



Machine-learning Inference of the Interior Structure of Low-mass Exoplanets

Philipp Baumeister¹ , Sebastiano Padovan² , Nicola Tosi^{1,2} , Grégoire Montavon³, Nadine Nettelmann² ,
Jasmine MacKenzie¹, and Mareike Godolt¹ 

¹ Centre of Astronomy and Astrophysics, Technische Universität Berlin, Hardenbergstraße 36, D-10623 Berlin, Germany; philipp.baumeister@tu-berlin.de

² Institute of Planetary Research, German Aerospace Center (DLR), Rutherfordstraße 2, D-12489 Berlin, Germany

³ Institute of Software Engineering and Theoretical Computer Science, Technische Universität Berlin, Marchstr. 23, D-10587 Berlin, Germany

Received 2019 August 30; revised 2019 November 15; accepted 2019 November 27; published 2020 January 23

Abstract

We explore the application of machine-learning based on mixture density neural networks (MDNs) to the interior characterization of low-mass exoplanets up to 25 Earth masses constrained by mass, radius, and fluid Love number, k_2 . We create a data set of 900,000 synthetic planets, consisting of an iron-rich core, a silicate mantle, a high-pressure ice shell, and a gaseous H/He envelope, to train a MDN using planetary mass and radius as inputs to the network. For this layered structure, we show that the MDN is able to infer the distribution of possible thicknesses of each planetary layer from mass and radius of the planet. This approach obviates the time-consuming task of calculating such distributions with a dedicated set of forward models for each individual planet. While gas-rich planets may be characterized by compositional gradients rather than distinct layers, the method presented here can be easily extended to any interior structure model. The fluid Love number k_2 bears constraints on the mass distribution in the planets' interiors and will be measured for an increasing number of exoplanets in the future. Adding k_2 as an input to the MDN significantly decreases the degeneracy of the possible interior structures. In an open repository, we provide the trained MDN to be used through a Python Notebook.

Unified Astronomy Thesaurus concepts: [Exoplanet structure \(495\)](#); [Exoplanets \(498\)](#); [Neural networks \(1933\)](#); [Planetary interior \(1248\)](#); [Computational methods \(1965\)](#); [Planetary theory \(1258\)](#)

1. Introduction

The characterization of the interior of observed exoplanets is one of the main goals in current exoplanetary science. With the large number of newly discovered exoplanets expected in the next 10 years by ground-based surveys such as the Wide Angle Search for Planets (WASP; Pollacco et al. 2006), the Next-Generation Transit Survey (NGTS; Wheatley et al. 2017), and the Hungarian Automated Telescope Network/Hungarian Automated Telescope Network-South (HATNet/HATSouth; Hartman et al. 2004; Bakos et al. 2013), as well as the ongoing *Transiting Exoplanet Survey Satellite (TESS)* space survey (Ricker et al. 2014) and the upcoming *PLAnetary Transits and Oscillations (PLATO)* mission (Rauer et al. 2014), a rapid characterization scheme of the interior structure of these planets will become increasingly necessary to further our understanding of planetary populations. The vast majority of the confirmed exoplanets has been identified either through transits or radial velocities surveys. Planets identified with both techniques are characterized by their mass and radius, which, combined, provide a first indication of the bulk composition through comparison with theoretical mass–radius curves (e.g., Valencia et al. 2006; Sotin et al. 2007; Zeng & Sasselov 2013). A common approach to the interior characterization of exoplanets is the use of numerical models to compute interior structures that comply with the measured mass and radius of the planet (e.g., Fortney et al. 2007; Sotin et al. 2007; Valencia et al. 2007; Wagner et al. 2011; Zeng & Sasselov 2013; Unterborn & Panero 2019). As this is an inverse problem, it requires the calculation of a large number of interior models to obtain an overview over possible interior structures (Rogers & Seager 2010a; Brugger et al. 2017; Dorn et al. 2017). If other observables are used in addition to mass and radius, the number of samples needed for an accurate inference of possible interior structures increases drastically, due to the increase in dimensionality (e.g., James et al. 2013). Thus, the inference can

quickly become computationally expensive. Moreover, with only mass and radius, possible solutions tend to be highly degenerate, with multiple, qualitatively different interior compositions that can match the observations equally well (e.g., Rogers & Seager 2010a, 2010b).

In this work, we explore a new approach to the interior characterization of exoplanets by employing a deep learning neural network to treat this inverse problem. In recent years, deep neural networks (DNNs) have been used in a number of exoplanetary science studies, with applications ranging from the detections of planetary transits (Pearson et al. 2018; Shallue & Vanderburg 2018; Chaushev et al. 2019) to atmospheric composition retrieval from measured planet spectra (Márquez-Neila et al. 2018; Zingales & Waldmann 2018), and the computation of critical core and envelope masses of forming planets (Alibert & Venturini 2019). We approach the characterization of exoplanets by first creating a large data set of synthetic planets and then training a multitask (Caruana 1997) mixture density network (MDN; Bishop 1994) to infer plausible thicknesses of the planetary layers using observables like mass and radius. Mixture density networks are a special case of neural networks, which predict the parameters of a Gaussian mixture distribution instead of single output values, enabling us to approximate the entire posterior probability density function at once. This property allows them to work with inverse problems that conventional neural networks often cannot address well. Neural networks are an optimal choice to approach the interior characterization of planets. First, as neural networks can essentially approximate any arbitrary nonlinear relations between some input and output values (Cybenko 1989; Goodfellow et al. 2017; Lu et al. 2017), they are very well suited to model the relation between observables and interior structure for an arbitrary number of observable parameters, as the mapping can be difficult to quantify, especially for a large number of model parameters (Bishop 1995). Second, the computation time needed by a well-trained neural

network to make a prediction is in the order of milliseconds. Compared with conventional methods, such as Markov Chain Monte Carlo, this significantly reduces the computational time for finding valid interior structures, allowing for rapid classification of exoplanets. Additionally, with the fast training times of neural networks (less than 30 minutes with our setup; see the end of Section 2.2), the effect of different combinations of observables can be tested rapidly (although validating the trained MDNs and optimizing the network architecture may still take a considerable amount of time).

Beyond mass and radius, additional observables that may be informative for the interior structure include elemental abundances of the host star, which may be representative of those of the planet (e.g., Dorn et al. 2015, 2017; Brugger et al. 2017), atmospheric composition (e.g., Madhusudhan et al. 2016), and the planetary fluid Love numbers (Padovan et al. 2018). Of these, the fluid Love numbers represent the most direct constraint on the interior structure, given that they only depend on the radial density distribution in the interior (Kramm et al. 2011).

Fluid Love numbers are a subset of the more general set of Love numbers, which are parameters describing how a planet responds to a perturbation (e.g., tidal and rotational deformations). They depend on the timescale of the perturbation and on the rheological properties of the interior (i.e., radial density profile, elastic properties, viscosity). For example, the rotational figure of the Earth is relaxed, i.e., it attained hydrostatic equilibrium and accordingly, its shape can be parameterized using fluid Love numbers (Padovan et al. 2018). At the same time, on the shorter tidal lunar and solar timescales, the response is described by tidal Love numbers (e.g., Petit & Luzum 2010).

The deformed shape of a transiting extrasolar planet will modify its transit light curve, introducing a second-order effect that can, in principle, be observed (Akınsanmi et al. 2019; Hellard et al. 2019). Similarly, for some particular configurations (e.g., a close-in planet with a distant companion, Mardling 2007), apsidal precession is mainly driven by the tidal interaction of the planet with its host star (Batygin et al. 2009), and a value of the Love number of the planet can be inferred (Csizmadia et al. 2019). Thus, radial velocity measurements, as well as transit observations, can be used under certain circumstances to infer the value of the Love number k_2 , a task that is currently possible only for close-in gas giants. The ever-increasing temporal baseline of exoplanet observations and sensitivity of ground-based and space telescopes will make the measurement of the Love numbers of extrasolar planets increasingly possible. For example, using the new Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations mounted on the Very Large Telescope facility of the European Southern Observatory (Pepe et al. 2010), on a 10 yr long observing baseline, one may infer the k_2 of a body with a mass of $3 M_{\oplus}$ (Sz. Csizmadia 2019, personal communication). The advent of larger telescopes such as the Thirty Meter Telescope (e.g., Skidmore et al. 2015) will further increase the sensitivity for smaller planets. In the following, we assume that measurements of the fluid Love numbers become available in the future.

