73th International Astronautical Congress, Paris, France.

IAC-22,C1,8,2,x70686

Preliminary Analysis of Small-Body Gravity Field Estimation using Physics-Informed Neural Networks and Kalman Filters

John R. Martin [a,*] and Hanspeter Schaub[b]

[a] *National Science Foundation Graduate Research Fellow, Aerospace Engineering Sciences, University of Colorado, Boulder, 431 UCB, Colorado Center for Astrodynamics Research, Boulder, CO, 80309.*, john.r.martin@colorado.edu

[b] *Professor and Department Chair, Schaden Leadership Chair, Ann and H.J. Smead Department of Aerospace Engineering Sciences, University of Colorado, Boulder, 431 UCB, Colorado Center for Astrodynamics Research, Boulder, CO, 80309. AAS Fellow, AIAA Fellow.*, hanspeter.schaub@colorado.edu

* Corresponding Author

Abstract

Gravity field estimation is a problem of increasing importance for small-body exploration. Specifically, the gravity fields of celestial bodies like asteroids and comets are often poorly characterized upon first approach, and engineers and scientists require lengthy mission phases dedicated to the construction of a more accurate gravity field model before attempting increasingly ambitious orbits or manuevers. Traditionally these gravity characterization periods rely on advanced orbit determination pipelines which are used to estimate the spacecraft state as well as low degree and order coefficients of a spherical harmonic gravity model. This research proposes an alternative gravity estimation framework which leverages a powerful and new gravity model: the Physics-Informed Neural Network (PINN) gravity model. PINN gravity models alleviate many of the inconveniences of past gravity representations such as divergence at the surface of the body (spherical harmonics), assumptions for the internal density profile of the body (polyhedral), and computational inefficiency (many). Moreover, the PINN gravity model has demonstrated greater sample efficiency than spherical harmonics, generating more accurate gravity representations from fewer and noisier measurements. Such characteristics make the PINN gravity model a desirable representation to be used online in small-body contexts, however no literature exists which discusses how this novel representation can be embedded into traditional filtering frameworks flown today. This work aims to fill this implementation hole, specifically exploring how the PINN gravity model can be used in conjunction with or embedded directly within Kalman filters. This work leverages simulated spacecraft data to train a PINN in an online manner compatible with current filtering techniques. This research also investigates various design choices including how frequently the PINN gravity model should be updated to ensure safe and productive learning, as well as highlights advantages and tradeoffs of transitioning to this new representation.

1. Introduction

In the past two decades, small-body exploration has blossomed into a major research focus for interplanetary exploration. Missions like Hayabusa2, Psyche, DAWN, OSIRIS-REx, Janus, and others have demonstrated this priority and the corresponding need enhanced tools to enable the rapid and safe exploration of these interesting targets [1–5]. For each past and future mission, it is imperative that the scientists and engineers prioritize the determination of a reliable gravity field model. These gravity models form the basis from which planetary scientists can build an intuition about the body's density distribution and surface properties, and by which dynamicists and mission designers can build trajectories that leverage the natural dynamics of the system [6–8].

Traditionally these gravity models are constructed by first assuming the body in question can be modeled as a point-mass, i.e. the body in question is infinitely small, perfectly spherical, and homogenous in density. These assumptions are often adequate for placing a spacecraft in an initial, high-altitude orbit around the body, but they quickly become problematic as more nuanced mission operations begin. Specifically, as spacecraft enter lower-altitude orbits, each of the assumptions begin to break down. The body is often not spherical – consider the asteroids Eros or Itokawa, both of which have highly irregular geometries – and their densities are not homogenous [9]. Taken together, these irregular shapes and inhomogenous densities produce non-uniform gravity fields which, in turn, yield highly non-keplerian motion.

Naturally it is important to capture these gravitational perturbations, particularly before the spacecraft enters lower orbits or attempts a landing. To accomplish this, dynamicists often turn to other, alternative gravity models. For ground-based simulation, dynamicists can use the polyhedral gravity model[10]. If the spacecraft is able to

capture many images of the body, teams on the ground can generate a shape model of the body. Using this shape model — comprised of many vertices and facets — dynamicists can compute the exact gravitational potentials and accelerations of that shape assuming the body has constant density. This approach provides a considerably more accurate gravity model than the prior point mass model, but it with two caveats. First, it is likely that the body violates the constant density assumption. While the polyhedral gravity model can accommodate inhomogeneous density profiles, these profiles are notoriously challenging to estimate uniquely [11]. Failing to account for density variations can, inturn, cause erroneous spacecraft dynamics which can prove to be dangerous when trajectories are propagated for extended periods of time. The second, and arguably larger, disadvantage of the polyhedral gravity model is its computational requirements[12]. High-fidelity shape models of asteroids can contain hundreds of thousands of vertices and facets which must be looped over at each propagation timestep. This can make it extremely expensive to compute accelerations both in ground based simulations and on-board.

