

Interplanetary Transfers via Deep Representations of the Optimal Policy and/or of the Value Function

Dario Izzo
European Space Agency
Noordwijk, The Netherlands
dario.izzo@esa.int

Ekin Öztürk
European Space Agency
Noordwijk, The Netherlands
ekin.ozturk@esa.int

Marcus Mörtens
European Space Agency
Noordwijk, The Netherlands
marcus.maertens@esa.int

ABSTRACT

A number of applications to interplanetary trajectories have been recently proposed based on deep networks. These approaches often rely on the availability of a large number of optimal trajectories to learn from. In this paper we introduce a new method to quickly create millions of optimal spacecraft trajectories from a single nominal trajectory. Apart from the generation of the nominal trajectory, no additional optimal control problems need to be solved as all the trajectories, by construction, satisfy Pontryagin's minimum principle and the relevant transversality conditions. We then consider deep feed forward neural networks and benchmark three learning methods on the created dataset: policy imitation, value function learning and value function gradient learning. Our results are shown for the case of the interplanetary trajectory optimization problem of reaching Venus orbit, with the nominal trajectory starting from the Earth. We find that both policy imitation and value function gradient learning are able to learn the optimal state feedback, while in the case of value function learning the optimal policy is not captured, only the final value of the optimal propellant mass is.

KEYWORDS

artificial neural networks, optimization, deep learning, trajectory design

ACM Reference Format:

Dario Izzo, Ekin Öztürk, and Marcus Mörtens. 2019. Interplanetary Transfers via Deep Representations of the Optimal Policy and/or of the Value Function. In *Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion)*, July 13–17, 2019, Prague, Czech Republic. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3319619.3326834>

1 INTRODUCTION

The use of deep neural networks (DNNs) for the guidance navigation and control of space systems (spacecraft, landers, etc.) is a prolific area of research as witnessed by the increasing number of results that appeared recently on these topics [3, 4, 6, 9, 11, 12, 14]. Most of this body of work can be divided into DNNs approximating the optimal policy u (e.g. the optimal thrust profile), and DNNs

approximating the value function v (e.g. optimal propellant consumption). The final goal of this line of research is to develop DNNs that can be part of the on-board software to steer the spacecraft and substitute the guidance and control loops currently used. Furthermore, accurate predictions of important quantities (e.g. propellant mass) for optimal transfers can provide advantages during preliminary design phases of trajectories where many options are typically screened and evaluated.

While the deployment of DNNs for trajectory design is promising, it has been shown over a large range of different application domains that the performance of DNNs depends heavily on the availability of suitable datasets for training. While most research literature is concerned with improving network architectures (number of layers, regularization, activation functions, etc.) or training procedures, in practice, it is often simply the lack of data that puts a limit on what any DNN architecture is capable of achieving.

The computation of a large-scale dataset of trajectories, including the optimal control profile and the value function is challenging, as a huge number of optimal control problems need to be solved. Since numerical instabilities and convergence issues are dominant, solution procedures are generally hard to automate and demanding in terms of computational resources.

The contribution of this paper is two-fold: First, we show how to create a large dataset of optimal trajectories without the need to solve more than one optimal control problem. This decreases the dataset generation time, in comparison to previous methods, by orders of magnitude. Second, we use the dataset to learn several control policies for an interplanetary mission to reach the orbit of Venus.

Building on the work of [9, 10, 12], we propose three different learning tasks, which we call *policy imitation*, *value function learning* and *value function gradients learning*. We are able to approximate, in the former case, the optimal policy (i.e. the optimal control profile) and, in the latter two cases, the value function (i.e. the solution to the corresponding Hamilton Jacobi Bellman equation). In the last case we find that the learned value function can also be satisfactorily used to compute the optimal policy.

This work is structured as follows: Section 2 provides the details of the nominal transfer to Venus designed as a mass-optimal trajectory by solving the two point boundary value problem resulting from Pontryagin's minimum principle. Based on this nominal trajectory, Section 3 describes the new method for the generation of a large dataset of different trajectories to the Venus orbit whose mass-optimality is still ensured by Pontryagin's minimum principle. Section 4 describes different strategies for the training of DNNs on such a dataset. The resulting controllers are evaluated experimentally in Section 5 and we conclude in Section 6.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '19 Companion, July 13–17, 2019, Prague, Czech Republic

© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6748-6/19/07...\$15.00

<https://doi.org/10.1145/3319619.3326834>

2 BACKGROUND

We consider, in the International Celestial Reference Frame (ICRF) heliocentric frame, the motion of a spacecraft of mass m whose velocity and position we indicate with $\mathbf{r}$ and $\mathbf{v}$. The spacecraft is equipped with an ion thruster having specific impulse I_{sp} and capable of delivering a maximum thrust c_1 regardless of the spacecraft distance from the Sun. We describe the spacecraft state via its mass and the modified equinoctial elements $\mathbf{x} = [p, f, g, h, k, L]^T$ as originally defined by Walker et al. [13].

