

PINION: physics-informed neural network for accelerating radiative transfer simulations for cosmic reionization

Damien Korber ,^{1,2}★ Michele Bianco ,¹ Emma Tolley ¹ and Jean-Paul Kneib ¹

¹*Institute of Physics, Laboratory of Astrophysics, École Polytechnique Fédérale de Lausanne (EPFL), Chemin Pegasi 51, 1290 Sauverny, Switzerland*

²*Observatoire de Genève, Université de Genève, Chemin Pegasi 51, 1290 Versoix, Switzerland*

Accepted 2023 February 22. Received 2023 February 13; in original form 2022 August 26

ABSTRACT

With the advent of the Square Kilometre Array Observatory (SKAO), scientists will be able to directly observe the Epoch of Reionization by mapping the distribution of neutral hydrogen at different redshifts. While physically motivated results can be simulated with radiative transfer codes, these simulations are computationally expensive and cannot readily produce the required scale and resolution simultaneously. Here we introduce the Physics-Informed neural Network for reIONization (PINION), which can accurately and swiftly predict the complete 4D hydrogen fraction evolution from the smoothed gas and mass density fields from pre-computed N-body simulation. We trained PINION on the C^2 -Ray simulation outputs and a physics constraint on the reionization chemistry equation is enforced. With only five redshift snapshots, PINION can accurately predict the entire reionization history between $z = 6$ and 12. We evaluate the accuracy of our predictions by analyzing the dimensionless power spectra and morphology statistics estimations against C^2 -Ray results. We show that while the network's predictions are in very good agreement with simulation to redshift $z > 7$, the network's accuracy suffers for $z < 7$. We motivate how PINION performance could be improved using additional inputs and potentially generalized to large-scale simulations.

Key words: radiative transfer – software: simulations – cosmology: dark ages, reionization, first stars..

1 INTRODUCTION

Immediately after the Epoch of Recombination, no significant radiation sources existed in the Universe. This period is called the Dark Ages. However, the evolution of the matter density field led to the generation of collapsed structures, leading to the formation of the first stars and galaxies. At $z \sim 30$, these first luminous sources started to emit ultraviolet light, which successively ionized the intergalactic medium (IGM) (Furlanetto, Oh & Briggs 2006; Ferrara & Pandolfi 2014). Indirect constraints such as high-redshift quasar spectra (Fan et al. 2006; McGreer, Mesinger & Fan 2011; McGreer, Mesinger & D'Odorico 2014), the decline of Lyman α emitting galaxies (Schenker et al. 2011; Stark, Ellis & Ouchi 2011; Pentericci et al. 2014; Tilvi et al. 2014) suggest that reionization was complete for $z < 6$. This change from a dark universe filled with neutral hydrogen to the fully ionized IGM is known as the Epoch of Reionization (EoR) (Gorbunov & Rubakov 2011; Dayal & Ferrara 2018). The study of the EoR is essential for cosmology. The evolution of reionization is driven by the matter–energy content and geometry of the Universe, and the distribution of ionized regions reflects how galaxies and black holes formed.

The presence of neutral hydrogen during the EoR will be evident by measuring the emerging brightness temperature from a gas cloud at a given redshift against the CMB temperature (Zaroubi 2013).

Modern telescopes such as LOFAR¹ are able to measure the global statistics of the 21-cm signal (Paciga et al. 2013; Yatawatta et al. 2013; Jelić et al. 2014; Parsons et al. 2014; Jacobs et al. 2015; Patil et al. 2017; Mertens et al. 2020; Trott et al. 2020). However, during reionization the signal fluctuations are expected to be non-Gaussian (Mondal, Bharadwaj & Majumdar 2017; Shimabukuro et al. 2017; Majumdar et al. 2018; Gorce & Pritchard 2019; Hutter et al. 2019; Watkinson, Greig & Mesinger 2022), and therefore cannot be fully depicted by the power spectra only. Next-generation telescopes such as the SKA² will be able not only to detect the 21-cm signal but also to produce maps of its fluctuations on the sky (Mellema et al. 2013; Iliiev et al. 2014; Dillon et al. 2015; Koopmans et al. 2015; Wyithe, Geil & Kim 2015; Bacon et al. 2020). These observations at multiple frequencies will allow us to directly measure the neutral hydrogen density distribution throughout the Universe evolution. The expected precision of SKA-Low will allow a groundbreaking comparison between observations and simulations. Future high-fidelity observations of the EoR demand equivalently high-fidelity simulations. A physically motivated simulation of reionization needs to include three main components: gravitational evolution of dark matter particles via N-body simulations, gas dynamics via hydrodynamical simulations, and radiative transfer (RT) simulations to compute the propagation of photons and the evolution of the ionized atoms. Although it is possible to do *on-the-fly* simulations that combine these three steps

¹<http://lofar.org>

²<https://skatelescope.org>

* E-mail: damien.korber@protonmail.ch

(Rosdahl et al. 2013; Ocvirk et al. 2016, 2020, 2021; Lewis et al. 2022), these simulations are very computationally expensive as they require small time-steps and large number of simulated particles (Abel, Norman & Madau 1999; Bolton, Meiksin & White 2004; Shapiro, Ilev & Raga 2004). It is also possible to *post-process* the RT step on to existing N-body and hydrodynamical simulations in order to lighten its computational footprint. The state-of-the-art post-process RT simulation is the C^2 -Ray code (Mellema et al. 2006). This method uses a 3D ray tracing method that relaxes the constraints on the time-step size. While C^2 -Ray can produce high-accuracy EoR simulations at large scale, is extremely computationally expensive to run, requiring several millions of core-hours to run across thousands of nodes.

Another post-process RT code is Grizzly (Ghara et al. 2018), which assumes spherical density profiles around each source, thus simplifying the problem to a 1D RT process and speeding up the computation by several orders of magnitude. There is no simulation, which is able to produce RT on large volumes while still accounting for small-scale physics.

Deep neural networks are well-known as universal function approximators (Hornik 1991), and have been successfully used to accelerate and improve simulations of physical processes in many domains (He et al. 2019; Liu et al. 2021). Deep learning has previously been applied in the context of reionization simulations. In particular, Chardin et al. (2019) obtained some encouraging results for the use of neural networks in predicting the 3D ionization map at $z = 6$.

Despite their success, neural networks suffer from interpretability and generalization problems (Nguyen, Yosinski & Clune 2015; Ribeiro, Singh & Guestrin 2016). While physical phenomena are often described by some theoretical or empirical model, neural networks attempt to learn these laws from statistically small and biased data. However, in recent years, a new class of neural networks with physics constraints has emerged. These networks are called *Physics-Informed Neural Networks* (PINNs) and encapsulate additional physics laws into the behaviour of the network (Karniadakis et al. 2021).

In this paper, we present PINION, a PINN that can accurately predict the reionization map at any redshift and volume. This paper is organized as follows. In Section 2, we present the numerical codes employed to simulate the reionization process and create the network training data set. In Section 3, we present the model used to study the evolution of ionization during the EoR and data pre-processing by smoothing the source distribution field to encapsulate the expanding behaviour of the ionizing front. Then in Section 4 we present PINION, the neural network that we developed for this project, with its physics constraint. Section 5 gives the results and their analysis, and Section 6 concludes with a summary of the results and the future prospects for this project.

2 REIONIZATION SIMULATIONS

The simulation used in this paper is briefly described here below, but we refer to the reader the recent works (e.g. Dixon et al. 2016; Ross et al. 2017; Bianco et al. 2021) for further details. The cosmological parameters used in the simulation are based on the WMPA 3-yr data observation and the results are consistent for a Λ CDM cosmology with the following parameters, $\Omega_\Lambda = 0.724$, $\Omega_m = 0.276$, $\Omega_b = 0.454$, $H_0 = 70.1 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.784$, and $n_s = 0.946$ (Spergel et al. 2007). The CUBEP³M code (Harnois-Déraps et al. 2013) is employed for the N-body simulation and it is run on a box of physical size 714 Mpc containing 6912^3 particles, with a spatial resolution of

5.17 kpc and particles mass of $5.12 \times 10^3 M_\odot$. The initial conditions are generated using the Zel'dovich approximation, while the power spectrum of the linear fluctuations is given by the CAMB code (Lewis, Challinor & Lasenby 2000). The simulation then starts at redshift $z = 150$, which gives enough time to significantly reduce the non-linear decaying modes (Crocce, Pueblas & Scoccimarro 2006). An on-the-fly spherical overdensity halo finder algorithm (Harnois-Déraps et al. 2013; Watson et al. 2013) with overdensity parameter $\Delta = 130$ creates a halo list catalogues at each redshift step, while the remaining particles are considered to be part of the IGM. In this study, we only consider haloes with mass $M_h \geq 10^9 M_\odot$ as these sources are not affected by radiative feedback and are considered to be the primary driver of reionization (Ilev et al. 2012; Dixon et al. 2016).