2. Methods

2.1. Interior Structure Model

We use a newly developed interior structure model named Tool for ATmospheres, Outgassing and Optimal INteriors of

Exoplanets (TATOOINE) to construct the planets used for the training data. A modeled planet consists of four compositionally distinct layers: an iron-rich core, a silicate mantle, an ice shell, and a H/He envelope. Each layer is characterized by its mass fraction, with the sum of the mass of all layers constrained to add up to the planet’s total mass, M_p .

The model calculates radial profiles of mass (m), density (ρ), and pressure (P) by solving the following three coupled equations:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r), \quad (1)$$

$$\frac{dP(r)}{dr} = -\frac{Gm(r)\rho(r)}{r^2}, \quad (2)$$

$$P(r) = f(\rho(r), T(r), c(r)). \quad (3)$$

Equation (1) expresses the mass of a spherical shell of thickness dr ; Equation (2) is the condition of hydrostatic equilibrium, where G is the gravitational constant; Equation (3) is an equation of state (EoS) where f indicates a function that is specific to the material of each model layer and depends on the density, temperature $T(r)$, and composition $c(r)$ profiles.

The model setup and parameters closely follow the modeling approach of Sotin et al. (2007), with the addition of a zero-temperature H/He envelope to account for gas-rich planets (Salpeter & Zapolsky 1967; Seager et al. 2007). Using a zero-temperature EoS implies that the envelope thickness inferred from mass and radius of the planet will tend to be too large (Section 2.1.4). Below the H/He envelope, temperatures in the planet follow an adiabatic profile starting with a temperature of 300 K at the envelope-ice interface in all models. We further assume that all planets share the same Earth-like mineral composition of the mantle and core. The interior structure of gas-rich planets could be better described by a compositional gradient rather than by chemically distinct layers, an approach that is best suited for predominantly rocky bodies. Nevertheless, the machine-learning based inference approach that we present below is general and can be easily adapted to any interior structure model (Section 2.2). We focus on planets below 25 Earth masses, as these show the greatest diversity in bulk density among the discovered exoplanets (e.g., Lopez & Fortney 2014), hinting at a wide range of interior structures.

2.1.1. Iron-rich Core

Density profiles derived from seismic data show that the density of Earth’s core is about 5%–10% smaller than the density of pure hcp-iron, suggesting the presence of lighter elements such as sulfur, oxygen, silicon, carbon, and hydrogen (e.g., Rubie et al. 2015). We chose an EoS for the core, where iron is alloyed with sulfur, which is depleted in Earth’s mantle compared with solar abundances (Murthy & Hall 1970; Poirier 1994). Several experimental studies suggest that the melting temperature of the Fe–FeS alloy is low enough to allow Earth’s outer core to be molten under the probable thermodynamic conditions (e.g., Yoo et al. 1993; Fei et al. 1997; Morard et al. 2008). We assume our model planet to have a liquid core composed of a Fe–FeS alloy with a molar fraction of 13% FeS (Poirier 1994; Sotin et al. 2007). For planets more massive than Earth, Stixrude (2014) suggests that the iron cores of planets in the entire Super-Earth mass range (up to 10 Earth masses) are at least partially liquid, and mantle convection is vigorous enough to sustain a dynamo. We neglect an inner solid core, as

the total core radius is not expected to be affected significantly, as the density difference between the solid and liquid phase is small (Sotin et al. 2007). The temperature- and pressure-dependent density is calculated using a Mie–Grüneisen–Debye EoS (e.g., Sotin et al. 2007; Fei et al. 2016).

2.1.2. Silicate Mantle

The silicate mantle in our model is divided into an upper and a lower mantle, separated by a phase transition at 23 GPa. The lower mantle is composed of bridgmanite and magnesiowüstite. The upper mantle is composed of olivine and orthopyroxene enstatite. The ratio of mantle components is assumed to be Earth-like, with a molar Mg/Si ratio of 1.131 (Sotin et al. 2007). Additionally, all mantle silicates form a solid-state solution of their respective Mg and Fe endmembers with the ratio determined by the Mg number (Mg#), which is defined as the mole fraction Mg/(Mg+Fe) of the mantle silicates. We use an Earth-like value of 0.9 for the Mg# in all modeled planets. Density profiles are calculated using a temperature-dependent, third-order Birch–Murnaghan EoS.

2.1.3. High-pressure Ice Layer

Many exoplanets have intermediate bulk densities, which may be indicative of large amounts of water in the interior (e.g., Léger et al. 2004; Adams et al. 2008; Barr et al. 2018; Zeng et al. 2019). We allow for the presence of water by using a layer of ice VII, which is the most stable water ice high-pressure phase over a wide range of pressures (Hemley et al. 1987). We neglect other water phases, as the density differences between solid ice phases are small (Hemley et al. 1987) and do not have a major effect on the thickness of this layer. By doing this, we neglect liquid and superionic phase effects, which would lead to more degenerate solutions of interior structures (Nettelmann et al. 2011). Also, it must be noted that higher temperatures in the ice shell may significantly change the thickness of the layer (Thomas & Madhusudhan 2016).

2.1.4. H/He Gas Envelope

Many massive ($>10 M_{\oplus}$) exoplanets probably have an extended primordial envelope consisting of hydrogen and helium (e.g., Lopez & Fortney 2014). To model this outer gaseous layer, we assume an envelope of solar-like composition with 71% hydrogen and 29% helium (by mass), and use a zero-temperature Thomas–Fermi–Dirac EoS from Salpeter & Zapolsky (1967) to model the pressure-dependent density. Using a zero-temperature envelope overestimates the density of the gas mixture compared with more sophisticated temperature-dependent EoS. As a consequence, more of the planet needs to be taken up by the envelope to account for the planet’s bulk density when modeling the interior. Especially for planet models with extended atmospheres, this means that, when inferring the thickness of the envelope from mass and radius of the planet, the distribution of possible solutions is shifted toward higher envelope thicknesses. We define the planet’s transit radius at a pressure of 100 mbar.

2.2. Mixture Density Networks

Feedforward neural networks are a fundamental part of machine-learning algorithms. Their goal is to learn the mapping between inputs x and corresponding outputs y (see e.g., Bishop 2006; Goodfellow et al. 2017). A neural network consists of

several layers of neurons. For a conventional feedforward neural network, each layer receives the outputs from the previous layer as input and computes a new set of outputs. This calculation is done by the neurons in the layer, where each neuron receives input from other neurons of the previous layer and computes an output value, using the so-called activation function. These functions are typically nonlinear, and thus enable the neural network to approximate any arbitrary, nonlinear function. The first layer in a network receives data features as input and is therefore called the input layer, and the final layer is called the output layer because it outputs the predictions of the network. All layers between input and output layer are referred to as “hidden.” Feedforward networks are so named because information flows only in one direction from the input layer through the hidden layers to the output layer, without loops or feedback.

A neural network is trained with a set of training data that provides examples of which output values correspond to each set of input values. During training, these data are passed through the network, where the weighting of each neuron is iteratively adjusted so that the predicted value approaches the actual target value. This difference is measured by the so-called loss function. The goal of training is to minimize the loss function. Therefore, the gradient of the loss function at the current iteration is calculated by backpropagating the predicted values through the network (Bishop 2006; Goodfellow et al. 2017), and then using a gradient descent algorithm to optimize the model weights.

The conventional approach in neural network training is to minimize the mean squared error between the known values from the training data and the values predicted by the network, which results in these models approximating the average of the training data, conditioned on the input parameters. For forward problems, where a set of input parameters maps to a single set of output values, this approach usually yields the desired solution. For inverse problems, where there may be multiple sets of output values for each feature, these models are inadequate in many cases (Bishop 1994). MDNs overcome this problem by combining conventional neural networks with a probability mixture model. While a conventional neural network maps inputs to deterministic output variables, a MDN predicts the full conditional probability distribution for each input. We use a Gaussian mixture model for our MDN, where the total distribution is the normalized linear combination of weighted normal distributions, as Gaussian mixtures can represent a wide range of data distributions (Bishop 1994, 2006).

