

EGU2020-16162

<https://doi.org/10.5194/egusphere-egu2020-16162>

EGU General Assembly 2020

© Author(s) 2021. This work is distributed under the Creative Commons Attribution 4.0 License.



Mars' thermal evolution from machine-learning-based 1D surrogate modelling

Siddhant Agarwal^{1,2}, Nicola Tosi¹, Doris Breuer¹, Sebastiano Padovan¹, and Pan Kessel²

¹German Aerospace Center, DLR, Planetary Physics, Germany (siddhant.agarwal@dlr.de)

²Machine Learning/Intelligent Data Analysis, TU Berlin, Berlin, Germany

The parameters and initial conditions governing mantle convection in terrestrial planets like Mars are poorly known meaning that one often needs to randomly vary several parameters to test which ones satisfy observational constraints. However, running forward models in 2D or 3D is computationally intensive to the point that it might prohibit a thorough scan of the entire parameter space. We propose using Machine Learning to find a low-dimensional mapping from input parameters to outputs. We use about 10,000 thermal evolution simulations with Mars-like parameters run on a 2D quarter cylindrical grid to train a fully-connected Neural Network (NN). We use the code GAIA (Hüttig et al., 2013) to solve the conservation equations of mantle convection for a fluid with Newtonian rheology and infinite Prandtl number under the Extended Boussinesq Approximation. The viscosity is calculated according to the Arrhenius law of diffusion creep (Hirth & Kohlstedt, 2003). The model also considers the effects of partial melting on the energy balance, including mantle depletion of heat producing-elements (Padovan et., 2017), as well as major phase transitions in the olivine system.

To generate the dataset, we randomly vary 5 different parameters with respect to each other: thermal Rayleigh number, internal heating Rayleigh number, activation energy, activation volume and a depletion factor for heat-producing elements in the mantle. In order to train in time, we take the simplest possible approach, i.e., we treat time as another variable in our input vector. 80% of the dataset is used to train our NN, 10% is used to test different architectures and to avoid over-fitting, and the remaining 10% is used as test set to evaluate the error of the predictions. For given values of the five parameters, our NN can predict the resulting horizontally-averaged temperature profile at any time in the evolution, spanning 4.5 Ga with an average error under 0.3% on the test set. Tests indicate that with as few as 5% of the training samples (= simulations x time steps), one can achieve a test-error below 0.5%, suggesting that for this setup, one can potentially learn the mapping from fewer simulations.

Finally, we ran a fourth batch of GAIA simulations and compared them to the output of our NN. In almost all cases, the instantaneous predictions of the 1D temperature profiles from the NN match those of the computationally expensive simulations extremely well, with an error below 0.5%.