The N-body particle list and the halo catalogues are then interpolated into a 300^3 grid with a smoothed-particle-hydrodynamic-like method (Shapiro et al. 1996; Mao et al. 2012), with a corresponding cell size of 2.381 Mpc. This dark matter density field and cumulative source mass field are used as inputs to the C^2 -Ray code, which computes the post-processed RT to generate the photoionization rates and ionization fractions. A fixed time-step $\Delta t = 11.54 \text{ Myr}$ is considered for redshifts $z = 6-50$. We refer the reader to, e.g. (Dixon et al. 2016; Majumdar et al. 2016; Bianco et al. 2021) for a more general overview of the RT method and N-body method employed in this paper. In our case, snapshots of C^2 -Ray for the redshifts range $z = 6-12$ were kept for training the network, which corresponds to the time when at least one source is present at each pixel of the smoothed grid and therefore where the reionization process is rapidly evolving.

The matter density field and source mass field used by C^2 -Ray are also provided as inputs to the neural network.

3 EVOLUTION OF REIONIZATION

In the EoR, we are interested in studying how the neutral hydrogen density n_{HI} is converted into the ionized hydrogen density n_{HII} . Reionization is driven by two different processes: *ionization*, by which high-energy photons ionize neutral hydrogen, and *recombination*, by which ionized hydrogen recombines with the local electron density n_e into neutral hydrogen. The evolution of ionized hydrogen is given by (Choudhury 2009; Pritchard & Loeb 2012):

$$\frac{dn_{\text{HII}}}{dt} = n_{\text{HI}}\Gamma - \mathcal{C}\alpha_B n_{\text{HII}}n_e, \quad (1)$$

where Γ is the reionization rate, \mathcal{C} is the clumping factor, which accounts for the inhomogeneity in the simulation, and $\alpha_B \equiv \alpha_B(T) = 2.59 \times 10^{-13} (T/10^4 \text{ K})^{-0.7} \text{ cm}^3 \text{ s}^{-1}$ is the temperature dependent case B recombination coefficient (Furlanetto et al. 2006). In this work, the clumping factor is set to $\mathcal{C} = 1$. This scenario corresponds to the case where sub-grid inhomogeneities in the IGM are ignored (Mao et al. 2019; Bianco et al. 2021). Moreover, we set the temperature to $T = 10^4 \text{ K}$, which assumes that gas in halo is able to radiatively cool through hydrogen and helium atomic lines (Wise et al. 2014). For a fixed density of hydrogen $n_{\text{H}} \equiv n_{\text{HI}} + n_{\text{HII}}$, we can define the neutral hydrogen fraction $x_{\text{HI}} \equiv n_{\text{HI}}/n_{\text{H}}$ and ionized hydrogen fraction $x_{\text{HII}} \equiv n_{\text{HII}}/n_{\text{H}}$. For simplicity, we assume pure hydrogen gas, $n_e \sim n_{\text{HII}}$. We can then rewrite equation (1) as

$$\frac{dx_{\text{HII}}}{dt} = (1 - x_{\text{HII}})\Gamma - \mathcal{C}\alpha_B n_{\text{H}} x_{\text{HII}}^2. \quad (2)$$

Note that in the following of this paper, we will refer to x_{HII} as the volume-average fraction of ionized hydrogen.

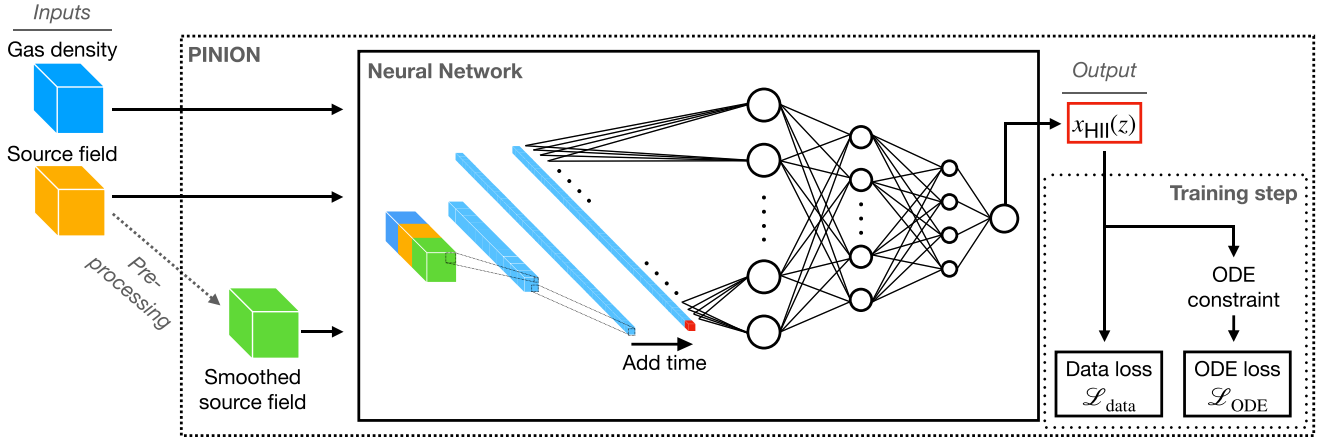


Figure 1. Diagram of PINION’s pipeline. The input data consist in two input subvolumes, the gas density and the source field that were produced by the N-body simulation. The source field is then processed to obtain a smoothed field. The network outputs one value for the subvolume central pixel x_{HII} . During the training steps, the output is then compared to the data to obtain the data loss $\mathcal{L}_{\text{data}}$ and the physics constraint to obtain the the ODE loss \mathcal{L}_{ODE} .

4 PINION

PINION³ is a 3D CNN (Convolutional Neural Network) that reduces the information from two 7^3 input cubes of gas density and mass of sources to predict x_{HII} at the cube centre for a single value of redshift. The input size corresponds to a volume of $(16.67 \text{ Mpc})^3$. We initially chose a size scale based on the physical hard limit imposed by Lyman limit system (Shukla et al. 2016), corresponding to an input mesh-grid size of 49^3 . However, we did not observe significant improvement in the result for this larger subvolume, and decided to reduce the input size to reduce the training and predicting computation time. We considered periodic boundary conditions on the full data set to be able to predict the border pixels. This is justified by the fact that the C^2 -Ray simulation we are training on also uses periodic boundary conditions.

4.1 PINION pipeline and architecture

PINION has two input channels: the source mass field and the gas density. An additional third input is derived internally by smoothing the source distribution field, described in Section 4.2. These three fields are then internally given as inputs to the network, which then returns the ionized fraction x_{HII} of the subvolume central pixel as the output. The network is composed of three convolutional layers with 3^3 kernel and doubling number of feature maps: 64, 128, and 256. The convolution results in 256 feature maps with channel size 1, to which the normalized look-back time is concatenated. This tensor is then used as input to a fully connected network with three hidden layers of 64, 16, and 4 nodes, and a single output node. Each layer uses the PReLU activation function (He et al. 2015) followed by a 3D batch normalization and a 10 per cent dropout. The output node uses a sigmoid activation function to truncate the output values between 0 and 1. In total, this network has 1130 252 trainable parameters. In Fig. 1, we show a simplified overview of PINION pipeline, the pre-process step and the neural network architecture.

We use the Adam optimization algorithm (Kingma & Ba 2014) during the network training to minimize our data loss function $\mathcal{L}_{\text{data}}$. In PINION, the data loss is the adjusted R^2 loss, also referred to the

coefficient of determination in statistics. This quantity is defined as

$$\mathcal{L}_{\text{data}}(x_{\text{true}}, x_{\text{pred}}) = \frac{\sum (x_{\text{pred}} - x_{\text{true}})^2}{\sum (x_{\text{true}} - \bar{x}_{\text{true}})^2} \equiv 1 - R^2, \quad (3)$$

where x_{pred} is the prediction from the network, x_{true} is the ground truth, and \bar{x}_{true} is its mean value (Gillet et al. 2019).

We train with the ODE loss as defined in Section 4.3. The results from this network will be discussed in Section 5 and show promising results.

4.2 Data pre-processing

During initial studies, we found that it was difficult for the network to learn the relationship between the input fields from the N-body simulation and the reionization across all redshifts. This is because, as shown in Fig. 2, while the gas density and source mass fields do not change drastically over the EoR, the reionization rate and fraction rapidly change as ionized regions progressively grow with time. Without any other constraints, the network must learn to predict extremely different outputs from very similar inputs across all redshifts.