The set of differential equations describing the spacecraft motion are then:

$$\begin{aligned} \dot{p} &= \sqrt{\frac{p}{\mu}} \frac{2p}{w} f_t \\ \dot{f} &= \frac{1}{m} \sqrt{\frac{p}{\mu}} \left\{ f_r \sin L + [(1+w) \cos L + f] \frac{f_t}{w} - (h \sin L - k \cos L) \frac{g \cdot f_n}{w} \right\} \\ \dot{g} &= \frac{1}{m} \sqrt{\frac{p}{\mu}} \left\{ -f_r \cos L + [(1+w) \sin L + g] \frac{f_t}{w} + (h \sin L - k \cos L) \frac{f \cdot f_n}{w} \right\} \\ \dot{h} &= \sqrt{\frac{p}{\mu}} \frac{s^2 f_n}{2mw} \cos L \\ \dot{k} &= \sqrt{\frac{p}{\mu}} \frac{s^2 f_n}{2mw} \sin L \\ \dot{L} &= \sqrt{\frac{p}{\mu}} \left\{ \mu \left(\frac{w}{p} \right)^2 + \frac{1}{w} (h \sin L - k \cos L) \frac{f_n}{m} \right\} \\ \dot{m} &= -\frac{\sqrt{f_r^2 + f_t^2 + f_n^2}}{I_{sp} g_0} \end{aligned} \quad (1)$$

where, $w = 1 + f \cos L + g \sin L$ and $s^2 = 1 + h^2 + k^2$ and f_r, f_t, f_n are the radial, tangential and normal components of the force generated by the spacecraft propulsion system. The gravity parameter is denoted with μ and the gravitational acceleration at sea level with g_0 .

It is useful to rewrite the equations above using a more concise notation. We thus introduce the matrices $\mathbf{B}$ and $\mathbf{D}$ defined as:

$$\begin{aligned} \sqrt{\frac{\mu}{p}} \mathbf{B}(\mathbf{x}) &= \begin{bmatrix} 0 & \frac{2p}{w} & 0 \\ \sin L & [(1+w) \cos L + f] \frac{1}{w} & -\frac{g}{w} (h \sin L - k \cos L) \\ -\cos L & [(1+w) \sin L + g] \frac{1}{w} & \frac{f}{w} (h \sin L - k \cos L) \\ 0 & 0 & \frac{1}{w} \frac{s^2}{2} \cos L \\ 0 & 0 & \frac{1}{w} \frac{s^2}{2} \sin L \\ 0 & 0 & \frac{1}{w} (h \sin L - k \cos L) \end{bmatrix} \\ &= \end{aligned} \quad (2)$$

and

$$\mathbf{D}(\mathbf{x}) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & \sqrt{\mu p} \left(\frac{w}{p} \right)^2 \end{bmatrix} \quad (3)$$

and rewrite the equations of motion in the form:

$$\begin{cases} \dot{\mathbf{x}} = \frac{c_1 u(t)}{m} \mathbf{B}(\mathbf{x}) \hat{\mathbf{i}}_\tau + \mathbf{D}(\mathbf{x}) \\ \dot{m} = -c_2 u(t) \end{cases} \quad (4)$$

where the spacecraft thrust is now indicated by $c_1 u \hat{\mathbf{i}}_\tau = [f_r, f_t, f_n]^T$ and bounded by $|u(t)| \leq 1$ and $|\hat{\mathbf{i}}_\tau(t)| = 1$. We will use the notation $\mathbf{u} \in \mathcal{U}$ to indicate the feasible control space. Furthermore, we define $c_2 = 1/(I_{sp} g_0)$. Note that in order to steer the spacecraft, at each instant we control the throttle magnitude $u(t) \in [0, 1]$ and the thrust direction $\hat{\mathbf{i}}_\tau$. We refer to these control variables also with the single symbol $\mathbf{u} = [u, \hat{\mathbf{i}}_\tau]$.

2.1 The Low-Thrust Problem

We consider here a free time orbital transfer problem, that is finding the controls $u(t)$ and $\hat{\mathbf{i}}_\tau(t)$ defined in $[0, t_f]$ and the transfer time t_f so that the functional:

$$J(u(t), t_f) = \int_0^{t_f} \{u - \epsilon \log[u(1-u)]\} dt \quad (5)$$

is minimized, and the spacecraft is steered from its initial mass m_0 and some initial point $\mathbf{x}_0$ to some final mass m_f and some final point $\mathbf{x}_f \in \mathcal{S}_f \subset \mathbb{R}^6$. The functional J , following the work of Bertrand and Epenoy [2], is parameterised by a continuation parameter $\epsilon \in [0, 1]$ which activates a logarithmic barrier smoothing the problem and ensuring that the constraint $u(t) \in [0, 1]$ is always satisfied. Clearly for $\epsilon \rightarrow 0$ the functional becomes $J = (m_0 - m_f)/c_2$, equivalent to minimizing the propellant mass.

2.2 Consequences of Pontryagin's Minimum Principle

Following the work of Pontryagin [8], we can infer the necessary conditions for an optimal solution of this problem by applying Pontryagin's minimum principle.¹

Let us introduce the co-states λ, λ_m as continuous functions defined in $[0, t_f]$ and define the Hamiltonian:

$$\begin{aligned} \mathcal{H}(\mathbf{x}, m, \lambda, \lambda_m, \mathbf{u}) &= \\ &= \frac{c_1 u}{m} \lambda^T \mathbf{B}(\mathbf{x}) \hat{\mathbf{i}}_\tau + \lambda_L \sqrt{\frac{\mu}{p^3}} w^2 - c_2 \lambda_m u + \{u - \epsilon \log[u(1-u)]\} \end{aligned} \quad (6)$$

and the system of equations:

$$\begin{cases} \dot{\mathbf{x}} = \frac{\partial \mathcal{H}}{\partial \lambda} = \frac{c_1 u(t)}{m} \mathbf{B} \hat{\mathbf{i}}_\tau(t) + \mathbf{D} \\ \dot{m} = \frac{\partial \mathcal{H}}{\partial \lambda_m} = -c_2 u(t) \\ \dot{\lambda} = -\frac{\partial \mathcal{H}}{\partial \mathbf{x}} \\ \dot{\lambda}_m = -\frac{\partial \mathcal{H}}{\partial m} \end{cases} \quad (7)$$

The explicit form of the various derivatives appearing in the equations above is reported in Appendix A.