The probability density of the target data y (the radius fractions of each planetary layer), conditioned on the input x (mass and radius, or mass, radius, and Love number k_2 of the planet), is represented by a linear combination of Gaussian functions $\phi_i(y | \mu_i(x), \sigma_i^2(x))$:

$$p(y | x) = \sum_{i=1}^m \alpha_i(x) \phi_i(y | \mu_i(x), \sigma_i^2(x)), \quad (4)$$

$$1 = \sum_{i=1}^m \alpha_i(x), \quad (5)$$

where m is the number of mixture components (i.e., the number of Gaussians), α_i are the normalized mixture weights, and μ_i and σ_i^2 are the means and variances of the individual Gaussian distributions. As such, the output layer of the MDN has $3 \times m$ neurons: m neurons, respectively, for the means, variances, and

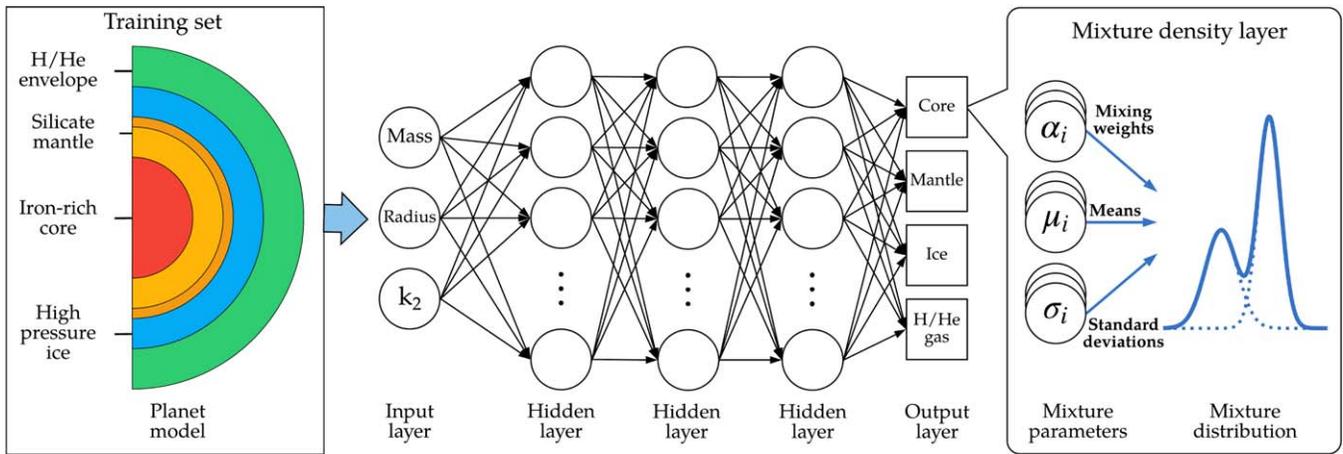


Figure 1. Schematic of the mixture density network architecture used in this work. A large training set of synthetic planets is used to train a MDN with three hidden layers of 512 nodes each. For each interior layer, the MDN outputs the parameters of a linear combination of Gaussians (mean, standard deviation, and weight of each Gaussian) to fit the distribution of planet interior structures constrained by the input parameters (mass, radius, fluid Love number k_2).

weights of the Gaussian mixture, in each case with their specific activation function.

Our MDN is built from a feedforward neural network with three hidden layers having 512 nodes each. For the output layer, we employ a multitask-learning approach (Caruana 1997), where the output layer consists of four separate MDN layers sharing the same hidden layer (see Figure 1). The MDN layers, corresponding to the radius fractions of the four (planetary) layers to be predicted, are built as specified above using $m = 20$ mixture components each. We found this setup to be among the best performing MDN architectures by carrying out a grid search of the hyperparameters of the MDN model, i.e., the amount of hidden layers, the amount of neurons per layer, as well as the learning rate of the optimizer (see below).

The hidden layers are activated with a Rectified Linear Unit (ReLU) function (Nair & Hinton 2010; Goodfellow et al. 2017):

$$\text{ReLU}(x) = \max(0, x). \quad (6)$$

The ReLU is a commonly used nonlinear activation function. Because of their close similarity to linear functions, they are easy to optimize and retain many of the advantages of linear models (Goodfellow et al. 2017).

The MDN layer uses a softmax activation for the α_i nodes, which ensures that the outputs of those nodes are between zero and one and add up to one. The μ_i and σ_i^2 output nodes are activated with a nonnegative exponential linear unit (NNELU) to ensure that predicted means and variances are always positive (Brando 2017). The NNELU is a regular ELU (Clevert et al. 2015) with an additional offset of +1 to keep it positive:

$$\text{NNELU}(x) = 1 + \text{ELU}(x) = 1 + \begin{cases} x & x \geq 0 \\ \exp(x) - 1 & x < 0. \end{cases} \quad (7)$$

We choose the loss function of every output layer to be the average negative log-likelihood (NLL)

$$\mathcal{L}(\Theta) = -\frac{1}{N} \sum_{(x,y) \in \mathcal{D}} \log p(y | x), \quad (8)$$

between predictions and input values, where N is the size of the training data set \mathcal{D} and Θ represents the model parameters. The

model is trained by minimizing the sum of the four individual output losses. Minimization of the loss function is done using the Adam optimizer (Kingma & Ba 2014) with a learning rate of 0.001. The learning rate defines how much the weights are updated when performing the gradient descent.

We include a dropout layer after each ReLU layer with a dropout rate of 0.5 (Srivastava et al. 2014). At each training iteration, a dropout layer randomly disables a number of neurons specified by the dropout rate (in our case half of the neurons) in the subsequent layer. Dropout is a common technique to prevent a neural network from overfitting the training data, because the network will be less sensitive to the weights of individual neurons. This forces neurons to be able to perform well regardless of other neurons in the network, which increases the network’s ability to generalize and improves the general robustness of the model. To further help against overfitting, we apply early stopping (Prechelt 2012) and stop training once the loss from the validation data does not improve for 15 consecutive training epochs.

The MDN is trained on a NVIDIA Quadro P5000 graphics card using the Keras library (Chollet et al. 2015), running on top of the TensorFlow library (Abadi et al. 2015). The code for the MDN layers is adapted from Martin & Duhaime (2019).

2.3. Training Set

We employ our interior structure code to construct each planet model of our training data set. We use a Monte Carlo sampling approach to create a data set of about $N = 900,000$ unique planet models with random interior structures. For each synthetic planet, the planet mass is set to a random value between 0.01 and 25 Earth masses on a logarithmic scale. The mass fractions of the core, ice, and gas layer are also set to a random value, with the mantle layer taking up the remaining mass. We sample core and ice layers linearly, while the gas mass fraction is sampled logarithmically between 10^{-5} and 1. A new random draw is performed if the mass fractions add up to more than 1. In Section 3.5 we discuss our choice of sampling distribution. From planet mass and mass fractions, the planet interior structure and planet radius are calculated. In addition, we calculate the fluid Love number k_2 for each planet following the matrix propagator approach of Padovan et al. (2018). We train the neural network

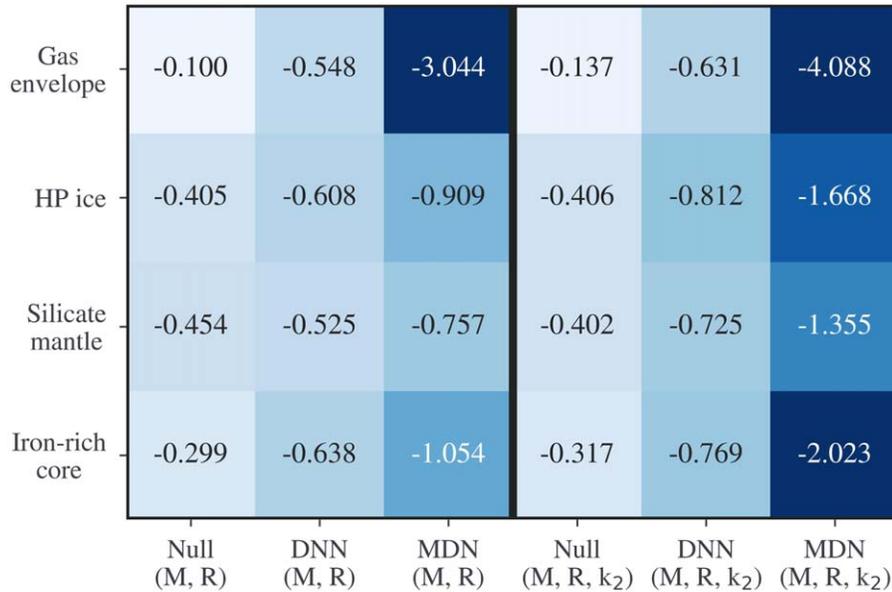


Figure 2. Comparison of the negative log-likelihood $\mathcal{L}(\Theta)$ (NLL) of the MDN and validation models for mass and radius inputs only (left panel) and mass, radius, and k_2 inputs (right panel). Lower values (darker shades of blue) correspond to a better approximation of the validation data by the model.

with 70% of the data set and validate the model results with the remaining 30%. We test two input scenarios: mass and radius only (case 1); mass, radius, and k_2 (case 2).