To ameliorate this problem, we provide the network with the ionizing front propagation as an additional input. This approach was inspired by He et al. (2019), which used a neural network to predict cosmological structure formation starting from the Zel’dovich approximation (Zel’dovich 1970) and in our case this input should give a crude approximation of the photoionization rate. We can calculate the mean free path of ionizing photons at a given redshift $\lambda_{\text{vHII}}(z)$ using the approximated form of equation (23) in Choudhury (2009) as follows:

$$\lambda_{\text{vHII}}(z) \approx \frac{c}{H(z)} \times 0.1 \left(\frac{1+z}{4} \right)^{-2.55}. \quad (4)$$

This equation is the empirical fit from simulations valid for redshifts $z > 3$. While the source model in Choudhury (2009) does not exactly correspond to our simulation, it provides a good starting point for the network. We can then approximate the size of the ionized bubbles using $\lambda_{\text{vHII}}(z)$, and use the smoothed sources mass field to infer the photoionization rate. The algorithm is the following:

- (i) Calculate $\lambda_{\text{vHII}}(z)$ for a given redshift z .

³<https://github.com/epfl-radio-astro/PINION>

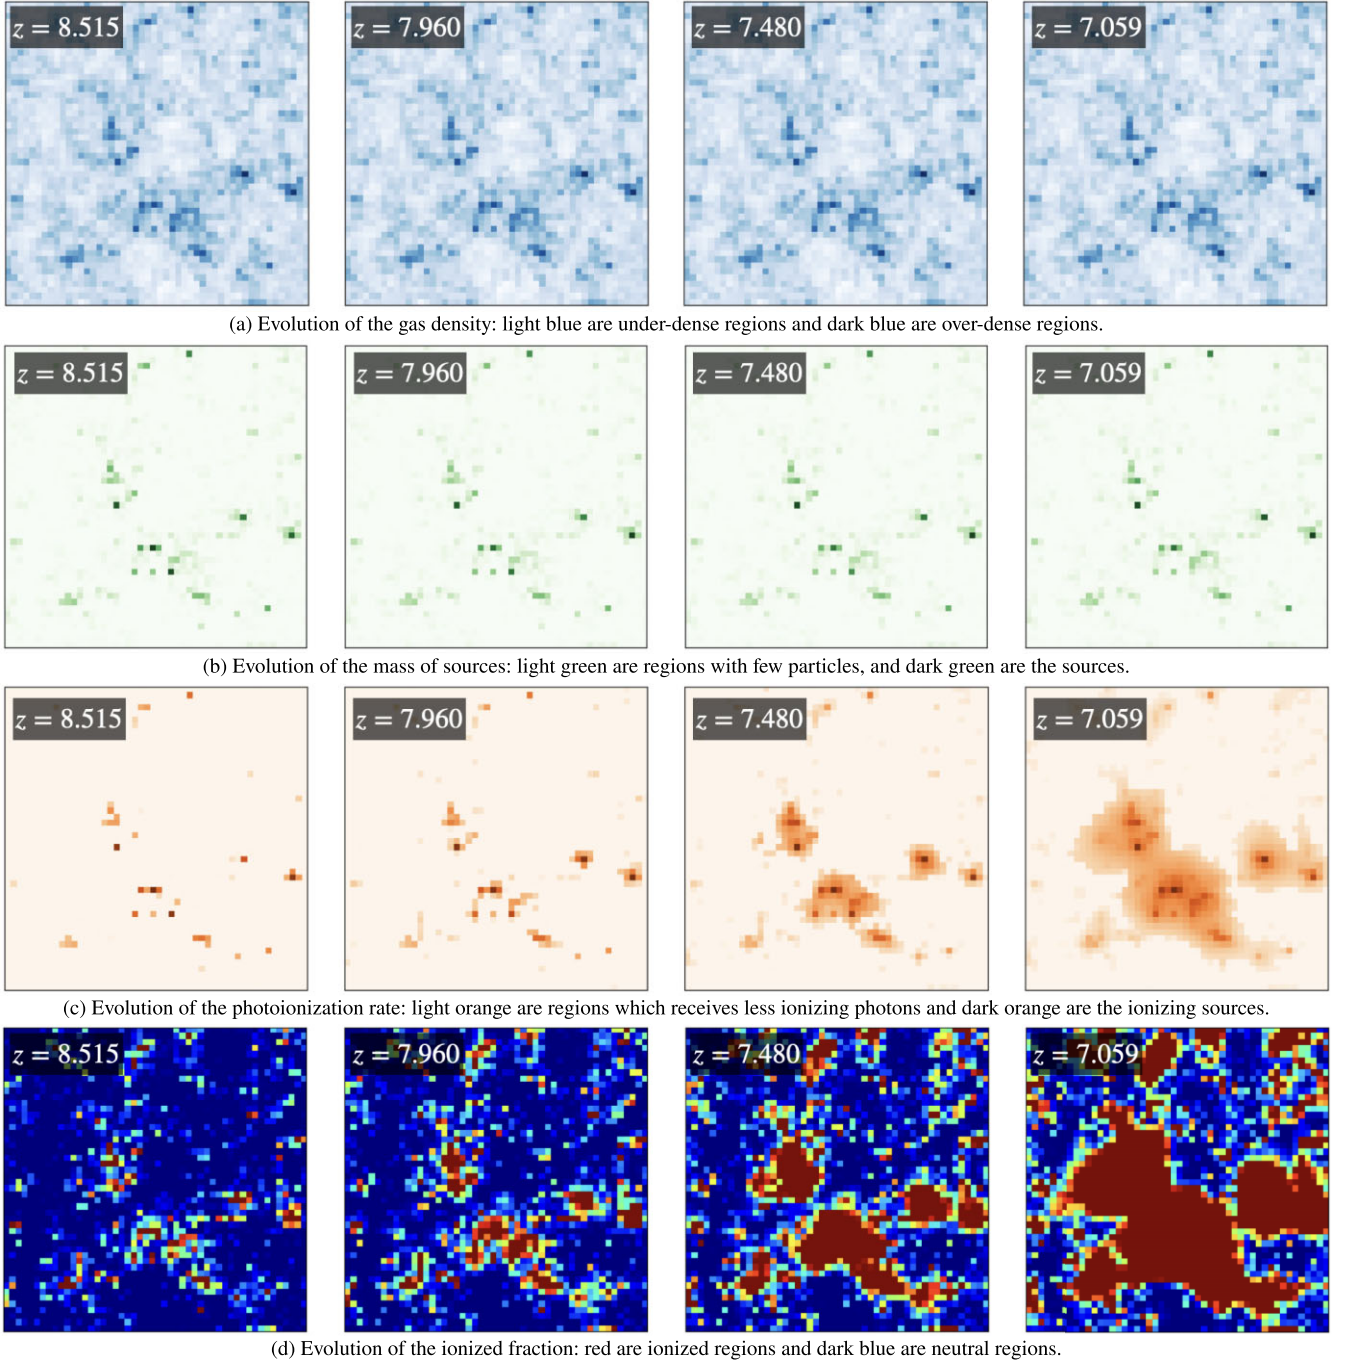


Figure 2. Evolution of the gas density, mass of sources, photoionization rate, and ionized fraction for four different redshifts. Note that the scale is different for every plot but the evolution shows that Fig. 2(a) and (b) are not expanding like Fig. 2(c) and (d).

(ii) Construct a spherical kernel of size $\lambda_{\text{vHI}}(z)$ by converting $\lambda_{\text{vHI}}(z)$ from Mpc to number of pixels given a resolution of 2.381 Mpc per pixel, setting every pixel to 0 when further than distance $\lambda_{\text{vHI}}(z)$ to the centre and scale the other pixel value between 0 and 1, where 1 is the centre pixel and 0 the borders.

(iii) Finally, convolve cube of mass of sources with the spherical kernel.

This algorithm creates a smoothed source distribution field that behaves approximately like the photoionization rate. Knowing that $\lambda_{\text{vHI}}(z)$ increases with redshift, the ionization front expand as well.

The purpose of this smoothed field is to provide an additional input to the network, and outside this machine learning context, it should not be interpreted as an estimate of the photoionization rate.

4.3 Physics constraint

One of the most popular ways to encode physics constraints in a deep neural network is through learning bias (Karniadakis et al. 2021). We have incorporated this in our network by using the ODE

from equation (2) as an additional loss:

$$\mathcal{L}_{\text{ODE}} = \frac{dx_{\text{HII}}}{dt} - (1 - x_{\text{HII}})\Gamma + \alpha_{\text{BH}}n_{\text{H}}x_{\text{HII}}^2. \quad (5)$$

This ODE depends on the time derivative of x_{HII} . While automatic differentiation can be used to obtain the exact value of the derivative (Baydin et al. 2018), it is unfortunately not feasible for a network with millions of parameters.