The alternative to the polyhedral gravity model is the popular spherical harmonics gravity model[13] (or its close cousin the ellipsoidal harmonics gravity model[14]). These gravity models provide slightly more forgiving assumptions about the body in question — namely no assumptions about the density of the body — by representing the gravity field as the superposition of harmonic basis functions, or the three-dimensional analogs to a Fourier series. These harmonics can provide a more representative model of the true gravity field than a point mass model, and they are most commonly expanded to relatively low degree and order. These truncated low-order models make spherical harmonics considerably easier to include within an orbit determination pipeline where the harmonic coefficients can be directly estimated over the mission lifetime.

Despite this, these harmonic gravity models are not without their own disadvantages. For one, these harmonics models rely on the assumption that the spacecraft will remain in orbit outside of the bounding sphere or ellipsoid. For missions that seek to land on the surface, or merely attempt closer flybys of the object, this assumption can be operationally limiting. Moreover, these harmonic models are extremely inefficient at capturing discontinuity. Large gravitationally perturbing features like craters, boulders, mountain ranges, etc. can require hundreds-of-thousands of harmonics superimposed together before they are represented accurately [12]. Moreover, the harmonic coefficients can be difficult to regress, requiring strict sampling requirements for observability.

Recently, the Physics-Informed Neural Network gravity model (PINN-GM) was proposed which attempts to bypass each of the limitations of past gravity models[12]. By learning, rather than prescribing, basis functions, the PINN-GM is able to produce high-fidelity gravity models without making any assumptions or imposing operational limits on the mission. While the PINN-GM offers a compelling solution to the gravity modeling problem offline, little research has been conducted to investigate how the PINN-GM can be trained online or included within an orbit determination pipeline. Early results demonstrate that the PINN-GM can produce high-fidelity models of the asteroids Eros and Bennu using relatively sparse and noisy measurements [15]; however the literature fails to demonstrate how the acceleration training data is generated in-situ. This paper attempts to fill these holes, focusing specifically on how the orbit determination pipeline can be constructed to leverage a PINN-GM trained in-situ over a mission gravity characterization phase.

2. Methods

2.1 Spherical Harmonics Gravity Model

Spherical harmonics are a popular way to represent the gravitational potential of an arbitrary body through:

$$U_l(r) = \frac{\mu}{r}\sum_{l=0}^{l}\sum_{m=0}^{l}\left(\frac{R}{r}\right)^l P_{l,m}[\sin(\phi)] \left(C'_{l,m}\cos(m\lambda) + S'_{l,m}\sin(m\lambda)\right) \tag{1}$$

where r is the field point, μ is the gravitational parameter of the body, R is the circumscribing radius of the celestial body, l is the degree of the model, m is the order of the model, $C'_{l,m}$ and $S'_{l,m}$ are the regressed spherical harmonic coefficients, λ is the longitude, ϕ is the latitude, and $P_{l,m}$ are the associated Legendre polynomials [13].

When expanded to low-degree and order, this gravity model is particularly well-suited to capture gravitational perturbations like planetary oblateness and obliquity. This ability make it an effective basis set for gravity field estimation to first-order — particularly for large, near-spherical bodies. This representation, however, does come with a variety of pitfalls.

First, the spherical harmonic gravity model assumes that spacecraft motion will be conducted outside of a bounding sphere with Brillouin radius R. If a spacecraft enters within the bounding sphere, the representation is prone to diverging numerics due to the $\left(\frac{R}{r}\right)^l$ term in the expansion. This makes the model unreliable when attempting to use it in small-body landing or touch-and-go operations where high-fidelity models are required (high values of l) and the field points may be at a considerably lower altitude then the bounding Brillouin sphere ($r << R$).

In addition to the numerically diverging representation, the spherical harmonic model is also inefficient at capturing discontinuity. Consider the Earth's gravitationally

perturbing features. Listed in order of significance, the Earth's field deviates from the point mass assumption due to things like planetary oblateness, mountain ranges, tectonic plate boundaries, and regional hotspots. Of these four perturbing features, three are highly discontinuous in nature. Harmonic basis are very poorly suited at capturing such discontinuity (as best demonstrated through the infamous Gibbs phenomenon[16]). Infinitely many harmonics must be superimposed together before spherical harmonics can represent a discontinuous or localized feature. This poses a problem both for large planetary objects and small-bodies, making spherical harmonics (or ellipsoidal harmonics) an inefficient basis set for capturing high-order gravitational perturbations.

