

Emulating the global 21-cm signal from Cosmic Dawn and Reionization

Aviad Cohen,¹ Anastasia Fialkov,^{2★} Rennan Barkana¹ and Raul A. Monsalve^{3,4,5}

¹*Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv 69978, Israel*

²*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

³*Department of Physics and McGill Space Institute, 3600 Rue University, QC H3A 2T8, Canada*

⁴*School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA*

⁵*Facultad de Ingeniería, Universidad Católica de la Santísima Concepción, Alonso de Ribera 2850, Concepción, Chile*

Accepted 2020 May 28. Received 2020 May 28; in original form 2019 October 14

ABSTRACT

The 21-cm signal of neutral hydrogen is a sensitive probe of the Epoch of Reionization (EoR), Cosmic Dawn, and the Dark Ages. Currently, operating radio telescopes have ushered in a data-driven era of 21-cm cosmology, providing the first constraints on the astrophysical properties of sources that drive this signal. However, extracting astrophysical information from the data is highly non-trivial and requires the rapid generation of theoretical templates over a wide range of astrophysical parameters. To this end emulators are often employed, with previous efforts focused on predicting the power spectrum. In this work, we introduce 21CMGEM– the first emulator of the global 21-cm signal from Cosmic Dawn and the EoR. The smoothness of the output signal is guaranteed by design. We train neural networks to predict the cosmological signal using a database of $\sim 30\,000$ simulated signals which were created by varying seven astrophysical parameters: the star formation efficiency and the minimal mass of star-forming haloes; the efficiency of the first X-ray sources and their spectrum parametrized by spectral index and the low-energy cut-off; the mean-free path of ionizing photons, and the cosmic microwave background optical depth. We test the performance with a set of ~ 2000 simulated signals, showing that the relative error in the prediction has an rms of 0.0159. The algorithm is efficient, with a running time per parameter set of 0.16 s. Finally, we use the database of models to check the robustness of relations between the features of the global signal and the astrophysical parameters that we previously reported.

Key words: software: development – galaxies: high-redshift – dark ages, reionization, first stars – cosmology: theory.

1 INTRODUCTION

The exploration of the Universe out to times earlier than the point of complete reionization is rapidly advancing. One of the most informative probes of these epochs is the 21-cm line produced by hydrogen atoms in the neutral intergalactic medium (IGM) at redshifts $z > 6$. This line redshifts to frequencies below 200 MHz and can be detected by low-frequency radio telescopes. Global 21-cm experiments measure the spectrum of this line averaged over the sky. The first tentative detection of the Cosmic Dawn signal was recently made by the Low-Band implementation of the Experiment to Detect the Global EoR Signature (EDGES, Bowman et al. 2018). Other global 21-cm experiments, including the Large-Aperture Experiment to Detect the Dark Ages (Bernardi et al. 2016; Price et al. 2018), the EDGES High-Band (Bowman & Rogers 2010; Monsalve et al. 2017, 2018, 2019), and the Shaped Antenna measurement of the background RAdio Spectrum (Singh et al. 2017,

2018), provide upper limits on the signal from Cosmic Dawn and the Epoch of Reionization (EoR), ruling out some extreme astrophysical scenarios. A parallel effort is being made by interferometric radio arrays that are placing upper limits on the fluctuations of the 21-cm signal, including the Donald C. Backer Precision Array for Probing the Epoch of Reionization (Kolopanis et al. 2019), the Low Frequency Array (LOFAR, Patil et al. 2017; Gehlot et al. 2019; Mertens et al. 2020), the Giant Metrewave Radio Telescope (Paciga et al. 2013), the Murchison Widefield Array (Beardsley et al. 2016; Barry et al. 2019; Li et al. 2019; Trott et al. 2020), and the Owens Valley Radio Observatory Long Wavelength Array (Eastwood et al. 2019). The most recent upper limit reported by LOFAR (Mertens et al. 2020) made it possible to place (weak) upper limits on the temperature of the neutral gas and ionization state of the Universe at $z = 9.1$ (Ghara et al. 2020; Mondal et al. 2020). Upcoming arrays, including the Hydrogen Epoch of Reionization Array (DeBoer et al. 2017), the Square Kilometer Array (Koopmans et al. 2015), and the New Extension in Nancay Upgrading LOFAR (Zarka et al. 2012), will provide measurements of the power spectrum over a wide range of scales and redshifts.

★ E-mail: afialkov@ast.cam.ac.uk

The 21-cm signal is driven by both astrophysical and cosmological processes and is thus a unique probe of the early Universe. The amplitude of the 21-cm line observed against the radio background radiation, normally assumed to be the cosmic microwave background (CMB; however, see Feng & Holder 2018; Ewall-Wice et al. 2018; Ewall-Wice, Chang & Lazio 2020), depends on the abundance of neutral hydrogen atoms as well as on the contrast between the spin temperature, T_S (the excitation temperature of the 21-cm transition), and the temperature of the background, T_{rad} . The former is driven to the kinetic temperature of the gas, T_K , by collisions as well as via absorption and re-emission of stellar Ly α photons (the Wouthuysen–Field, WF, coupling, Wouthuysen 1952; Field 1958). In the absence of collisions and/or Ly α radiation the spin temperature is driven to the temperature of the background. The gas is seen in absorption against the background if $T_S < T_{\text{rad}}$ (usually during the Dark Ages and Cosmic Dawn). Once the population of the first X-ray sources builds up and heats the IGM above the temperature of the background, the gas is seen in emission. In the course of reionization the abundance of neutral hydrogen atoms decreases and the IGM signal gradually vanishes. Overall, the signal measures properties of star formation, the abundance and luminosity of ultraviolet and X-ray sources, and possibly, properties of dark matter if the latter has an effect on the thermal and ionization histories of the gas (e.g. Evoli, Mesinger & Ferrara 2014; Tashiro, Kadota & Silk 2014; Muñoz, Kovetz & Ali-Haïmoud 2015; Barkana 2018; Fialkov, Barkana & Cohen 2018; D’Amico, Panci & Strumia 2018).

Our currently limited knowledge about primordial star and black hole formation translates into large uncertainties in the predicted 21-cm signal. As a result, a wide space of astrophysical parameters should be explored when predicting the 21-cm signature. Because full-scale numerical simulations are prohibitively expensive, alternative techniques, such as fast algorithms, emulators, or machine learning methods, are often employed to walk through the allowed space of astrophysical signals (e.g. Greig & Mesinger 2015, 2017; Shimabukuro & Semelin 2017; Schmit & Pritchard 2018). The effort has so far focused on the power spectrum of the 21-cm signal: Greig & Mesinger (2015) presented 21CMMC – a Monte Carlo Markov Chain (MCMC) tool which returns three reionization parameters (the mean-free path of ionizing photons, the minimum temperature of star forming haloes and the ionizing efficiency of sources) given power spectrum measurements (similar efforts include works by Liu et al. 2016; Hassan et al. 2017). Because X-ray heating might play an important role during the EoR (Mesinger, Ferrara & Spiegel 2013; Fialkov, Barkana & Visbal 2014), 21CMGEM was recently extended to include three heating parameters: the bolometric luminosity of X-ray sources per unit star formation rate, as well as the low-energy cut-off and the slope of the X-ray spectral energy distribution (SED, Greig & Mesinger 2017). Shimabukuro & Semelin (2017) took a different approach to find the best fit reionization parameters given power spectrum measurements: artificial neural networks (NNs) were trained on the data from 70 EoR simulations performed using the 21cmFAST code (Mesinger, Furlanetto & Cen 2011). The performance of the algorithm was tested on an additional set of 54 simulations. Schmit & Pritchard (2018) used NNs to emulate the power spectra generated by 21cmFAST and found a good agreement with 21CMMC. Jennings et al. (2019) compared five different machine learning techniques for emulating the power spectrum of models generated with the code SIMFAST21 (Santos et al. 2010). Finally, Kern et al. (2017) presented a more sophisticated emulator based on Gaussian processes, which could be applied to a broad range of problems. They demonstrate the performance on a six-parameter

model for the 21-cm signal including reionization and heating parameters as well as five additional cosmological parameters. With the exception of Kern et al. (2017), all the above-mentioned tools are designed to reconstruct the parameters from a 21-cm power spectrum measurement. Similar tools specifically designed for the global 21-cm signal are lacking.

The recently reported results from EDGES Low-Band (Bowman et al. 2018) revealed an anomalously strong and narrow absorption feature at ~ 78 MHz which, if truly of cosmological origin, cannot be explained by the standard astrophysical model outlined above. Even though concerns about the signal being of cosmological origin have been expressed in the literature (it could be a result of an uncompensated systematic error or be imprinted by the Galactic foregrounds, Hills et al. 2018; Singh & Subrahmanyan 2019; Spinelli, Bernardi & Santos 2019; Sims & Pober 2020), several exotic theories have been suggested to explain this signal. One possible explanation is that dark matter scattered off baryons, draining energy from the gas and leading to its overcooling (e.g. Muñoz et al. 2015; Barkana 2018; Fialkov et al. 2018; Muñoz & Loeb 2018). Another explanation invoked in the literature requires the existence of a strong radio background in addition to the CMB. Such an excess could be created by an anomalously bright population of high-redshift black holes at $z \sim 20$ (Bowman et al. 2018; Feng & Holder 2018; Ewall-Wice et al. 2018, 2020). As we await independent observational confirmation of the intriguing EDGES result, it is important to keep studying both the standard picture and exotic scenarios. In this paper, we explore a wide range of standard astrophysical scenarios. We use a large data set of models, which cover the widest astrophysical parameter space (see the next section), to develop a 21-cm global emulator (21CMGEM) for the first time. Given a set of seven astrophysical parameters, the emulator makes a prediction for the global 21-cm signal over a wide-redshift range ($z = 5\text{--}50$) that includes both the EoR and Cosmic Dawn. Although our models do not capture the EDGES absorption feature, the algorithm developed here could be applied to a revised set of models with additional physics. 21CMGEM, along with the global signals that were used to create the emulator, is publicly available at <https://www.ast.cam.ac.uk/~afialkov/Publications.html>. The tool has recently been employed to derive constraints on astrophysical parameters using the EDGES High-Band spectrum¹ (90–190 MHz, Monsalve et al. 2019).

This paper is organized as follows. In Section 2, we detail our seven-parameter astrophysical model. In Section 3, we describe the simulation, the limits on the astrophysical parameter space, and the database of $\sim 30\,000$ models. We also re-examine consistency relations between the astrophysical parameters and the features of the global signal first derived by Cohen et al. (2017). The design of the emulator is outlined in Section 4, and its performance assessed. Finally, we summarize our results in Section 5.

2 THE HIGH-REDSHIFT UNIVERSE

The 21-cm signal from Cosmic Dawn and the EoR is driven by several astrophysical processes including star formation, heating, and ionization. To produce the 21-cm signal, we use our seminumerical method (e.g. Visbal et al. 2012; Fialkov et al. 2013; Fialkov et al. 2014) which generates realizations of the universe in large cosmological volumes (384^3 comoving Mpc³) and over

¹ Note that we previously referred to the emulator as GLOBAL21CM (Monsalve et al. 2019).

a large redshift range ($z = 5\text{--}60$). The simulation follows the hierarchical growth of structure (including effects of the relative streaming velocity between dark matter and gas, Tseliakhovich & Hirata 2010), tracks star formation (averaged over a 3 Mpc scale) and follows the evolution of X-ray, Ly α , Lyman–Werner (LW, 11.2–13.6 eV), and ionizing radiative backgrounds. The simulation takes into account the effect of the relative streaming velocity on star formation, as well as the effect of the LW radiation and of the photoheating feedback on star formation (see details below). We parametrize the high-redshift astrophysics using seven key parameters: the star formation efficiency (f_*), the minimum circular velocity of star-forming haloes (V_c), the X-ray radiation efficiency (f_X), power-law slope (α) and low energy cut-off (ν_{\min}) of the X-ray SED, the mean-free path of ionizing photons (R_{mfp}), and the CMB optical depth (τ).