3. Results

3.1. Validation with Baseline Models

We establish a baseline model to evaluate the training performance of the MDN. We use a DNN with the same parameters and training routine as the MDN, with the exception that the output layer consists of four nodes relating to the radius fraction of the four planet layers. This layer is additionally activated with a softmax activation function (Goodfellow et al. 2017) to ensure that all output values add up to one, introducing an additional constraint. This DNN is in effect a MDN with only one mixture component. We also define a null model that always predicts the mean and standard deviation of layer thicknesses in the training data, thus representing the average interior structure in the data set. The null model defines the minimum requirements a model must fulfill, in that the thickness of each planet layer must be predicted better than the global average of the training data. A model performing equal or worse than the null model is basically guessing random values. For an indicator of the goodness of the training of the model, we compare the NLL (Equation (8)) of all three models as shown in Figure 2. The three models on the left were trained on mass and radius only (case 1), the three models on the right were trained with mass, radius, and k_2 (case 2). Lower values (corresponding to darker shades of blue) indicate a better performance of the respective models.

Both the DNN and MDN perform better than the null model for all planetary interior layers. The MDN shows significant improvement for all interior layer predictions compared to the DNN due to its ability to accommodate multiple interior solutions for one mass and radius input.

In Figure 3, we demonstrate the dependence of the NLL loss on the size of the training data set. For each training size, we changed the random seed to sample a different set of training data. For both case 1 and case 2, the loss improves with increasing

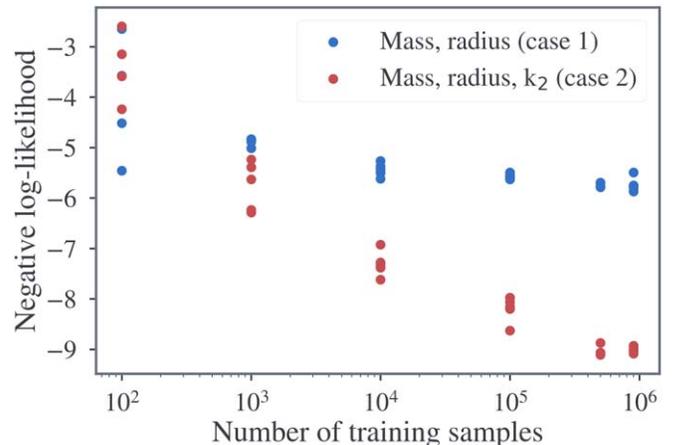


Figure 3. Dependence of the NLL on the size of the training set for the MDN using only mass and radius inputs (blue dots), and for the MDN using mass, radius, and k_2 (red dots). For each training set size, we changed the random seed to generate a different set of training samples.

sample numbers, and the spread of loss values is reduced. The NLL flattens for higher sample numbers. We conclude that a training set size of $N \approx 10^5$ is sufficient to ensure a good fit of the MDN. Case 2, using mass, radius, and k_2 , needs more training samples than case 1, due to the increased number of observables.

3.2. MDN Accuracy and Degeneracies

The ability of the network to constrain the interior structure of each planet based on given input parameters is displayed in Figure 4, which shows the distributions predicted by the MDN from the validation data set of each interior layer plotted against the actual relative thickness of the layer. A MDN model that perfectly predicts the actual data set would plot along the diagonal.

With only mass and radius as inputs, the MDN can infer the relative thickness of the gaseous envelope well. This result is similar to the findings of Lopez & Fortney (2014), who showed

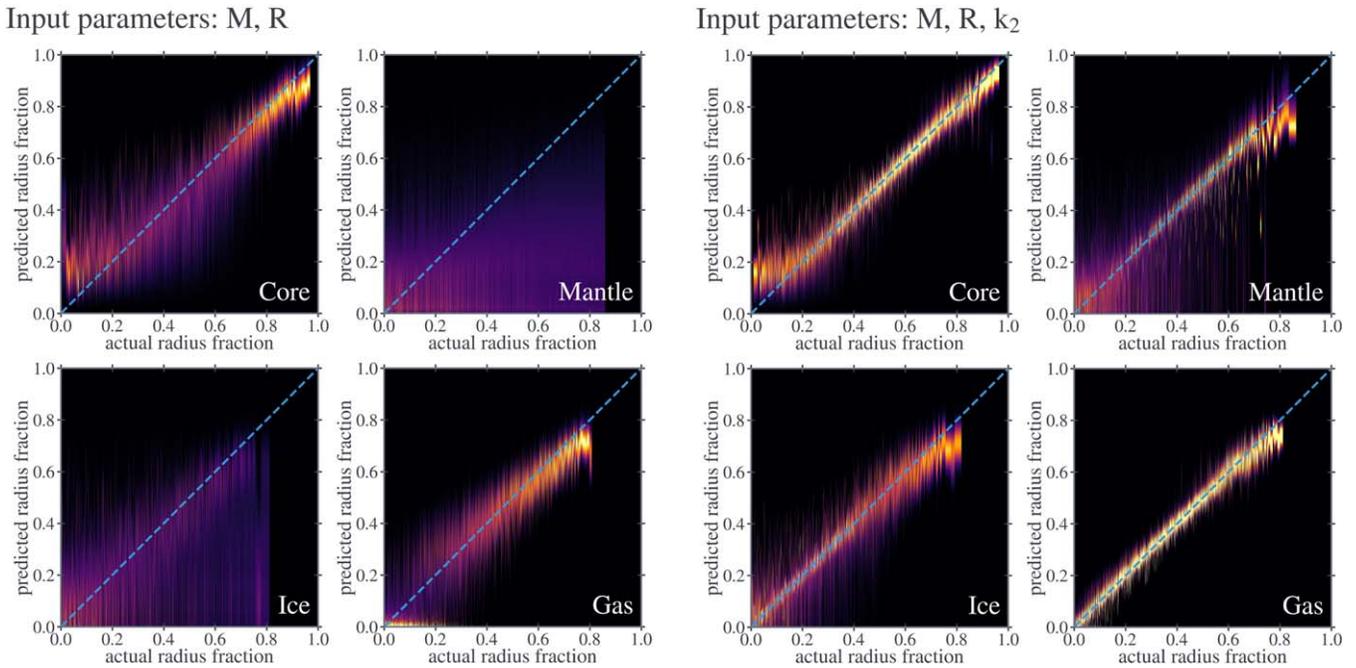


Figure 4. Predicted distributions of the thickness of core, mantle, ice, and gas-envelope layers as a function of actual values from the validation data set. The dashed diagonal line indicates a perfect prediction of the layer thickness. The four panels on the left show predictions for a MDN trained on mass and radius inputs only. The four panels on the right show the predicted distributions for a MDN trained additionally with the fluid Love number k_2 . The predicted distributions are colored according to the local probability density.

that the planetary radius can be used as a first-order proxy for the mass of the H/He envelope. Likewise, the iron-rich core is relatively well constrained. For model planets with large cores, the predictions are quite accurate, as the high density of these planets does not allow for much variation in possible interior structures, given that we do not allow for planets denser than a naked core. The same applies to planets with thick gas envelopes, as there is nothing less dense than the gas. The MDN tends to overestimate the core thickness when the core is small. The reason for this behavior lies in the nature of the mass–radius diagram: for planets with small cores, a slight increase of the core mass leads to a larger relative increase in core size than for planets with massive cores, leading to a higher degeneracy of the interior structures. As such, very small core sizes are undersampled in the training data set (see also Figure 7). This makes the core size of small core planets more difficult for the MDN to predict. Similarly, the thickness of large layers tends to be underestimated due to less training samples available.

The high-pressure ice layer thicknesses are constrained relatively well for the most part, but the MDN tends to underestimate the thickness of the ice layer in some parts of the data set. The silicate mantle is barely constrained. The MDN mostly predicts a broad distribution with constant mean values, but performs slightly better for planets with small mantles. This is not an unexpected result, as it is a consequence of the inherent degeneracy in the mantle and ice layer thickness due to both layers having intermediate densities in comparison with other layers in the planet. As such, their masses and radii can easily be compensated by those other layers. Figure 4 shows that the thickness of mantle and ice layers is barely constrainable for most planets.

