GENERATING DARK MATTER SUBHALO POPULATIONS USING NORMALIZING FLOWS

JACK LONERGAN[©] Department of Physics and Astronomy, University of Southern California

> ANDREW BENSON[®] The Observatories of the Carnegie Institution for Science

DANIEL GILMAN^O Department of Astronomy and Astrophysics, University of Chicago Version April 23, 2025

ABSTRACT

Strong gravitational lensing is a powerful tool for probing the nature of dark matter, as lensing signals are sensitive to the dark matter substructure within the lensing galaxy. We present a comparative analysis of strong gravitational lensing signatures generated by dark matter subhalo populations using two different approaches. The first approach models subhalos using an empirical model, while the second employs the GALACTICUS semi-analytic model of subhalo evolution. To date, only empirical approaches have been practical in the analysis of lensing systems, as incorporating fully physical models was computationally infeasible. To circumvent this, we utilize a generative machine learning algorithm, known as a normalizing flow, to learn and reproduce the subhalo populations generated by GALACTICUS subhalo distribution while significantly reducing computation time compared to direct simulation. Moreover, we find that subhalo populations from GALACTICUS produce comparable results to the empirical model in replicating observed lensing signals under the fiducial dark matter model. This work highlights the potential of machine learning techniques in accelerating astrophysical simulations and improving model comparisons of dark matter properties.

Subject headings: First, second

1. INTRODUCTION

Dark matter is an unknown form of non-baryonic matter that composes up to 85% of all matter in the universe (Ade et al. 2016; Aghanim et al. 2020). Given that dark matter constitutes the majority of matter in the universe, it plays a significant role in the formation and evolution of structure throughout the universe's history. Quantum fluctuations in the early universe produced slight overdensities in an otherwise homogeneous matter density field, and these overdensities eventually formed gravitationally bound structures known as dark matter halos (Spergel et al. 2003; Hinshaw et al. 2013). These halos merged over time, resulting in a present-day hierarchical distribution of matter (Davis et al. 1985; Cole & Lacey 1996; Makishima 1998). The current cosmological model is the Lambda cold dark matter (Λ CDM) which accurately models the large-scale structure of the universe (Frenk et al. 1985; Peebles 2020). The fiducial model begins to experience tensions with observations on smaller, galactic scales, however, and many studies have proposed models to accurately replicate the distribution of dark matter to match observations on these scales (see Bul-

jacklone@usc.edu abenson@carnegiescience.edu gilmanda@uchicago.edu lock & Boylan-Kolchin 2017, for a review of problems and proposed solutions).

A wide range of techniques for simulating the formation and evolution of dark matter substructure in galaxies have been developed, ranging from analytic models (Zentner et al. 2005; Giocoli et al. 2008; Yang et al. 2011) to N-body simulations (Moore et al. 1999; Springel et al. 2008; Nadler et al. 2023). Analytic models offer an efficient way to model a broad range of halos at the cost of making simplifying physical assumptions to allow the model to be solved analytically. N-body simulations, on the other hand, numerically solve the equations governing an evolving matter density field in an expanding universe. The development of such simulations has allowed us to replicate the distribution of matter at high resolutions over large time intervals, providing a detailed picture of the evolution of the universe (Mikkola & Aarseth 1993; Bertschinger 1998; Bagla 2005). Solving these equations numerically, however, demands significant computational resources, making the computational cost one of the main drawbacks of N-body simulations (Dehnen & Read 2011). One middle-ground approach between analytic models and simulations is the class of semi-analytic models (SAMs) of subhalo evolution. As with N-body simulations, these models perform simulations with the modification that, in SAMs, the fundamental objects are halos themselves, rather than individual mass particles. These halos take on simple geometries (i.e., spherical halos), which allows for some solutions to become analytic and greatly reduces the computational cost where analytic solutions are unavailable (Taylor & Babul 2001; Benson et al. 2004; Zentner et al. 2005; Pullen et al. 2014; Jiang et al. 2021; Dekker et al. 2022). Each method presents a trade-off between computational cost and accuracy, so different models are chosen depending on the specific requirements of their application.

In addition to a wide variety of approaches to modeling dark matter, there are also many ways in which we can detect dark matter, despite its lack of interactions with electromagnetic radiation. Here, we specifically focus on one such method-strong gravitational lensing of quasar images—which can be used to constrain dark matter particle properties (Gilman et al. 2020; Vegetti & Vogelsberger 2014; Hezaveh et al. 2016). As light passes by a galaxy and its dark matter halo, its path is deflected by the gravitational potential of that galaxy, its halo, and the population of subhalos. The resulting distortion of the quasar image can be quantified to constrain the underlying population of subhalos, their distribution of masses, spatial distribution, and density profiles. While strong gravitational lensing of galaxies is merely one method for constraining dark matter models, it is considered advantageous over other methods such as analyzing the Lyman- α forest or counting of galaxies because the effect of gravitational lensing only depends on the total gravitational potential, which, in the low mass (sub)halos to which this method is sensitive, is dominated by dark matter and is not expected to be significantly altered by baryonic processes (Dutton et al. 2016). As such, strong gravitational lensing bypasses the need to understand the complex, baryonic physics of the system (Treu 2010; Gannon et al. 2025). Additionally, there are typically high degrees of uncertainty when attempting to distinguish the gravitational effects of dark matter when also accounting for baryonic matter (Fitts et al. 2019), making it convenient to work with measurements that are largely unaffected by baryonic physics.

Previous studies have shown that dark matter models can be constrained using imaging of gravitationally strongly lensed quasars (Dalal & Kochanek 2002; Nierenberg et al. 2014; Xu et al. 2015; Hsueh et al. 2020; Gilman et al. 2021; Nadler et al. 2021)x. For example, Gilman et al. (2020) used a parameterized model to generate dark matter halo populations within the lensing galaxy (which going forward will be referred to as the "empirical model"), repeating this process of creating halo populations between 300,000 to 1,200,000 times for a given lens system. This was done to obtain a statistically representative distribution of subhalo populations responsible for the lensing signature and thereby assess the relative likelihood of a given dark matter model reproducing the observed lensing statistics. While this model accurately reproduces certain summary statistics of the subhalo population, such as the subhalo mass function and distribution of infall redshifts, it also makes some substantial simplifying assumptions. Specifically, it does not fully account for correlations between the orbital properties of a subhalo and tidal stripping (instead it

connects tidal stripping to only the current orbital radius), growth of the host halo, and the radial distribution of subhalos. Work has been done to account for the effects of tidal stripping in analytic models (e.g. Han et al. 2016), yet underlying physical assumptions (such as a mass-independent tidal stripping function) remain present. The impact of such assumptions made in analytic models could be examined if halo populations used in the analysis accounted for these correlations (i.e. halo populations from SAM's or N-body simulations). The primary obstacle preventing such populations from being implemented is that it is computationally infeasible to generate $\mathcal{O}(10^6)$ realizations of dark matter halos directly from either SAMs or N-body simulations.