2.1 Star formation

The simulation takes into account the effect of radiative and mechanical feedback processes on star formation. Star formation is possible in dark matter haloes that are massive enough to enable efficient cooling of the in-falling gas (e.g. Tegmark et al. 1997). We use the threshold mass, or, equivalently (at a given redshift), the minimum circular velocity of star-forming haloes, as one of the free parameters. The lowest temperature coolant in the early universe is molecular hydrogen, which allows stars to form in haloes more massive than $M_{\min}^{\text{molecular}} \sim 10^5 M_\odot$, or with circular velocity larger than $V_c = 4.2 \text{ km s}^{-1}$ (e.g. Tegmark et al. 1997; Barkana & Loeb 2001; Abel, Bryan & Norman 2002; Bromm, Coppi & Larson 2002; Yoshida et al. 2003). LW radiation produced by the first stars eventually halts star formation in molecular cooling haloes (Haiman, Rees & Loeb 1997), shifting it to more massive atomic cooling haloes of $M_{\min}^{\text{atomic}} \sim 10^7 M_\odot$ ($V_c = 16.5 \text{ km s}^{-1}$, Haiman, Abel & Rees 2000; Machacek, Bryan & Abel 2001; Wise & Abel 2007; O’Shea & Norman 2008). The timing and duration of this transition is affected by uncertainties in the efficiency of the LW feedback (Visbal et al. 2014; Schauer et al. 2015). In addition, star formation in low-mass haloes is modulated by the relative streaming velocity between dark matter and baryons (e.g. Tseliakhovich & Hirata 2010; Dalal, Pen & Seljak 2010; Fialkov et al. 2012; Schauer et al. 2019). On the other hand, the minimum cooling mass can rise above the atomic cooling threshold via feedback mechanisms such as supernova explosions (e.g. Wyithe & Loeb 2013). At lower redshifts, when reionization becomes significant and the gas in the IGM is heated above 10^4 K , photoheating feedback becomes important. This feedback mechanism prevents further accretion of gas on to haloes below $10^8\text{--}10^9 M_\odot$ (V_c up to $\sim 75 \text{ km s}^{-1}$, e.g. Rees 1986; Weinberg, Hernquist & Katz 1997; Navarro & Steinmetz 2000; Sobacchi & Mesinger 2013; Cohen, Fialkov & Barkana 2016). In our parameter study, we explored the values of V_c between 4.2 and 100 km s^{-1} . However, the emulator was optimized for the starting value of V_c (before additional feedbacks are imposed) in the $4.2\text{--}76.5 \text{ km s}^{-1}$ range.

Another free parameter in the simulation is the fraction of gas in dark matter haloes that is converted into stars, referred to as the star formation efficiency. In general, this quantity depends on the halo mass. At low redshifts, observations find a mass-dependent star formation efficiency, e.g. Behroozi et al. (2019) show the evolution of stellar mass in haloes above $M \gtrsim 10^{10} M_\odot$ and at $0 \lesssim z \lesssim 10$. The observed star formation efficiency peaks at a value of a few per cent in haloes of $\sim 2.8 \times 10^{11} M_\odot$ (e.g. Mirocha, Furlanetto & Sun 2017; Behroozi et al. 2019), as it is regulated

by feedback mechanisms, and the process is less efficient in both higher and lower mass haloes. Such trends were recently incorporated in simulations of reionization with applications to synergies between the 21-cm signal and galaxy surveys with *James Webb Space Telescope* at $z \lesssim 10$ (Mirocha et al. 2017; Park et al. 2020). However, due to the lack of observations at higher redshifts and lower halo masses, applying such models to our work would require considerable extrapolations. This is because the Cosmic Dawn signal is driven by dark matter haloes of $10^5\text{--}10^8 M_\odot$ which can start forming stars as early as $z \sim 40$.

Star formation in the low-mass haloes characteristic of the high-redshift Universe is virtually unconstrained by observations, while numerical simulations yield a large scatter, finding an efficiency of a few per cent or much lower (Jeon et al. 2014; Wise et al. 2014; O’Shea et al. 2015). Therefore, to parametrize the process of star formation in our simulations we assume constant star formation efficiency in haloes heavier than the atomic cooling mass (and this value we designate f_*), while in lower mass haloes a logarithmic cutoff in the efficiency is employed

$$f_*(M) = \begin{cases} f_* & M_{\min}^{\text{atomic}} < M, \\ f_* \frac{\log(M/M_{\min})}{\log(M_{\min}^{\text{atomic}}/M_{\min})} & M_{\min} < M < M_{\min}^{\text{atomic}}, \\ 0 & \text{otherwise,} \end{cases}$$

where M_{\min} corresponds to the cut-off circular velocity V_c (see Cohen et al. 2017, for more details). We vary f_* between 0.0001 and 0.5.

2.2 Heating

The least constrained component of the modelling is the set of properties of the first X-ray sources that heat up the cosmic gas. The most plausible sources that dominate the X-ray radiative background at high redshifts are X-ray binaries (XRBs, Mirabel et al. 2011; Fragos et al. 2013); however, other candidates have also been discussed in the literature, including hot gas in galaxies, mini-quasars (Madau et al. 2004), X-rays produced via inverse Compton scattering of the CMB photons off electrons accelerated by supernovae (Oh 2001), or more exotic scenarios such as dark matter annihilation (e.g. Cirelli, Iocco & Panci 2009; Liu, Ridgway & Slatyer 2020).

The SED of the early X-ray sources is a key astrophysical parameter (Fialkov et al. 2014) and might strongly affects the 21-cm signal from both the EoR and Cosmic Dawn. The effect of hard X-ray sources with energy around 2 keV on the thermal and ionization histories (and, thus, on the resulting 21-cm signal) is significantly different from that of soft sources with energies of $\sim 0.5 \text{ keV}$: soft sources generate strong fluctuations on relatively small scales (up to a few tens of comoving Mpc) in the gas temperature and, subsequently, in the 21-cm intensity; on the other hand, hard sources produce a more homogeneous and less efficient heating, generating mild fluctuations on larger scales (> 100 comoving Mpc). XRBs, as well as mini-quasars, have a hard SED (see the discussion in Fialkov et al. 2016) that peaks at a few keV, while other sources can have softer SEDs. Absorption of soft X-rays with energy lower than ν_{\min} (typically of $\sim 0.1\text{--}0.5 \text{ keV}$) by dust in the host galaxy could contribute to effective hardening of X-ray SEDs (Fragos et al. 2013). We parametrize the X-ray SED by a power law of the slope α (i.e. $d \log(E_X)/d \log(\nu) = -\alpha$) and a low-frequency cut-off ν_{\min} . Since there is significant degeneracy between these two parameters, we vary α only slightly (in the range: $\alpha = 1\text{--}1.5$), and ν_{\min} in the wide range of $0.1\text{--}3 \text{ keV}$.




Figure 1. Left: histogram of the relative error in the prediction of ζ based on seven model parameters $[f_*, V_c, f_X, \alpha, \nu_{\min}, R_{\text{mfp}}, \tau]$. The total number of test cases was 2186. We find that 76 per cent of the cases have a relative error smaller than 1 per cent. Right: histogram of the relative error in the prediction of ν_{16} per cent, the frequency (which is a measure of redshift) at which the neutral fraction reaches 16 per cent. 92 per cent of cases have a relative error smaller than 1 per cent.

In addition to the shape of the SED, the total X-ray luminosity of sources is the other important parameter. Here, we adopt the standard expression for the luminosity per star formation rate (L_X –SFR relation, see Fialkov et al. 2014; Cohen et al. 2017, for more details) inferred from low-redshift observations of nearby starburst galaxies and XRBs (Grimm, Gilfanov & Sunyaev 2003; Gilfanov, Grimm & Sunyaev 2004; Mineo, Gilfanov & Sunyaev 2012):

$$\frac{L_X}{\text{SFR}} = 3 \times 10^{40} f_X \text{ erg s}^{-1} \text{ M}_{\odot}^{-1} \text{ yr.} \quad (1)$$

In the above expression, L_X is the bolometric luminosity, and f_X is the (constant) X-ray efficiency of sources, which we use as the third X-ray parameter. The standard normalization for XRBs (with $f_X = 1$) takes into account an order of magnitude increase in the L_X –SFR relation in the low-metallicity environments expected at high redshifts (Fragos et al. 2013). The high-redshift f_X is poorly constrained: A model-dependent upper limit of $f_X \sim 10$ –1000 can be derived using the measurement of the unresolved cosmic X-ray background (Fialkov et al. 2016); a lower (also model-dependent) limit of $f_X \sim 0.001$ is hinted at by 21-cm experiments (Singh et al. 2017, 2018; Monsalve et al. 2018, 2019; Mondal et al. 2020). To explore the parameter space, we vary f_X between 0 and 1000. However, the emulator was optimized in the range of f_X between 0 and 10.

2.3 Reionization

We parametrize the process of reionization with two parameters: the first parameter is the mean-free path of ionizing photons, R_{mfp} , which we vary between 10 and 50 comoving Mpc (Alvarez & Abel 2012; Greig & Mesinger 2015). This parameter approximately quantifies the effect of dense small-scale absorption systems in that it is the mean-free path of ionizing photons in a large-scale ionized bubble. In practice, it is set as an upper limit on the distance to sources that can participate in the reionization of a given cell.

The second EoR parameter is the ionizing efficiency of sources, defined as

$$\zeta = f_* f_{\text{esc}} N_{\text{ion}} \frac{1}{1 + \bar{n}_{\text{rec}}}, \quad (2)$$

where f_{esc} is the fraction of ionizing photons that escape into the IGM, \bar{n}_{rec} is the mean number of recombinations per ionized hydrogen atom, and N_{ion} is the number of ionizing photons produced per stellar baryon. Given a star formation history (i.e. fixing all the rest of the parameters $[f_*, V_c, f_X, \alpha, \nu_{\min}, R_{\text{mfp}}]$), and assuming a mass-independent ionizing efficiency, there is a one-to-one correspondence between ζ and the CMB optical depth. Because τ (rather than ζ) is directly probed by the CMB experiments (specifically by the *Planck* satellite, Planck Collaboration XIII 2016), we choose to work with τ as the free parameter. In our parameter exploration, we varied τ between ~ 0.04 and ~ 0.2 . However, high values of τ are ruled out (e.g. Planck Collaboration XIII 2016), and we find it difficult to produce τ below 0.055 and still be consistent with observational constraints (see below). Therefore, the emulator has been optimized for τ in the range between 0.055 and 0.1.

The non-linear mapping between τ and ζ , which is a function of the other input astrophysical parameters, is carried out using an NN which was trained on a set of 27 455 cases and tested with 2186 cases. This NN has seven input model parameters $[f_*, V_c, f_X, \alpha, \nu_{\min}, R_{\text{mfp}}, \tau]$ (and, thus, seven input neurons), one hidden layer of 40 neurons and 1 output, ζ . The Levenberg–Marquardt algorithm (Levenberg 1944; Marquardt 1963) was used to minimize the mean-square error between the true value provided by the training data set and the value predicted by the network. We evaluate the performance of the NN by quantifying its accuracy in predicting ζ . We find that 76 per cent of the cases have a relative error smaller than 1 per cent and the mean relative error is 0.77 per cent. The histogram of the relative errors is shown in Fig. 1 (left-hand panel).