Along an optimal trajectory, the Hamiltonian must be zero (free terminal time problem) and minimal w.r.t. the choices of u and $\hat{\mathbf{i}}_\tau$. For the optimal thrust direction $\hat{\mathbf{i}}_\tau^*$ it follows that

$$\hat{\mathbf{i}}_\tau = \hat{\mathbf{i}}_\tau^*(t) = -\frac{\mathbf{B}^T \lambda}{|\mathbf{B}^T \lambda|} \quad (8)$$

where the time dependence on the right hand side follows both from the co-states and from $\mathbf{B}$ and has not been indicated explicitly for brevity. While for the optimal throttle u^* , necessarily:

$$u(t) = u^*(t) = \frac{2\epsilon}{2\epsilon + SF(t) + \sqrt{4\epsilon^2 + SF(t)^2}} \quad (9)$$

where we introduce a switching function

$$SF(t) = 1 - \frac{c_1}{m} |\mathbf{B}^T \lambda| - c_2 \lambda_m. \quad (10)$$

¹Note that we have stated a minimization problem, hence the conditions are actually slightly different from the ones originally derived in Pontryagin's work.




Figure 1: Solution of the two point boundary value problem for the throttle magnitude $u(t)$ and decreasing values of ϵ .

2.2.1 The two point boundary value problem. The test case considered in this work is that of a transfer from any $\mathbf{x}_0, m_0$ to Venus' orbit (not a rendezvous). In this case the final value for the mass and the final value for the anomaly L are left free so that some transversality conditions apply ($\lambda_{L_f} = 0, \lambda_{m_f} = 0$). For any starting state $\mathbf{x}_0$ and m_0 , we need to choose the values for λ_0 and λ_{m_0} and the time t_f such that solving the initial value problem for Eq.(7) results in a final state that matches the arrival conditions at Venus orbit, the transversality conditions and the free-time condition on the Hamiltonian. Formally, we introduce the shooting function:

$$\begin{aligned} \phi(\lambda_0, \lambda_{m_0}, t_f) = \\ = [p_f - p_V, f_f - f_V, g_f - g_V, h_f - h_V, k_f - k_V, \lambda_{L_f}, \lambda_{m_f}, \mathcal{H}_f] \end{aligned} \quad (11)$$

where we indicate with a subscript V the modified equinoctial elements of Venus orbit and with the subscript f the final values of the modified equinoctial elements resulting from numerically integrating Eq.(7) for a time t_f and from the initial conditions $\mathbf{x}_0, m_0, \lambda_0, \lambda_{m_0}$. We have thus transformed our optimal control problem into the problem: $\phi(\lambda_0, \lambda_{m_0}, t_f) = 0$.

2.2.2 The nominal trajectory. We consider a spacecraft with mass $m_0 = 1500$ [kg], a nuclear electric propulsion system specified by $I_{sp} = 3800$ [sec] and $c_1 = 0.3$ [N]. We look for the optimal transfer from the Earth to Venus orbit with a launch on the 7th of May 2005. The planet ephemerides are computed using JPL low-precision ephemerides. We solve the optimal control problem by solving the problem $\phi(\lambda_0, \lambda_{m_0}, t_f) = 0$. Note that this is, essentially, a system of eight equations in eight unknowns and can be solved by root finding methods (e.g. Powell, Levenberg-Marquardt) as well as by SQP or interior point methods (e.g. SNOPT or IPOPT). As it is well known (see [5] for example), the convergence radius for this problem can get rather small, to the point that if we were to try to directly solve the mass optimal problem (i.e. $\epsilon = 0$) we




Figure 2: Nominal trajectory to Venus orbit. Thrust arcs are indicated in red.

would fail consistently as almost any initial guess on the co-state would not converge. However, solving the problem for $\epsilon = 0.1$ is reasonably simple as convergence is frequently achieved when starting with random co-states (we sampled them from a uniform distribution with a standard deviation of 10). Note that we use non dimensional units for the state so that the astronomical unit AU is used for length, the spacecraft initial mass for mass, and the rest is set as to get $\mu = 1$. Gradually decreasing ϵ from 0.1 down to 10^{-6} (as visualized in Figure 1), allows us to obtain the final mass optimal trajectory which is visualized in Figure 2. We refer to the final trajectory (with $\epsilon = 10^{-6}$) as the *nominal trajectory* in the following. The nominal trajectory reaches the orbit of Venus after $t_f^* = 1.376$ [years] and spends $m_p^* = 210.47$ [kg] of propellant.

2.3 Consequences of Bellman's Principle of Optimality

For $\epsilon = 0$, we can apply Bellman's principle of optimality [1] to the optimal control problem that we have stated in the previous section. We indicate with $v(\mathbf{x}, m)$ the value function, i.e. the optimal value of the functional defined by Eq.(5) for any spacecraft state $\mathbf{x}, m_0$. Since the value function is time-independent, the Hamilton Jacobi Bellman (HJB) equations can be written as:

$$0 = \min_{\mathbf{u} \in \mathcal{U}} (u + \nabla v \cdot \mathbf{f}) \quad (12)$$

$$\mathbf{u} = \arg \min_{\mathbf{u} \in \mathcal{U}} (u + \nabla v \cdot \mathbf{f}). \quad (13)$$

These equations hold in all points where $v(\mathbf{x}, m)$ is differentiable. We use $\mathbf{f}$ to denote the right hand side of Eq.(4) including the mass equation. Comparing Eq.(12) and Eq.(6), it follows that $\nabla v = [\lambda, \lambda_m]$ and thus $\mathcal{H} = u + \nabla v \cdot \mathbf{f}$. Thus, the co-states of Pontryagin's theory are the gradients of the value function introduced in Bellman's theory. This fact, albeit rarely exploited in interplanetary trajectory optimization research, provides a convenient basis for the design

of new learning procedures, which we introduce and evaluate later in Section 4.