Finally, harmonic gravity models also suffer from the drawback of sample inefficiency. In order to regress high-degree harmonics, the field in question must sampled with a homogenous distribution of positions to avoid aliasing other harmonics [17]. These sampling requirements can prove difficult for trajectory and mission designers for which it may be fuel inefficient to manuever the spacecraft to these carefully distributed points.

Despite this, spherical harmonics remain a popular gravity model to use in the gravity estimation problem. Part of the reason for this is that the model is analytically differentiable. When collecting measurements of the spacecraft in an orbit determination pipeline, engineers can differentiate the accelerations produced by low-degree harmonics with respect to the spacecraft state to then be used in estimating the low-degree spherical harmonic coefficients. This analytic differentiability is also necessary when propagating state covariances.

2.2 Physics-Informed Neural Network Gravity Model

In 2022, the second generation of the Physics-Informed Neural Network Gravity Model (PINN-GM-II) was introduced [15]. This gravity model bypasses many of the failings of the prior analytic models by allowing a neural network to learn an efficient basis set that is naturally amenable to the gravity field the dynamicist is attempting to model. The training of the neural network simultaneously prioritizes predicting accurate predictions, while also enforcing that the learned model satisfies relevant differential equations such as Laplace's equation ($\nabla^2 U = 0$), relationships between potentials and accelerations ($\mathbf{a} = -\nabla U$), and conservative vector field properties ($\nabla \times \nabla U = 0$). This augmenting of a traditional cost function with differential constraints is what makes the network "Physics-Informed" [18].

For the second generation PINN-GM the cost function is defined as

$$J_{\text{ALC}}(\Theta) = \frac{1}{N_f}\sum_{i=1}^{N_f}\left|\boldsymbol{a}_i + \overset{\text{AD}}{\nabla} f(\boldsymbol{x}_i|\Theta)\right|^2 + \left|\overset{\text{AD}}{\nabla^2} f(\boldsymbol{x}_i|\Theta)\right|^2 + \left|\overset{\text{AD}}{\nabla} \times \overset{\text{AD}}{\nabla} f(\boldsymbol{x}_i|\Theta)\right|^2 \tag{2}$$

where $\overset{\text{AD}}{\nabla}$ refers to the gradient taken using automatic differentiation, and $f(\boldsymbol{x}_i|\Theta)$ is the gravitational potential as learned by the neural network.

2.2.1 Automatic Differentiation

Automatic differentiation is a way to compute the exact derivative of any algorithm [19]. Because all algorithms are comprised of sequences of elementary functions (+, -, /, *, exp, log, sin, cos, etc.) — to which there exist exact derivatives — a generalized form of chain rule can be applied to compute the exact derivative of the output of the algorithm with respect to the input (or any other intermediate value within said algorithm). Automatic differentiation is most famously applied in stochastic gradient descent algorithms — the same algorithms used to train neural networks [20]. Physics-informed neural networks simply take the same tool to compute the derivatives of the network to be used in the differential constraints of the cost function.

2.2.2 Advantages of the PINN-GM

A number of advantages are acquired by approaching the gravity field modeling problem with a PINN-GM. Foremost, the PINN-GM is extremely sample efficient. Precise and carefully distributed collections of data are not required to regress an accurate model. Instead, the PINN-GM can be trained with arbitrarily distributed data; iteratively improving its model based on where new information is gathered. In addition, the physics-informed constraints act as a form of regularization, making the PINN-GM robust to noise in the training data set. The PINN-GM is also efficient to evaluate — taking orders-of-magnitude less time to produce acceleration estimates than equivalent accuracy spherical harmonic or polyhedral gravity models [12]. Lastly, the PINN-GM is also exactly differentiable. This is a key advantage of the PINN-GM for orbit determination. Just like spherical harmonics, the Jacobian of the accelerations produced by the gravity model can be evaluated exactly for used in propagation of the STM and covariance updates.