The inclusion of k_2 dramatically improves the performance of the MDN, which now adequately predicts the correct

thickness of each interior layer (right panels of Figure 4). The biggest improvement is in the prediction of the mantle layer. However, the mantle remains the least well-predicted layer as already shown by the scores of Figure 2. The MDN still tends to overestimate the thickness of the core, mantle and ice layer when these layers are thin due to undersampling in the training data set. The previous underestimation of the ice layer is no longer observed. The gas layer thickness is predicted with a high accuracy, even for thin layers. All in all, the inclusion of k_2 , by providing a measure of the degree of mass concentration in the interior, significantly constrains the solution space and helps reduce the degeneracy of interior solutions.

3.3. Application to Solar System and Extrasolar Planets

We investigate the predicted interior structures for four representative planets. In the solar system, we consider Earth, being the archetypal terrestrial planet, and Neptune as a gas and ice-rich planet at the upper end of our mass range (with a mass of $17.1 M_{\oplus}$ and a radius of $3.865 R_{\oplus}$). For exoplanets, we investigate Kepler-10b ($3.72 M_{\oplus}$, $1.47 R_{\oplus}$, Weiss et al. 2016), which with its high density is a likely Super-Earth (Valencia 2010; Batalha et al. 2011; Wagner et al. 2012), and GJ 1214b (Charbonneau et al. 2009; Rogers & Seager 2010b; Nettelmann et al. 2011), a well-studied, relatively low-density planet ($6.26 M_{\oplus}$, $2.85 R_{\oplus}$, Højsøe et al. 2013), likely with a volatile-rich interior and extended gas envelope.

Figure 5 shows the distributions of predicted solutions of the MDN for each interior layer for the four planets using only mass and radius as inputs for the network. The histograms behind each prediction curves show all interior solutions within a 5% margin of mass and radius, assuming a uniform error distribution where every input parameter is equally likely. To obtain these samples, we modeled an independent data set by Monte Carlo sampling from the same prior distribution of mass

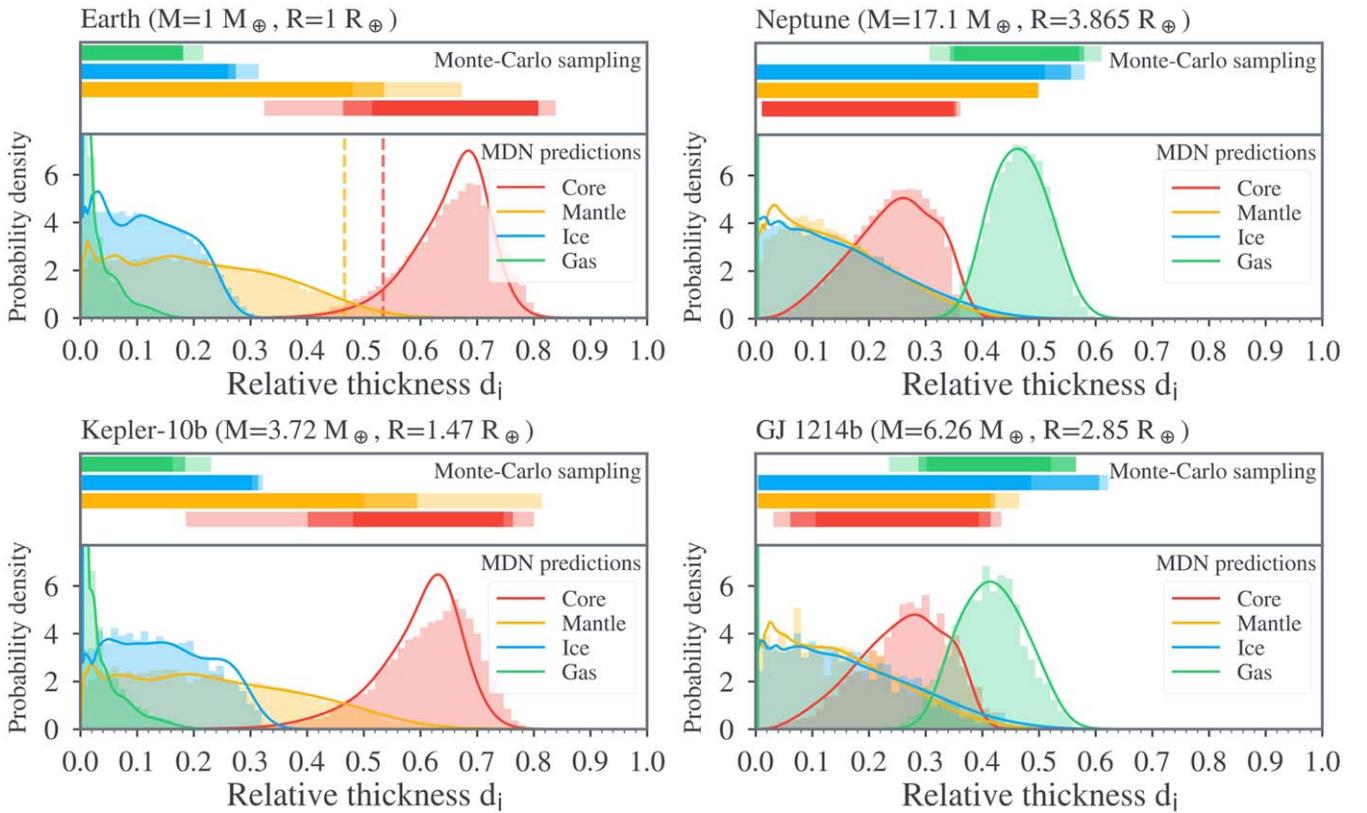


Figure 5. MDN predictions of the relative interior layer thickness for four representative planet cases using only mass and radius as inputs to the network. The colored lines show the combined Gaussian mixture prediction of the MDN. The area under each curve is normalized to one. The histograms show possible interior solutions within 5% of the input values obtained from an independent Monte Carlo sampling (see the main text for more details). The colored bars in the upper plots represent the space of solutions of valid interior structures obtained from the same Monte Carlo sampling. The shade of each bar corresponds to an uncertainty of 1%, 2%, and 5% in the observable parameters (from darkest to lightest, respectively). The vertical dashed lines in the Earth plot show the actual thickness of Earth’s core and mantle from PREM (Dziewonski & Anderson 1981). Data for Kepler-10b from Weiss et al. (2016). Data for GJ 1214b from Harpsøe et al. (2013).

fractions that was used to create the training data set. Additionally, the bars on top of each plot show the spread of possible solutions obtained from this Monte Carlo sampling, where the color shade shows all samples within the 1%, 2%, and 5% error margins in input parameters (from darkest to lightest). The relative error δ of each Monte Carlo sample is calculated as the Euclidean distance from the input parameters

$$\delta = \sqrt{\left(\frac{\Delta M}{M}\right)^2 + \left(\frac{\Delta R}{R}\right)^2}, \quad (9)$$

where M and R are the true mass and radius values, and ΔM and ΔR the absolute errors between the Monte Carlo sample and true value. The comparison of the MDN predictions (solid lines) and the Monte Carlo distributions (histograms) in Figure 5 show that the MDN accurately replicates the distributions of possible interior solutions for all planets.

In the case of the Kepler-10b, the predicted core size is consistent with values from other studies (e.g., Wagner et al. 2012), showing it to be most likely a terrestrial-type planet. In contrast, the predicted distributions of GJ 1214b are more similar to Neptune’s interior. This compares well with the study of Nettelmann et al. (2011), who conclude that the interior of GJ 1214b is consistent with having a large gaseous, water-enriched envelope. Other possible interior structures in their study include the planet being mostly rocky with a large H/He envelope taking up 40% of the planet’s radius.

For the Earth, the dashed lines in Figure 5 indicate the actual core and mantle thickness as known from the Preliminary Reference Earth Model (Dziewonski & Anderson 1981). These do not align with the peaks of the predicted MDN distributions, which reproduce the correct values only near their tails (Figure 5). This is a result of the large degeneracy in interior solutions. A planet like the Earth, consisting of iron core and silicate mantle only, is only one solution (in terms of possible solutions of the interior structure) compared to a generic planet having ice and gas layers. The situation would partly improve with additional constraints on the planet, such as transmission spectra or detailed escape modeling ruling out the presence of a substantial atmosphere (e.g., Owen & Wu 2017; Dorn & Heng 2018; Kubyskhina et al. 2018).