3 BUILDING A DATASET OF OPTIMAL TRAJECTORIES

3.1 Reducing the Cost of Computation

The availability of large-scale datasets is essential to train DNNs. This section shows how to create a large number of optimal trajectories without the need to solve any particular optimal control problem except for the nominal trajectory. The brute-force approach for computing a dataset of optimal trajectories is to decide on a number of interesting initial states and then solve the corresponding optimal control problem for each (e.g. following all the steps outlined in Section 2). This approach scales extremely poorly because of the known computational difficulties associated with solving optimal control problems, both using direct and indirect methods.

Previous work (e.g. Sanchez and Izzo [9, 10]) deployed a continuation (homotopy) approach to reduce some of the complexity involved by eliminating the need to search for a new initial guess all the time: by perturbing the initial state of a nominal trajectory, the unperturbed co-state provides (most of the time) a good initial guess to solve the newly created two points boundary value problem. However, assuming that no convergence issues occur, it is still necessary to solve Eq.(11) for the new initial conditions considered, which incurs a significant computational cost.

In the following, we describe a more efficient way to obtain a similar result avoiding entirely convergence issues, guaranteeing optimality and reducing the computational costs down to the cost of a single numerical propagation of Eq.(7) for a new trajectory. The idea, surprisingly simple and applicable more generally than this Venus orbit acquisition task, is to perform a backward in time propagation of Eq.(7) starting from suitably perturbed final values of the state and co-states of the nominal trajectory. This perturbation needs to be chosen such that the transversality conditions and the condition on the Hamiltonian are still satisfied.

Formally, consider the nominal optimal trajectory and indicate with $\mathbf{x}_f, m_f, \lambda_f, \lambda_{m_f}$ the final values (i.e. at t_f^*) of the states and the co-states. Consider the new set of final co-state values:

$$\lambda'_f = \lambda_f + \delta\lambda \quad (14)$$

where the perturbation $\delta\lambda$ is in some ball $B_\rho \in \mathbb{R}^7$ of size ρ such that the transversality condition on the anomaly ($\delta\lambda_L = 0, \delta\lambda_m = 0$) is satisfied. Note that the trajectory that results from propagating backward in time Eq.(7) from the new final conditions $\mathbf{x}_f, m_f, \lambda'_f, \lambda_{m_f}$ is fulfilling all of Pontryagin's necessary conditions for optimality, except $\mathcal{H}_f = 0$. It is possible to find new final values $m'_f = m_f + \delta m$ and $L'_f = L_f + \delta L$ (e.g. with a simple root finding algorithm) so that the condition on the Hamiltonian is fulfilled. Thus the new trajectory resulting from propagating backward in time Eq.(7) from the conditions $\mathbf{x}'_f, m'_f, \lambda'_f, \lambda_{m'_f}$ is now fulfilling all of Pontryagin's necessary conditions for optimality. Such a trajectory will neither end where the nominal trajectory ends, nor will it start from where the nominal trajectory starts, but it is nevertheless optimal (with respect to Eq. 5) and represents a valid sample to




Figure 3: Overlapping optimal dataset trajectories. Thrust arcs are indicated in red.

learn from. This procedure reduces the cost of computing one more valid training sample to that of a backward in time integration of Eq.(7) plus the (negligible) cost of a single root finding call to solve $\mathcal{H}_f(m'_f, L'_f) = 0$.

As we only modify m_f and L , our final state is always constrained to be on the orbit of Venus, and since our perturbations for the final co-states are small the timing and direction of thrust maneuvers are close to the nominal trajectory during the backward integration. Together, these imply that it is unlikely, if not impossible, for the final state of the backward integration to go significantly beyond the orbit of Earth during an integration period of $t_f^* = 1.376$ [years]. This is confirmed when looking at the dataset in the next subsection.

3.2 The dataset

To build the particular dataset used in this work, we start from the nominal trajectory of Figure 1 and generate 500 000 perturbations of the arrival state $\mathbf{x}'_f, m'_f, \lambda'_f$ using $\rho = 0.1$ as a perturbation radius. Only perturbations for which the root finder for $\mathcal{H}_f(L'_f, m'_f) = 0$ converged are kept. These valid final states are then integrated backwards in time for $t_f = 1.376$ [years]. At 100 equi-spaced points in time, the state $\mathbf{x}', m', \lambda'$ and the optimal policy u^*, i_t^* are sampled and stored in the dataset. Figure 3 shows an overlapping plot of these trajectories.

During the dataset generation, 453 212 feasible perturbations were found. Thus the final dataset consists of 45 321 200 samples which took approx. 12 hours to generate using 20 threads and 25 000 perturbations per thread on an Intel Xeon E5-2650L v4 processor at 1.70 GHz. Thus, we were able to generate a new optimal trajectory approximately every 1.7 seconds. The bulk of the computational time is spent to integrate Eq.(7) numerically with a relative error tolerance of 10^{-13} and an absolute error tolerance of 10^{-13} using the LSODA integrator. In Figure 3 we visualize the trajectory dataset in its entirety. Note that all trajectories we will learn from are thus time free, optimal control trajectories.