2.4 Observational constraints

The parameter space outlined above is constrained by the available observations of the EoR. In addition to the limits on τ from the CMB experiments, we consider two other types of constraints when developing the global signal emulator:

(i) Stellar models indicate that for the extreme case of massive population III stars, $N_{\text{ion}} = 40\,000$ (Bromm, Kudritzki & Loeb

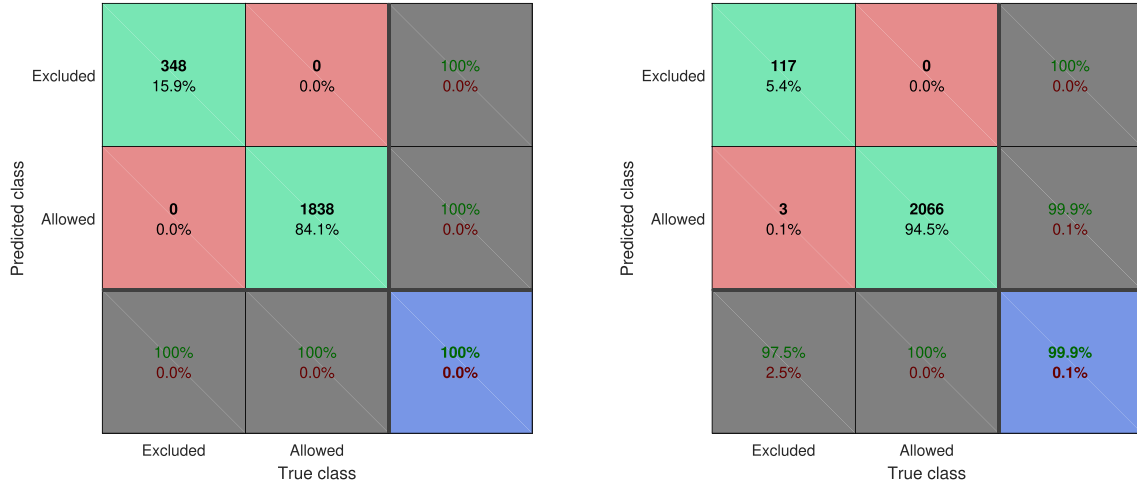


Figure 2. Left-hand panel: the confusion matrix of the exclusion process $\zeta > \zeta_{\max}$. Right-hand panel: the confusion matrix of the exclusion criterion $\nu_{16 \text{ per cent}} > 205.85 \text{ MHz}$. The structure of each confusion matrix is as follows: each matrix has nine fields with the green squares showing the number of excluded and allowed cases which were correctly classified, red squares showing the number of miss-classified cases, grey showing the percentage of the correct predictions for each row/column, and blue showing the total accuracy defined as the ratio of the total number of correctly identified cases (both excluded and allowed) to the total number of considered cases. The classification is done correctly in 100 per cent of cases for ζ , and in 99.9 per cent of cases for $\nu_{16 \text{ per cent}}$.

2001); therefore, we set an upper limit of $\zeta_{\max} = 40\,000 f_*$ based on equation (2). Hence, our first requirement for a parameter set to be valid is that $\zeta < \zeta_{\max}$.

(ii) Absorption seen in the spectra of high-redshift quasars measures the neutral fraction of the Universe (e.g. Bañados et al. 2018). A 2σ upper limit of $x_{\text{HI}, \max} = 16$ per cent on the neutral fraction at $z = 5.9$ ($\nu = 205.85 \text{ MHz}$) was derived from quasar absorption troughs (McGreer, Mesinger & D’Odorico 2015). Our second requirement is, thus, $x_{\text{HI}}(z = 5.9) < 16$ per cent. In this paper, we do not take into account the latest constraints from the Ly- α emitting galaxies (Mason et al. 2019) and high-redshift quasars (e.g. Bañados et al. 2018) as they became available when our paper was close to being completed (however, see Monsalve et al. 2019). To incorporate the neutral fraction constraint in our modelling, we train an NN to predict at which frequency, denoted by $\nu_{16 \text{ per cent}}$, the neutral fraction reaches 16 per cent for the given set of astrophysical parameters, $\nu_{16 \text{ per cent}} = \nu(x_{\text{HI}} = 0.16)$. The reionization history is considered valid if this frequency is lower than 205.85 MHz (i.e. the redshift is higher than 5.9). Because for many cases $x_{\text{HI}}(z = 5.9)$ is zero, $\nu_{16 \text{ per cent}}$ can be more easily inferred with high accuracy than the neutral fraction at $z = 5.9$. Technical details of this NN are discussed in Section 4.1.3. The performance of the NN in predicting $\nu_{16 \text{ per cent}}$ is evaluated in Fig. 1 (right-hand panel) where we show the histogram of relative errors. We find that the mean relative error is 0.47 per cent and 92 per cent of cases have a relative error smaller than 1 per cent.

As a part of the global signal emulator, described in detail in Section 4, the code checks whether or not an input parameter set renders a valid EoR history, that is given the generated values of ζ and $\nu_{16 \text{ per cent}}$, whether $\zeta < \zeta_{\max}$ and $\nu_{16 \text{ per cent}} < 205.85 \text{ MHz}$. The success/failure rates of the validation process is summarized in the form of confusion matrices shown in Fig. 2. Out of 2186 tested cases, 348 are excluded based on their values of ζ and 117 are excluded based on the values of $\nu_{16 \text{ per cent}}$ (22 overlap, i.e. are inconsistent with either constraint). The classification is done correctly in 100 per cent of cases for ζ , and in 99.9 per cent of cases for $\nu_{16 \text{ per cent}}$.

The above-mentioned constraints on ζ and $\nu_{16 \text{ per cent}}$ condition the entire parameter space, because the reionization history depends on several astrophysical parameters simultaneously. Top panel of Fig. 3 illustrates the mapping between the EoR constraints and the allowed regions in the $f_* - V_c$ plane for the specific choice of the EoR parameters, $\tau = 0.055$ and $R_{\text{mfp}} = 50 \text{ Mpc}$. For each combination of f_* and V_c , we use the trained NNs to check whether the reionization history is valid or not. In the figure, the area where the two exclusion criteria overlap is painted in black, the excluded region with $\zeta > \zeta_{\max}$ is shown in blue, while the region with $x_{\text{HI}}(z = 5.9) > 16$ per cent is red. The white regions have valid reionization histories.

The shape of the excluded and allowed regions is easy to understand. Consider first the requirement $\zeta < \zeta_{\max}$. For given values of τ and R_{mfp} (as well as the fixed heating parameters of $f_X = 1$, $\alpha = 1.5$ and $\nu_{\min} = 0.2 \text{ keV}$), models with a low star formation efficiency require high values of ζ that exceed the upper limit. Therefore, cases with low f_* are excluded. Now, the lower the value of V_c is, the more star forming haloes there are, making it easier to reionize (and match the required value of τ) without needing to exceed ζ_{\max} . Therefore, the maximum excluded f_* is a monotonically growing function of V_c . This function grows rapidly at the highest V_c due to the exponential dependence of the halo abundance on V_c , while it changes slowly at $V_c < 16.5 \text{ km s}^{-1}$ because in this mass range the number of stars at a given f_* is regulated by the LW feedback by the time of the bulk of cosmic reionization.

Consider the second requirement, $x_{\text{HI}}(z = 5.9) < 16$ per cent. Since τ has fixed the average timing of reionization, the neutral fraction constraint rules out cases with low V_c , since those are characterized by a more gradual evolution of reionization and, thus, a higher remaining neutral fraction at $z = 5.9$, regardless of the values of the other parameters. Therefore, this requirement rules out the left-hand portion of the $f_* - V_c$ plane (for a fixed R_{mfp}).

The exclusion contours for several choices of τ and R_{mfp} at the fixed values of the heating parameters are shown in the middle panel of Fig. 3, with the previously examined (reference) case of $\tau = 0.055$ and $R_{\text{mfp}} = 50 \text{ Mpc}$ shown in solid blue. A lower R_{mfp} implies a more gradual end to reionization (thus raising the residual x_{HI} at $z = 5.9$)




Figure 3. Constraints on V_c and f_* imposed by the constrained reionization history. Top: we show allowed (white) and excluded (blue for $\xi > \xi_{\max}$, red for $x_{\text{HI}}(z = 5.9) > 16$ per cent, black for both) regions in the $f_* - V_c$ plane for $\tau = 0.055$ and $R_{\text{mfp}} = 50$ Mpc. Also assumed are $f_X = 1$, $\alpha = 1.5$, and $\nu_{\min} = 0.2$ keV. Middle panel: the total exclusion contours (the excluded regions are under and to the left of the curves) for $\tau = 0.055$ (solid lines), $\tau = 0.064$ (dashed), $R_{\text{mfp}} = 10$ Mpc (red), and $R_{\text{mfp}} = 50$ Mpc (blue). The same X-ray parameters were assumed. Bottom panel: the total exclusion contours are shown for $\tau = 0.049$ (red), 0.052 (orange), 0.055 (green), 0.060 (cyan), and 0.064 (blue); these are the regions that are excluded for all values of R_{mfp} and X-ray parameters that we consider (i.e. these are *not* averaged over those parameter regions). The highest τ which is completely excluded within our parameter space is $\tau = 0.046$.

since sources then cannot contribute ionizing photons beyond this shorter distance. A higher τ moves the bulk of reionization towards higher redshifts, making it more compatible with the observational constraint at the fixed redshift of 5.9. With a lower R_{mfp} of 10 Mpc (solid red), the excluded area is larger with a lowest allowed value of $V_c = 52 \text{ km s}^{-1}$ ($\sim 28 \times M_{\text{min}}^{\text{atomic}}$), compared to 29 km s^{-1} ($\sim 5 \times M_{\text{min}}^{\text{atomic}}$) for the reference case. Increasing τ (the dashed lines correspond to $\tau = 0.064$) allows a wider range of V_c . In that case, if $R_{\text{mfp}} = 10$ Mpc (dashed red) then only $V_c < 6.3 \text{ km s}^{-1}$ is excluded, and even that is only if $f_* \lesssim 0.25$. At the highest star formation efficiencies, partial ionization by X-rays becomes significant, speeding up the process of reionization.

After searching over the full range of R_{mfp} and the heating parameters, we show the absolutely excluded regions for various values of τ in the bottom panel of Fig. 3; that is these are regions that are always excluded, there is no averaging here. We found a lower limit for the optical depth of $\tau = 0.046$. For the best-fitting *Planck* value of $\tau = 0.055$, we found lower limits on the circular velocity of $V_c \sim 26 \text{ km s}^{-1}$ ($\sim 4 \times M_{\text{min}}^{\text{atomic}}$) and on the star formation efficiency of $f_* \sim 0.0004$. However, for $\tau = 0.064$ (1σ away from the best-fitting *Planck* measurement) no values of V_c are excluded and the absolute minimum on the star formation efficiency is $f_* \sim 0.0002$.