Instead, we encode the ODE loss as a finite-difference, using the Runge–Kutta 4 (RK4) method (Scherer 2013) to calculate the derivative. RK4 requires an intermediate step to perform derivations, so we simply take a step size two times larger than the time increment of the data. Let $x_n \equiv x_{\text{HII}}(t_n)$ the ionization fraction at the n th time-step and $\Gamma_n \equiv \Gamma(t_n)$ the photoionization rate at the n th time-step. Then, defining $h_n = \Delta t_n + \Delta t_{n+1}$ as the new time-step, where Δt_n is the time between snapshots n and $n+1$, and $D_n = \alpha_{\text{BH}}n_{\text{H},n}$, we can define the RK4 terms as

$$k_1 = (1 - x_n)\Gamma_n - D_n x_n^2, \quad (6a)$$

$$k_2 = (1 - x_n - \frac{h_n}{2}k_1)\Gamma_{n+1} - D_{n+1}(x_n + \frac{h_n}{2}k_1)^2, \quad (6b)$$

$$k_3 = (1 - x_n - \frac{h_n}{2}k_2)\Gamma_{n+1} - D_{n+1}(x_n + \frac{h_n}{2}k_2)^2, \quad (6c)$$

$$k_4 = (1 - x_n - h_n k_3)\Gamma_{n+2} - D_{n+2}(x_n + h_n k_3)^2, \quad (6d)$$

$$x_{n+2}^* = x_n + \frac{h_n}{6} [k_1 + 2k_2 + 2k_3 + k_4] \equiv x_n + \frac{h_n}{6} \mathcal{K}, \quad (6e)$$

where x_n is the predicted mean fraction of ionized hydrogen, obtained as the output of the neural network, x_{n+2}^* is the step obtained by computing the evolution of the ODE from x_n and Γ_n is the true photoionization rate from $\text{C}^2\text{-RAY}$. The large time-step makes it difficult to keep the solution stable. Indeed, at higher redshifts, most of the Universe is neutral, but at lower redshifts the ionization quickly increases. With large time-steps, RK4 has no issues tracking the evolution of the ODE at earlier times, as changes are slow. However, at later times, the photoionization rate is rapidly changing and the RK4 solution might drastically increase. To fix this issue, the equation (6e) is clipped between 0 and 1 to guarantee that the solution does not diverge. One can now write the physics loss as follows:

$$\Delta x = x_{n+2} - \text{MinMax} \left(\left| x_n + \frac{h_n}{6} \mathcal{K} \right|, 0, 1 \right), \quad (7a)$$

$$\mathcal{L}_{\text{ODE}}(x_n) = \frac{1}{N} \sum_{n=0}^N \left(\frac{\Delta x}{h_n} - (1 - x_n)\Gamma_n + D_n x_n^2 \right)^2, \quad (7b)$$

where $\text{MinMax}(x, 0, 1)$ ensures that x is encapsulated between 0 and 1. In Fig. 3, we compute the residual to check that the RK4 loss accurately describe the behaviour of the $\text{C}^2\text{-RAY}$ simulation. We see good agreement between the data and the ODE as the standard deviation is within 2 per cent for all redshifts. When the reionization fraction is not rapidly evolving for $z > 9$, we see excellent agreement within 0.5 per cent.

Although we use 3D convolutional layers, we account for the time evolution constraint by using the batch dimension of the network input. For a given spatial position of the input fields, its evolution for the 46 redshift snapshots is added as adjacent inputs along the batch dimension. The truth information is fed to the network in the same manner. This way, we can use spatial convolution and study the evolution of the predicted quantity with the time evolution constraint from the ODE loss. This additional loss constrains the network's behaviour and improves performance.

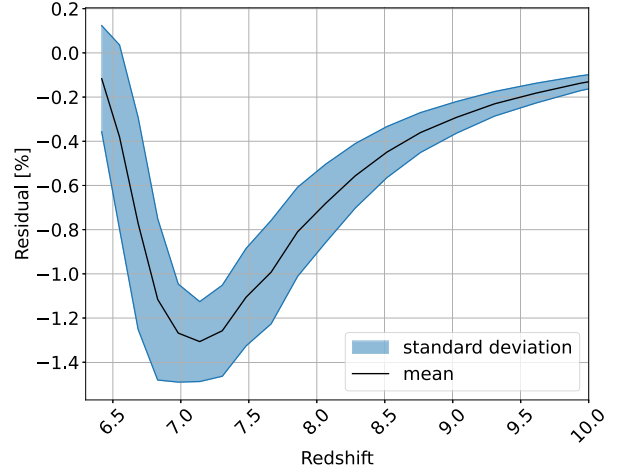


Figure 3. Residual between the RK4 ODE evolution of x_{HII} and the simulated x_{HII} . This plot is obtained by averaging the evaluated ODE over 1000 random subvolumes of size 47.62 Mpc. The redshift is cut at $z = 10$ to highlight the lower redshifts, where the residual is the biggest.

4.4 PINION and approximate RT solvers

Compared to other fast approximations of RT, PINION offers a more physically motivated but less interpretable solution.

For example, RT codes such as GRIZZLY assume an uniform density profile around each isolated source (Ghara et al. 2018). This approximation reframes the reionization process as a 1D problem using rotational symmetries, which considerably simplifies the computations. Therefore, this approach computes the ionization profile around isolated sources, thus quantifying the radius of the spherical ionization front as a function of time and optical depth (Shapiro et al. 2006) and then redistributes photons in overlapping regions in a second step (Ghara, Choudhury & Datta 2015). Recent works use 1D RT codes to study global statistics of the 21-cm signal for EoR and Cosmic Dawn (Ghara et al. 2020, 2021; Datta et al. 2022; Kamran et al. 2022).

PINION does not make any assumptions or approximations about the source or density field. However, PINION must infer the dynamics of RT during training, learning from a subvolume of size $(16.67 \text{ Mpc})^3$, which we consider the *region of influence*. We assume that this region contains all the nearby sources that contribute to the ionization state of the central pixel. During training, the ODE constraint equation (7) enforces realistic behaviour of the learned dynamics. However, as in all deep learning solutions, PINION results will have less interpretability compared to approximate 1D RT codes.

5 RESULTS AND ANALYSIS

In order to characterize the impact of the ODE loss, we consider three different network training scenarios:

- (i) NP (*no physics*): The training contains the entire evolution of pixels, but does not have a physics constraint: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}}$. This scenario is noted by dashdot lines ($\cdot -$) in figures.
- (ii) PFD (*physics and full data*): The training uses the ODE loss and the data loss for the entire time evolution of pixels: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \mathcal{L}_{\text{ODE}}$. This scenario is noted by dashed lines ($- -$) in figures.
- (iii) LD (*low data*): The training uses the ODE loss, but the data loss is only evaluated at five different snapshots (limited) of the




Figure 4. Visualization of a small cutout of the ionized hydrogen map at redshift $z = 7.305$ simulated by C^2 -Ray and the three PINION scenarios. The mean fraction of ionized hydrogen over the whole cube at this redshift is $\bar{x}_{\text{HII}} = 0.389$ for C^2 -Ray, $\bar{x}_{\text{HII}} = 0.398$ for PFD, $\bar{x}_{\text{HII}} = 0.364$ for NP, and $\bar{x}_{\text{HII}} = 0.391$ for LD. The residual is computed as $R = x_{\text{HII}}^{\text{PINION}} - x_{\text{HII}}^{C^2\text{-Ray}}$. Axes and colour bars are shared for readability.

simulation instead of the full reionization history: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{limited}} + \mathcal{L}_{\text{ODE}}$. This scenario is noted by dotted lines ($\cdot \cdot$) in figures.

During the training, the three scenarios do not differ greatly in convergence time. Training each scenario took 1.5 GPU-hours on one node for 400 epochs. The training was performed on a node from CSCS’s GPU cluster Piz Daint⁴ that is equipped with NVIDIA Tesla P100 16GB GPU. In each scenario, only the network weights with the lowest validation loss is kept.

To predict the reionization fraction at a single pixel location, the network needs the surrounding 7^3 pixel subvolume, corresponding to a volume of $(16.667 \text{ Mpc})^3$, for the three internal input fields. These subvolumes are sliced out of the full cubes and are concatenated to form a higher dimensional tensor for the training data set. For the PFD and NP scenarios, 4000 subvolumes over 46 redshifts are provided to the network (184 000 training subvolumes in total) with a batch size of 3680. For the LD training, 4000 subvolumes over 46 redshifts with a batch size of 3680 are also considered, but the data loss only had the ground truth for five fixed redshifts evenly spaced over the EoR evolution: ($z = 12.048, 9.938, 8.515, 7.480,$ and 6.686) in the evolution, meaning that it compares the prediction to the truth data on 20 000 subvolumes instead of 184 000 for $\mathcal{L}_{\text{data}}$. To validate the training, 500 validation subvolumes are considered, and the final network weights are selected from the training snapshot with the smallest total validation loss. Note that all the subvolumes are randomly chosen from the full data set.

5.1 Network output

The network results for each PINION training scenario are compared to the C^2 -Ray ground truth in Fig. 4 for a single redshift slice. In general, all the scenarios are able to accurately reproduce the ionization map. However, small differences can be observed at the borders of large and small ionized regions.