4 EXPERIMENTS

Using the dataset constructed in the previous section, our goal is to evaluate if it is possible to learn a control policy and/or the value function by a deep neural network.

4.1 Ideas Behind Our Experiments

Since this goal can take many forms, we are particularly interested in the following learning tasks:

- (1) *Policy imitation*: In the first experiment we apply the imitation learning pipeline described in [7, 9, 11] to our training data. The result is a network $\mathcal{N}_u(\mathbf{x}, m)$ predicting the optimal u and i_τ directly.
- (2) *Value function learning*: In the second experiment we train a network $\mathcal{N}_v(\mathbf{x}, m)$ to predict the value function $v(\mathbf{x}, m) = m_p/c_2$ (i.e. the optimal propellant mass required to get to Venus' orbit). The loss function used during training and evaluated on a batch of size b is $\ell = \sum_{i=1}^b |v(\mathbf{x}_i, m_i) - \mathcal{N}_v(\mathbf{x}_i, m_i)|^2$. Once the network is trained, the Hamilton Jacobi Bellman equation Eq.(13) is used to derive the policy $\mathbf{u}(\mathbf{x}, m)$. Besides providing a way to compute $\mathbf{u}(\mathbf{x}, m)$, this network can also be used for predicting the propellant consumption in preliminary design phases.
- (3) *Value function gradients learning*: In the third experiment we train a network $\mathcal{N}_{\nabla v}(\mathbf{x}, m)$ to predict the value function $v(\mathbf{x}, m)$ and enforce the constraint that $[\partial_{\mathbf{x}} \mathcal{N}_{\nabla v}, \partial_m \mathcal{N}_{\nabla v}] = \nabla v = [\lambda, \lambda_m]$. The idea being that gains are to be achieved, with respect to the previous case, as Eq.(13) contains the value function gradient which is here directly learned. The loss function used during training and evaluated on a batch of size b is thus:

$$\ell = \sum_{i=1}^b \left\{ |v(\mathbf{x}_i, m_i) - \mathcal{N}_v(\mathbf{x}_i, m_i)|^2 + \left| \lambda_i - \frac{\partial \mathcal{N}_{\nabla v}}{\partial \mathbf{x}}(\mathbf{x}_i, m_i) \right|^2 + \left| \lambda_{m_i} - \frac{\partial \mathcal{N}_{\nabla v}}{\partial m}(\mathbf{x}_i, m_i) \right|^2 \right\} \quad (15)$$

4.2 Network Architecture and Training

The network architectures used for the three experiments are mostly similar except for the depth and the output layers. In all cases, the networks take as input the state $[\mathbf{x}, m]$ and consist of n hidden layers with softplus activations between each layer (following the results in [12]).

The policy imitation (experiment 1) network consists of $n = 4$ hidden layers with 100 units. The output layer has a sigmoid activation and 4 units (throttle and thrust vector).

The value function networks (experiment 2 and 3) consist of $n = 9$ hidden layers with 100 units. The output layer has a linear activation and 1 unit. The reason for a linear activation is to avoid normalizing the output, which would also effect the computed gradients used in the loss function.

The network output can then be differentiated with respect to its inputs to estimate the value function gradients necessary to derive the control policy from Eq.(13). In the case of the value function network (experiment 2), these gradients are only used after training to compute the policy. For the value function gradients

network (experiment 3), the mean square error of the estimated value function gradients with respect to the co-states is added to the loss-function of the network. Thus the training for the value function gradient network is minimizing both the error of the value-function and its corresponding gradients.

In terms of the training set, the policy network was trained on a normalized dataset where the normalization consists of a standard scaling (i.e. removing the mean and scaling to unit variance). The targets u and i_τ have been scaled to $[0, 1]$ in order to match the output of the last sigmoid activations. During preliminary experiments, we found that for the value function network and the value function gradient network, a better performance was achieved by omitting any normalization of the data.

We split the data into a train, validation and test set by a ratio of 0.8/0.1/0.1. All networks were trained on mini-batches of size $b = 8192$ using the Amsgrad optimizer with an initial learning rate of 10^{-5} for 300 epochs. For each experiment, we select the network with the lowest mean square error on the validation set for all further computations presented in this work.

5 RESULTS

We evaluate the quality of the trained networks by two different methods. First, we test the network performance on the test set, containing optimal state control pairs from trajectories never seen during training. Second, and more importantly, we test how well the network performs when used to control the spacecraft dynamics, as computed by integrating the equations of motion in Eq.(4) forward in time starting from the initial state of the nominal trajectory (which was excluded from the training, validation and test sets) using the policies resulting from the network for $\mathbf{u}$.

5.1 Evaluation of Control Predictions on the Test Set

For each sample of the test set we compute the optimal policy prediction from the networks (u and i_τ) and thus the mean and standard deviation of the absolute error. In Table 1, the values for each of the control components are shown. The policy network displays the best performance in terms of the u , i_r and i_t components. On the other hand, the value function gradients network outperforms the policy network in terms of the i_n component with regards to the spread in errors across the test set despite an overall higher mean absolute error. The value function network without gradient learning, while approximating the value function reasonably well, fails to reconstruct the optimal policy. This suggests that learning the gradients in addition to the value function is necessary for the reconstruction of an optimal control profile.