To bypass these limitations, we propose a more computationally efficient method to produce large numbers of realizations of subhalo populations with statistical properties as predicted from simulations. Specifically, we make use of a generative machine learning algorithm known as a normalizing flow. We train the normalizing flow on a set of simulation data, and then, once trained, we can sample new subhalo populations in a fraction of the time it would take to generate populations of subhalos directly from simulations. In principle, this methodology can be applied to any relevant simulation data set, but here we make use of simulated data from the GALACTICUS¹ semi-analytic model (Benson 2012) which has also been used to guide the construction of the empirical model in Gilman et al. (2020). In addition to generating GALACTICUS halo populations using the normalizing flows algorithm, we show an example application in which we apply the generated subhalo realizations to the forward modeling approach of Gilman et al. (2020) for a single lensed system.

This paper is organized as follows: In section 2, we introduce the empirical and GALACTICUS halo population models, as well as the normalizing flows algorithm. In section 3, we detail the lensing analysis conducted in (Gilman et al. 2020) and describe the implementation of emulator subhalo populations into this analysis. In section 4, we cover the emulator's efficiency and present the results of the forward modeling application. In section 5, we discuss the limitations of this machine learning approach and possibilities for future work. Finally, section 6 presents our general conclusions.

2. Methods

In this section, we first describe two different models for subhalo populations. We summarize the empirical model and then describe in detail the properties of GALACTICUS and how it generates a realization of a subhalo population. We then describe the specifics of the normalizing flows based algorithm used to replicate GALACTICUS realizations

2.1. Models

2.1.1. The Empirical Model

The dark matter halos in the empirical model are constructed from a series of analytic models—describing the number, spatial distribution, and density profiles—for

 $^{^1}$ https://github.com/galacticusorg/galacticus, we use revision 4313a4e2c7b68abbec547107e118ad22b33fc556.

both the cold dark matter and warm dark matter models. and for both subhalos and halos along the line of sight of the lensing galaxy. This paper, however, will focus specifically on CDM subhalo populations. The reason for this is that the properties of line-of-sight halos are relatively simple as they do not undergo the strong tidal evolution experienced by subhalos and already have accurately fitting equations (Sheth & Tormen 2002; Navarro 1996). Subhalos, on the other hand, are influenced by complicated, non-linear evolutionary processes such as tidal stripping, tidal heating, tidal truncation, dynamical friction, etc., whose intricate details are not fully accounted for in the empirical model. So it is these halos whose properties will be most impacted by analytic approximations. The details of the analytic model can be found in Gilman et al. (2020), with tidal physics updated as described in Keelev et al. (2024), but we summarize here the main points relevant to CDM subhalos:

• Subhalo density profiles are modeled as truncated Navarro-Frenk-White (NFW Navarro 1996) profiles:

$$\rho(r) = \frac{\rho_{\rm s}}{x(1+x)^2} \frac{\tau^2}{x^2 + \tau^2},\tag{1}$$

where $x = r/r_{\rm s}$, $\tau = r/r_{\rm t}$, $\rho_{\rm s}$ is the scale density, $r_{\rm s}$ is the scale radius, and $r_{\rm t}$ is the truncation radius. Halo masses are defined as the total mass with a spherical region of radius r_{200} within which the mean density is 200 times the critical density of the universe at the current redshift.

- The truncation radius of a given subhalo is defined by $r_t = r_{50}$, where r_{50} denotes the radius which encloses 50 times the critical density (Keeley et al. 2024).
- Subhalos in the empirical model are spatially distributed in 3D inside the virial radius of the host halo. Within this volume, subhalos are assumed to trace the mass profile of the host outside of an inner radius which is set equal to half the host scale radius. Inside that inner radius, subhalos are distributed uniformly in three dimensions, motivated by tidal disruption simulations around lens galaxies (Jiang & van den Bosch 2017). This results in a distribution that is approximately uniform in projection in the inner region of the halo that is relevant to lensing, in agreement with N-body simulations (Xu et al. 2015).
- The subhalo mass function is given by:

$$\frac{\mathrm{d}^2 N_{\mathrm{sub}}}{\mathrm{d}m \mathrm{d}A} = \frac{\Sigma_{\mathrm{sub}}}{m_0} \left(\frac{m}{m_0}\right)^{\alpha} \mathcal{F}(M_{\mathrm{halo}}, z), \qquad (2)$$

where $\Sigma_{\rm sub}$ is a normalization parameter characterizing the amplitude of the mass function, the pivot is denoted by $m_0 = 10^8 M_{\odot}$, and $\mathcal{F}(M_{\rm halo}, z)$ accounts for how the number density of halos depends on both the host halo mass and redshift (Gannon et al. 2025).

2.1.2. GALACTICUS Model

The second model that we use to generate subhalo populations in this work is the semi-analytic galaxy formation model known as GALACTICUS (Benson 2012). GALACTICUS is a highly modular code that simulates the evolution of galaxies and their dark matter halos by representing them as individual nodes in an overall merger tree. As these nodes evolve, they undergo merging events resulting in the hierarchical assembly of dark matter halos. Constructing a set of evolved dark matter halos in GALACTICUS requires two main steps: First, a merger tree is constructed, which represents, at each point in time, the set of progenitor halos from which the final halo is formed. Then, each node is evolved forward in time by numerically solving a series of differential equations that describe the relevant physics.

The initial step in setting up a halo population is to generate the merger tree. This is done using a Monte Carlo (MC) method to match predictions from the extended Press-Schechter (EPS) formalism (Bower 1991; Lacey & Cole 1994). The details of how GALACTICUS merger trees are constructed can be found in Cole et al. (2000), with merger rates updated to match the calibrated results of Parkinson et al. (2008). These calibrations were shown to be in good agreement with the measurements of progenitor mass functions from N-body simulations. For this study, merger trees were simulated for a host halo of mass $M_{\rm lens} = 10^{13.3} {\rm M}_{\odot}$ at $z_{\rm lens} = 0.5$ with a mass resolution of $10^6 {\rm M}_{\odot}$. This host mass and lens redshift were chosen to match the conditions of a typical strong lens galaxy. The mass resolution was chosen to be $10^6 M_{\odot}$ as halos below this mass have negligible contributions to the overall lensing signal. As discussed by Gilman et al. (2020), the lensing signal of a subhalo becomes undetectable somewhere in the mass range 10^{6} – $10^7 M_{\odot}$, so $10^6 M_{\odot}$ acts as a conservative lower limit. Subhalos were also restricted to lie within a 20 kpc annulus from the center of the host, as it is subhalos within the vicinity of the Einstein radius that primarily contribute to the observed lensing signature (Gannon et al. 2025).

