep physical models. *MNRAS* **495**, 4227–4236 (2020).
- Y. P. Jing, Correcting for the alias effect when measuring the power spectrum using a fast Fourier transform. *Acta Pathol. Jpn.* **620**, 559–563 (2005).
- M. Davis, G. Efstathiou, C. S. Frenk, S. D. M. White, The evolution of large-scale structure in a universe dominated by cold dark matter. *Acta Pathol. Jpn.* **292**, 371–394 (1985).
- F. List, G. F. Lewis, A unified framework for 21 cm tomography sample generation and parameter inference with progressively growing GANs. *MNRAS* **493**, 5913–5927 (2020).
- Y. Li, W. Hu, M. Takada, Super-sample covariance in simulations. *Phys. Rev. D* **8**, 083519 (2014).
- A. Barreira *et al.*, Separate universe simulations with IllustrisTNG: Baryonic effects on power spectrum responses and higher-order statistics. *MNRAS* **488**, 2079–2092 (2019).
- K.-W. Huang, Y. Ni, Y. Feng, T. Di Matteo, The early growth of supermassive black holes in cosmological hydrodynamic simulations with constrained Gaussian realizations. *MNRAS* **496**, 1–12 (2020).
- S. Lin *et al.*, The completed SDSS-IV extended baryon oscillation spectroscopic survey: GLAM-QPM mock galaxy catalogs for the emission line galaxy sample. *MNRAS* **498**, 5251–5262 (2020).
- S. Satpathy *et al.*, On the possibility of baryon acoustic oscillation measurements at redshift $z > 7.6$ with the roman space telescope. *MNRAS* **498**, 4955–4970 (2020).
- S. Agarwal, R. Davé, B. A. Bassett, Painting galaxies into dark matter haloes using machine learning. *MNRAS* **478**, 3410–3422 (2018).
- M. A. Marshall *et al.*, “The host galaxies of $z=7$ quasars: Predictions from the BlueTides simulation” in *Monthly Notices of the Royal Astronomical Society* (Oxford University Press, 2020), vol. 499, pp. 3819–3836.
- C. Brummel-Smith *et al.*, ENZO: An adaptive mesh refinement code for astrophysics (version 2.6). *J. Open Source Softw.* **4**, 1636 (2019).
- Y. Feng, M.-Y. Chu, U. Seljak, P. McDonald, FASTPM: A new scheme for fast simulations of dark matter and haloes. *MNRAS* **463**, 2273–2286 (2016).
- C. Modi, F. Lanusse, M. Mustafa, U. Seljak, Simulating the universe in tensorflow (2020). <https://blog.tensorflow.org/2020/03/simulating-universe-in-tensorflow.html> (Accessed 1 April 2020).
- U. Seljak, G. Aslanyan, Y. Feng, C. Modi, Towards optimal extraction of cosmological information from nonlinear data. *JCAP* **2017**, 009 (2017).
- F. Leclercq, W. Enzi, J. Jasche, A. Heavens, Primordial power spectrum and cosmology from black-box galaxy surveys. *MNRAS* **490**, 4237–4253 (2019).
- C. Dong, C. C. Loy, K. He, X. Tang, Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**, 295–307 (2015).
- C. Ledig *et al.*, “Photo-realistic single image super-resolution using a generative adversarial network” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Conference Publishing Services, Los Alamitos, CA, 2017), pp. 4681–4690.
- J. Johnson, A. Alahi, L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution” in *European Conference on Computer Vision: ECCV 2016*, B. Leibe, J. Matas, N. Sebe, M. Welling, eds. (Springer, Cham, Switzerland, 2016), pp. 694–711.
- M. Mirza, S. Osindero, Conditional generative adversarial nets. arXiv [Preprint] (2014). <https://arxiv.org/abs/1411.1784> (Accessed 7 November 2014).
- P. Isola, J.-Y. Zhu, T. Zhou, A. A. Efros, “Image-to-image translation with conditional adversarial networks” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2017), pp. 1125–1134.
- T. Karras, S. Laine, T. Aila, “A style-based generator architecture for generative adversarial networks” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2019), pp. 4401–4410.
- T. Karras *et al.*, “Analyzing and improving the image quality of StyleGAN” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2020).
- K. He, X. Zhang, S. Ren, J. Sun, “Deep residual learning for image recognition” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2016), pp. 770–778.
- M. Arjovsky, S. Chintala, L. Bottou, “Wasserstein generative adversarial networks” in *Proceedings of the 34th International Conference on Machine Learning*, D. Precup, T. Precup, Y. Whye, Eds. (Proceedings of Machine Learning Research, ICML, 2017), vol. 70, pp. 214–223.
- I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A. C. Courville, “Improved training of Wasserstein GANs” in *NIPS’17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, I. Guyon *et al.*, Eds. (Advances in Neural Information Processing Systems, Curran Associates, Inc., Red Hook, NY, 2017), vol. 30, pp. 5767–5777.
- Y. Feng, T. Di-Matteo, R. A. Croft, S. Bird, N. Battaglia, S. Wilkins, The BlueTides simulation: First galaxies and reionization. *MNRAS* **455**, 2778–2791 (2016).