A more detailed view of the error distribution is shown in Figure 6 for the throttle and Figure 7 for the thrust direction, confirming that the value function gradients network mis-throttles (i.e. throttling when it should not and not throttling when it should) on average more frequently than the policy network. It is interesting to note, again, that in terms of the distribution of errors in the thrust direction, the value function gradients network achieves a narrower distribution of errors centered around 0 for the u_n component of the thrust.




Figure 4: Absolute error for predicting the optimal thrust direction $\hat{i}_\tau$ of the nominal trajectory (log-scale).

controller	u	i_r	i_t	i_n
$\mathcal{N}_u$	$4.0633 \cdot 10^{-3} \pm 1.6114 \cdot 10^{-3}$	$9.4601 \cdot 10^{-4} \pm 2.0558 \cdot 10^{-6}$	$5.4004 \cdot 10^{-4} \pm 8.7714 \cdot 10^{-7}$	$2.1191 \cdot 10^{-3} \pm 5.0105 \cdot 10^{-5}$
$\mathcal{N}_v$	$3.8448 \cdot 10^{-1} \pm 2.3425 \cdot 10^{-1}$	$1.7492 \cdot 10^{-2} \pm 1.7631 \cdot 10^{-4}$	$3.6827 \cdot 10^{-2} \pm 7.3182 \cdot 10^{-4}$	$1.5677 \cdot 10^{-1} \pm 5.6198 \cdot 10^{-3}$
$\mathcal{N}_{\nabla v}$	$6.7513 \cdot 10^{-2} \pm 6.0727 \cdot 10^{-2}$	$2.9607 \cdot 10^{-3} \pm 6.2466 \cdot 10^{-6}$	$8.4705 \cdot 10^{-4} \pm 7.3286 \cdot 10^{-7}$	$2.6342 \cdot 10^{-3} \pm 4.6612 \cdot 10^{-6}$

Table 1: Mean and Standard Deviation of the Absolute Error of network predictions for the three network types.




Figure 5: Absolute error for predicting the optimal throttle u^* of the nominal trajectory (log-scale).

controller	m_f [kg]	Δm [kg]	$m_f + \Delta m$ [kg]
optimal	210.4742	–	210.4742 (+ 0.0000)
$\mathcal{N}_u$	210.3197	0.1858	210.5055 (+ 0.0313)
$\mathcal{N}_v$	151.1732	129.7239	280.8971 (+ 80.7503)
$\mathcal{N}_{\nabla v}$	205.4894	5.6123	211.1017 (+ 0.6275)

Table 2: Propellant mass spent to reach Venus orbit from the nominal initial condition. This mass consists of the propellant mass m_f spent using the considered controller to steer the spacecraft for a time t_f^* and the mass (Δm) needed for an additional (optimal) corrective maneuver able to match Venus orbit exactly.

5.2 Closing the Gap to Venus Orbit

Next, we compare how well the trained neural networks are able to reproduce the nominal trajectory (which was not part of the dataset

used for training) during a forward in time numerical integration of the equation of motion stated in Eq.(4). Starting from the nominal initial conditions, we thus predict the control actions using each network and integrate forward in time for a period of $t_f^* = 1.376$ [years] (the time needed by the nominal trajectory to reach the orbit of Venus). Figure 8 shows the distance from the orbit of Venus obtained in the three cases. We observe that $\mathcal{N}_u$ and $\mathcal{N}_{\nabla v}$ come closest to Venus orbit while $\mathcal{N}_v$ is comparatively far away.

For a fair comparison between all three strategies with respect to the nominal trajectory, we correct all orbits (in a mass-optimal way) to match the exact orbit of Venus and add the corresponding propellant mass (Δm) needed for this corrective maneuver to the propellant mass (m_f) spent by the controller. Consequently, we solve the optimal control problem for reaching Venus from each of the final orbits reached by the networks. The results are reported in Table 2. We note that the additional mass needed by $\mathcal{N}_u$ and $\mathcal{N}_{\nabla v}$ to close the gap to Venus orbit is below 1 [kg].

6 CONCLUSIONS AND DISCUSSION

The results presented show that our strategy for the generation of optimal trajectories is able to provide datasets of suitable size for supervised training of deep neural networks. The major benefit of this technique is the low computational cost and its reliability, as the usage of complex solvers for optimal control problems can be avoided in favor of a much simpler numerical integration procedure. For the test case of an interplanetary transfer to Venus orbit, we find that the small disturbances ρ , which we applied to the final co-states, allowed us to generate enough variety within the dataset to learn reliable control policies by imitation learning. While the policy network is structurally simple and comparatively small in size, the achieved precision is promising and provides a first step towards the goal of achieving on-board optimal guidance and control by a such a neural network.

Additionally to the prediction of the optimal control policy, the value networks are able to estimate the value function up to high




Figure 6: Distribution across the test set of the absolute error introduced when predicting the throttle using the three trained networks.




Figure 7: Distribution of thrust errors for the radial (top), tangential (middle) and normal (bottom) thrust components with a logarithmic y-scale.

precision and thus provide an approximation to the mass-budget needed for optimal transfers from different starting points. If the error of the gradients of the value function is incorporated during training, a control profile equally good as the imitation learning can be obtained.

While these preliminary results are encouraging, we plan to further study the potential of this approach for designing more ambitious transfers. In particular, the perturbation parameter of the final co-states ρ is the key for this approach as it significantly

impacts the resulting trajectory dataset, its needed size and, subsequently, the quality and usability of the trained networks. Limiting ρ to a comparatively small value allowed the trajectories in our dataset to resemble the nominal trajectory closely. However, the investigation of larger values of ρ is of great interest and remains for future work. The dataset size and its relation to ρ is also of great interest, since in this work they were both selected in a somewhat arbitrary fashion.