2.5 Data set

Using the modelling outlined above, we created a data set of 29 641 global 21-cm signals that cover a very wide range of possible values of the seven astrophysical parameters, $V_c = 4.2\text{--}100 \text{ km s}^{-1}$, $f_* = 0.0001\text{--}0.50$, $\alpha = 1\text{--}1.5$, $\nu_{\min} = 0.1\text{--}3 \text{ keV}$, $f_X = 0\text{--}1000$, $\tau = 0.04\text{--}0.2$, $R_{\text{mfp}} = 10\text{--}50 \text{ Mpc}$, and verifying whether or not the ionization history complies with the EoR constraints (Section 2.4). The sampling of the parameter space was done randomly with uniform priors on $\log_{10}(V_c)$, $\log_{10}(f_*)$, $\log_{10}(f_X)$, R_{mfp} , and τ . The SED was randomly chosen with $\alpha = 1, 1.3$, or 1.5 , and $\nu_{\min} = 0.1, 0.2, 1$, or 3 keV . The 21-cm spectra are created over the redshift range $z = 5\text{--}50$ and are sampled at $\Delta z = 0.1$. The set of models was (randomly) split into the training and testing sets. The parameters of the testing set are restricted to $V_c = 4.2\text{--}76.5 \text{ km s}^{-1}$, $f_* = 0.0001\text{--}0.50$, $\alpha = 1\text{--}1.5$, $\nu_{\min} = 0.1\text{--}3 \text{ keV}$, $f_X = 0\text{--}10$, $\tau = 0.055\text{--}0.1$, and $R_{\text{mfp}} = 10\text{--}50 \text{ Mpc}$ over which ranges the performance of 21CMGEM was optimized. The training and testing data sets are available online at <https://www.ast.cam.ac.uk/~afalkov/Publications.html>.

3 CONSISTENCY RELATIONS

Cohen et al. (2017) derived universal relations between astrophysical quantities (such as the heating rate, ϵ_X , and the intensity of the Ly α background, J_α), and the three key points of the global signal, including the high- z maximum at the redshift labelled $z_{\text{max}}^{\text{hi}}$ (at matching frequency $\nu_{\text{max}}^{\text{hi}}$) and the brightness temperature $T_{\text{max}}^{\text{hi}}$; the absorption trough located at z_{min} (or ν_{min}) and reaching T_{min} ; and the low- z maximum at $z_{\text{max}}^{\text{lo}}$ (or $\nu_{\text{max}}^{\text{lo}}$) with $T_{\text{max}}^{\text{lo}}$. That work was based on a data set of 193 signals generated using a five-parameter model (V_c , f_* , τ , f_X , and either a hard or soft X-ray SED) with the parameters sampled on a grid (see Cohen et al. 2017, for a detailed description of the sampling). Here, we verify the validity of the above-mentioned relations in the context of our extended seven-parameter model and using a subset of 1948 randomly drawn combinations of the parameters. We find a good general agreement between this work and the previous study. However, compared to the previous study, we find significantly larger scatter owing to the larger explored astrophysical parameter space.





Figure 4. Brightness temperature at the high-redshift maximum point as a function of the observed frequency $\nu_{\max}^{\text{hi}} = 1420 \text{ MHz}/(1 + z_{\max}^{\text{hi}})$. The colours indicate the value of V_c as indicated on the colour bar: dark blue corresponds to the lowest value of V_c (4.2 km s^{-1}), and dark red corresponds to its highest value (76.5 km s^{-1}). Also shown is a fitting function (equation 3, solid) along with our older fit from equation (8) of Cohen et al. (2017, dashed) for comparison. Black \times 's show models that were excluded by our observational constraints. We observe a tight correlation between T_{\max}^{hi} and ν_{\max}^{hi} .

At the onset of Cosmic Dawn, the 21-cm signal is driven by atomic physics and the early process of Ly α coupling due to star formation, which results in a close relation between z_{\max}^{hi} and T_{\max}^{hi} as shown in Fig. 4. There is low scatter relative to a relation that can be fitted with a quadratic function of the form

$$T_{\max}^{\text{hi}} = a(1 + z_{\max}^{\text{hi}})^2 + b(1 + z_{\max}^{\text{hi}}) + c. \quad (3)$$

Using the extended data set, we find a similar relation to the one reported by Cohen et al. (2017, equation 8 and fig. 2 in that paper), with the best-fitting parameters changed by 10–20 percent. The new best-fitting values are $[a, b, c] = [-0.02925, 1.053, -9.667]$.

The value of z_{\max}^{hi} (and, hence, the value of the brightness temperature at this redshift) directly depends on the intensity of the Ly α background that drives the WF coupling. Therefore, it is



natural to expect that the intensity of the Ly α background can be inferred from the high-redshift maximum of the signal. Following Cohen et al. (2017, see their equations 9 and 10 and Fig. 3), we examine the relationship between z_{\max}^{hi} and the mean Ly α intensity measured at this redshift, as well as its derivative with respect to the scale factor $a = 1/(1 + z)$, and show the new results in Fig. 5. The best fits to the new data are:

$$\log(J_\alpha) = a_1 \log^2(1 + z_{\max}^{\text{hi}}) + b_1 \log(1 + z_{\max}^{\text{hi}}) + c_1, \quad (4)$$

and

$$\log\left(\frac{J_\alpha}{da}\right) = a_2 \log^2(1 + z_{\max}^{\text{hi}}) + b_2 \log(1 + z_{\max}^{\text{hi}}) + c_2, \quad (5)$$

where $[a_1, b_1, c_1] = [-10.64, 37.15, -54.31]$ and $[a_2, b_2, c_2] = [-7.851, 30.34, -47.73]$. On average we find a good agreement between the new study and the results of Cohen et al. (2017). However, the scatter in J_α is now substantially larger due to the contribution of X-rays to the Ly α background via X-ray excitation of neutral hydrogen.

The complexity of the universe increases as the population of the first heating sources forms. The location and the amplitude of the absorption trough show a very large scatter (left-hand panel of Fig. 6, see also fig. 4 of Cohen et al. 2017) due to the dependence of the signal on both the parameters of heating and of star formation. The latter regulates the strength of the WF coupling: for an efficient WF coupling, T_S is close to the kinetic temperature of the gas (and the absorption trough is deeper); while for a very inefficient coupling T_S moves towards the temperature of the background radiation (and the absorption trough is shallower). On the other hand, the role of the X-ray sources is to heat up the gas: the weaker the heating is, the more time the universe has to cool down as a result of the adiabatic expansion. Therefore, we get a lower limit given by the strongest possible absorption in the case of a fully coupled, adiabatically cooled gas:

$$T_{\min} \geq 26.8 \left(\frac{1 + z_{\min}}{10}\right)^{1/2} \left(1 - \frac{1 + z_{\text{dec}}}{1 + z_{\min}}\right) \text{ mK}, \quad (6)$$

where $z_{\text{dec}} = 137$ (Cohen et al. 2017). The depth of the absorption trough as a function of ν_{\min} is shown in the left-hand panel of Fig. 6 colour coded as a function of f_* .




Figure 5. The Ly α intensity in units of $\text{erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1} \text{ sr}^{-1}$ (left) and its derivative with respect to the scale factor $1 + z_{\max}^{\text{hi}}$ (right) as a function of $1 + z_{\max}^{\text{hi}}$. The colours indicate the value of V_c in accordance with the colour bar. Also shown are the fitting function for the present data set (solid, equations 4 and 5 on the left- and right-hand panels, respectively); and the fits from Cohen et al. (2017, dashed, equations 9 and 10 on the left- and right-hand panels, respectively) are shown for comparison. Black \times 's show models that were excluded by our observational constraints. The large scatter in J_α is a result of neutral hydrogen excitation by X-rays.

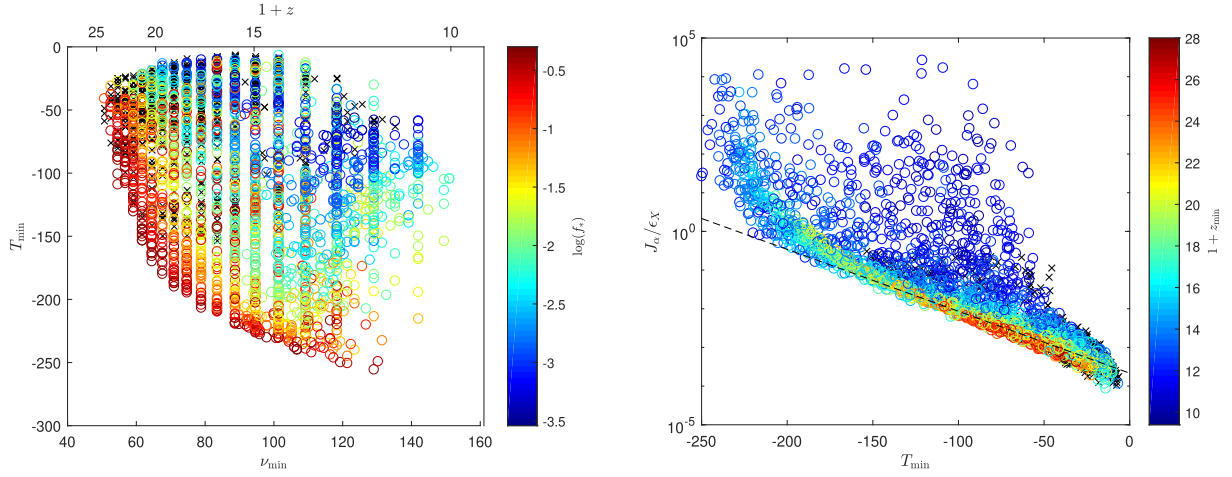


Figure 6. Left-hand panel: brightness temperature at the minimum point as a function of observed frequency of this point (bottom axis) or the equivalent one plus redshift (top axis). The colours indicate the value of the star formation efficiency (see the colour bar on the right). High values of f_* are needed for efficient WF coupling and deep absorption troughs. Right-hand panel: the ratio between the Ly- α intensity (in units of $\text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1} \text{sr}^{-1}$) and the X-ray heating rate (in units of $\text{eV s}^{-1} \text{baryon}^{-1}$) measured at z_{\min} as a function of the brightness temperature at the minimum point. The colours indicate the redshift of the minimum point (see the colour bar). Also shown is the fitting function from Cohen et al. (2017, dashed, equation 13). Black \times 's show models that were excluded by our observational constraints. Large variation in the properties of X-ray sources explored in this work contributes to the larger scatter in the J_α/ϵ_X relation.

Cohen et al. (2017) suggested that the ratio between the Ly α intensity and the X-ray heating rate can be inferred from the value of the brightness temperature at the minimum point. The right-hand panel of Fig. 6 shows that this does not entirely persist. The larger variation in the properties of X-ray sources employed here compared to what was implemented by Cohen et al. (2017), results in a large scatter in the J_α/ϵ_X relation. In particular, models with very low f_X (values that are unusually low compared with low-redshift galaxies, but are still possible) break this relation. As shown in the figure, this is the case only for models for which the measured value of z_{\min} would be low.

Finally, we examine the emission peak of the 21-cm signal during the EoR. As was pointed out by Fialkov et al. (2014), for a large part of the astrophysical parameter space X-ray heating plays a major role in the 21-cm signal during the EoR. Specifically, in cases of extremely inefficient heating there is no transition of the 21-cm signal into emission and the signal is seen in absorption throughout cosmic history. Therefore, the location and the amplitude of the emission feature depend not only on the EoR parameters but also on the heating rate (as well as on the parameters of star formation). Because of the complex dependence, one would expect to find a large scatter in the values of $(z_{\max}^{\text{lo}}, T_{\max}^{\text{lo}})$. However, as can be seen from the left-hand panel of Fig. 7 (see also equation 15 and fig. 7 of Cohen et al. 2017), the scatter is relatively low because the EoR history is significantly constrained by current observations (Section 2.4). The location and amplitude of the emission peak for the present data set are in good agreement with our previous results. The relation can be fitted with:

$$T_{\max}^{\text{lo}} = \begin{cases} a \frac{1}{1+z_{\max}^{\text{lo}}} + b & \text{if } 1+z_{\max}^{\text{lo}} > \frac{-a}{b} \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where $[a, b] = [-500.1, 59.05]$.