Once a merger tree is constructed in GALACTICUS, halos in the tree are evolved forward in time by numerically solving differential equations to produce a population of dark matter halos at the epoch of interest (in this case, the redshift of the lens galaxy). Specifically, GALACTI-CUS tracks the evolution of the total mass of a given halo following the mass history along each branch of the merger tree. Orbital properties of subhalos, such as their 3D position, velocity, bound mass, and density profile, are tracked by incorporating the effects of the host halo potential, along with dynamical friction, tidal stripping, and tidal heating (Pullen et al. 2014; Yang et al. 2020; Benson & Du 2022). GALACTICUS also computes the evolution of the dark matter profile scale length, determined from the halo's concentration (Diemer & Joyce 2019).

When generating subhalo realizations using GALACTI-CUS directly, numerically solving differential equations for the evolution of each halo takes substantial amounts of time. In this work, we generated 300 merger trees in total, with each individual tree taking between 6–8 CPUhours of computation time to evolve to z_{lens} . It would be computationally impractical to generate $\mathcal{O}(10^6)$ of these trees for each lens to be studied, which motivates the search for more computationally efficient alternatives.

2.2. The Normalizing Flows Algorithm

While GALACTICUS is capable of simulating populations of dark matter subhalos directly, we instead propose training a normalizing flow emulator on data generated by GALACTICUS, and then using that emulator to generate populations of subhalos statistically consistent with those produced by GALACTICUS. Such algorithms take in an input data distribution and learn a mapping from a simple, latent space to the input data space. The mapping between latent and data spaces has some convenient transformation properties such as being invertible and at least once differentiable. The invertible mapping allows the flow-based algorithm to perform both sampling and likelihood estimation (Papamakarios et al. 2021).

The flows-based algorithm used in this work takes in a data distribution $p(\mathbf{z})$ and generates an invertible map F between the data distribution and a latent distribution $\tilde{p}(\mathbf{u})$. The relationship between the prior and likelihood distributions is modeled using the change of variables formula:

$$p(\mathbf{z}) = \tilde{p}\left(F^{-1}(\mathbf{z})\right) \left| \det\left(\frac{\partial F^{-1}}{\partial \mathbf{z}}\right) \right|.$$
 (3)

Here, the final term is the determinant of the Jacobian of F^{-1} evaluated at each component of \mathbf{z} . This term is included to account for the change in volume when transforming between the latent and data spaces. The mapping F is composed of n individual transformations $F = f_1 \circ f_2 \circ \cdots \circ f_n$, each of which is independent, differentiable, and invertible. As the data passes through the k^{th} individual transformation in the mapping f_k , it satisfies $\mathbf{w}_k = f_k(\mathbf{w}_{k-1})$. Here, \mathbf{w}_k represents the data at stage k of the transformation, so $\mathbf{w}_0 = \mathbf{u}$ and $\mathbf{w}_n = \mathbf{z}$.

In this work, our normalizing flow is composed of n = 12 affine coupling layers, which are individual bijective functions that construct the map. An individual affine transformation f acting on the data \mathbf{w}_k is defined by

$$f(\mathbf{w}_k; \mathbf{s}, \mathbf{t}) = \exp(\mathbf{s}) \odot \mathbf{w}_k + \mathbf{t}, \tag{4}$$

where \odot represents element-wise multiplication, and **s** and **t** are the vectors that contain the scale and translation coefficients for the affine transformation. The **s**, **t** coefficient vectors are obtained by passing previous iterations of the data $\mathbf{w}_{1:k-1}$ through a neural network of 5 dense layers: four intermediate ReLU activation functions, followed by a tanh activation function. Using 5 dense layers to generate the scaling and translation coefficients of the affine transformation adds complexity and non-linearity to the model. To prevent overfitting, each individual scaling and transformation layer has an L2 regularization metric with a regularization parameter of 0.01. The overall model is compiled using the Adam optimizer with a learning rate of 0.0001 (Loshchilov 2017).

The normalizing flows algorithm is specifically applied to an input data distribution that parametrizes GALACTICUS subhalos. Each subhalo is characterized by an infall² mass $M_{\rm infall}$, concentration c, bound

mass M_{bound} , infall redshift z_{infall} , truncation radius r_{t} and two dimensional projected radius³ r_{2D} taken from GALACTICUS realizations where these parameters were computed by solving the differential equations governing subhalo evolution, and together fully characterize the properties of a subhalo that are needed for lensing calculations. The collection of these six parameters for each generated subhalo forms a six-dimensional parameter space for the population of subhalos. Before this 6D distribution is input into the flows algorithm, the unnormalized data $\mathbf{x} = (M_{\text{infall}}, c, M_{\text{bound}}, z_{\text{infall}}, r_{\text{t}}, r_{2D})$ is first normalized to a new vector \mathbf{y} whose components are defined through the following set of equations:

$$y_1 = \log_{10}(M_{\text{infall}}/M_{\text{host}}),\tag{5}$$

$$=c,$$
 (6)

(0)

$$y_3 = \log_{10}(M_{\text{bound}}/M_{\text{infall}}),\tag{7}$$

$$y_4 = z_{\text{infall}},\tag{8}$$

$$y_7 = \log_1\left(r_1/R_{1,1}, r_2\right)$$

$$y_5 = \log_{10}(r_t/R_{\rm vir,\ host}),$$
 (9)

$$y_6 = \log_{10}(r_{2\rm D}/R_{\rm vir,\ host}).$$
 (10)

In these equations, the variables $M_{\rm host}$ and $R_{\rm vir,\ host}$ refer to the virial mass and radius, respectively, of the host halo as computed from GALACTICUS. These virial quantities are defined in terms of the virial density contrast $\Delta_{\rm vir}$, which in this work is based on the spherical collapse model in a universe that contains collisionless matter and a cosmological constant (Bryan & Norman 1998). Only subhalos along the line of sight within a double cone with opening angle $\theta = 3R_{\rm Ein}$, where $R_{\rm Ein}$ is the Einstein radius, contribute to the lensing signal (Gilman et al. 2020).