Figure 8: Distance from Venus during trajectory for value function (blue), value function gradients (orange) and policy networks (green).

Sampling of the trajectories at equally spaced points in time resulted in a bias of the dataset, as it leads to introduce more sample points at larger distances from the target orbit. Although this did not pose a problem during our experiments, a different sampling strategy could easily reduce the bias and would be most likely beneficial for training on challenging transfers. Furthermore, foregoing the computation of a nominal trajectory seems entirely possible and worth investigating, for example by generating the final co-states at random (and not perturbing some initial values).

Lastly we note that for the value function gradient learning as presented here, the co-states need to be known (i.e. part of the dataset). While our approach is able to deliver the co-states for each trajectory, a direct method for solving the optimal control problem would not. In such cases we hypothesize that adding Eq.(12) to the loss function might improve the quality of the controller even if information on the co-states is not available.

A APPENDIX

This appendix contains the explicit forms of all the derivatives necessary to write Eq.(7) explicitly.

The $\dot{\lambda}_p$ equation

We get:

$$\begin{aligned}\dot{\lambda}_p &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial p} \mathbf{i}_r - w^2 \lambda_L \frac{\partial}{\partial p} \sqrt{\frac{\mu}{p^3}} \\ &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial p} \mathbf{i}_r + \frac{3}{2} w^2 \lambda_L \sqrt{\frac{\mu}{p^5}}\end{aligned}\quad (16)$$

$$\begin{aligned}2\sqrt{\mu p} \frac{\partial \mathbf{B}(\mathbf{x})}{\partial p} &= \\ &= \begin{bmatrix} 0 & \frac{6p}{w} & 0 \\ \sin L & [(1+w)\cos L + f] \frac{1}{w} & -\frac{g}{w} (h \sin L - k \cos L) \\ -\cos L & [(1+w)\sin L + g] \frac{1}{w} & \frac{f}{w} (h \sin L - k \cos L) \\ 0 & 0 & \frac{1}{w} \frac{s^2}{2} \cos L \\ 0 & 0 & \frac{1}{w} \frac{s^2}{2} \sin L \\ 0 & 0 & \frac{1}{w} (h \sin L - k \cos L) \end{bmatrix}\end{aligned}\quad (17)$$

The $\dot{\lambda}_f$ equation

We get:

$$\begin{aligned}\dot{\lambda}_f &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial f} \mathbf{i}_r - 2\lambda_L w \sqrt{\frac{\mu}{p^3}} \frac{\partial w}{\partial f} \\ &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial f} \mathbf{i}_r - 2\lambda_L w \sqrt{\frac{\mu}{p^3}} \cos L\end{aligned}\quad (18)$$

$$\begin{aligned}w^2 \sqrt{\frac{\mu}{p}} \frac{\partial \mathbf{B}(\mathbf{x})}{\partial f} &= \\ &= \begin{bmatrix} 0 & -2p \cos L & 0 \\ 0 & w - (\cos L + f) \cos L & g \cos L (h \sin L - k \cos L) \\ 0 & -(\sin L + g) \cos L & (w - f \cos L) (h \sin L - k \cos L) \\ 0 & 0 & -\frac{s^2}{2} \cos^2 L \\ 0 & 0 & -\frac{s^2}{2} \sin L \cos L \\ 0 & 0 & -(h \sin L - k \cos L) \cos L \end{bmatrix}\end{aligned}\quad (19)$$

The $\dot{\lambda}_g$ equation

We get:

$$\begin{aligned}\dot{\lambda}_g &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial g} \mathbf{i}_r - 2\lambda_L w \sqrt{\frac{\mu}{p^3}} \frac{\partial w}{\partial g} \\ &= -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial g} \mathbf{i}_r - 2\lambda_L w \sqrt{\frac{\mu}{p^3}} \sin L\end{aligned}\quad (20)$$

$$\begin{aligned}w^2 \sqrt{\frac{\mu}{p}} \frac{\partial \mathbf{B}(\mathbf{x})}{\partial g} &= \\ &= \begin{bmatrix} 0 & -2p \sin L & 0 \\ 0 & -(\cos L + f) \sin L & -(w - g \sin L) (h \sin L - k \cos L) \\ 0 & w - (\sin L + g) \sin L & -f \sin L (h \sin L - k \cos L) \\ 0 & 0 & -\frac{s^2}{2} \cos L \sin L \\ 0 & 0 & -\frac{s^2}{2} \sin^2 L \\ 0 & 0 & -(h \sin L - k \cos L) \sin L \end{bmatrix}\end{aligned}\quad (21)$$

The $\dot{\lambda}_h$ equation

We get:

$$\dot{\lambda}_h = -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial h} \mathbf{i}_r\quad (22)$$

$$\frac{\partial \mathbf{B}(\mathbf{x})}{\partial h} = \sqrt{\frac{p}{\mu}} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -\frac{g}{w} \sin L \\ 0 & 0 & \frac{f}{w} \sin L \\ 0 & 0 & \frac{f}{w} \cos L \\ 0 & 0 & \frac{h}{w} \sin L \\ 0 & 0 & \frac{h}{w} \cos L \end{bmatrix}$$