The right-hand panel of Fig. 7 shows the relation between the amplitude of the emission feature and the heating rate at z_{\max}^{lo} , which can be fitted with

$$\log(\epsilon_X) = aT_{\max}^{\text{lo}} + b, \quad (8)$$

with $[a, b] = [0.07026, -17.95]$. While there is significant scatter, this dependence can be used to constrain the properties of X-ray sources directly from the measurement of the global 21-cm signal.

4 21CMGEM: THE GLOBAL SIGNAL EMULATOR

The main product of this work is the global signal emulator which, given a set of seven input astrophysical parameters, outputs a realization of the global 21-cm signal sampled at $\Delta z = 0.1$ over the redshift range $z = 5-50$. In addition to the 21-cm spectrum, 21CMGEM outputs frequencies at which the neutral fraction is 0.16 per cent and 11 per cent along with values of the neutral fraction at $z = 5.9, 7.08$, and 7.54 . The values of the neutral fraction can be compared to the observational constraints on the reionization history at these redshifts: McGreer et al. (2015) published the upper limit $\bar{x}_{\text{HI}} < 0.06 + 0.05$ (at 68 per cent confidence) at $z = 5.9$, Greig et al. (2017) find $\bar{x}_{\text{HI}} = 0.40^{+0.21}_{-0.19}$ (68 per cent) from the damping wing analysis of a quasar at $z = 7.08$; while Bañados et al. (2018) find $\bar{x}_{\text{HI}} = 0.65^{+0.15}_{-0.32}$ (68 per cent) at $z = 7.54$ using the spectrum of ULASJ1342+0928, the highest redshift quasar detected so far. This auxiliary information can be used to apply external constraints to the models (see Monsalve et al. 2019).

Designed to detect features in the global 21-cm signal, the total-power experiments are very sensitive to steps and wiggles in the data. To avoid spurious apparent detections, the smoothness of the mock 21-cm signal over the entire observed frequency band is one of the major requirements. Predicting the signal in each frequency bin separately (as is done with the power spectrum emulators) is not sufficient as it leads to discontinuities in the spectrum. Instead, our approach here is to decompose the signals on to a new basis of smooth functions that span the entire simulated data set. Principal component analysis (PCA, Pearson 1901) is employed to find the basis of such functions. Dividing the entire database into training and testing sets, we train NNs to predict the PCA coefficients, along with the key points of the global signal, for any input set of astrophysical parameters. This information is then used to generate




Figure 7. Left-hand panel: brightness temperature as a function of observed frequency of the low-redshift maximum point. The colours indicate the CMB optical depth as is indicated on the colour bar. We also show the fitting function for the present data set (solid, equation 7) and the fit from Cohen et al. (2017, dashed, equation 15). Owing to the constrained EoR history, the scatter in the $T_{\max}^{\text{lo}} - z_{\max}^{\text{lo}}$ relation is low. Right-hand panel: X-ray heating rate (in units of $\text{eV s}^{-1} \text{baryon}^{-1}$) as a function of the brightness temperature at the low- z maximum. The colours indicate the CMB optical depth as is indicated on the colour bar. We also show the fitting function for the present data set (solid, equation 8) and the fit from Cohen et al. (2017, dashed, equation 16). Black \times 's show models that were excluded by our observational constraints. The nature of X-ray sources can be directly constrained from the measurement of ϵ_X .

the output 21-cm signal. The main steps of the emulation process, as well as details of the training and optimization of the algorithm, are described in the rest of this section.

4.1 Design

4.1.1 Classification

As our parameter study shows, all the analysed global signals have a universal shape featuring a high-redshift maximum and an absorption trough (note: for a reminder of the overall shape of the global 21-cm signal, see the examples shown in Fig. 13). The only non-universal feature is the emission signal during the EoR which is either present (we refer to this type of signal as *positive*) or not (*negative* signals) depending on the astrophysics. Because of this fundamental difference in the shape of the signals, our algorithm is twofold and treats the two types of signals separately. The classification into positive and negative cases is an essential part of the training and the prediction processes. If the signal is positive, it has four key points: the high-redshift maximum, absorption trough, low-redshift maximum, and the redshift of complete reionization (ν_{\max}^{hi} , ν_{\min} , ν_{\max}^{lo} , and ν_{reion} , respectively). A negative signal has only three key points (ν_{\max}^{hi} , ν_{\min} , and ν_{reion} , respectively). The key points divide each positive (negative) case into 3 (2) segments. As we detail in the next subsection, each segment is analysed separately using the PCA.

The *bagged trees* algorithm (Breiman 1996; Loh & Shih 1997) was used to determine whether a case is negative or positive. This algorithm fits many decision trees, each time using a different subset of the training set, and the decision is made by voting. After optimization, we chose to use bagged trees with 30 tree learners and tree size chosen using fivefold cross-validation. The classification was tested on 1014 negative and 580 positive cases. We first tested the accuracy of the classification process against each of the test cases and visually compared the results to assess the performance. Knowing the location of the emission feature (low- z maximum point) compared to the timing of the other key points helps to

improve the quality of the classification.² The success rate of the algorithm is 99.9 per cent as is demonstrated by the confusion matrix (Fig. 8). Note that for this test we only used cases with $T_{\max}^{\text{lo}} > 0.2$ mK. This is because cases with a lower (but still positive) emission peak are really neither positive nor negative, and misclassification in this case does not lead to an inaccurate prediction of the 21-cm signal itself.

4.1.2 PCA

The core of our emulator is PCA which, given a database, finds an orthogonal basis that spans the data. Eigenfunctions (or eigenvectors) of this basis are smooth functions found using the covariance matrix of the data; while eigenvalues are a measure of the variance of the data along each particular eigenvector. The basis is constructed so that the first principal component (the eigenvector with the largest eigenvalue) has the largest possible variance, the component with the second greatest variance is the second principal component, and so on. Using the basis of smooth functions to represent the 21-cm signal guarantees the smoothness of the outcome.

The astrophysical key points divide each positive/negative signal into 3/2 distinct frequency segments. For a positive signal, the segments are $s_1 \in [\nu_{\max}^{\text{hi}}, \nu_{\min}]$, $s_2^p \in [\nu_{\min}, \nu_{\max}^{\text{lo}}]$, and $s_3^p \in [\nu_{\max}^{\text{lo}}, \nu_{\text{reion}}]$; while for a negative signal, the segments are $s_1 \in [\nu_{\max}^{\text{hi}}, \nu_{\min}]$ and $s_2^n \in [\nu_{\min}, \nu_{\text{reion}}]$. We find it best to split the data into these segments and analyse them separately. Note that, because the first segment, s_1 , is defined identically for both positive and negative cases, over s_1 all signals are analysed together; while over other segments the positive

²We found that the original algorithm misclassified 5.5 per cent of negative cases as positive. For almost all of these cases, at the output of the algorithm the order of the predicted redshifts of the key points was wrong (e.g. the redshift of the low- z maximum point was predicted to be higher than the redshift of the absorption trough, which is unphysical). We used the ill-ordered key points as a diagnostic and for such cases changed the classification of the model from positive to negative. After this procedure, the algorithm returned the correct answer in 99.9 per cent of cases.

Predicted class	Negative	1014 63.5%	0 0.0%	100% 0.0%
	Positive	2 0.1%	580 36.3%	99.7% 0.3%
		99.8% 0.2%	100% 0.0%	99.9% 0.1%
		Negative	Positive	True class

Figure 8. Confusion matrix of the classifier. The green squares show the number of correctly classified positive and negative cases, the red squares show the number of misclassified cases, grey shows the percentage of the correct predictions for each row/column, and blue shows the total accuracy defined as the ratio of the total number of correctly identified cases (both positive and negative) to the total number of considered cases. The classification is correct in 99.9 per cent of cases.

and negative cases are treated separately. In order to uniformly normalize the signals within each segment of the data, we perform a coordinate transformation into a new coordinate system x_s, y_s in which each signal varies in the range $x_s \in [0, 1]$ and $y_s \in [-1, 1]$. For instance, on s_1 the following coordinate transformation from the ν - T_{21} plane to the x_s - y_s plane is performed:

$$x_s = \frac{\nu - \nu_{\max}^{\text{hi}}}{\nu_{\min} - \nu_{\max}^{\text{hi}}}, \quad y_s = \frac{T - T_{\max}^{\text{hi}}}{T_{\min} - T_{\max}^{\text{hi}}}. \quad (9)$$

In other words, for both the negative and positive signals, s_1 is chosen so that $(\nu_{\max}^{\text{hi}}, T_{\max}^{\text{hi}})$ is mapped to $(x_s, y_s) = (0, 0)$, and the minimum point (ν_{\min}, T_{\min}) is mapped to $(x_s, y_s) = (1, -1)$. Each re-normalized segment is separately analysed using PCA.

In principle, for a perfect reconstruction of the signal via PCA decomposition the number of coefficients should be the same as the size of the database (i.e. $\sim 30\,000$ in our case). However, for our data set, the first four eigenvalues strongly dominate, allowing us to truncate the basis and use only the first four eigenfunctions to represent the signal in each segment. Fig. 9 is an illustration of the PCA decomposition for a signal over s_1 . We show a re-normalized signal (blue) and the first four eigenfunctions of the basis (shades of brown). The red curve shows the sum of the first four PCA components (each with its corresponding coefficient), reproducing the true signal nearly perfectly. We quantify the accuracy of the reconstruction process along each segment by calculating the rms of the error defined as

$$\text{rms} = \sqrt{\text{mean} \left[(y_{s,\text{sim}}(x_s) - y_{s,\text{pred}}(x_s))^2 \right]}. \quad (10)$$

The mean rms error across all the reconstruction cases is 0.0020 on s_1 , 0.0058 and 0.0075 on s_2^p and s_3^p , respectively, and 0.0045 for s_2^n .

4.1.3 Training of the NNs

Having created the PCA decomposition for each set of astrophysical parameters from the training data set of 27 455 cases, we tabulated




Figure 9. Example of the PCA decomposition for a signal in the x_s - y_s space on s_1 . The four PCA components used to reconstruct this segment are shown in shades of brown: the first component is the darkest and the fourth component is the lightest. The sum of the first four PCA eigenfunctions each with the corresponding coefficient, is shown in red, while the original signal is in blue (the reconstruction is so good that the lines overlap). The signal reconstructed by NNs using the predicted PCA coefficients is shown as dashed black. The first four PCA components are enough to represent the signal in each segment.

the values of the PCA coefficients along with the key points (both the frequency and the corresponding brightness temperature of each key point). Using this library, NNs were trained to retrieve the PCA coefficients along with the values of the key points given an input set of the astrophysical parameters. Architecturally, all the NNs described in this section are identical having one hidden layer of 40 neurons and employing the Levenberg–Marquardt algorithm to minimize the mean-squared error between the true value provided by the training data set and the value predicted by the network.