To evaluate the truncation radii of GALACTICUS subhalos, the ratio $\rho(r)/\rho_{\rm NFW}(r)$ of a given subhalo is evaluated at a range of radii ranging from the innermost regions of the halo out to the virial radius. Here, $\rho(r)$ is the actual density profile generated by GALACTICUS including the effects of tidal stripping and heating, and $\rho_{\rm NFW}(r)$ is the NFW density profile of the halo prior to any tidal effects. A truncation fraction $f_{\rm trunc}(r, r_{\rm t})$ is then defined in terms of the truncation radius, $r_{\rm t}$:

$$f_{\rm trunc}(r, r_{\rm t}) = \frac{1}{1 + (r/r_{\rm t})^2}.$$
 (11)

To find the best-fit truncation radius we minimize the objective function $\chi^2(r_t)$:

$$\chi^2(r_{\rm t}) = \sum_{r \le R_{\rm vir}} \left(\log_{10} \frac{\rho(r)/\rho_{\rm NFW}(r)}{f_{\rm trunc}(r, r_{\rm t})} \right)^2.$$
(12)

Once the GALACTICUS data has been normalized, it is then subsequently shifted so that it lies within a hypercube of side length 2, centered at the origin, with each coordinate in the range [-1, 1]. This newly transformed data will be referred to as the hypercube data \mathbf{z} , and its components are defined in terms of the normalized data

 $^{^2\,}$ "Infall" here means the time at which a halo first became a subhalo, by falling inside the virial radius of a larger halo.

³ For the purposes of gravitational lensing, only the projected radius is relevant as all subhalos effectively lie in plane—the extent of the subhalo population along the line of sight (of order the virial radius of the host halo) is negligible compared to the distance from observer to lens, and lens to source.

y through the following definition:

$$z_i = 2\sigma_i - 1, \quad \sigma_i = \frac{y_i - y_{\min,i}}{y_{\max,i} - y_{\min,i}}.$$
 (13)

Here, y_i refers to the *i*th component of the normalized data, and $y_{\min,i}, y_{\max,i}$ are the minimum and maximum values of the *i*th parameter, respectively. It should be noted that equation 13 is invertible, so it is straightforward to transform the data out of the hypercube coordinates back to their starting, normalized values. It is this hypercube data, \mathbf{z} , that is fed into the emulator as input data. The purpose of this additional shift from \mathbf{y} to \mathbf{z} is so that data points do not exponentially grow as they pass through layers of the emulator, leading to potential numerical instability. Keeping the data restricted to a [-1, 1] interval makes it easy to transform as it propagates through the neural network. From the total input data, an 80/20 split is implemented between a validation and training set for the model.

In Figure 1, we show the results for the emulator's subhalo population in comparison to the population from GALACTICUS. The left column shows normalized 2D density plots of GALACTICUS subhalos. The right column is the emulator's attempt to replicate the distribution. Both columns show a total of 49,633 subhalos sampled from their respective distributions. Overall, the similarities between the two columns indicate that the emulator can accurately learn the normalized GALACTICUS subhalo distribution. A more detailed discussion quantifying the accuracy of the emulator is in Appendix A. It should be noted that the emulator's invertible mapping, plus the fact that the emulator knows nothing about the physics of subhalos (it just views the input as a distribution of data points), will have some difficulties modeling sharp edges at the boundaries of the distribution. Therefore, the emulator plots are clipped to ensure that each of the subhalos obeys the following physical and model constraints:

- Subhalos must have infall masses greater than twice the mass resolution, specifically $M_{\rm infall} > 2 \times 10^6 {\rm M}_{\odot}$. This constraint is a result of how GALACTICUS generates halos in a merger tree. A given halo will have two progenitor halos of smaller masses, and a halo with $M_{\rm infall} < 2 \times 10^6 {\rm M}_{\odot}$ would necessarily imply that at least one of its progenitors is below the mass resolution. Since subhalos below the mass resolution (by definition) are not included, GALACTICUS's population of (sub)halos will start to become incomplete below this limit.
- Emulator subhalos must have bound masses within the range $[10^6, 10^9] M_{\odot}$. The lower bound is enforced as, during the evolution of subhalos in GALACTICUS, if the bound mass falls below $10^6 M_{\odot}$, the subhalo is removed from the calculation. The upper bound is imposed to match the range of subhalo masses considered in the empirical model.
- The bound masses of a subhalo must be less than its corresponding infall mass, since, in GALACTI-CUS, subhalos can only lose mass after infall due to tidal stripping.

- The infall redshifts of emulated subhalos must be greater than or equal to the redshift of the lensing galaxy, specifically $z_{\text{lens}} \ge 0.5$ —by definition, for a halo to be a subhalo by the time it is observed at z_{lens} it must have fallen in at some earlier time, $z_{\text{infall}} > z_{\text{lens}}$.
- We restrict ourselves to looking at subhalos in a narrow region around the Einstein radius, as these are the subhalos that predominantly contribute to the strong gravitational lensing signal. Specifically, this is enforced by requiring $r_{2D} \leq 20$ kpc. A similar choice is made in the empirical model.

3. CONSTRAINING DARK MATTER FROM THE LENSING SIGNATURE

We have shown that the normalizing flows algorithm is able to accurately replicate the distribution of GALACTI-CUS dark matter subhalos. The power of the normalizing flows algorithm, however, is not limited to its ability to reproduce subhalo populations. The broader motivation for implementing a flows-based algorithm is that the emulated subhalos, with the same physical accuracy as GALACTICUS subhalos, can be efficiently implemented into any analysis that utilizes subhalo populations.

In this section, we describe how subhalos emulated from GALACTICUS can be implemented into an analysis code that measures gravitational lensing signatures. Specifically, emulator subhalos are input into the code PYHALO⁴ which outputs summary statistics describing how well input halo realizations match observational data from observed lensed quasar systems (Gilman et al. 2024). We extended PYHALO to support subhalos with the emulator's parameterization, which differs in detail from that of the empirical model. A more detailed description of the analysis process is depicted in Gilman et al. (2020), but we will outline the general analysis process here.

The overall goal is to utilize a forward-generative model that samples the target posterior distribution using an Approximate Bayesian Computing (ABC) technique. Specifically, the posterior distribution of interest is the distribution of subhalo population parameters \mathbf{q}_{s} given the data \mathbf{D} , which comes in the form of quasar image positions and flux ratios. Mathematically, this distribution can be modeled by the relation:

$$p(\mathbf{q}_{s}|\mathbf{D}) \propto \pi(\mathbf{q}_{s}) \prod_{n=1}^{N} \mathcal{L}(\mathbf{d}_{n}|\mathbf{q}_{s}).$$
 (14)

Here, N refers to the number of lenses in the system, $\pi(\mathbf{q}_s)$ is the prior distribution of the dark matter parameters, and \mathbf{d}_n is the data from the n^{th} lens. In this work, the target data we use is the set of image positions and flux ratios for a smooth simulated mock lens from an elliptical power law mass profile with Einstein radius of 1 arc second, axis ratio of 0.73, a logarithmic power law slope of 2.05, and an external shear strength of 0.057. The term "smooth" here refers to a lens system with no substructure. While equation 14 provides a straightforward approach to posterior sampling in theory, the likelihood is difficult to evaluate in practice due to its large

 $^{^4~}$ https://github.com/dangilman/pyHalo, we use revision 9a5d619f3feef297fce8b0a302862991dcb2bd74.