2.3 Kalman Filter

Kalman filters are online algorithms used to estimate relevant spacecraft state and environmental parameters given uncertain measurements [21]. The general framework under which the state and parameters are estimated is as follows:

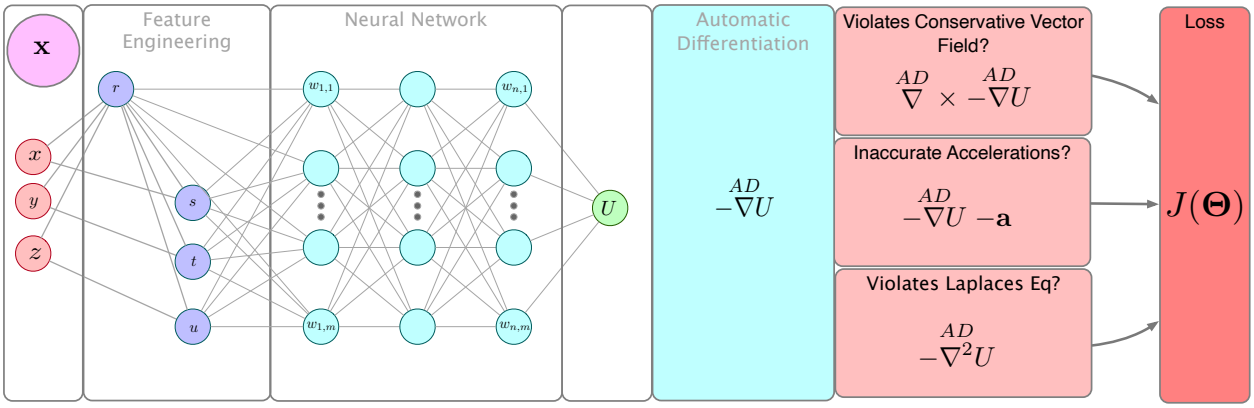

Fig. 1: Second generation of the PINN-GM [12]

1. Initialize the filter with an initial reference state $\boldsymbol{x}_0$, the initial state deviation / error $\Delta\boldsymbol{x}_0$, the initial state covariance P_0, process noise covariance matrix Q, measurement noise covariance matrix R.

2. Obtain a measurement $\boldsymbol{y}_i$

3. Propagate the reference state and state transition matrix (STM), $\Phi(t_i, t_{i-1})$ to the time of the measurement through the differential equations

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix}, \quad \dot{\Phi} = \frac{\partial \dot{\boldsymbol{x}}}{\partial \boldsymbol{x}} \Phi \tag{3}$$

   where $\ddot{x}, \ddot{y}, \ddot{z}$ correspond to the perturbing accelerations like gravity, solar radiation pressure, and/or thrust vectors.

4. Perform the time update:

$$\Delta x_i^- = \Phi \Delta x_{i-1}^+ \tag{4}$$
$$P_i^- = \Phi P_{i-1}^+ \Phi^T + Q_i \tag{5}$$

5. Perform the measurement update:

$$r_i = y_i - \hat{y}_i \tag{6}$$
$$H_i = \frac{\partial h}{\partial x}\Big|_{x_i} \tag{7}$$
$$K_i = P_i^- H_i^T (H_i P_i^- H_i^T + R_i)^{-1} \tag{8}$$
$$\Delta x_i^+ = \Delta x_i^- + K_i(r_i - H_i \Delta x_i^-) \tag{9}$$
$$P_i^+ = (\mathbb{I} - K_i H_i) P_i^- (\mathbb{I} - K_i H_i)^T \tag{10}$$

   where h is the measurement function.

6. Repeat for each incoming measurement

For a more comprehensive derivation of these equations, see [22].

A few important notes regarding the algorithm: first, the Kalman filter relies heavily on the assumption that there exists an analytic form of the perturbing accelerations such that the partial $\frac{\partial \dot{x}}{\partial x}$ can be evaluated for the propagation of the STM. This is precisely why low-degree spherical harmonics gravity model are popular. The partials of the acceleration produced by the model can be derived analytically, and once programmed, can be evaluated efficiently for the propagation of the STM.

While this analytic form is convenient, it is not explicitly required. If the gravity model is of sufficiently high fidelity, evaluating the jacobian may be too computationally cumbersome. Alternatively, an analytic form may not exist. In these cases the jacobian can be computed numerically through finite differencing which introduces its own computational expense and accuracy limitations.

The second interesting note is that the filter, in its current form, captures dynamical uncertainty through the process noise matrix Q_i. By formulating the filter in this way, these dynamical uncertainties will never be known explicitly – it simply ensures that the filter does not grow overconfident in its state estimate given the limited knowledge of the dynamics. This is known as stochastic noise compensation (SNC).

While SNC is a popular way to inject noise into the filter to ensure conservative state estimates, it is not the only way to capture dynamical uncertainty. An alternative approach is to use dynamic model compensation (DMC).