We found that the accuracy of the prediction is improved if we add to the modelling combinations of the astrophysical parameters that we expect to map more directly to the 21-cm signal. We made use of the fact that we know the simulated cosmology. Assuming the standard collisionless cold dark matter scenario and hierarchical structure formation, we can infer the mean collapse fraction at every redshift ($f_{\text{coll}}(z)$), the fraction of mass that is contained in haloes of mass above the minimum cooling threshold, Barkana & Loeb 2004). We appended five more parameters to the original set of the seven astrophysical parameters, bringing the total number of input parameters of each NN to 12. The auxiliary parameters include: $f_{\alpha} f_{\text{coll}}(20)$ which is proportional to the intensity of the Ly α radiation before the Ly α coupling; $f_{\alpha} f_{\text{coll}}(15)$ which scales as the intensity of X-ray radiation before the heating saturation; and $\zeta f_{\text{coll}}(10)$ which is a measure of the ionizing radiation at the onset of reionization. In addition, we added the fraction of X-ray energy above 1 keV and the fraction of X-ray energy above 2 keV ($f_{\text{XR} > 1 \text{ keV}}$ and $f_{\text{XR} > 2 \text{ keV}}$, respectively) to characterize the X-ray SED. Lastly, we applied physical cuts on the predicted signal to assist the NNs. In particular, an upper limit of $T_{\max}^{\text{hi}} = 0$ was imposed because T_{\max}^{hi} is expected to always be negative in the range of scenarios considered here. We also set a lower limit on the signal at the minimum point, T_{\min} , in accordance with equation (6).

In total, predicting a positive/negative signal requires generating 19/13 parameters: four PCA coefficients for each of the three/two

Table 1. Accuracy in prediction of the key points ν_{\max}^{hi} , T_{\max}^{hi} , ν_{\min} , T_{\min} , ν_{\max}^{lo} , T_{\max}^{lo} , and ν_{reion} . The percentage of cases with a relative error below 2 per cent (5 per cent) in the prediction is shown in the second (third) row.

Key point	ν_{\max}^{hi}	T_{\max}^{hi}	ν_{\min}	T_{\min}	ν_{\max}^{lo}	T_{\max}^{lo}	ν_{reion}
Prediction error below 2 per cent [per cent]	97.44	99.59	77.02	78.88	66.88	38.54	96.68
Prediction error below 5 per cent [per cent]	100	100	99.42	98.28	98.57	59.87	99.88

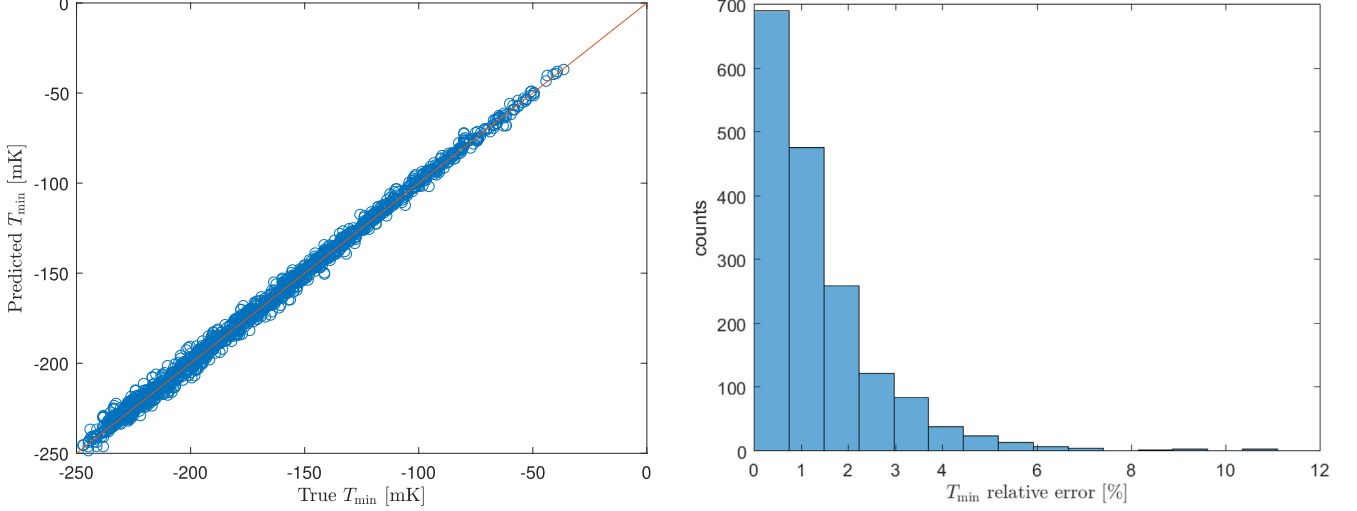


Figure 10. Left-hand panel: predicted versus true value of the amplitude of the absorption feature, shown for 1743 cases. Also shown is the perfect prediction ($Y = X$, solid red line). Right-hand panel: histogram of the relative error in the predicted amplitude of the absorption trough. 98.28 per cent of cases have a relative error of less than 5 per cent.

segments plus four/three key points each having two coordinates (frequency and brightness temperature), minus one degree of freedom because the value of the brightness temperature at ν_{reion} is by definition zero. As part of the optimization process, we had to choose between using one network which would predict all the 19/13 output parameters, 19/13 networks each of which would return a single parameter, or a few NNs predicting groups of the parameters. We found that predicting several parameters with a single network sometimes decreases the error in the predicted signal. However, it also can result in outliers, that is a few cases with very large error. To minimize the frequency of outliers while preserving the overall accuracy, we decided to group correlated parameters within the same network. For instance, the four PCA components of a given segment are correlated and were computed with a single NN that has 12 input parameters (and, thus, 12 input neurons), one hidden layer of 40 neurons and four outputs (the PCA coefficients). The outcome of this prediction is demonstrated in Fig. 9 (dashed black line). Using equation (10), we assessed the performance of the prediction process and found an r.m.s. error of 0.023 on s_1 , 0.055 and 0.136 on s_2^p and s_3^p respectively, and 0.031 on s_2^q .

Two NNs were trained to predict the coordinates of the critical points in the ν - T_{21} space for positive/negative cases. These NNs have 12 input parameters and 7/5 outputs (4/3 temperature values and 3/2 frequency coordinates). The accuracy of the reconstruction of these coordinates is summarized in Table 1. We found that the algorithm is well tuned to predict the signal from Cosmic Dawn, with 100 per cent of cases having better than 5 per cent accuracy in the amplitude and the location of the high-redshift maximum, and more than 98 per cent of cases having better than 5 per cent accuracy in the prediction of the depth and location of the absorption

trough. The low-redshift maximum point is the hardest to predict since it is affected by all the astrophysical parameters and also the amplitude of the signal at this point is quite small. In addition, because this feature does not exist for negative cases, the training data set which could be used for $(\nu_{\max}^{\text{lo}}, T_{\max}^{\text{lo}})$ was smaller. The maximal absolute error obtained when predicting T_{\max}^{lo} was 2.3 mK, with 59.87 per cent of cases returning relative errors smaller than 5 per cent. In 98.57 per cent of cases ν_{\max}^{lo} was found to better than a 5 per cent error. Finally, the success rate for the prediction of the timing of reionization was close to 100 per cent.

Separate, but architecturally identical (with 12 input parameters, one hidden layer of 40 neurons, one output), NNs were trained to predict frequencies at which the neutral fraction is 0.16 per cent (ν_{16} per cent, which we used in Section 2.4) and 11 per cent along with values of the neutral fraction at $z = 5.9, 7.08$, and 7.54 .

As an illustration, in Fig. 10 we show the accuracy of the algorithm in reconstructing the amplitude of the absorption feature for 1743 cases (all our test cases that were not excluded by the observational constraints in Section 2.4). The line $Y = X$ corresponds to a perfect prediction. The scatter shows the error in this prediction, which is also quantified in the histogram (right-hand panel). We find that 98.28 per cent of cases have a relative error of less than 5 per cent, while 78.88 per cent of cases have an error less than 2 per cent (as indicated in Table 1).

4.2 Prediction pipeline

Using the trained NNs, the global 21-cm signal is predicted given a set of seven input parameters. The complete prediction algorithm is summarized in Fig. 11 and contains the following steps:

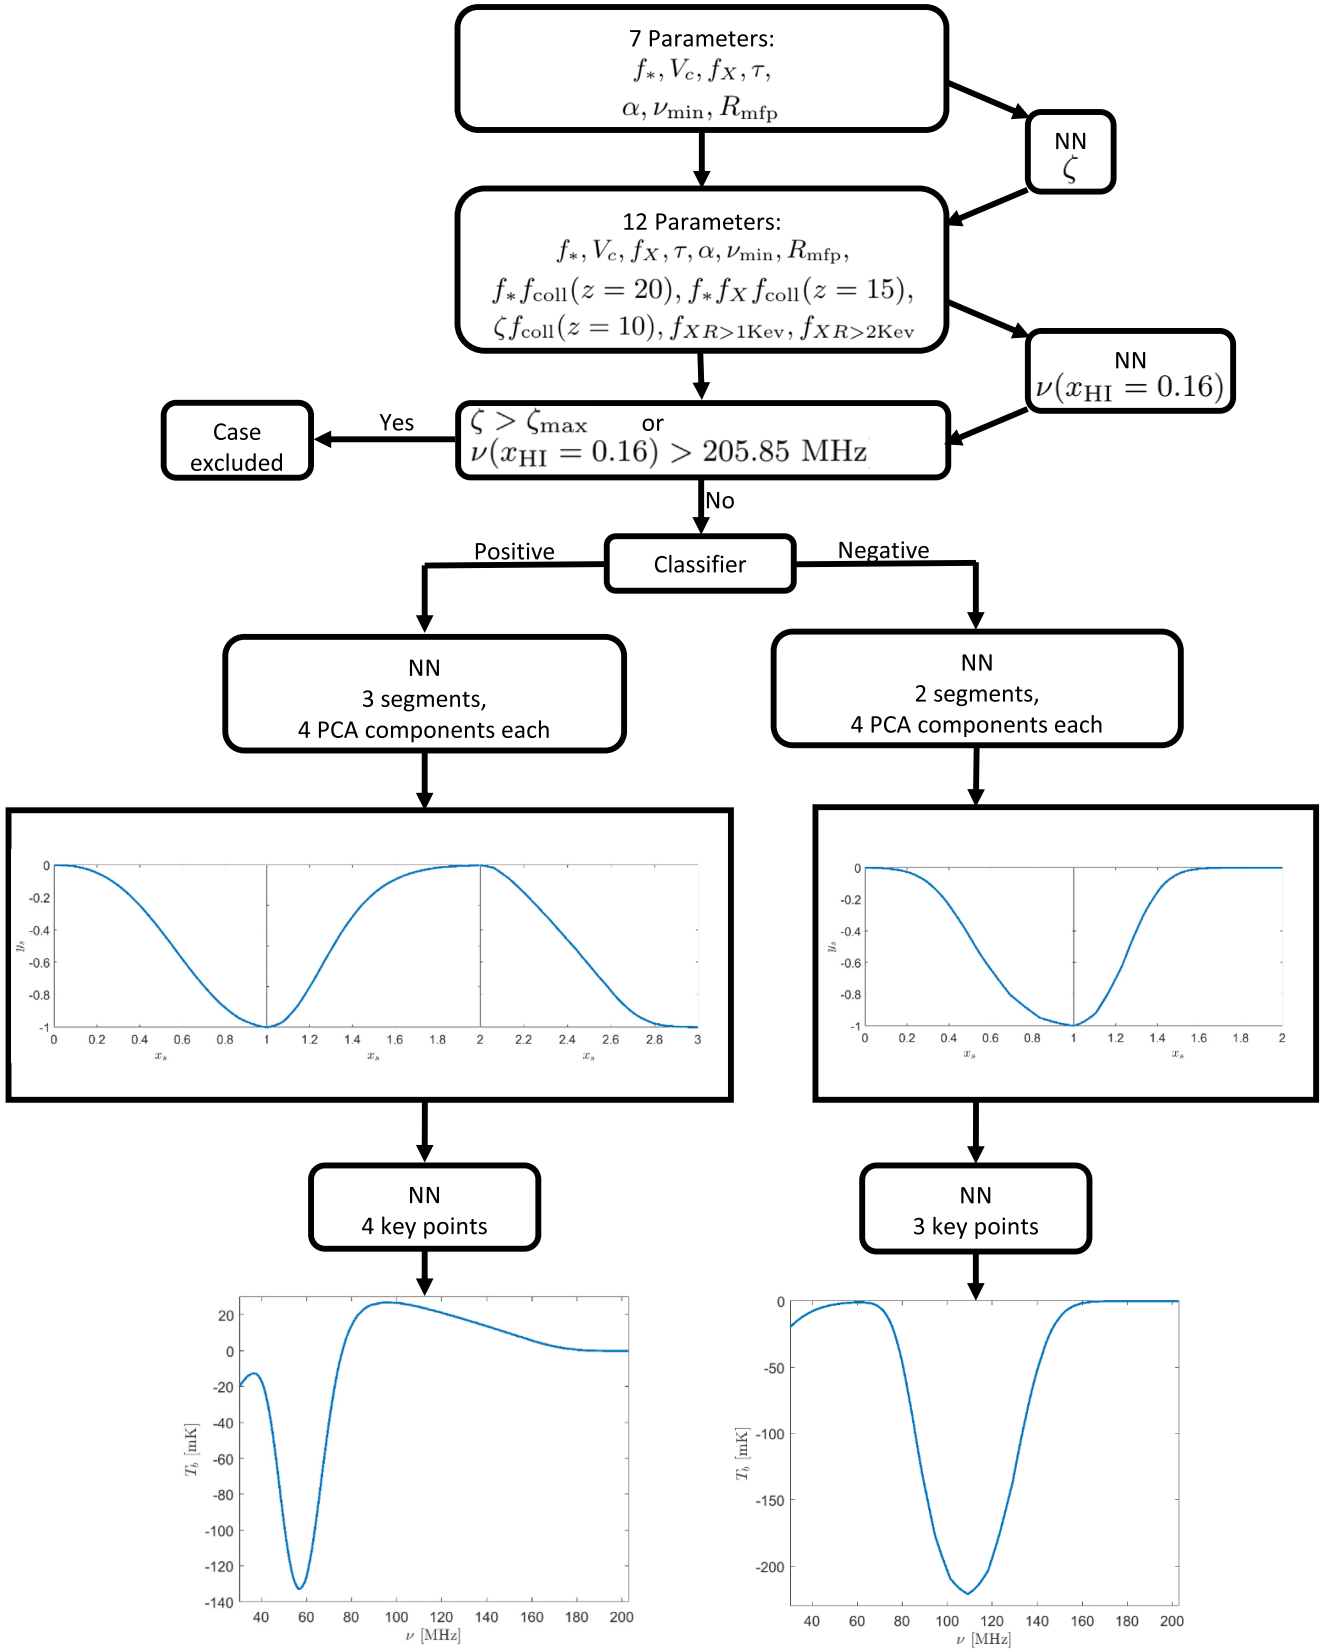


Figure 11. Flowchart of the prediction process as outlined in Section 4.2.




Figure 12. Histogram of the errors in the predicted global 21-cm signals, normalized by their maximum amplitude, as defined in equation (11). The results are shown for the entire test set of 1743 cases. The error is lower than 0.05 for 98.9 per cent of cases.

(i) Given the seven input astrophysical parameters $[f_*, V_c, f_X, \tau, \alpha, \nu_{\min}, R_{\text{mfp}}]$ trained NN (Section 2.3) is used to infer the value of the ionizing efficiency ζ . The algorithm calculates five auxiliary parameters $f_* f_{\text{coll}}(20), f_* f_X f_{\text{coll}}(15), \zeta f_{\text{coll}}(10), f_{\text{XR} > 1 \text{ keV}}, f_{\text{XR} > 2 \text{ keV}}$.

(ii) Using the full set of parameters and an NN that predicts $\nu_{16 \text{ per cent}}$ the algorithm verifies whether this case has a valid EoR history or not (as described in Section 2.4).

(iii) If the case is valid, the algorithm uses decision trees to classify the case and determines if it is expected to have an emission feature or not (i.e. whether the case is negative or positive, Section 4.1.1).

(iv) Based on the input parameters, NNs (Section 4.1.3) predict the PCA coefficients for each of the segments as well as the coordinates of the key astrophysical points (as explained in Section 4.1.2).

(v) A coordinate transformation (inverse of Equation 9) is performed to return the signal in physical units of mK as a function of frequency in MHz.

4.3 Performance analysis

In this section, we test the overall performance of the emulator, assessing its accuracy in predicting the global signal for each of the test cases in the set of 1743 signals.

We define the error in the predicted signal, $T_{\text{pred}}(\nu)$, compared to the signal generated by the full simulation for the same parameter set, $T_{\text{sim}}(\nu)$, as the rms value of the difference between the two signals, normalized by the maximal amplitude of the true signal:

$$\text{Error} = \frac{\sqrt{\text{mean} \left[(T_{\text{sim}}(\nu) - T_{\text{pred}}(\nu))^2 \right]}}{\max |T_{\text{sim}}(\nu)|}. \quad (11)$$

Over the entire test set the mean value for the error is 0.0159 and the median is 0.0130. The histogram of the errors for all the tested cases is shown in Fig. 12. We find that the error is lower than 0.05 for 98.9 per cent of cases.

To illustrate the performance of the emulator, we show several specific cases in Fig. 13: (a) the case with a 10th percentile error (i.e. 10 per cent of the cases have smaller error), with an error of 0.0072, (b) the median error of 0.013, (c) the mean error of 0.0159, (d) 90th

percentile error of 0.0271, (e) 95th percentile error of 0.0349, and (f) the largest error of 0.1055. Visually, the cases with the mean and median errors (top right and middle left panels) are in excellent agreement with the simulated signal.

4.4 Limitations

21CMGEM is designed to cover a wide range of redshifts (5–50) and was optimized to return a small mean error over the entire range. This is both an advantage and a disadvantage. The reionization parameters τ and R_{mfp} only affect the low-redshift portion of the signal (at $z \lesssim 10$) where the amplitude of the signal is very low (compared to the deep absorption trough at Cosmic Dawn). Therefore, if 21CMGEM were used as a part of MCMC to recover these parameters from data, large errors would be expected. For instance, $R_{\text{mfp}} = 70$ Mpc results in a slightly faster end to reionization, compared to 50 Mpc. However, the difference between the global signals with $R_{\text{mfp}} = 70$ and 50 Mpc is very small. For example, for a model with $V_c = 16.5 \text{ km s}^{-1}$, $f_* = 0.05$, $f_X = 1$, hard SED, and $\tau = 0.073$, the error is 9.7×10^{-4} , which is much smaller than the typical precision of 21CMGEM with the median value of rms of 0.01, as discussed in Section 4.3 of this paper. A related shortcoming is the precision of 21CMGEM in reconstructing the emission feature and the large error on $T_{\text{max}}^{\text{lo}}$ (as we discussed in Section 4.1.3).

5 SUMMARY AND CONCLUSIONS

In this paper, we have presented a database of 29 641 global 21-cm signals generated over the widest possible space of seven astrophysical parameters that include the star formation efficiency, minimum cooling mass, X-ray radiation efficiency, the slope and the low-energy cut-off of the X-ray spectrum, the mean-free path of the ionizing photons, and the CMB optical depth. The parameter space is constrained by the observations of the CMB and quasar absorption lines as well as by the maximum possible ionizing efficiency (corresponding to massive metal-free stars).

We used this data set to verify the consistency relations between the astrophysical parameters and the properties of the global 21-cm signal first reported in our previous paper (Cohen et al. 2017), finding a good agreement in all relations except for the value of J_α/ϵ at z_{\min} which shows much larger scatter due to the wider selection of X-ray spectra considered here. In particular, there remains a tight predicted relationship between the brightness temperature and the observed frequency of the high-redshift maximum point (Fig. 4); a measurement of this point can be used to infer the Ly α intensity at that time, though with significant scatter (Fig. 5). Also, the brightness temperature and observed frequency of the low-redshift maximum point follow a tight relation, which can be used to estimate the X-ray intensity (Fig. 7).

We utilized the database to develop and test 21CMGEM which, given a set of astrophysical parameters, predicts the global 21-cm signal. Additional outputs include values of the neutral fraction at $z = 5.9, 7.08$, and 7.54 along with frequencies at which the neutral fraction is 0.16 per cent and 11 per cent. The crucial elements of the emulator are:

- (i) Smoothness of the output signal is guaranteed by construction.
- (ii) The database can be divided into two categories: signals that have an emission feature and signals that are only seen in absorption. The classification is done using bagged decision trees.




Figure 13. Comparison between the simulated signal (blue) and the predicted signal (red). The error and parameters [f_* , V_c , f_X , τ , α , ν_{\min} , R_{mfp}] of each panel are: (a) 10th percentile error (error = 0.0072; [0.050,76.5,0.001,0.0781,1.3,1,20]); (b) median error (error = 0.0130; [0.5,4.2,0.1,0.0779,1.0,5,20]); (c) mean error (error = 0.0159; [0.0285,5.41,3.01,0.0738,1.0,5,34]); (d) 90th percentile error (error = 0.0271; [0.2573,21.50,0.0001,0.0701,1.5,0.7,14]); (e) 95th percentile error (error = 0.0349; [0.5,24.2,0.02,0.0666,1.5,0.1,35]); and (f) the largest error (error = 0.1055; [0.3835,7.79,0.0045,0.1,1.5,0.2,48.30]).

(iii) We train NNs to predict the cosmological signal based on the seven astrophysical parameters. Each signal is broken into a few segments separated by the key astrophysical points and is decomposed into a basis of smooth orthogonal functions using PCA. Because PCA ranks the eigenfunctions by variance in decreasing order, most of the information is encoded in the first few terms. This allows us to reduce the dimensionality and use only the first four functions of the basis. NNs are used to predict the PCA coefficients as well as the two ends of each segment given a set of astrophysical parameters.

(iv) The algorithm also checks whether the case satisfies current constraints on reionization. The constraints that are taken into account include limits on the total CMB optical depth, the upper limit on the ionization efficiency of stars, and the upper limit on the neutral fraction at $z \sim 5.9$ derived from the absorption profile of high-redshift quasars. Using these conditions, the minimum circular velocity of star-forming haloes as well as the star formation efficiency can be constrained. We find a lower limit of $V_c \sim 26 \text{ km s}^{-1}$ ($\sim 4 \times M_{\text{min}}^{\text{atomic}}$) for an optical depth of 0.055.