FIG. 1.— Two dimensional density distributions of normalized subhalo parameter spaces generated from GALACTICUS (left column) versus the emulator (right column). Yellow regions correspond to regions of higher number density. The parameters y_1, \ldots, y_6 denote the normalized infall masses, concentrations, bound masses, infall redshifts, truncation radii, and projected radii respectively.

dimensionality. Therefore, an alternative approach can be taken where the data itself is forward modeled.

Instead of directly sampling the likelihood, an alternative approach is to first create simulated data which has lenses \mathbf{d}'_n . This data set is composed of quasar image positions and flux ratios produced from a simulated realization of subhalos \mathbf{m}_{sub} generated from model parameters \mathbf{q}_s . Then, in the case where $\mathbf{d}_n = \mathbf{d}'_n$ for some n, we will have found values for \mathbf{q}_s which have correctly reproduced the observational data and thus have found a single data point in the posterior distribution. By generating many different subhalo realizations \mathbf{m}_{sub} for different \mathbf{q}_s and finding which realizations most accurately reproduce the data, we can effectively sample from the posterior distribution without needing to calculate the likelihood directly.

To quantify how well a given realization of subhalos reproduces observed data, a summary statistic S_{lens} as follows:

$$S_{\text{lens}}(\mathbf{f}, \mathbf{f}') \equiv \sqrt{\sum_{i=1}^{3} \left(f'_{i} - f_{\text{obs}(i)}\right)^{2}}.$$
 (15)

The vector **f** represents the flux ratios from the target data, and the vector **f'** represents flux ratios produced from the set of simulated subhalos. Given that the object being observed is a quadruply imaged quasar, there will naturally be four lensed images and, thus, three flux ratios. The statistic S_{lens} represents a discrepancy between the observed and simulated flux ratios, such that the most accurate realizations will have the smallest values of S_{lens} . A single S_{lens} value corresponds to a given subhalo population, and the distribution of S_{lens} values over many realizations provides insight into a model's relative ability to reproduce target data.

4. RESULTS AND ANALYSIS

In this section, we present two sets of results from our study. First, we give a summary of run times needed to generate subhalo populations and perform lensing analysis. We then show the results of implementing emulator subhalos into the gravitational lensing analysis from the previous section.

4.1. Summary of Run Times

Table 1 presents the run times needed to produce a subhalo population for the empirical model, direct generation from GALACTICUS, and the emulator trained on GALACTICUS data.

 $\label{eq:table_table_table} \begin{array}{c} {\rm TABLE \ 1} \\ {\rm Run \ times \ for \ generating \ a \ single \ subhalo \ population}. \end{array}$

Model	CPU time [s]
Empirical model	2
GALACTICUS	2.6×10^4
Emulator	2

Each value measures the time taken to generate a single population of subhalos (for the specific host halo and resolutions described above). Here, the "Empirical model" row refers to the amount of time needed to generate a realization of subhalos from the empirical model, while the "Galacticus" row refers to how much time it takes to generate a realization of subhalos from GALACTICUS directly. The "Emulator" row refers to how much time it takes the normalizing flows algorithm to generate a single realization of normalized subhalos (and subsequently unnormalize them) once trained. Although not mentioned in the table, training the emulator takes around we found that it took around 1.8×10^4 s to train the emulator on GALACTICUS data. Of course, the emulator needs to be trained only once on an input data set, after which any number of subhalo populations can be generated.

We see that the time to train the emulator is less than the time it takes to generate a single subhalo population directly from GALACTICUS, and note that the training dataset consisted of 300 GALACTICUS subhalo population realizations. This highlights the computational efficiency of using the emulator to generate many populations of subhalos, as once the emulator is trained it takes only a couple of seconds to generate subsequent realizations. Moreover, we see that the time to generate realizations from the emulator (once trained) is similar to the time it takes to generate realizations from the empirical model⁵. Therefore, once the emulator learns the mapping between the latent and data spaces we can produce realizations of GALACTICUS subhalo populations on timescales comparable to that of the empirical model.

Once subhalo populations are created either from either the empirical model or the emulator, it takes similar amounts of time for the forward modeling analysis to be applied to these realizations. This is to be expected since regardless of whether the empirical or emulator model is being applied, an $N \times 6$ array is fed into SAMANA⁶ (where N is the number of subhalos, each of which is described by 6 parameters). It takes around 80s to run SAMANA and produce a single S_{lens} value. Therefore, producing 1,000,000 S_{lens} values takes ~22,000 CPU-hours. Generating 300 GALACTICUS trees and training the emulator takes only ~2,100 CPU-hours. The process of generating training data and training the emulator is, therefore, not the bottleneck in the overall pipeline for producing S_{lens} distributions⁷.

4.2. Forward Modeling the Data

As an example of emulated subhalos being applied in a lensing analysis, we construct S_{lens} distributions from 10,000 realizations of subhalo populations from both empirical and emulated subhalo models. Figure 2 shows the distributions of flux ratios between the empirical and emulator models. We see that there is very good agreement between the distribution of flux ratios, indicating that the emulator produces subhalo populations with lensing properties comparable to those in the empirical model. Additionally, the red bins show the case where there is no dark matter substructure. We see that in the f_2/f_1 and

 $^{^5}$ Furthermore, this time is small compared to the time taken to perform lensing calculations for a given subhalo population, such that subhalo population generation is not a significant contributor to the total time taken to analyze lensing data.

 $^{^{6}}$ https://github.com/dangilman/samana, we use revision 30 cdf 83609 e 50 bdc 4a 0a 7 d 89 877 cf 61 39 51 b 7115.

 $^{^7}$ Both GALACTICUS and SAMANA can be run in parallel, greatly reducing the wall-time needed to complete these calculations, but this does not change the overall number of CPU-hours required in each case.

 f_3/f_1 ratios that there are noticeable differences between the red bins and the other two models.

Looking at the f_4/f_1 ratios however, it appears that the distribution of flux ratios is nearly identical, regardless of if the model contains substructure or not. We investigate this further in Figure 3, which shows two dimensional distributions of the flux ratios. We see that the density distributions for the no subhalos model appears noticeably different from both the empirical and emulator models. This suggests that the emulator reproduces flux ratios more closely aligned with the empirical model than with a model devoid of dark matter subhalos. Additionally, this suggests that the similar f_4/f_1 histograms in Figure 2 are simply a consequence of marginalization, and do not imply that the underlying flux ratio distributions are nearly identical.