2.3.1 Dynamic Model Compensation

Dynamical model compensation assumes that process noise is correlated in time, and this correlation can be exploited to reconstruct unmodelled dynamics. By assuming there is an underlying differential equation governing the evolution of the process noise, the noise can be treated as part of the dynamical system and estimated as part of the state. This is often done by assuming the process noise

evolves according to a first-order Gauss-Markov process

$$\dot{\mathbf{w}} = -\frac{1}{\tau}\mathbf{w} \tag{11}$$

Using the approximation, the unknown perturbations, w, can be augmented the state vector and estimated directly to produce an approximation of any unmodelled accelerations like those produced by a mis-aligned thrust vector or high-order gravitational accelerations. More information regarding DMC can be found in [23].

2.4 PINN-GM Kalman Filter

As discussed, traditional filters in small-body settings tend to leverage spherical harmonic gravity models due to their analytic convenience and differentiability; however, these conveniences come at the cost of operational limitations and challenging estimation requirements to regress the harmonic coefficients. This paper provides the first implementation of the novel PINN-GM into the orbit determination pipeline.

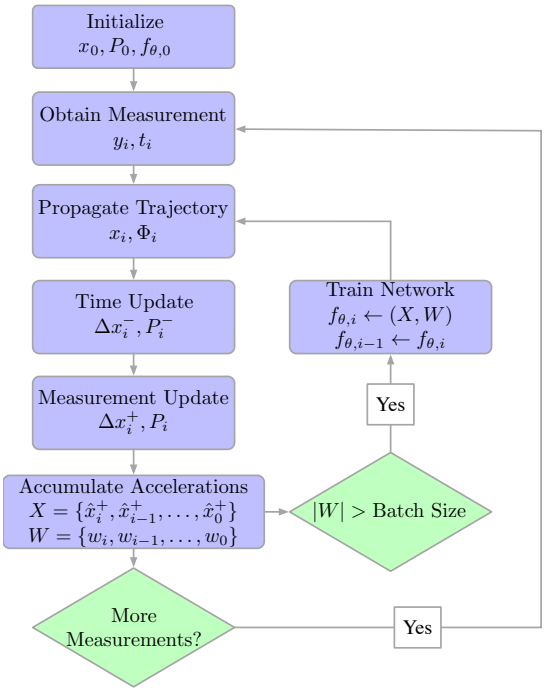

Fig. 2: PINN-GM Kalman Filter

The proposed filtering framework in shown in Figure 2. This new filter uses a PINN-GM as the dynamics backend to the estimation pipeline, and uses DMC to generate training data to then update the PINN-GM over time. Explicitly, the PINN-GM can be trained if there exists measurements of the spacecraft position and the corresponding gravitational acceleration at that position. In the past, these measurements have been generated offline using previous constructed high-fidelity gravity models (EGM-2008 for Earth, or high-fidelity shape models for asteroids). In most circumstances, however, these offline models do not exist, and measurements must be taken in-situ.

DMC offers a compelling way to provide these training data. By estimating the unmodelled accelerations as part of the state, the filter can accumulate "observations" of the gravitational acceleration at an estimated spacecraft position. Once a sufficient number of these observations exist, the neural network can be updated with traditional stochastic gradient descent using that data. Moreover, in the recent PINN-GM-II paper it has been shown that the model is robust to noisy training data, so even if the estimates are erroneous, the PINN-GM is able to converge on a reliable solution.

One of the advantages of approaching the gravity field estimation problem in this way is that there is no longer a strict sampling requirement to update the model. Spherical harmonics is prone to aliasing other harmonics if the measurements are not properly distributed. The PINN-GM can accept data regardless of distribution, and use that data to improve the gravity model in whatever region it was collected.

The other advantage to using a PINN-GM within the orbit determination pipeline is that it allows for the rapid differentiation of the dynamics with respect to the spacecraft state $\frac{\partial \dot{x}}{\partial x}$. Thanks to automatic differentiation, a PINN-GM of arbitrary fidelity can be differentiated and used to propagate the state and update the covariance. Moreover, this approach is general such that as the gravity model improves it can immediately be leveraged within the pipeline without having to reformulate the state.