The $\dot{\lambda}_k$ equation

We get:

$$\dot{\lambda}_k = -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial k} \mathbf{i}_r\quad (23)$$

$$\frac{\partial \mathbf{B}(\mathbf{x})}{\partial k} = \sqrt{\frac{p}{\mu}} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & \frac{g}{w} \cos L \\ 0 & 0 & -\frac{f}{w} \cos L \\ 0 & 0 & \frac{k}{w} \cos L \\ 0 & 0 & \frac{k}{w} \sin L \\ 0 & 0 & -\frac{1}{w} \cos L \end{bmatrix}$$

The $\dot{\lambda}_L$ equation

We get:

$$\dot{\lambda}_L = -\frac{c_1 u}{m} \lambda^T \frac{\partial \mathbf{B}(\mathbf{x})}{\partial L} \mathbf{i}_r - 2w \sqrt{\frac{\mu}{p^3}} \lambda_L w_L\quad (24)$$

where,

$$\mathbf{w}^2 \sqrt{\frac{\mu}{p}} \frac{\partial \mathbf{B}(\mathbf{x})}{\partial L} = \begin{bmatrix} 0 & -2p(g \cos L - f \sin L) & 0 \\ \mathbf{w}^2 \cos L & -(1 + \mathbf{w})\mathbf{w} \sin L - \mathbf{w}_L(\cos L + f) & ((\mathbf{w}h + \mathbf{w}_L k) \cos L + (\mathbf{w}k - \mathbf{w}_L h) \sin L)g \\ \mathbf{w}^2 \sin L & (1 + \mathbf{w})\mathbf{w} \cos L - \mathbf{w}_L(\sin L + g) & ((\mathbf{w}h + \mathbf{w}_L k) \cos L + (\mathbf{w}k - \mathbf{w}_L h) \sin L)f \\ 0 & 0 & -\frac{s^2}{2}(\mathbf{w} \sin L + \mathbf{w}_L \cos L) \\ 0 & 0 & \frac{s^2}{2}(\mathbf{w} \cos L - \mathbf{w}_L \sin L) \\ 0 & 0 & (\mathbf{w}h + \mathbf{w}_L k) \cos L + (\mathbf{w}k - \mathbf{w}_L h) \sin L \end{bmatrix} \quad (25)$$

where,

$$\mathbf{w}_L = \frac{\partial \mathbf{w}}{\partial L} = g \cos L - f \sin L \quad (26)$$

The $\dot{\lambda}_m$ equation

We get:

$$\dot{\lambda}_m = -\frac{c_1 u}{m^2} |\lambda^T \mathbf{B}(\mathbf{x})| \quad (27)$$

REFERENCES

- [1] Richard Bellman. 1966. Dynamic programming. *Science* 153, 3731 (1966), 34–37.
- [2] Régis Bertrand and Richard Epenoy. 2002. New smoothing techniques for solving bang–bang optimal control problems—numerical results and statistical interpretation. *Optimal Control Applications and Methods* 23, 4 (2002), 171–197.
- [3] Lin Cheng, Zhenbo Wang, Yu Song, and Fanghua Jiang. 2019. Real-Time Optimal Control for Irregular Asteroid Landings Using Deep Neural Networks. *arXiv preprint arXiv:1901.02210* (2019).
- [4] Roberto Furfaro, Ilaria Bloise, Marcello Orlandelli, and Pierluigi Di Lizia. 2018. Deep Learning for Autonomous Lunar Landing. In *2018 AAS/AIAA Astrodynamics Specialist Conference*. 1–22.
- [5] Thomas Haberkorn, Pierre Martinon, and Joseph Gergaud. 2004. Low thrust minimum-fuel orbital transfer: a homotopic approach. *Journal of Guidance, Control, and Dynamics* 27, 6 (2004), 1046–1060.
- [6] Dario Izzo, Marcus Märtens, and Binfeng Pan. 2018. A Survey on Artificial Intelligence Trends in Spacecraft Guidance Dynamics and Control. *arXiv preprint arXiv:1812.02948* (2018).
- [7] Dario Izzo, Christopher Sprague, and Dharmesh Tailor. 2018. Machine learning and evolutionary techniques in interplanetary trajectory design. *arXiv preprint arXiv:1802.00180* (2018).
- [8] L.S. Pontryagin, V.G. Boltyanskii, R.V. Gamkrelidze, and E.F. Mishchenko. 1962. *The Mathematical theory of optimal processes*. Wiley & Sons.
- [9] Carlos Sánchez-Sánchez and Dario Izzo. 2018. Real-time optimal control via Deep Neural Networks: study on landing problems. *Journal of Guidance, Control, and Dynamics* 41, 5 (2018), 1122–1135.
- [10] Carlos Sánchez-Sánchez, Dario Izzo, and Daniel Hennes. 2016. Learning the optimal state-feedback using deep networks. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 1–8.
- [11] Christopher Iliffe Sprague, Dario Izzo, and Petter Ögren. 2019. Learning a Family of Optimal State Feedback Controllers. *arXiv preprint arXiv:1902.10139* (2019).
- [12] Dharmesh Tailor and Dario Izzo. 2019. Learning the optimal state-feedback via supervised imitation learning. *arXiv preprint arXiv:1901.02369* (2019).
- [13] M. J. H. Walker, B. Ireland, and J. Owens. 1985. A set of modified equinoctial orbit elements. *Celestial Mechanics* 36 (Aug. 1985), 409–419. <https://doi.org/10.1007/BF01227493>
- [14] Yue-he Zhu and Ya-zhong Luo. 2019. Fast Approximation of Optimal Perturbed Many-Revolution Multiple-Impulse Transfers via Deep Neural Networks. *arXiv preprint arXiv:1902.03741* (2019).