Figure 4 shows the cumulative distribution of S_{lens} values from both models. The plot on the left shows the cumulative distribution function (CDF) of the S_{lens} summary statistic. The plot on the right shows a zoomed-in version of the CDFs for $S_{\text{lens}} \leq 0.1$. Subhalo realizations in the right-hand plot are especially of interest, as these populations produce flux ratios that most closely resemble the target data. We see that emulated subhalos when implemented into the forward modeling analysis, produce comparable results to the empirical model. This is expected (assuming that our emulation procedure accurately captures the distribution of subhalo properties predicted by GALACTICUS) as the empirical model itself was constructed to approximately match the results of GALACTICUS simulations.

One notable difference between the empirical and emulator models is the average number of subhalos per realization. On average, the empirical model produces around 23,000 subhalos per realization, compared to only 1,200 subhalos per realization coming from the emulator. The reason for this difference is the choice of how the subhalo mass resolution is implemented in each model. For GALACTICUS, the mass resolution is applied to the bound mass of subhalos—that is, any subhalo that falls below a bound mass of $10^6 M_{\odot}$ is removed from the population⁸. Conversely, in the empirical model, the mass resolution is applied to infall masses, such that a subhalo may have a bound mass much smaller than $10^{6} M_{\odot}$ —such halos are *not* discarded in the empirical model. We find that ~90% of empirical subhalos have $M_{\rm bound} < 10^6 {\rm M}_{\odot}$, which is to be expected (Du et al. 2025).

A natural question to ask is whether or not the additional subhalos from the empirical model that are excluded in the emulator model have a significant impact on the resulting S_{lens} distribution curves. In Figure 5, we plot the S_{lens} distributions for the empirical model in two cases: One instance where we include subhalos with infall masses above $10^6 M_{\odot}$ but bound masses below $10^6 M_{\odot}$ (black), and one instance where we exclude these halos (green). By excluding subhalos whose bound mass falls below the resolution, we remove the vast majority of subhalos from a given population. However despite removing such a large portion of the subhalo population, the distribution curves are essentially overlapping. This further supports the work from Gilman et al. (2020), which shows that these low-mass subhalos do not significantly affect the observed lensing signal.

5. LIMITATIONS/FUTURE WORK

Here, we highlight the limitations of our emulator approach as well as possibilities for future work. One limitation of the emulator approach arises from how the normalizing flow is constructed. Requiring that the map between latent and data spaces be differentiable and invertible restricts the possible forms that the map can take, and thus certain distributions become difficult to replicate (i.e. distributions with sharp edges, multimodal distributions, regions with complex topologies, etc.). An example of this limitation is discussed in Figure 6 in Appendix A. Another limitation of this current approach is that the emulator knows nothing about sub-subhalos, as this information is not provided to it during training, so it is unable to capture correlations between these halos and their host subhalos. While this is expected to have only a small effect on the resulting subhalo populations (sub-subhalos typically make up less than 10% of the total population in GALACTICUS), future work could avoid this issue by using a second emulator to populate each subhalo with sub-subhalos for example.

There are several other directions that could be explored in future work to improve upon the results of this study. For example, in this work, we used the emulator to generate a six-dimensional joint distribution of subhalo parameters for a specific host halo mass and redshift and for the case of cold dark matter. Previous studies have demonstrated how normalizing flow algorithms can be used to learn conditional distributions (Winkler et al. 2019; Friedman & Hassan 2022; Abbasi et al. 2023)—often referred to as conditional normalizing flow (CNF) algorithms.

By implementing a CNF algorithm we could greatly increase the flexibility of this model. This generalization can, in principle, be applied to any set of interdependent parameters, but two specific areas of interest are as follows. First, the emulator could be conditioned on host halo mass and redshift, allowing a single emulator to be trained and applied to all observed lenses. Another potentially powerful application of a CNF would be to construct an emulator conditioned on the physics of a given dark matter model. For example, performing the same lensing analysis for a warm dark matter (WDM) model would require knowledge of the WDM particle mass. Current studies estimating $m_{\rm WDM} \gtrsim 7$ keV (Nadler et al. 2021; Nierenberg et al. 2023; Keeley et al. 2024). With our current approach, one would need to train a separate emulator for each WDM particle mass to be considered. If a CNF algorithm was implemented, however, the same 6D parameter space could be learned as a function of WDM particle mass, allowing rapid exploration of WDM models, and for constraints on WDM particle mass to be derived. This would, of course, require the generation of GALACTICUS realization of subhalo populations for WDM models of varying masses in order to train the CNF. Generalization to other, beyond-CDM dark matter models (e.g., conditioning on any selfinteraction cross-section of the dark matter particle) may also be possible.

 $^{^8}$ This choice is made for computational efficiency and because, even if such subhalos were retained, the population of subhalos with $M_{\rm bound} < 10^6 {\rm M}_{\odot}$ would be incomplete due to the missing contribution from subhalos with $M_{\rm infall} < 10^6 {\rm M}_{\odot}$.



FIG. 2.— Histograms for the three flux ratios are plotted where f_i refers to the flux of the *i*th lensed image. The blue bins depict the distributions of subhalos from the empirical model, the orange bins show the subhalo distributions for the emulator model, and the red bins show flux ratios when the lensing galaxy contains no subhalos.

6. CONCLUSIONS

In this work, we examined the strong gravitational lensing perturbations to flux ratios of a quadruply lensed quasar resulting from a population of subhalos produced by the empirical model of Gilman et al. (2020) as updated by Keeley et al. (2024), and compared the results against same analysis conducted using subhalo population generated from GALACTICUS via a normalizing flow emulator. A statistical approach was used in which many subhalo populations were generated, but we only kept the realizations that most closely resembled observational data for both the empirical and emulator models. This emulation approach makes such analyses computationally feasible—generating sufficient realizations directly from GALACTICUS would be computationally impractical.

We found that the normalizing flows algorithm accurately replicates subhalo populations, achieving this in a fraction of the time required for direct generation from GALACTICUS. The distributions of summary statistics describing the perturbations are similar (although not identical) from the empirical model and emulated GALACTICUS—this is as expected as the empirical model was constructed to approximately match GALACTICUS





FIG. 3.— Two dimensional density distributions of flux ratios for the empirical model (left column), emulator model (middle column) and model without substructure (right column).



FIG. 4.— Left panel: inverted CDFs of subhalo populations from both the empirical model (black) and the emulator model (red). Right panel: a zoom-in on the left plot for the lowest S_{lens} values.



simulations, and serves as a demonstration that results derived using the empirical model are not substantially affected by the simplifications made in that approach.

In future work, we aim to extend this approach by generalizing the normalizing flows algorithm used to generate subhalo populations. We currently train the algorithm to learn a joint distribution of the subhalo parameter space. In future work, we aim to use the emulator to learn a conditional distribution, allowing subhalo populations to be generated as a function of, for example, the mass of a warm dark matter particle. Doing this will allow us to test the strong lensing signature for many different WDM models, and thereby place accurate constraints on the mass of a WDM particle.