3. Problem Setup

To investigate the performance of the PINN-GM enabled Kalman filter, a simple scenario is developed. Consider a spacecraft orbiting a non-rotating asteroid with the shape of 433-Eros. The spacecraft is subject only to the natural gravitational dynamics of this artificial asteroid, produced using the polyhedral gravity model. The initial orbital elements and the position propagated for approximately 3 orbit periods are shown in Table 1 and Figure 3 respectively. In this scenario, there exists perfect range and range-rate measurements of the spacecraft taken from an inertially-fixed observer at a distance of approximately 0.1 AU. It is also assumed that the bulk density / gravitational parameter for the asteroid in question is known a-priori. The state, covariance, process noise matrix, and measurement noise matrix used to initialize the filter are listed in Table 2. This scenario, while not representative of a true system, provides a testbench to characterize the preliminary performance of the PINN-GM applied to orbit determination and gravity field estimation.

Table 1: Initial orbital elements

Element	Value	Unit
a	30,000	[m]
e	0.2	[-]
i	0.1	[rad]
ω	0.0	[rad]
Ω	0.0	[rad]
M	0.0	[rad]

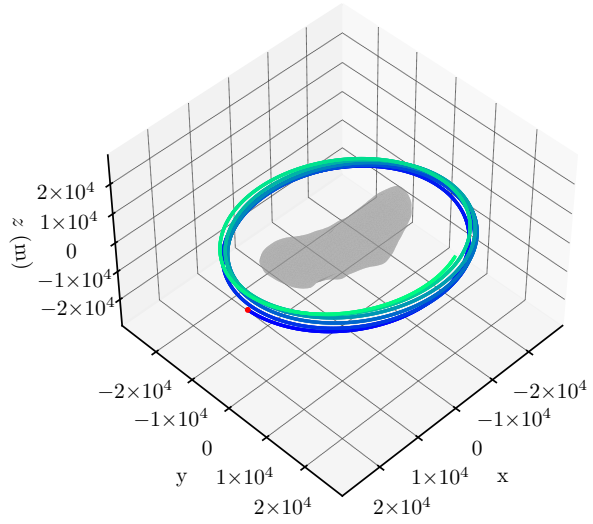

Fig. 3: Spacecraft trajectory about the asteroid

An experiment is designed to first characterize the PINN-GM enabled filter with respect to certain hyperparameters. Specifically, the PINN-GM is first trained online using a point-mass reference model (i.e. the PINN-GM is trained to initially mimic point mass behavior). This filter uses dynamic model compensation to estimate unmodelled accelerations, and once a sufficient number of state estimates and acceleration estimates are produced, the PINN-GM is updated using stochastic gradient decent.

The online training of the PINN-GM using the position and DMC estimates is a nuanced process which is sensitive to multiple parameters. For one, stochastic gradient decent relies on a user-defined learning rate, η, which weights how heavily the measurements should influence the model update step. If this learning rate is too large, the network will disproportionally favor correctly modeling the incoming data, potentially at the expense of other regions that were trained using past data. Alternatively, if the the learning rate is too small, the network will require an excessive amount of data before learning the proper dynamics.

Table 2: Initial state, covariance, process noise matrix, and measurement noise matrix

State	Value	Unit
x_0	[24,000, 0, 0]	[m]
v_0	[0, 4.7, 4.716]	[m/s]
w_0	[0 , 0, 0]	[m/s^2]
P_{x_0}	diag([100, 100, 100])	[m^2]
P_{v_0}	diag([0.01, 0.01, 0.01])	[m^2/s^2]
P_{w_0}	diag([1E-14 , 0, 0])	[m^2/s^4]
Q	diag([1E-22, 1E-22, 1E-22])	[m^2/s^2]
R	diag([1, 1E-6])	[m^2] and [m^2/s^4]

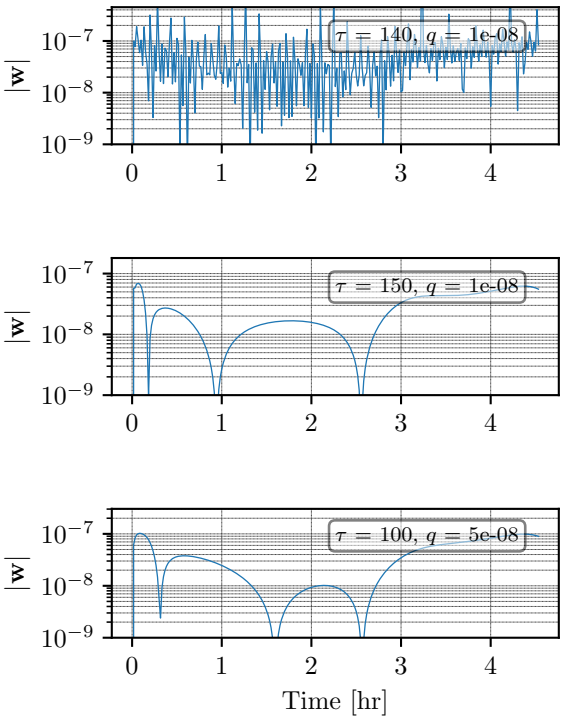

Fig. 4: DMC with different values of q and τ

The second parameter is how many samples should be used in a stochastic gradient update step, i.e. what is the mini-batch size? Large mini-batches can produce the most representative estimates of the cost landscape, but this can cause the networks to converge into local minimums and intrinsically requires more measurements. Small batch sizes have much noisier gradient descent steps, which can cause the network to learn more slowly albeit with greater chance of launching out of a local minimum into a more desirable part of the cost landscape.