The x_{HII} redshift evolution in each scenario is shown next to C^2 -Ray in Fig. 5. The three scenarios evolve similarly and show good

agreement with the ground truth. All the PINION scenarios are able to find the first ionizing bubbles ($z > 8$). See for example the bubble at $(z, L) = (9, 610 \text{ Mpc})$ and $(8.1, 610 \text{ Mpc})$ in Fig. 5. However, for redshift $z < 7$, we start to see large differences between C^2 -Ray and the PINION predictions. For instance, the neutral island at $(z, L) = (6.4, 600 \text{ Mpc})$ or $(6.6, 490 \text{ Mpc})$ on the same figure. Between the different training scenarios, the NP scenario exhibits much more noise at low redshift ($z \approx 6.5$) compared to the physics-constrained scenarios. Furthermore, we observe that neutral islands are smaller at this redshift in all predictions. Nevertheless, NP scenario appears to predict the most accurate neutral island size.

5.2 Evolution of the volume-averaged \bar{x}_{HII}

We then evaluate the evolution of the mean fraction of ionized hydrogen over several distinct subvolumes. Fig. 6, top panel, shows the evolution of the ionized fraction for four different subvolumes, which each have different ionization history. Generally, we see very good agreement for $z > 8$. For early ionizing subvolumes (volumes 3 and 4), the PINION reionization fraction predictions are always postponed for the NP scenario. The physically constrained scenarios (PFD and LD) ionize earlier for $z > 7.5$, and later otherwise. For late ionizing subvolumes (volumes 1 and 2), as well as the global average (black line), when $z \gtrsim 7$ we observe good agreement for the LD scenario, overprediction for the PFD scenario, and underprediction for the NP scenario. For $z < 7$ all scenarios ionize earlier than the simulation. In general, we observe that for redshifts $z \gtrsim 7$ the LD scenario gives better results. For $z \lesssim 7$, the NP scenario is better at reproducing later ionizations, while LD remains better for earlier ionizations. In the residual plot, we observe a higher standard deviation at lower redshift ($z \lesssim 7.5$) in general. For volumes 1 and 2, the deviation peaks at $z \approx 6.6$ while volumes 3 and 4 have a negative peak earlier at $z \approx 7.1$. The global average shows both behaviours, with a flex point at $z \approx 7$, but less residual error. This is correlated with what observed in Section 5.1 for the same redshift range. In Fig. 6 bottom panel, we have the residual error of our prediction. We can see for $z \gtrsim 7$ and all subvolumes that the LD scenario follow well the simulation, while the NP scenario underpredicts with

⁴<https://www.cscs.ch/computers/piz-daint/>




Figure 5. Slices through the redshift axis of the ionized hydrogen light-cone, simulated by C^2 -Ray and the three PINION scenarios. The colour bar indicates the ionization fraction. For readability, we cut at redshift $z = 10$ and position $L \in [400, 600]$ Mpc.

residual under 2.5 per cent for $z \gtrsim 7.5$ while the PFD overpredicts it with less than 1 per cent difference for $z \gtrsim 7.5$. In Table 1, we show the corresponding redshift of four reionization milestones, when $\bar{x}_{\text{HII}} = 0.3, 0.5, 0.7$, and towards reionization completion at $\bar{x}_{\text{HII}} = 0.95$, interpolated from the results in Fig. 6. In Appendix C, we show a comparison between mass-averaged and volume-averaged global ionization fraction.

5.3 Morphology of the ionization bubbles

In the study of the EoR, of particular interest are the size and volume distributions of the ionized regions in the simulation (Furlanetto, Zaldarriaga & Hernquist 2004). There are several methods to measure these statistics (Friedrich et al. 2011; Lin et al. 2016; Giri et al. 2018), but in this work we consider the mean free path method and the friend-of-friend method. The former indicates the fraction of ionized bubbles in a given spherical-averaged size range, while the latter indicates the fraction of the total ionized volume that is contained in regions of a given volume. For both methods, we employ the `Tools21cm`⁵ python package for EoR simulations analysis (Giri, Mellema & Jensen 2020).

Fig. 7 shows the mean free path R of the different PINION scenarios and the simulation plotted against the probability distribution $P(R)$. We chose to study redshifts for which $\bar{x}_{\text{HII}} \approx 0.3, 0.5, 0.7$, which approximately corresponded to redshifts $z = 7.480, 7.059$, and 6.830. All three PINION scenarios exhibit similar morphology. In particular, in each of them, we observe an excess around the peaks,

meaning that we have a more uniform distribution of bubble size. The fact that we observe this for all three redshift means that the universe evolves more uniformly with the PINION prediction than with the C^2 -Ray simulation.

For redshifts $z = 7.480$ (blue) and $z = 7.059$ (orange), PINION predicts an excess of smaller scale bubbles below the median, and a lack of larger scale bubbles above the median compared to C^2 -Ray. For redshift $z = 7.480$ (blue), only the NP scenario lacks large-scale bubbles ($R \gtrsim 20$ Mpc), which means that the NP scenario produces smaller bubbles on average than the PFD and LD scenarios. For redshift $z = 7.059$ (orange), however, we observe that all three scenarios have a turning point above which the prediction lacks large-scale bubbles. For large- and small-scale bubbles, the LD scenario gives a more accurate distribution, but around each peak, the PFD scenario tend to be closer to the distribution in the simulation.

For redshift $z = 6.830$ (green), we observe a deficit of small-scale bubbles for all scenarios. In general, we observe that the predicted distributions are shifted to the larger scale, meaning that the bubbles are larger on average than in the simulation.

We also study the volume distribution with the friend-of-friend method. Contrary to the mean free path, which measures the size of the bubbles, the friend-of-friend algorithm evaluates the volume distribution of the bubbles using an island-finding algorithm. Pixels are recursively grouped into islands based on a threshold criterion, in our case of $x_{\text{HII}} > 0.5$. The average volume of each island V is then plotted against the cumulative probability distribution $V P(V)$. Fig. 8 shows the result for the friend-of-friend analysis. First, for redshift $z = 7.480$ (blue), one can see that both the PFD and LD scenarios are able to accurately reconstruct the bubble volume distribution

⁵<https://github.com/sambit-giri/tools21cm>




Figure 6. Top panel: Evolution of the mean fraction of ionized hydrogen for subvolumes of size 47.61 Mpc. Four volumes with different end of ionization were hand-picked and are displayed with different colours. Black lines represent the global average for the whole volume. Solid lines are C^2 -Ray, dashed lines are PFD training, dashdot lines are NP training, and dotted lines are LD training. The x -axis is common to both plot and was cropped at redshift $z = 10$ for readability. Some characteristic redshifts can be found in Table 1. Bottom panel: The residual between the PINION scenarios and the C^2 -Ray simulation for the above evolutions are displayed with the same colours. In blue, the standard deviation of the residual for at each redshift for 1000 different subvolumes of the same size is shown.

and the *percolation cluster*.⁶ On the other hand, the NP scenario tends to overpredict the number of non-overlapping bubbles, as one can observe an overprediction of the bubbles distribution and an underestimation of the total volume of the percolation cluster. For redshifts $z = 7.059$ (orange) and $z = 6.830$ (green) all three scenarios are able to reproduce the volume distribution of both the bubbles and the percolation cluster. In general, we observe that the volume size distribution is well-reproduced by PINION regardless of the scenario, aside from NP at higher redshifts.

5.4 HI dimensionless power spectrum

We also evaluate the dimensionless power spectrum of the neutral hydrogen field $\Delta_{\text{HI}} = P(k) \frac{k^3}{2\pi^2}$ for both C^2 -Ray and PINION,

⁶The volume formed by the overlapping bubbles.

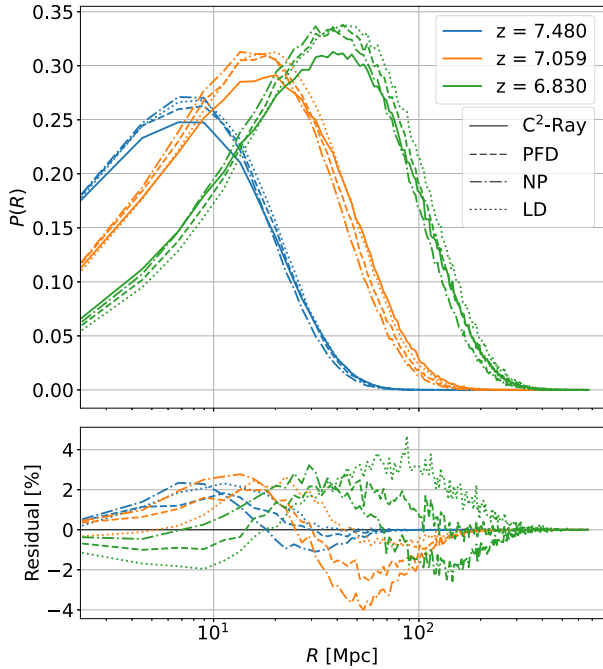
shown in Fig. 9. We observe almost perfect agreement in large- and small-scale structure for $z = 7.480$ (blue). However, as ionization progresses, we begin to see tension between PINION and C^2 -Ray. At larger scales, $k < 10^{-1} \text{Mpc}^{-1}$, the power spectrum is underestimated by all PINION predictions at all redshifts, meaning a statistically more ionized field than the ground truth. We observe that at these scales the discrepancy between the C^2 -Ray simulation and the PINION prediction increases with decreasing redshift, from almost perfect agreement at $z = 7.480$ (blue) to up to a factor of 2 difference for $z = 6.830$ (green).