ACKNOWLEDGMENTS

FIG. 5.— Distribution of $S_{\rm lens}$ values for the empirical model. The black curve includes subhalos which have infall masses above $10^6 {\rm M}_{\odot}$ and bound masses below $10^6 {\rm M}_{\odot}$. The green curve excludes these halos.

We thank Charles Gannon for their helpful discussions and their assistance in calibrating the empirical model.

REFERENCES

Abbasi, R., Ackermann, M., Adams, J., et al. 2023, Conditional normalizing flows for IceCube event reconstruction, Tech. rep. Ade, P. A., Aghanim, N., Arnaud, M., et al. 2016, Astronomy & Astrophysics, 594, A13

- Aghanim, N., Akrami, Y., Ashdown, M., et al. 2020, Astronomy & Astrophysics, 641, A6
- Bagla, J. S. 2005, Current science, 1088
- Benson, A. J. 2012, New Astronomy, 17, 175
- Benson, A. J., & Du, X. 2022, Monthly Notices of the Royal Astronomical Society, 517, 1398
- Benson, A. J., Lacey, C. G., Frenk, C. S., Baugh, C. M., & Cole, S. 2004, MNRAS, 351, 1215,
 - doi: 10.1111/j.1365-2966.2004.07870.x
- Bertschinger, E. 1998, Annual Review of Astronomy and Astrophysics, 36, 599
- Bower, R. G. 1991, MNRAS, 248, 332,
- doi: 10.1093/mnras/248.2.332
- Bryan, G. L., & Norman, M. L. 1998, The Astrophysical Journal, 495, 80
- Bullock, J. S., & Boylan-Kolchin, M. 2017, Annual Review of Astronomy and Astrophysics, 55, 343
- Cole, S., & Lacey, C. 1996, Monthly Notices of the Royal Astronomical Society, 281, 716
- Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000,
- Monthly Notices of the Royal Astronomical Society, 319, 168 Dalal, N., & Kochanek, C. 2002, The Astrophysical Journal, 572, 25
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, ApJ, 292, 371, doi: 10.1086/163168
- Dehnen, W., & Read, J. I. 2011, The European Physical Journal Plus, 126, 55
- Dekker, A., Ando, S., Correa, C. A., & Ng, K. C. Y. 2022,
- Phys. Rev. D, 106, 123026, doi: 10.1103/PhysRevD.106.123026
- Diemer, B., & Joyce, M. 2019, The Astrophysical Journal, 871, 168
- Du, X., Gilman, D., Treu, T., Benson, A., & Gannon, C. 2025, arXiv e-prints, arXiv:2503.07728, doi: 10.48550/arXiv.2503.07728
- Dutton, A. A., Macciò, A. V., Frings, J., et al. 2016, MNRAS, 457, L74, doi: 10.1093/mnrasl/slv193
- Fitts, A., Boylan-Kolchin, M., Bozek, B., et al. 2019, MNRAS, 490, 962, doi: 10.1093/mnras/stz2613
- Frenk, C. S., White, S. D., Efstathiou, G., & Davis, M. 1985, Nature, 317, 595
- Friedman, R., & Hassan, S. 2022, arXiv preprint arXiv:2211.12724 Gannon, C., Nierenberg, A., Benson, A., et al. 2025, arXiv e-prints, arXiv:2501.17362, doi: 10.48550/arXiv.2501.17362
- Gannon, C., Nierenberg, A., Benson, A., et al. 2025, arXiv
- preprint arXiv:2501.17362 Gilman, D., Birrer, S., Nierenberg, A., & Oh, M. S. 2024, Monthly Notices of the Royal Astronomical Society, 533, 1687
- Gilman, D., Birrer, S., Nierenberg, A., et al. 2020, Monthly Notices of the Royal Astronomical Society, 491, 6077
- Gilman, D., Bovy, J., Treu, T., et al. 2021, Monthly Notices of the Royal Astronomical Society, 507, 2432
- Giocoli, C., Pieri, L., & Tormen, G. 2008, MNRAS, 387, 689, doi: 10.1111/j.1365-2966.2008.13283.x
- Han, J., Cole, S., Frenk, C. S., & Jing, Y. 2016, Monthly Notices of the Royal Astronomical Society, 457, 1208
- Hezaveh, Y. D., Dalal, N., Marrone, D. P., et al. 2016, The Astrophysical Journal, 823, 37
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, The
- Astrophysical Journal Supplement Series, 208, 19
- Hsueh, J.-W., Enzi, W., Vegetti, S., et al. 2020, Monthly Notices of the Royal Astronomical Society, 492, 3047

APPENDIX

A. ACCURACY OF THE EMULATOR

Here, we describe tests that quantify the accuracy of the emulator in its ability to replicate GALACTICUS subhalo populations. We first conducted a two-sample Kolmogorov-Smirnov (KS) test, where a test statistic $D_{m,n}$ is defined by:

$$D_{m,n} = \sup_{x} |F_m(x) - G_n(x)|,$$
 (A1)