The third parameter is the number of epochs used during training. An epoch corresponds to the total number of times the entire data set is iterated over. I.e. if a mini-batch has size 32, but the data set has 128 samples then 4 gradient descent updates are taken in a single epoch. The number of epochs and the learning rate have coupled effects on one another. A small learning rate would require more epochs to achieve the same loss value as a large learning rate and a small number of epochs. The former would produce a more stable, but slow, learning process; whereas the later would produce a fast and less stable descent.

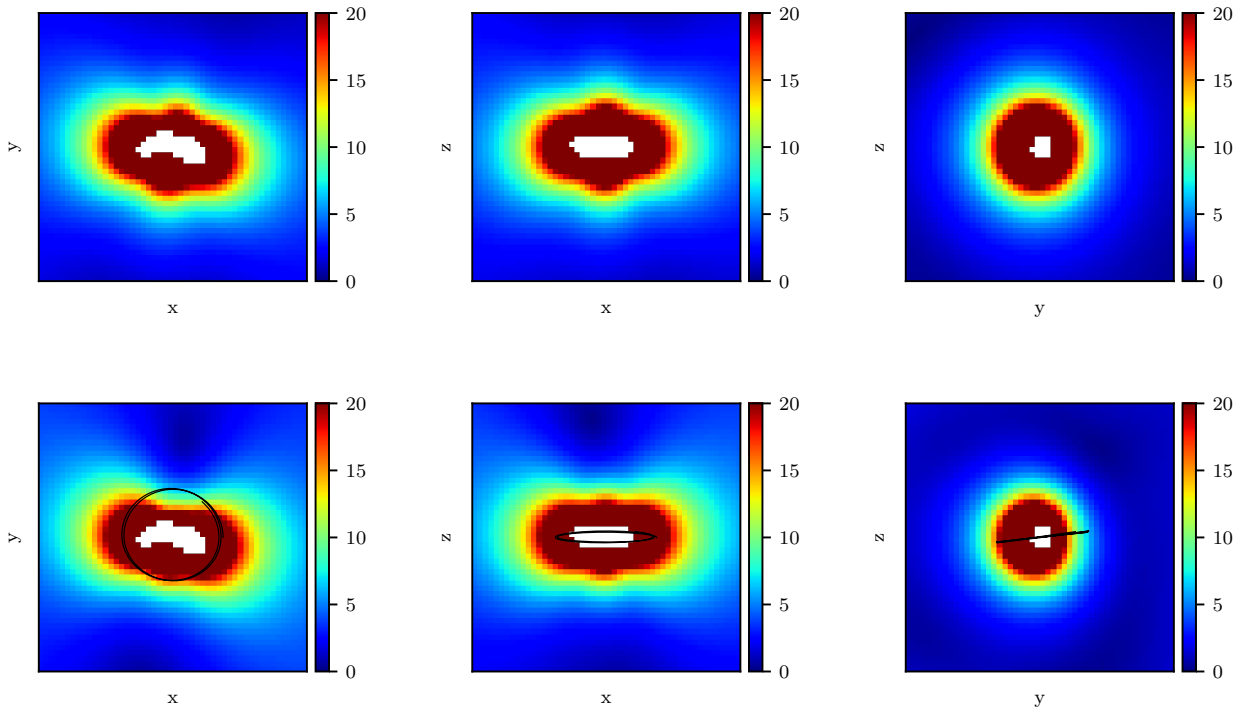

Fig. 5: Top Row: Point mass gravitational acceleration error. Bottom Row: PINN-GM acceleration error after trained with in-situ DMC estimates and best performing hyperparameters. Color bar corresponds to percent error in the gravitational acceleration. Black line is trajectory of spacecraft for 1/5th of the mission duration.