5.5 Reionization redshift



From the predicted reionization maps, we can calculate the reionization redshift z_{reion} . Defined as the redshift at which a pixel becomes ionized fraction is $x_{\text{HII}} \geq 0.5$. In Fig. 10, we compare the true and reconstructed z_{reion} for our three training scenarios.

Table 1. Linearly interpolated redshifts for a given volume-averaged \bar{x}_{HII} . The values are derived from the results from Fig. 6.

	$z_{0.3}$	$z_{0.5}$	$z_{0.7}$	$z_{0.95}$	
Global	7.519	7.108	6.847	6.510	C^2 -Ray
	7.545	7.120	6.891	6.562	PFD
	7.451	7.066	6.857	6.512	NP
	7.513	7.120	6.903	6.560	LD
Vol. 1	7.031	6.685	6.535	6.369	C^2 -Ray
	7.073	6.817	6.670	6.432	PFD
	7.023	6.780	6.633	6.385	NP
	7.059	6.824	6.676	6.421	LD
Vol. 2	7.299	6.923	6.711	6.484	C^2 -Ray
	7.316	6.974	6.780	6.533	PFD
	7.233	6.929	6.744	6.487	NP
	7.292	6.980	6.791	6.528	LD
Vol. 3	7.854	7.446	7.189	6.903	C^2 -Ray
	7.906	7.434	7.125	6.788	PFD
	7.810	7.377	7.085	6.752	NP
	7.878	7.428	7.137	6.792	LD
Vol. 4	8.037	7.592	7.307	6.965	C^2 -Ray
	8.095	7.609	7.282	6.843	PFD
	7.995	7.548	7.236	6.806	NP
	8.071	7.605	7.291	6.854	LD

**Figure 7.** Top panel: Study of the ionization bubbles morphology using the mean free path method. Solid lines are C^2 -Ray, dashed lines are PFD training, dashdot lines are NP training, and dotted lines are LD training. The lines with the same colours are picked at the same redshift. The blue line corresponds to approximately $\bar{x}_{\text{HII}} = 0.3$, the orange to $\bar{x}_{\text{HII}} = 0.5$, and the green to $\bar{x}_{\text{HII}} = 0.7$. Bottom panel: Residual plot for the above distribution. The residual is obtained by subtracting the predicted PINION distribution with C^2 -Ray's distribution.

At first glance, all three histograms look very similar, but some differences can be observed in the x projection. We observe a larger residual at higher redshifts, with -4.3 ± 0.5 per cent, -3.1 ± 0.6 per cent, and -3.3 ± 0.6 per cent for $z = 11.834$ in the PFD, NP, and LD scenarios, respectively. As shown in Sec-

**Figure 8.** Study of the ionization bubbles volume distribution using the friend-of-friend method. Solid lines are C^2 -Ray, dashed lines are PFD training, dashdot lines are NP training, and dotted lines are LD training. Both lines are picked at the same redshift. The blue line corresponds to approximately $\bar{x}_{\text{HII}} = 0.3$, the orange to $\bar{x}_{\text{HII}} = 0.5$, and the green to $\bar{x}_{\text{HII}} = 0.7$.**Figure 9.** Top panel: Dimensionless power spectrum of x_{HII} . Solid lines are C^2 -Ray, dashed lines are PFD training, dash-dot lines are NP training, and dotted lines are LD training. These lines are picked at the same redshifts. The blue line corresponds to approximately $\bar{x}_{\text{HII}} = 0.7$, the orange to $\bar{x}_{\text{HII}} = 0.5$, and the green to $\bar{x}_{\text{HII}} = 0.3$. Bottom panel: Ratio plot for the above dimensionless power spectrum. The ratio is obtained by dividing the predicted power spectrum with C^2 -Ray's power spectrum.

tion 5.4, this indicates that PINION ionizes earlier than C^2 -Ray at high redshift. However, this difference reduces with redshift and in the particular case of NP, will have a positive peak of 1.4 ± 1.5 per cent at $z = 8.133$ which this time indicates that at this ionization redshift, PINION tends to ionize later compared




Figure 10. Comparison between the PINION prediction and the C^2 -Ray prediction of the redshift of reionization z_{reion} . For each training scenario, the *central panel* shows the 2D histogram of the predicted and truth redshift at which a region is considered ionized. A region is considered ionized when it reaches $x_{\text{HII}} \geq 0.5$. We linearly interpolated the redshift to have the value at the threshold to avoid tessellation effects. The red line shows emphasis the 1:1 relationship, which would correspond to a perfect result. The *vertical plots* show the y projection of the histogram, with the expected value and standard deviation of the residual in blue. The red line emphasis the perfect prediction, which is a residual of 0. The value of the residual is divided by the central redshift of the bins of interest. The *horizontal plots* show the x projection. The axis and the colours gradient are shared between scenarios to improve readability.

to C^2 -Ray. For the y projection, we observe a large positive bias in both NP and LD with 4.1 ± 3.5 per cent for NP and 2.2 ± 3.1 per cent at $z = 9.414$ for LD. Near the end of the EoR, $z \leq 7$, we observe a sudden change in the residuals of every scenario. We believe this results from small variations on the borders of neutral islands, which are difficult for the network to predict precisely.

We can also use Fig. 10 to compare PINION to (Chardin et al. 2019), a network which predicts the 3D z_{reion} map. PINION shows much better agreement at early- and mid-redshift, likely due to the addition of the physics constraint. Additionally, because PINION outputs x_{HII} , while PINION can reproduce (Chardin et al. 2019), the inverse is not possible.

5.6 Interpretation

We serve a few consistent differences between C^2 -Ray and the PINION predictions in our analysis. PINION predictions are more uniform, having more uniform bubble sizing and less variety in mean ionization fraction. The NP scenario ionizes later at high redshift, whereas the physics-constrained scenarios PDF and LD ionize sooner. For redshifts with rapidly changing ionization fraction, at $z \lesssim 7$, there is tension between the PINION predictions and C^2 -Ray at small and particularly large scales. There are several possible explanations for these systematic errors.

Data bias

The C^2 -Ray simulation used to train PINION is an imbalanced data set depending on the redshift. For instance, the vast majority of pixels at high redshift, early reionization history, are neutral rather than ionized. As such, it is possible that this biases the network predictions towards more neutral predictions. Any such bias would be most evident in the NP scenario, and we see for $z > 7$ in Fig. 6 that the NP scenario always predicts delayed ionization, thus more neutral volumes compared to truth. Additionally, in comparison to the PFD and LD scenarios, the NP has larger neutral islands and smaller ionized bubbles.

Finite-difference instabilities

Instabilities coming from the finite-difference approximation used to calculate the derivative can be observed as a direct issue for the predictions. Although PINION uses a Runge–Kutta 4 scheme for differentiation, which has an error of the order $\mathcal{O}(h^5)$, the time-steps between each snapshot are large compared to what is required by the scheme. We expect the tension introduced by this bias to affect the LD and PFD scenarios, and we indeed observe that the NP scenario has the best agreement with C^2 -Ray for $z < 7$. The finite-difference instability issue is observed, as shown in Fig. 3, but the maximal observed difference lies within 2 per cent. This is lower than the residual that we observe in Fig. 6, where a standard deviation within 12 per cent is observed for $z < 7$. While instabilities of the derivative evaluation do introduce a slight bias, this does not account for the majority of the differences observed between C^2 -Ray and PINION.

Inaccuracies from smoothed source distribution field

The last main issue encountered in this project is the use of the smoothed source distribution field. As described in Section 4.2, it is a way to provide additional information to the network, given that the two outputs from the N-body simulation do not provide enough information about the growing structure of the ionized regions. However, as the smoothed field is generated by convolving uniform spheres with the input fields, it is much more uniform in structure compared to the behaviour of the ionization front, and therefore the photoionization rate. The morphology of the ionization rate map and the smoothed source distribution field are compared in Appendix A, and in Appendix B we test the influence of the smoothed source distribution field by adjusting the scale length of equation (4) on the reionization history.

6 CONCLUSIONS

PINION successfully predicts the reionization history from RT simulation standard inputs with high accuracy. PINION could be applied to even larger volume simulations to create RT simulations

of unprecedented volume and resolution. Additionally, the physics-informed \mathcal{L}_{ODE} drastically reduces the amount of data required to train the network, with almost identical performance between the LD and PFD scenarios.