The algorithm was trained on 27 455 simulated signals, and an additional 2186 cases were used as the test set. The predicted signal has an rms error of 0.0159, corresponding to 1.59 per cent of the signal amplitude, with 98.9 per cent of cases having errors lower than 0.05. The algorithm is efficient, with running time per parameter set of 0.16 s (while one full simulation run typically

takes a few hours on a single core). This tool can be used in the fitting process (e.g. MCMC) to constrain the high-redshift parameter space using the data of global signal experiments. We have used it recently with the data from EDGES High-Band (Monsalve et al. 2019). 21CMGEM and the training and testing data sets are available online at <https://www.ast.cam.ac.uk/~afialkov/Publications.html>.

ACKNOWLEDGEMENTS

We acknowledge the usage of the Harvard Odyssey cluster. Part of this project was supported by the Royal Society University Research Fellowship of AF. RB and AC acknowledge Israel Science Foundation grant 823/09 and the Ministry of Science and Technology, Israel. This project/publication was made possible for RB through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. RB was also supported by the ISF-NSFC joint research program (grant no. 2580/17). RAM was supported by the NASA Solar System Exploration Virtual Institute cooperative agreement 80ARC017M0006.

REFERENCES

- Abel T., Bryan G. L., Norman M. L., 2002, *Science*, 295, 93
 Alvarez M. A., Abel T., 2012, *ApJ*, 747, 126

- Bañados E. et al., 2018, *Nature*, 553, 473
- Barkana R., 2018, *Nature*, 555, 71
- Barkana R., Loeb A., 2001, *Phys. Rep.*, 349, 125
- Barkana R., Loeb A., 2004, *ApJ*, 609, 474
- Barry N. et al., 2019, *ApJ*, 884, 1
- Beardsley A. P. et al., 2016, *ApJ*, 833, 102
- Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
- Bernardi G. et al., 2016, *MNRAS*, 461, 2847
- Bowman J. D., Rogers A. E. E., 2010, *Nature*, 468, 796
- Bowman J. D., Rogers A. E. E., Monsalve R. A., Mozdzen T. J., Mahesh N., 2018, *Nature*, 555, 67
- Breiman L., 1996, *Mach. Learn.*, 24, 123
- Bromm V., Kudritzki R. P., Loeb A., 2001, *ApJ*, 552, 464
- Bromm V., Coppi P. S., Larson R. B., 2002, *ApJ*, 564, 23
- Cirelli M., Iocco F., Panci P., 2009, *J. Cosmol. Astropart. Phys.*, 10, 009
- Cohen A., Fialkov A., Barkana R., 2016, *MNRAS*, 459, L90
- Cohen A., Fialkov A., Barkana R., Lotem M., 2017, *MNRAS*, 472, 1915
- D'Amico G., Panci P., Strumia A., 2018, *Phys. Rev. Lett.*, 121, 1103
- Dalal N., Pen U.-L., Seljak U., 2010, *J. Cosmol. Astropart. Phys.*, 11, 007
- DeBoer D. R. et al., 2017, *PASP*, 129, 5001
- Eastwood M. W. et al., 2019, *AJ*, 158, 84
- Evoli C., Mesinger A., Ferrara A., 2014, *J. Cosmol. Astropart. Phys.*, 11, 024
- Ewall-Wice A., Chang T.-C., Lazio J., Dore O., Seiffert M., Monsalve R. A., 2018, *ApJ*, 868, 63
- Ewall-Wice A., Chang T.-C., Lazio T. J. W., 2020, *MNRAS*, 492, 6086
- Feng C., Holder G., 2018, *ApJ*, 858, 17
- Fialkov A., Barkana R., Tseliakhovich D., Hirata C. M., 2012, *MNRAS*, 424, 1335
- Fialkov A., Barkana R., Visbal E., Tseliakhovich D., Hirata C. M., 2013, *MNRAS*, 432, 2909
- Fialkov A., Barkana R., Visbal E., 2014, *Nature*, 506, 197
- Fialkov A., Cohen A., Barkana R., Silk J., 2016, *MNRAS*, 464, 3498
- Fialkov A., Barkana R., Cohen A., 2018, *Phys. Rev. Lett.*, 121, 1101
- Field G. B., 1958, *PIRE*, 46, 240
- Fragos T., Lehmer B. D., Naoz S., Zezas A., Basu-Zych A., 2013, *ApJ*, 776, L31
- Gehlot B. K. et al., 2019, *MNRAS*, 488, 4271
- Ghara R. et al., 2020, *MNRAS*, 493, 4728
- Gilfanov M., Grimm H.-J., Sunyaev R., 2004, *MNRAS*, 347, 57
- Greig B., Mesinger A., 2015, *MNRAS*, 449, 4246
- Greig B., Mesinger A., 2018, *MNRAS*, 477, 3217
- Greig B., Mesinger A., Haiman Z., Simcoe R., 2017, *MNRAS*, 466, 4239
- Grimm H.-J., Gilfanov M., Sunyaev R., 2003, *MNRAS*, 339, 793
- Haiman Z., Rees M. J., Loeb A., 1997, *ApJ*, 476, 458
- Haiman Z., Abel T., Rees M. J., 2000, *ApJ*, 534, 11
- Hassan S., Dave R., Finlator K., Santos M. G., 2017, *MNRAS*, 468, 22
- Hills R., Kulkarni G., Meerburg P. D., Puchwein E., 2018, *Nature*, 564, 32
- Jennings W. D., Watkinson C. A., Abdalla F. B., McEwen J. D., 2019, *MNRAS*, 483, 2907
- Jeon M., Pawlik A. H., Bromm V., Milosavljevic M., 2014, *MNRAS*, 444, 3288
- Kern N. S., Liu A., Parsons A. R., Mesinger A., Greig B., 2017, *ApJ*, 848, 23
- Kolopanis M. et al., 2019, *ApJ*, 883, 133
- Koopmans L. et al., 2015, *Proc. Sci.*, The Cosmic Dawn and Epoch of Reionisation with SKA. SISSA, Trieste, PoS(AASKA14)001
- Levenberg K., 1944, *Quart. Appl. Math.*, 2, 164
- Li W. et al., 2019, *ApJ*, 887, 141
- Liu A., Pritchard J. R., Allison R., Parsons A. R., Seljak U., Sherwin B. D., 2016, *Phys. Rev. D*, 93, 3013
- Liu H., Ridgway G. W., Slatyer T. R., 2020, *Phys. Rev. D*, 101, 353
- Loh W. Y., Shih Y. S., 1997, *Stat. Sin.*, 7, 815
- Machacek M. E., Bryan G. L., Abel T., 2001, *ApJ*, 548, 509
- Madau P., Rees M. J., Volonteri M., Haardt F., Oh S. P., 2004, *ApJ*, 604, 484
- Marquardt D., 1963, *SIAM J. Appl. Math.*, 11, 431
- Mason C. A., Naidu R. P., Tacchella S., Leja J., 2019, *MNRAS*, 489, 2669
- McGreer I. D., Mesinger A., D'Odorico V., 2015, *MNRAS*, 447, 499
- Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Mesinger A., Ferrara A., Spiegel D. S., 2013, *MNRAS*, 431, 621
- Mineo S., Gilfanov M., Sunyaev R., 2012, *MNRAS*, 419, 2095
- Mirabel I. F., Dijkstra M., Laurent P., Loeb A., Pritchard J. R., 2011, *A&A*, 528, 149
- Mirocha J., Furlanetto S. R., Sun G., 2017, *MNRAS*, 464, 1365
- Mondal R. et al., 2020, preprint ([arXiv:2004.00678](https://arxiv.org/abs/2004.00678))
- Monsalve R. A., Rogers A. E. E., Bowman J. D., Mozdzen T. J., 2017, *ApJ*, 847, 64
- Monsalve R. A., Greig B., Bowman J. D., Mesinger A., Rogers A. E. E., Mozdzen Th. J., Kern N. S., Mahesh N., 2018, *ApJ*, 863, 11
- Monsalve R. A., Fialkov A., Bowman J. D., Rogers A. E. E., Mozdzen T. J., Cohen A., Barkana R., Mahesh N., 2019, *ApJ*, 875, 67
- Muñoz J. B., Loeb A., 2018, *Nature*, 557, 684
- Muñoz J. B., Kovetz E. D., Ali-Haïmoud Y., 2015, *PRD*, 92, 3528
- Navarro J. F., Steinmetz M., 2000, *ApJ*, 538, 477
- O'Shea B. W., Norman M. L., 2008, *ApJ*, 673, 14
- O'Shea B. W., Wise J. H., Xu H., Norman M. L., 2015, *ApJ*, 807, 12
- Oh S. P., 2001, *ApJ*, 553, 499
- Paciga G. et al., 2013, *MNRAS*, 433, 639
- Park J., Gillet N., Mesinger A., Greig B., 2020, *MNRAS*, 491, 3891
- Patil A. H. et al., 2017, *ApJ*, 838, 65
- Pearson K., 1901, *Philos. Mag.*, 2, 559
- Planck Collaboration XIII, 2016, *A&A*, 596, A107
- Price D. C. et al., 2018, *MNRAS*, 478, 4193
- Rees M. J., 1986, *MNRAS*, 222, 27
- Santos M. G., Ferramacho L., Silva M. B., Amblard A., Cooray A., 2010, *MNRAS*, 406, 2421
- Schauer A. T. P., Whalen D. J., Glover S. C. O., Klessen R. S., 2015, *MNRAS*, 454, 2441
- Schauer A. T. P., Glover S. C. O., Klessen R. S., Ceverino D., 2019, *MNRAS*, 484, 3510
- Schmit C. J., Pritchard J. R., 2018, *MNRAS*, 475, 1213
- Shimabukuro H., Semelin B., 2017, *MNRAS*, 468, 3869
- Sims P. H., Pober J. C., 2020, *MNRAS*, 492, 22
- Singh S., Subrahmanyam R., 2019, *ApJ*, 880, 26
- Singh S. et al., 2017, *ApJ*, 845, L12
- Singh S. et al., 2018, *ApJ*, 858, 54
- Sobacchi E., Mesinger A., 2013, *MNRAS*, 432, 3340
- Spinelli M., Bernardi G., Santos M., 2019, *MNRAS*, 489, 4007
- Tashiro H., Kadota K., Silk J., 2014, *Phys. Rev. D*, 90, 3522
- Tegmark M., Silk J., Rees M., Blanchard A., Abel T., Palla F., 1997, *ApJ*, 474, 1
- Trott C. M. et al., 2020, *MNRAS*, 493, 4711
- Tselikhovich D., Hirata C. M., 2010, *Phys. Rev. D*, 82, 3520
- Visbal E., Barkana R., Fialkov A., Tseliakhovich D., Hirata C. M., 2012, *Nature*, 487, 70
- Visbal E., Zoltan H., Terrazas B., Bryan G. L., Barkana R., 2014, *MNRAS*, 445, 107
- Weinberg D. H., Hernquist L., Katz N., 1997, *ApJ*, 477, 8
- Wise J. H., Abel T., 2007, *ApJ*, 671, 1559
- Wise J. H., Demchenko V. G., Halicek M. T., Norman M. L., Turk M. J., Abel T., Smith B. D., 2014, *MNRAS*, 442, 2560
- Wouthuysen S. A., 1952, *AJ*, 57, 31
- Wyithe J. S. B., Loeb A., 2013, *MNRAS*, 428, 2741
- Yoshida N., Abel T., Hernquist L., Sugiyama N., 2003, *ApJ*, 592, 645
- Zarka P., Girard J. N., Tagger M., Denis L., 2012, in Boissier S., de Laverny P., Nardetto N., Samadi R., Valls-Gabaud D., Wozniak H., eds, *LSS/NenuFAR: The LOFAR Super Station project in Nançay*, p. 687