3.4. Influence of k_2 on the Predictability of the Interior Structure

As shown in Figure 4, adding the fluid Love number k_2 as a third observable helps to significantly constrain the interior structure of the planets. Figure 6 (left) shows the predictions of the neural network including the fluid Love number k_2 for the Earth (with $k_2 = 0.933$ from Lambeck 1980). For Neptune (right), the actual value of k_2 is 0.392 calculated from the first order of the expansion of the gravitational moment $J_2 = 3.40843 \times 10^{-3}$ (Jacobson 2009), using the volumetric mean radius of 24.622×10^6 m and a *Voyager* rotation period of 16.11 hr (Hubbard 1984). In this work, however, the assumption of the

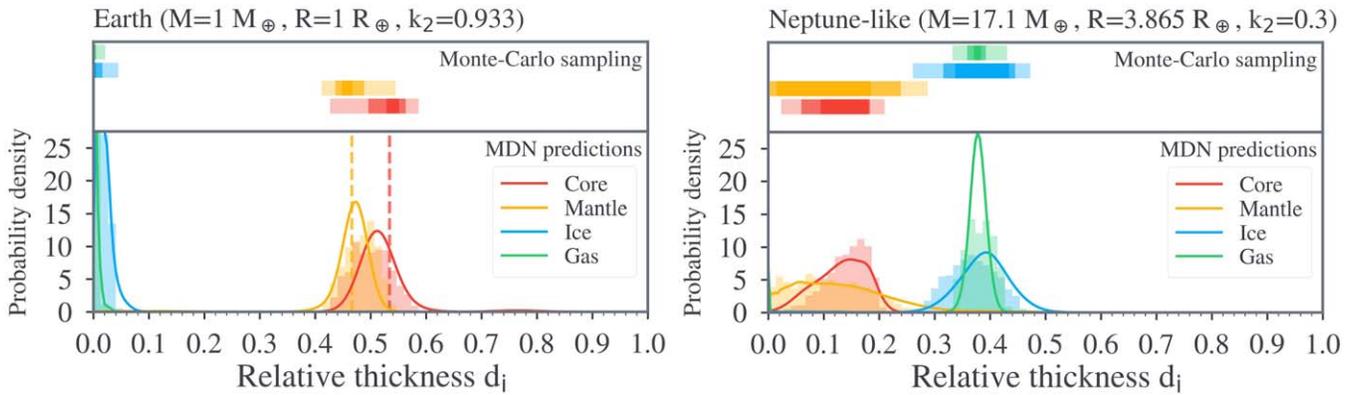


Figure 6. MDN predictions of the relative interior layer thickness for the Earth (left) and a Neptune-like planet with $k_2 = 0.3$ (right) as shown in Figure 5, but using mass, radius, and the fluid Love number k_2 as inputs to the network.

zero-temperature EoS for the gaseous envelope puts the real mass, radius, and k_2 triplet of Neptune outside of our sampled parameter space, preventing us from reaching the actual value of k_2 . As mentioned in Section 2.1.4, our model overestimates the thickness of the gas envelope when constraining the planet’s interior to mass and radius. As a result, planets with significant gas fractions are more centrally condensed in our model, resulting in low values of k_2 . For a planet with Neptune mass of $17.1 M_{\oplus}$ and radius of $3.865 R_{\oplus}$, possible values of k_2 lie between 0.2 and 0.325 with our model. As such, we chose to use a value of $k_2 = 0.3$, which allows the interior model to return a good spread of solutions while still being close to the actual value. We call this planet *Neptune-like*, noting that a more sophisticated, temperature-dependent EoS of the gas envelope would be needed to model Neptune in better detail and with the correct k_2 .

In both cases, the inclusion of k_2 significantly shrinks the solution space. For the Earth, the interior structure of the Earth is constrained well to the actual values. In the case of the Neptune-like planet, the MDN predictions clearly indicate this planet to be an ice-rich planet with an extended gas envelope and a small iron-silicate core. The predictions are also consistent with the planet having no iron-silicate core at all. The gas-envelope thickness, assuming a Love number of 0.3, can be well constrained by the MDN.

Compared to the 3-layer interior structures models of Neptune by Podolak et al. (1995) and more recent ones by Nettelmann et al. (2013), which show that Neptune’s ice-enriched gaseous envelope makes up 10%–30% of the planet’s radius, the MDN overestimates the thickness of the gas envelope to be 30%–60% of the planet’s radius when considering mass and radius only (Figure 5), and 30%–44% of the planet’s radius when considering mass, radius, and k_2 (Figure 6). There are several reasons for this finding. For one, we use a basic H/He atmosphere with no temperature-dependence of the envelope, and a pure water ice layer underneath. Compared to interior density profiles of Nettelmann et al. (2013), we overestimate the density of the ice layer by $\approx 25\%$ in the lower parts of the planet. This is a result of our choice of a high-pressure EoS neglecting ionic and super-ionic water phases. As a result the gas envelope needs to be larger to arrive at the same mass and radius, compensating the denser interior (see Section 2.1.4). Compared to more recent studies, the zero-temperature EoS that we use overestimates the density especially in the upper part of the atmosphere. The combination of these two factors causes the envelope to be comparatively larger than currently believed. In turn, this also leads to our planet models being more centrally

condensed, so that k_2 tends to be underestimated in all models with extended atmospheres.

Additionally, our interior model assumes strictly immiscible layers. Some layer phases, such as H_2 and H_2O , can actually be miscible under pressure and temperature conditions in gas giants. As such, the layered interior model cannot fully account for the interior of Neptune (see also Podolak et al. 1995).

The Neptune-like predictions of the MDN are however fully consistent with the Monte Carlo sampling results. As such, deviations from the actual interior structure are only due to simplifications in the interior model, reflected in the training data. Therefore, we expect that using a more appropriate atmosphere and an improved ice model will help us to improve the accuracy of our predictions.

3.5. Influence of Prior Distribution

As the MDN predicts the posterior distribution of possible planetary structures solely from the data it was given during training, the distribution of training data points can be a main source of prediction errors. In Figure 7, we show the effect of a different prior choice on the predicted interior structures of Earth and Neptune. Solid lines represent the baseline prior distributions in mass (blue) and radius (red), where mass sampling is linear except than for the gaseous envelope, where it is logarithmic. The dashed lines are the priors obtained assuming a linear mass sampling in each layer. This drastically reduces the amount of planets with a small envelope (red lines), shifting focus to planets with more extended envelopes. This slightly changes the priors of the other layers as well.

The bottom plots in Figure 7 show a comparison of the predictions of a MDN trained with this new resampled data set with the predictions of the original MDN. For the Earth case, the greatest difference lies in the prediction of the envelope thickness. Here, the MDN tends to predict a preferentially thicker envelope. This is due to the new prior giving less priority to very small envelopes. As such, there are fewer training samples available for small envelopes. As a result of the changed envelope prediction, the predictions of the other layers also change slightly. Mantle and ice predictions shift to slightly lower thicknesses, the core shifts to slightly larger thicknesses. However, the predicted interior structures still lie in a similar range, only the shape of the curves are different. For Neptune, there is very little difference between predictions for the new and old priors, as planets with larger envelope fractions are equally well sampled with both prior distributions.