PINION is relatively fast for a reionization simulation. Training the network for each scenario takes about 1.5 GPU-hours and the prediction for each scenario takes less than 84 GPU-hours. However, the prediction is an already highly parallelizable task that can be further improved. The data I/O is a major bottleneck in network training and prediction, and optimization was not explored as part of this work. For example, the prediction for a 300^3 volume is done in parallel on 1000 GPU nodes, each working on a smaller subvolume. In the current implementation, each node loads the entire data set and then predicts all the pixels in the subvolume in approximately 5 min. This could easily be reduced to less than 2 min by only reading the subvolume.

While PINION is already a viable candidate for reproducing the C^2 -Ray simulations, there is still room for improvement. First, no hyperparameter optimization was performed for the training, so tweaks to the optimization algorithm or network structure may improve performance. Second, the smoothed field likely introduces strong systematic biases into the network's predictions. Starting from a more accurate proxy for the ionization front or the photoionization rate would result in more accurate predictions. Using an approximate but more physically motivated RT simulation such as *Grizzly* as an input field may drastically improve performance while keeping computation times low. Finally, the data bias could be reduced by modifying the data loss to account for imbalanced data.

PINION is a promising candidate for almost unsupervised learning. The network requires very little input data to train, and the ODE loss only requires a realistic estimation of the photoionization rate for the calculation of equation (5), which in this work we took instead the Γ from C^2 -Ray simulations. Moreover, we demonstrated that by providing a simplistic approach in equation (4) for calculating the smoothed field, the network can retrieve information on the redshift evolution observed by the photoionization rate maps, Fig. 2c. In principle, any mean free path model for the UV radiation could be used as a proxy in the smoothed field.

ACKNOWLEDGEMENTS

MB acknowledges the financial support from the SNSF under the Sinergia Astrosignals grant (CRSII5_193826). We acknowledge access to Piz Daint at the Swiss National Supercomputing Centre, Switzerland under the SKA's share with the project ID sk02. This work has been done in partnership with the SKACH consortium through funding by SERI. The authors would like to thank Professor Ilian Iliev for valuable discussions and comments and for providing the C^2 -Ray simulation outputs employed in this work. The authors would like to thank the referee for their valuable comments and suggestions that have improved the quality of this manuscript.

DATA AVAILABILITY

The data underlying this article is available upon request, and can also be re-generated from scratch using the publicly available C^2 Ray, CUBE³M and `Tools21cm` codes. The PINION code (ascl:2209.008) and its trained network weights are available on the EPFL radio astronomy GitHub page: <https://github.com/epfl-radio-astro/PINION>.

REFERENCES

- Abel T., Norman M. L., Madau P., 1999, *ApJ*, 523, 66
- Bacon D. J. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e007
- Baydin A., Pearlmutter B., Radul A., Siskind J., 2018, *J. Mach. Learn. Res.*, 18, 1
- Bianco M., Iliev I. T., Ahn K., Giri S. K., Mao Y., Park H., Shapiro P. R., 2021, *MNRAS*, 504, 2443
- Bolton J., Meiksin A., White M., 2004, *MNRAS*, 348, L43
- Chardin J., Uhlrich G., Aubert D., Deparis N., Gillet N., Ocvirk P., Lewis J., 2019, *MNRAS*, 490, 1055
- Choudhury T. R., 2009, *Curr. Sci.*, 97, 841
- Croce M., Pueblas S., Scoccimarro R., 2006, *MNRAS*, 373, 369
- Datta K. K., Ghara R., Hoque A., Majumdar S., 2022, *MNRAS*, 509, 945
- Dayal P., Ferrara A., 2018, *Phys. Rep.*, 780–782, 1
- Dillon J., Tegmark M., Liu A., et al., 2015, *PRD*, 91, 23002
- Dixon K. L., Iliev I. T., Mellema G., Ahn K., Shapiro P. R., 2016, *MNRAS*, 456, 3011
- Fan X. et al., 2006, *ApJ*, 132, 117
- Ferrara A., Pandolfi S., 2014, *Proc. Int. Sch. Phys. Fermi*, 186, 1
- Friedrich M. M., Mellema G., Alvarez M. A., Shapiro P. R., Iliev I. T., 2011, *MNRAS*, 413, 1353
- Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
- Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, *ApJ*, 613, 1
- Ghara R. et al., 2020, *MNRAS*, 493, 4728
- Ghara R., Choudhury T. R., Datta K. K., 2015, *MNRAS*, 447, 1806
- Ghara R., Giri S. K., Ciardi B., Mellema G., Zaroubi S., 2021, *MNRAS*, 503, 4551
- Ghara R., Mellema G., Giri S. K., Choudhury T. R., Datta K. K., Majumdar S., 2018, *MNRAS*, 476, 1741
- Gillet N., Mesinger A., Greig B., Liu A., Ucci G., 2019, *MNRAS*, 484, 282
- Giri S. K., Mellema G., Dixon K. L., Iliev I. T., 2018, *MNRAS*, 473, 2949
- Giri S. K., Mellema G., Jensen H., 2020, *J. Open Source Soft.*, 5, 2363
- Gorbunov D. S., Rubakov V. A., 2011, *Introduction to the Theory of the Early Universe: Hot Big Bang Theory*, 2nd edn. World Scientific Publishing Company, Singapore
- Gorce A., Pritchard J. R., 2019, *MNRAS*, 489, 1321
- Harnois-Déraps J., Pen U. L., Iliev I. T., Merz H., Emberson J. D., Desjacques V., 2013, *MNRAS*, 436, 540
- He K., Zhang X., Ren S., Sun J., 2015, *Proc. 2015 IEEE ICCV*, Santiago, p. 1026
- He S., Li Y., Feng Y., Ho S., Ravanbakhsh S., Chen W., Póczos B., 2019, *Proc. Natl. Acad. Sci.*, 116, 13825
- Hornik K., 1991, *Neural Netw.*, 4, 251
- Hutter A., Watkinson C. A., Seiler J., Dayal P., Sinha M., Croton D. J., 2019, *MNRAS*, 492, 653
- Iliev I. T., Mellema G., Ahn K., Shapiro P. R., Mao Y., Pen U.-L., 2014, *MNRAS*, 439, 725
- Iliev I. T., Mellema G., Shapiro P. R., Pen U.-L., Mao Y., Koda J., Ahn K., 2012, *MNRAS*, 423, 2222
- Jacobs D. C. et al., 2015, *ApJ*, 801, 51
- Jelić V., de Bruyn A., Mevius M., et al., 2014, *AAP*, 568, A101
- Kamran M., Ghara R., Majumdar S., Mellema G., Bharadwaj S., Pritchard J. R., Mondal R., Iliev I. T., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 001
- Karniadakis G. E., Kevrekidis I. G., Lu L., Perdikaris P., Wang S., Yang L., 2021, *Nat. Rev. Phys.*, 3, 422
- Kingma D. P., Ba J., 2014, *Proc. 3rd ICLR conf.*, San Diego
- Koopmans L. V. E., et al., 2015, *PoS*, AASKA14, 001
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 2, 2
- Lewis J. S. W. et al., 2022, *MNRAS*, 516, 3389
- Lin Y., Oh S. P., Furlanetto S. R., Sutter P. M., 2016, *MNRAS*, 461, 3361
- Liu X., Tian S., Tao F., Yu W., 2021, *Compos. Part B Eng.*, 224, 109152
- McGreer I. D., Mesinger A., D'Odorico V., 2014, *MNRAS*, 447, 499
- McGreer I. D., Mesinger A., Fan X., 2011, *MNRAS*, 415, 3237
- Majumdar S. et al., 2016, *MNRAS*, 456, 2080
- Majumdar S., Pritchard J. R., Mondal R., Watkinson C. A., Bharadwaj S., Mellema G., 2018, *MNRAS*, 476, 4007

Mao Y., Koda J., Shapiro P. R., Iliiev I. T., Mellema G., Park H., Ahn K., Bianco M., 2019, *MNRAS*, 22, 1

Mao Y., Shapiro P. R., Mellema G., Iliiev I. T., Koda J., Ahn K., 2012, *MNRAS*, 422, 926

Mellema G. et al., 2013, *Exp. Astron.*, 36, 235

Mellema G., Iliiev I. T., Alvarez M. A., Shapiro P. R., 2006, *New Astron.*, 11, 374

Mertens F. G. et al., 2020, *MNRAS*, 26, 1

Mondal R., Bharadwaj S., Majumdar S., 2017, *MNRAS*, 464, 2992

Nguyen A., Yosinski J., Clune J., 2015, Proc. 2015 IEEE CVPR Conf., Boston, p. 427

Ocvirk P. et al., 2016, *MNRAS*, 463, 1462

Ocvirk P. et al., 2020, *MNRAS*, 496, 4087

Ocvirk P., Lewis J. S. W., Gillet N., Chardin J., Aubert D., Deparis N., Thélie É., 2021, *MNRAS*, 507, 6108