- Jiang, F., Dekel, A., Freundlich, J., et al. 2021, MNRAS, 502, 621, doi: 10.1093/mnras/staa4034
- Jiang, F., & van den Bosch, F. C. 2017, Monthly Notices of the Royal Astronomical Society, 472, 657
- Keeley, R. E., Nierenberg, A. M., Gilman, D., et al. 2024, MNRAS, 535, 1652, doi: 10.1093/mnras/stae2458
- Keeley, R. E., Nierenberg, A. M., Gilman, D., et al. 2024, arXiv preprint arXiv:2405.01620
- Lacey, C., & Cole, S. 1994, MNRAS, 271, 676, doi: 10.1093/mnras/271.3.676
- Loshchilov, I. 2017, arXiv preprint arXiv:1711.05101
- Makishima, K. 1998, in Symposium-International Astronomical Union, Vol. 188, Cambridge University Press, 181-184
- Mikkola, S., & Aarseth, S. J. 1993, Celestial Mechanics and Dynamical Astronomy, 57, 439
- Moore, B., Ghigna, S., Governato, F., et al. 1999, ApJ, 524, L19, doi: 10.1086/312287
- Nadler, E. O., Birrer, S., Gilman, D., et al. 2021, The
- Astrophysical Journal, 917, 7 Nadler, E. O., Mansfield, P., Wang, Y., et al. 2023, The Astrophysical Journal, 945, 159
- Navarro, J. F. 1996, in Symposium-international astronomical union, Vol. 171, Cambridge University Press, 255-258
- Nierenberg, A., Treu, T., Wright, S., Fassnacht, C., & Auger, M. 2014, Monthly Notices of the Royal Astronomical Society, 442, 2434
- Nierenberg, A., Keeley, R., Sluse, D., et al. 2023, arXiv preprint arXiv:2309.10101
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. 2021, Journal of Machine Learning Research, 22, 1
- Parkinson, H., Cole, S., & Helly, J. 2008, Monthly Notices of the Royal Astronomical Society, 383, 557
- Peebles, P. J. E. 2020, Principles of physical cosmology (Princeton university press)
- Pullen, A. R., Benson, A. J., & Moustakas, L. A. 2014, ApJ, 792, 24, doi: 10.1088/0004-637X/792/1/24
- Pullen, A. R., Benson, A. J., & Moustakas, L. A. 2014, The Astrophysical Journal, 792, 24
- Sheth, R. K., & Tormen, G. 2002, MNRAS, 329, 61, doi: 10.1046/j.1365-8711.2002.04950.x
- Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, The Astrophysical Journal Supplement Series, 148, 175
- Springel, V., Wang, J., Vogelsberger, M., et al. 2008, MNRAS, 391, 1685, doi: 10.1111/j.1365-2966.2008.14066.x
- Taylor, J. E., & Babul, A. 2001, ApJ, 559, 716, doi: 10.1086/322276
- Treu, T. 2010, Annual Review of Astronomy and Astrophysics, 48.87
- Vegetti, S., & Vogelsberger, M. 2014, Monthly Notices of the Royal Astronomical Society, 442, 3598
- Winkler, C., Worrall, D., Hoogeboom, E., & Welling, M. 2019, arXiv preprint arXiv:1912.00042
- Xu, D., Sluse, D., Gao, L., et al. 2015, Monthly Notices of the Royal Astronomical Society, 447, 3189
- Yang, S., Du, X., Benson, A. J., Pullen, A. R., & Peter, A. H. 2020, Monthly Notices of the Royal Astronomical Society, 498, 3902
- Yang, X., Mo, H. J., Zhang, Y., & van den Bosch, F. C. 2011, ApJ, 741, 13, doi: 10.1088/0004-637X/741/1/13
- Zentner, A. R., Berlind, A. A., Bullock, J. S., Kravtsov, A. V., & Wechsler, R. H. 2005, ApJ, 624, 505, doi: 10.1086/428898

where m, n are the sizes of the two data sets, and $D_{m,n}$ represents the maximum absolute difference between two distributions. Table 2 shows a summary of the $D_{m,n}$ statistics comparing the GALACTICUS and emulator (trained on GALACTICUS) distributions:

The $D_{m,n}$ values from Table 2 imply that the GALACTICUS and emulator distributions are not identical. While this may seem problematic, the goal is not necessarily for the emulator to exactly replicate the distribution of GALACTICUS subhalos. Rather, the goal is for the emulator to produce a strong lensing signal that

TABLE 2Results from the two-sample KS test.

Parameter	$D_{m,n}$
Infall mass	0.008
Concentration	0.023
Bound mass	0.016
Infall redshift	0.008
Truncation radius	0.013
2D projected radius	0.009



FIG. 6.— CDFs of the concentration parameter for both the GALACTICUS and emulator distributions, zoomed in to the region around the minimum allowed concentration in GALACTICUS, which is where the maximum discrepancy occurs between distributions. The vertical dashed line represents the value of *c* at which the maximum discrepancy between distributions occurs.

is sufficiently close to the strong lensing signature that would have been produced by GALACTICUS subhalos. Here, "sufficiently close" means that remaining differences should be small compared to other approximations made in the GALACTICUS modeling, which are at least of order 10% (Nadler et al. 2023).

We investigate the source of the discrepancy between distributions by comparing their CDFs. For example, Figure 6 shows the CDF of the concentration parameter, the parameter with the largest $D_{m,n}$ value. Figure 6 compares CDFs of the concentration parameter c between the GALACTICUS (black) and emulator (red) subhalo distributions, and is zoomed in on a region around c = 3.795 (grey dashed line) where $D_{m,n}$ is maximized. There is a relatively steep increase in the number of GALACTICUS subhalos around c = 3.795, which is the minimum concentration of subhalos in the simulation 9 . The emulator CDF does not replicate this sharp increase perfectly, which can be attributed to the inherent limitations of a normalizing flows algorithm. Such models tend to have difficulties accurately replicating sharp changes in distribution due to the requirements of invertibility and differentiability of the mapping between latent and data spaces.

To test if the emulator is able to produce a lensing



FIG. 7.— A comparison of the S_{lens} CDFs between the empirical model (black curve) and the emulator trained on the empirical model (red curve). For reference, the blue curve shows results if no subhalos are included.

signal statistically equivalent to the subhalo population it replicates, we performed a separate test in which we trained the emulator on subhalo realizations from the empirical model. Since the lensing signature, S_{lens} , from the empirical model can be computed directly (by simply generating larger numbers of subhalo populations from the empirical model—something we can not do for GALACTICUS because of the computational cost), we can assess how accurately the emulator is able to reproduce this signature. The results can be seen in Figure 7.

In Figure 7, the black curve shows the empirical model, while the red curve shows the emulator trained on subhalos generated by the empirical model. For comparison, the blue curve shows the result obtained if no subhalos are included. The blue curve lies notably lower than the red and black curves, implying that there is a higher fraction of realizations with low S_{lens} values. This is to be expected, as the smooth mock baseline model contains no substructure, so the model with no subhalos is the correct model in this case. There is a slight discrepancy between the red and black curves, indicating that the emulator does not perfectly capture the distribution of subhalos in the empirical model. However, the offset between red and black curves is substantially smaller than the offset between either of these curves and the no subhalos model (blue curve), indicating that the inclusion of subhalos produces a noticeable effect on the S_{lens} curves and that the magnitude of this effect is similar between empirical subhalos and emulated subhalos.

Lastly, we examine the model loss over the course of the emulator training in Figure 8. The model loss here is the quantity that the algorithm seeks to minimize over the course of the training period to produce an emulator that matches the input data.

In Figure 8, the blue and orange curves represent the model loss as a function of the total number of epochs for the training and validation data sets respectively. There is an 80/20 split across 300 input GALACTICUS trees between the training and validation sets respectively. The validation loss starts to flatten out around epochs 50–75, with only a slow decline at further epochs. Nevertheless, the validation model loss *does* still slowly decrease,

⁹ GALACTICUS imposes a minimum concentration of $c_{\rm vir} = 4$, defined relative to the spherical collapse virial radius. Here we show concentration, c, defined relative to the R_{200} definition of virial radius, which results in a small, redshift-dependent shift in the cut-off, causing the cut-off to be somewhat smoothed out.



FIG. 8.— The emulator model loss as a function of training epoch. The blue curve shows the model loss for the training data set, and the orange curve shows the model loss on the validation data set.

indicating that our emulator is not strongly affected by overfitting.

This paper was built using the Open Journal of Astrophysics IATEX template. The OJA is a journal which provides fast and easy peer review for new papers in the astro-ph section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at http://astro.theoj.org.