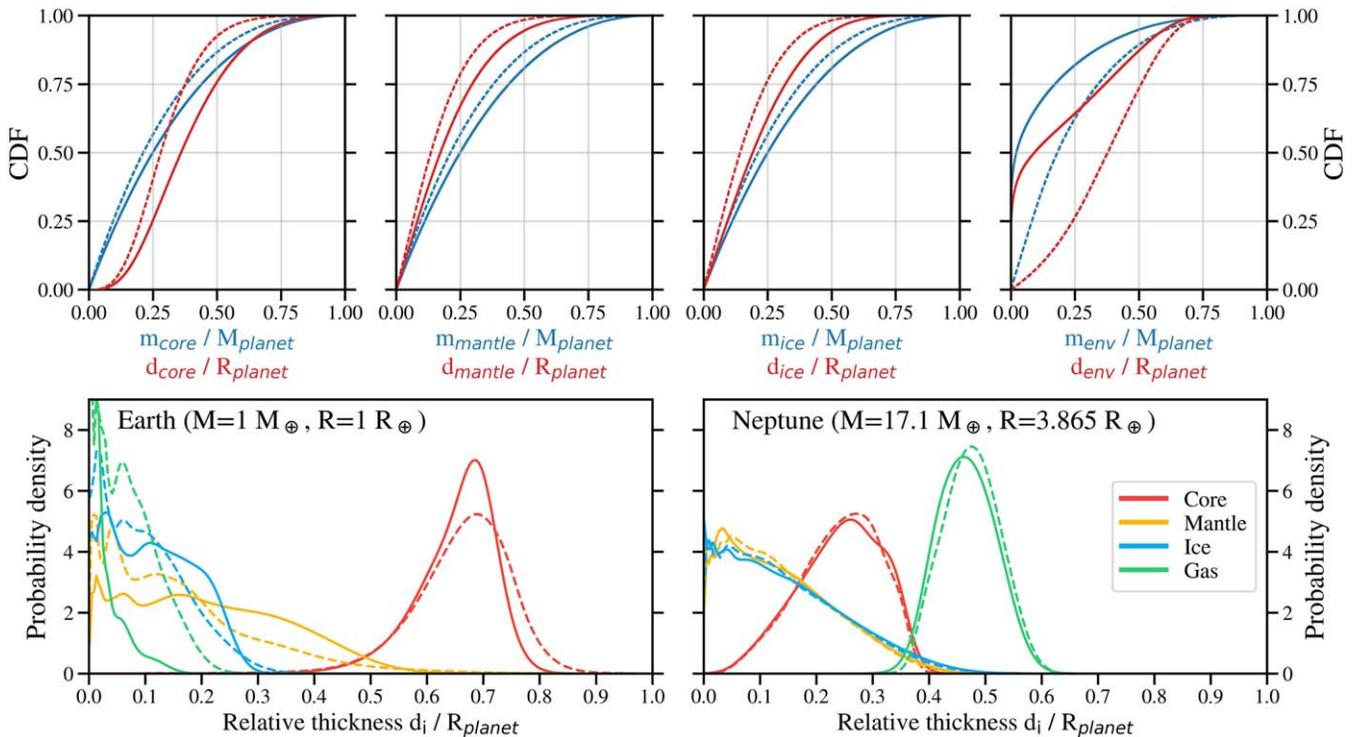


Figure 7. Top plots: cumulative density function (CDF) of the training data set for layer mass fractions (blue) and layer radius fractions (red). Solid lines represent the distribution of the full training set with a logarithmic sampling of the envelope mass fractions. For the dashed lines, the data set was resampled so that the distribution of gas mass fractions approximates the distribution of the other layers’ mass fractions (with linear sampling). Bottom plots: comparison of predictions for both prior distributions for Earth and Neptune (using mass and radius as network inputs).

We conclude that the exact shape of the prior distributions is not too important as long as there are enough samples in each region of the solution space.

4. Discussion and Conclusions

We show that MDNs (Figure 1) can accurately predict the space of interior solutions of planets with mass below 25 Earth masses (Figure 5). There are two main advantages in this machine-learning-based approach compared to conventional methods (e.g., Monte Carlo sampling): First, predicting the distribution of possible interior structures of a planet with a trained MDN only takes few milliseconds on a normal desktop CPU (we used an AMD EPYC 7351) compared to potentially several hours if the parameter space needs to be sampled with a forward model. Thus, trained MDNs can be used to provide a rapid first characterization of the space of admissible interior structures of observed planets. These can be then investigated with more sophisticated models that account for additional observables relevant to specific planets such as stellar irradiation, planet age, and atmospheric composition. Second, the trained MDN represents a stand-alone tool: while its predictions rely on the assumptions of the underlying interior model, the interior model itself as well as the training data set are not needed to operate the MDN. This renders MDNs a valuable tool for the exoplanetary science community because a fully trained MDN can be used to obtain an indication of possible interior structures without the need of a dedicated code for solving the forward problem. The file size of our MDN is ≈ 7.5 MB, making its dissemination easy. Our trained MDN models are available on GitHub.⁴

The generation of a data set of forward models of appropriate size can be computationally expensive, especially if a complex interior model is used or if additional observables are considered (Figure 3). Conversely, the training of the network itself only takes a few minutes, and can be significantly sped up by utilizing GPU acceleration (Shi et al. 2016). Therefore, several MDNs corresponding to different combinations of input parameters (i.e., observables) can be efficiently tested, making it possible to quickly identify combinations of parameters that better characterize the interior of a planet.

Similar to other inference methods, the predicted distributions depend to some extent on the choice of the prior distributions (see, e.g., Dorn et al. 2017). With better informed priors constrained by observations and physical considerations, the predictions will be more accurate. The MDN approach presented here reproduces the well-known degeneracy of interior structures based on mass–radius pairs only (Figures 4 and 5). Using k_2 as an additional input significantly reduces the set of possible interior structures (Figures 4 and 6).

In the solar system there are only few planetary objects, each characterized by several remote and in situ observations that provide a relatively accurate picture of their formation, history, and interior structure. The problem of understanding exoplanets is complementary, in that thousands of objects are known, each characterized by a severely limited number of observations. MDNs provide a fast alternative to other, often time-consuming, inference methods, allowing an efficient approach to exploiting exoplanetary data, which span a parameter space much wider than that of the planets of the solar system. As new type of data will become available with the advent of new observing facilities (e.g., atmospheric compositions from *James Webb Space Telescope*), models more complex that

⁴ https://github.com/philippbaumeister/MDN_exoplanets

those presented here, and larger data sets of forward models to train appropriate MDNs will be required. These are modifications that only depend on the available computing power, and can be readily incorporated in the method presented here as the need arises.

We thank the anonymous reviewer for insightful comments that helped improve the paper. We describe contributions to the paper using the taxonomy of Brand et al. (2015). For each entry, the first name is lead, the following are contributing. *Conceptualization*: S.P., N.T.; *Methodology*: P.B., S.P., N.T., G.M., M.G.

Software: P.B., S.P.; *Investigation*: P.B.; *Writing—Original Draft*: P.B., S.P., N.T.; *Writing—Review & Editing*: P.B., S.P., N.T., G.M., N.N., J.M., M.G.; *Visualization*: P.B.; *Supervision*: N.T., S.P., M.G.; *Funding Acquisition*: N.T., M.G. The authors acknowledge the support of the DFG Priority Program SPP 1992 “Exploring the Diversity of Extrasolar Planets” (TO 704/3-1, GO 2610/2-1) and the DFG Research Unit FOR 2440 “Matter under planetary interior conditions.” N.T. also acknowledges support by the Helmholtz Association (VH-NG-1017), and M.G. by the DFG (GO 2610/1-1).