Paciga G. et al., 2013, *MNRAS*, 433, 639

Parsons A. R. et al., 2014, *ApJ*, 788, 106

Patil A. et al., 2017, *ApJ*, 838, 65

Pentericci L. et al., 2014, *ApJ*, 793, 113

Pritchard J. R., Loeb A., 2012, *Rep. Prog. Phys.*, 75, 086901

Ribeiro M. T., Singh S., Guestrin C., 2016, Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining. KDD'16. Association for Computing Machinery, New York, p. 1135,

Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, *MNRAS*, 436, 2188

Ross H. E., Dixon K. L., Iliiev I. T., Mellema G., 2017, *MNRAS*, 468, 3785

Schenker M. A., Stark D. P., Ellis R. S., Robertson B. E., Dunlop J. S., McLure R. J., Kneib J.-P., Richard J., 2011, *ApJ*, 744, 179

Scherer P. O., 2013, Computational Physics: Simulation of Classical and Quantum Systems. Graduate Texts in Physics, Springer International Publishing, Heidelberg

Shapiro P. R., Iliiev I. T., Alvarez M. A., Scannapieco E., 2006, *ApJ*, 648, 922

Shapiro P. R., Iliiev I. T., Raga A. C., 2004, *MNRAS*, 348, 753

Shapiro P. R., Martel H., Villumsen J. V., Owen J. M., 1996, *ApJ*, 103, 270

Shimabukuro H., Yoshiura S., Takahashi K., Yokoyama S., Ichiki K., 2017, *MNRAS*, 468, 1542

Shukla H., Mellema G., Iliiev I. T., Shapiro P. R., 2016, *MNRAS*, 458, 135

Spergel D. N. et al., 2007, *ApJ Suppl. Ser.*, 170, 377

Stark D. P., Ellis R. S., Ouchi M., 2011, *ApJ*, 728, L2

Tilvi V. et al., 2014, *ApJ*, 794, 5

Trott C. M. et al., 2020, *MNRAS*, 493, 4711

Watkinson C. A., Greig B., Mesinger A., 2022, *MNRAS*, 510, 3838

Watson W. A., Iliiev I. T., D'Aloisio A., Knebe A., Shapiro P. R., Yepes G., 2013, *MNRAS*, 433, 1230

Wise J. H., Demchenko V. G., Halicek M. T., Norman M. L., Turk M. J., Abel T., Smith B. D., 2014, *MNRAS*, 442, 2560

Wyithe S., Geil P. M., Kim H., 2015, *PoS*, AASKA14, 015

Yatawatta S. et al., 2013, *A&A*, 550, 136

Zaroubi S., 2013, in Wiklind T., Mobasher B., Bromm V. eds, The First Galaxies. Vol. 396. Springer, Berlin, Heidelberg, p. 45

Zel'Dovich Y. B., 1970, *A&A*, 5, 84

APPENDIX A: VALIDITY OF THE DATA PRE-PROCESSING

As described in Section 4.2, to encode additional information about the expanding behaviour of the ionization rate map in the PINION input data, we pre-processed the source field to create the smoothed source distribution field. This field is a very rough approximation of this behaviour and should not be used outside PINION's context, as the CNN component of PINION is needed to correct for accurate physical behaviour.

However, large differences between this smoothed field and the ionization rate map might introduce issues in the reconstitution of the x_{HII} map, as seen in Section 5.6. We highlight the behavioural differences between the ionization rate and the smoothed field in




Figure A1. Dimensionless power spectrum of the photoionization rate and the smoothed source distribution field. As these are two different quantities (different scales and different units), these two fields are transformed into dimensionless quantities before computing the power spectrum.

Fig. A1. This figure shows the dimensionless power spectrum for the dimensionless quantity $\hat{x} = x/\bar{x} - 1$, where x corresponds to the photoionization rate field and the smoothed source distribution field, respectively. This quantity compares the scale dependence of structures caused by differences in the two fields, and any divergence can be fixed by adjusting a scaling factor. This figure shows that the dimensionless power spectrum is well-reproduced for redshifts around $z \approx 7.5$. However, for earlier redshifts ($z \gtrsim 8.5$) the smoothed field overpredicts large-scale structure $k < 10^{-1} \text{ Mpc}^{-1}$, while it is in relative good agreement with the small scale. For later redshifts ($z \lesssim 7$), the smoothed field underpredicts the small-scale structures, while overpredicting the large-scale structures.

APPENDIX B: MODIFYING THE SMOOTHED FIELD MEAN FREE PATH

To evaluate the impact of the mean free path for the smoothed fields, we recalculated the smoothed field with the mean free path reduced by a multiplicative factor $f=0.5$, physically equivalent to the ionizing photons travelling half the distance at a given redshift. The full analysis described in this paper was run under this modification, with and without network retraining (T and NT corresponding to *Trained* and *Not-Trained*, respectively). The resulting figures for redshift $z = 7.305$ are shown in Figs B1–B3. The mean fraction of ionized hydrogen over the whole cube at this redshift is $\bar{x}_{\text{HII}} = 0.389$ for $\text{C}^2\text{-RAY}$, $\bar{x}_{\text{HII}} = 0.398$ for $\text{PFD}(f=1)$, $\bar{x}_{\text{HII}} = 0.308$ for $\text{PFD-NT}(f=0.5)$, and $\bar{x}_{\text{HII}} = 0.370$ for $\text{PFD-T}(f=0.5)$. These plots show that the network outputs are agnostic to a change in the mean free path scale as long as the network is retrained using the modified smoothed field. Indeed, the PSD ($f=1$) case, presented in Section 5.1, and the corresponding retrained case results, $\text{PFD-T}(f=0.5)$, are very similar in these figures, indicating that the network can retrieve a reasonable ionization map as long as it is trained with some smoothed pre-processed field. For the predicted results without proper training, $\text{PFD-NT}(f=0.5)$, the reionization is, as expected, generally slower. The opposite behaviour was also observed for a factor $f=2$, but is omitted for brevity. This indicates the importance of proper training on the pre-processed inputs for PINION and the robustness of PINION to particular choices of pre-processing strategies.




Figure B1. Visualization of a small cutout of the ionized hydrogen map at redshift $z = 7.305$. The first two columns are the same as in Fig. 4, $f = 1$ indicating no modification to the mean free path. The third and fourth columns correspond to the prediction using the $f = 0.5$ smoothed field without modifying the training and with a retrained network. Axes and colour bars are shared for readability.




Figure B2. Slices through the redshift axis of the ionized hydrogen light-cone. The C^2 -Ray simulation was described in recent works (e.g. Dixon et al. 2016; Ross et al. 2017; Bianco et al. 2021) and the three other scenarios correspond to the usual PFD scenarios with modified scale factors and training. The colour bar indicates the ionization fraction. For readability, we cut at redshift $z = 10$ and position $L \in [400, 600]$ Mpc.




Figure B3. Top panel: Dimensionless power spectrum of x_{HI} . Solid lines are $\text{C}^2\text{-Ray}$, dashed lines are PFD ($f = 1$) scenario, dash-dot lines are PFD-NT ($f = 0.5$) scenario, and dotted lines are PFD-T ($f = 0.5$) scenario. These lines are chosen for the same redshifts. The blue line corresponds to approximately $\bar{x}_{\text{HI}} = 0.7$, the orange to $\bar{x}_{\text{HI}} = 0.5$, and the green to $\bar{x}_{\text{HI}} = 0.3$. Bottom panel: Ratio plot for the above dimensionless power spectrum. The ratio is obtained by dividing the predicted power spectrum with $\text{C}^2\text{-Ray}$'s power spectrum.

APPENDIX C: MASS-AVERAGED ANALYSIS

To complement the volume-averaged quantities analysis from Section 5, we compared the main results from Fig. 6 with their mass-averaged counterparts. The mass-averaged ionization fraction is defined as

$$\bar{x}_{\text{HI}}^m = 1 - \frac{\sum_i (1 - x_{\text{HI},i})(1 + \delta_i)}{\sum_i (1 + \delta_i)}, \quad (\text{C1})$$

where $x_{\text{HI},i}$ is the volume-averaged ionization fraction and $(1 + \delta_i)$ is the gas overdensity. The summation term i is computed on the entire simulated volume. The comparison between the mass and volume-averaged quantities is shown in Fig. C1. One can notice that the general behaviour of large volume simulation, where the mass-averaged tend to ionize earlier than the volume-averaged counterparts (Iliev et al. 2014).




Figure C1. Comparison between the mass-averaged and volume-averaged ionization fraction for Vol. 1, 4, and global average, as defined in Fig. 6. Each volume is separated into two panels, the top one showing the evolution of \bar{x}_{HI} over redshift the residual between the two for each scenario.

This paper has been typeset from a $\text{T}_{\text{E}}\text{X}/\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ file prepared by the author.