ORCID iDs

Philipp Baumeister  <https://orcid.org/0000-0001-9284-0143>

Sebastiano Padovan  <https://orcid.org/0000-0002-8652-3704>

Nicola Tosi  <https://orcid.org/0000-0002-4912-2848>

Nadine Nettelmann  <https://orcid.org/0000-0002-1608-7185>

Mareike Godolt  <https://orcid.org/0000-0003-4770-8551>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, <https://www.tensorflow.org/>
- Adams, E. R., Seager, S., & Elkins-Tanton, L. 2008, *ApJ*, **673**, 1160
- Akinsanmi, B., Barros, S. C. C., Santos, N. C., et al. 2019, *A&A*, **621**, A117
- Alibert, Y., & Venturini, J. 2019, *A&A*, **626**, A21
- Bakos, G. Á., Csabry, Z., Penev, K., et al. 2013, *PASP*, **125**, 154
- Barr, A. C., Dobos, V., & Kiss, L. L. 2018, *A&A*, **613**, A37
- Batalha, N. M., Borucki, W. J., Bryson, S. T., et al. 2011, *ApJ*, **729**, 27
- Batygin, K., Bodenheimer, P., & Laughlin, G. 2009, *ApJL*, **704**, L49
- Bishop, C. M. 1994, Mixture Density Networks, Tech. Rep. NCRG/94/004
- Bishop, C. M. 1995, Neural Networks for Pattern Recognition (New York: Oxford Univ. Press)
- Bishop, C. M. 2006, Pattern Recognition and Machine Learning (New York: Springer)
- Brand, A., Allen, L., Altman, M., Hlava, M., & Scott, J. 2015, *LePub*, **28**, 151
- Brando, A. 2017, Master’s thesis, Univ. de Barcelona
- Brugger, B., Mousis, O., Deleuil, M., & Deschamps, F. 2017, *ApJ*, **850**, 93
- Caruana, R. 1997, *Machine Learning*, **28**, 41
- Charbonneau, D., Berta, Z. K., Irwin, J., et al. 2009, *Natur*, **462**, 891
- Chaushev, A., Raynard, L., Goad, M. R., et al. 2019, *MNRAS*, **488**, 5232
- Chollet, F., et al. 2015, Keras: The Python Deep Learning library, <https://keras.io>
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. 2015, arXiv:1511.07289
- Csizmadia, S., Hellard, H., & Smith, A. M. S. 2019, *A&A*, **623**, A45
- Cybenko, G. 1989, *Mathematics of Control, Signals and Systems*, **2**, 303
- Dom, C., & Heng, K. 2018, *ApJ*, **853**, 64
- Dom, C., Khan, A., Heng, K., et al. 2015, *A&A*, **577**, A83
- Dom, C., Venturini, J., Khan, A., et al. 2017, *A&A*, **597**, A37
- Dziewonski, A. M., & Anderson, D. L. 1981, *PEPI*, **25**, 297
- Fei, Y., Bertka, C. M., & Finger, L. W. 1997, *Sci*, **275**, 1621
- Fei, Y., Murphy, C., Shibasaki, Y., Shahar, A., & Huang, H. 2016, *GeoRL*, **43**, 6837
- Fortney, J. J., Marley, M. S., & Barnes, J. W. 2007, *ApJ*, **659**, 1661
- Goodfellow, I., Bengio, Y., & Courville, A. 2017, Deep Learning (London: MIT Press)
- Harpsoe, K. B. W., Hardis, S., Hinse, T. C., et al. 2013, *A&A*, **549**, A10
- Hartman, J. D., Bakos, G., Stanek, K. Z., & Noyes, R. W. 2004, *AJ*, **128**, 1761
- Hellard, H., Csizmadia, S., Padovan, S., et al. 2019, *ApJ*, **878**, 119
- Hemley, R. J., Jephcoat, A. P., Mao, H. K., et al. 1987, *Natur*, **330**, 737
- Hubbard, W. B. 1984, Planetary Interiors (New York: Van Nostrand Reinhold)
- Jacobson, R. A. 2009, *AJ*, **137**, 4322
- James, G., Witten, D., Hastie, T., & Tibshirani, R. 2013, An Introduction to Statistical Learning (New York: Springer)
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kramm, U., Nettelmann, N., Redmer, R., & Stevenson, D. J. 2011, *A&A*, **528**, A18
- Kubyskhina, D., Fossati, L., Erkaev, N. V., et al. 2018, *A&A*, **619**, A151
- Lambeck, K. 1980, The Earth’s Variable Rotation: Geophysical Causes and Consequences (Cambridge: Cambridge Univ. Press)
- Léger, A., Selsis, F., Sotin, C., et al. 2004, *Icar*, **169**, 499
- Lopez, E. D., & Fortney, J. J. 2014, *ApJ*, **792**, 1
- Lu, Z., Pu, H., Wang, F., Hu, Z., & Wang, L. 2017, Advances in Neural Information Processing Systems, in press, arXiv:1709.02540
- Madhusudhan, N., Agúndez, M., Moses, J. I., & Hu, Y. 2016, *SSRv*, **205**, 285
- Mardling, R. A. 2007, *MNRAS*, **382**, 1768
- Márquez-Neila, P., Fisher, C., Sznitman, R., & Heng, K. 2018, *NatAs*, **2**, 719
- Martin, C., & Duhaime, D. 2019, cpmpercussion/keras-mdn-layer v0.2, Zenodo, doi:10.5281/zenodo.2578015
- Morard, G., Andrault, D., Guignot, N., et al. 2008, *E&PSL*, **272**, 620
- Murthy, V. R., & Hall, H. 1970, *PEPI*, **2**, 276
- Nair, V., & Hinton, G. E. 2010, in Proc. 27th Int. Conf. on Machine Learning, ICML’10 (Omnipress), 807, <http://dl.acm.org/citation.cfm?id=3104322.3104425>
- Nettelmann, N., Fortney, J. J., Kramm, U., & Redmer, R. 2011, *ApJ*, **733**, 2
- Nettelmann, N., Helled, R., Fortney, J. J., & Redmer, R. 2013, *P&SS*, **77**, 143
- Owen, J. E., & Wu, Y. 2017, *ApJ*, **847**, 29
- Padovan, S., Spohn, T., Baumeister, P., et al. 2018, *A&A*, **620**, A178
- Pearson, K. A., Palafox, L., & Griffith, C. A. 2018, *MNRAS*, **474**, 478
- Pepe, F. A., Cristiani, S., Rebolo Lopez, R., et al. 2010, *Proc. SPIE*, **7735**, 77350F
- Petit, G., & Luzum, B. 2010, IERS Technical Note No. 36, Frankfurt am Main: Verlag des Bundesamts für Kartographie und Geodäsie, <https://www.iers.org/IERS/EN/Publications/TechnicalNotes/tn36.html>
- Podolak, M., Weizman, A., & Marley, M. 1995, *P&SS*, **43**, 1517
- Poirier, J.-P. 1994, *PEPI*, **85**, 319
- Pollacco, D. L., Skillen, I., Collier Cameron, A., et al. 2006, *PASP*, **118**, 1407
- Prechelt, L. 2012, in Neural Networks: Tricks of the Trade, ed. G. Montavon, G. B. Orr, & K.-R. Müller (2nd ed.; Berlin: Springer), 53
- Rauer, H., Catala, C., Aerts, C., et al. 2014, *ExA*, **38**, 249
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, *JATIS*, **1**, 014003
- Rogers, L. A., & Seager, S. 2010a, *ApJ*, **712**, 974
- Rogers, L. A., & Seager, S. 2010b, *ApJ*, **716**, 1208
- Rubie, D. C., Nimmo, F., & Melosh, H. J. 2015, in Treatise on Geophysics, ed. G. Schubert (Amsterdam: Elsevier), 43
- Salpeter, E. E., & Zapolsky, H. S. 1967, *PhRv*, **158**, 876
- Seager, S., Kuchner, M., Hier-Majumder, C., & Militzer, B. 2007, *ApJ*, **669**, 1279
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, **155**, 94
- Shi, S., Wang, Q., Xu, P., & Chu, X. 2016, arXiv:1608.07249
- Skidmore, W., & TMT International Science Development Teams, & Science Advisory Committee, T. 2015, *RAA*, **15**, 1945
- Sotin, C., Grasset, O., & Mocquet, A. 2007, *Icar*, **191**, 337
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. 2014, Journal of Machine Learning Research, **15**, 1929, <http://jmlr.org/papers/v15/srivastava14a.html>
- Stixrude, L. 2014, *RSPTA*, **372**, 20130076
- Thomas, S. W., & Madhusudhan, N. 2016, *MNRAS*, **458**, 1330
- Unterborn, C. T., & Panero, W. R. 2019, *JGRE*, **124**, 1704
- Valencia, D. 2010, IAU Symp. 276, The Astrophysics of Planetary Systems: Formation, Structure, and Dynamical Evolution (Cambridge: Cambridge Univ. Press), 181
- Valencia, D., O’Connell, R. J., & Sasselov, D. D. 2006, *Icar*, **181**, 545
- Valencia, D., Sasselov, D. D., & O’Connell, R. J. 2007, *ApJ*, **665**, 1413
- Wagner, F., Sohl, F., Hussmann, H., Grott, M., & Rauer, H. 2011, *Icar*, **214**, 366
- Wagner, F. W., Tosi, N., Sohl, F., Rauer, H., & Spohn, T. 2012, *A&A*, **541**, A103
- Weiss, L. M., Rogers, L. A., Isaacson, H. T., et al. 2016, *ApJ*, **819**, 83
- Wheatley, P. J., West, R. G., Goad, M. R., et al. 2017, *MNRAS*, **475**, 4476
- Yoo, C. S., Holmes, N. C., Ross, M., Webb, D. J., & Pike, C. 1993, *PhRvL*, **70**, 3931
- Zeng, L., Jacobsen, S. B., Sasselov, D. D., et al. 2019, *PNAS*, **116**, 9723
- Zeng, L., & Sasselov, D. 2013, *PASP*, **125**, 227
- Zingales, T., & Waldmann, I. P. 2018, *AJ*, **156**, 268