The final parameters of note are the value of τ used within the dynamics of the noise model and q, the magnitude of the process noise matrix. Large values of τ assume that the unmodelled accelerations have greater correlation in time whereas small values assume the accelerations are more instantaneous. Given the irregular geometries and varying densities of the celestial body in question, one needs to carefully consider the time scales over which the high-order gravitational accelerations would be expected to change by a considerable amount before selecting τ. The parameter q is the remaining process noise left unmodelled after DMC. Large values produce more conservative state and acceleration estimates, where as too small q can produce an overconfidence in the solution.

A hyperparameter experiment is performed which tests a variety of combinations of the mentioned parameters for a subset of the mission duration; specifically the PINN-GM only learns from the DMC estimates collected during the first two orbits. The performance of the PINN-GM within the filter is estimated by measuring the average percent error of the gravitational accelerations produced along the XY, YZ, and XZ planes. For reference, the PINN-GM trained only on a point mass reference has an average error of 7.57%, therefore any PINN-GM model that achieves lower error is considered to be learning a better representation of the true gravitational dynamics.

The network hyperparameters include the learning rate, the number of epochs run per network update, the mini-batch size, and the type of PINN constraint used (see [15] for details). Additionally, different values for the dynamic model are tested (τ and q).

4. Results

First, the accuracy of the default PINN-GM trained on a point mass reference model is explored. Figure 5 shows the error along each cartesian plane of this model compared to the truth model (a 200k facet polyhedral model of the asteroid). At low altitudes, this model is in excess of 20% error, and at increasing altitudes transitions towards 5% error. These high errors can become large problems depending on the orbit of the spacecraft. Even if the spacecraft is not performing a landing or touch-and-go manuever, this point mass gravity model may be insufficiently representative of the true system such that the linearization applied in the filter will begin breaking down over extended periods of time when the state deviation grows to far from the true trajectory.

When averaged across all planes, the error of the PINN-GM trained on a point mass reference model is 7.57%. This average error serves as the benchmark metric by

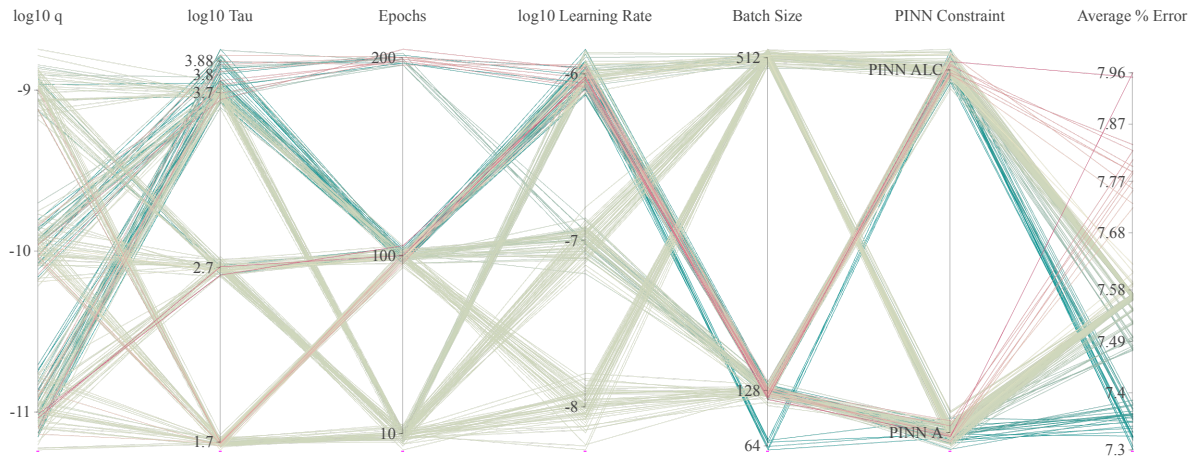

Fig. 6: Results of the hyperparameter search quantifying average percent error along the XY, YZ, and XZ planes.

which different hyperparameter configurations are evaluated. Figure 6 shows the effect of different hyperparameter combinations on this metric. The first observation from Figure 6 is that there exist hyperparameter combinations that yield solutions that produce error below the 7.57%. This suggests that the PINN-GM can learn more accurate dynamics using DMC acceleration estimates as training data. While the most dramatic improvement only achieves an average error of 7.30% (a 3.5% increase in accuracy compared to the baseline), this provides the first proof of concept that there is a viable means by which these models can be trained.

In contrast, there are also hyperparameter combinations which yield suboptimal performance, producing higher average error after the network updates than the original model. Generally these less successful models are the result of high learning rates, combined with long training times (epochs), and small mini-batch sizes. These results are somewhat expected as small batch sizes produce less reliable estimates of the loss landscape, and the high learning rate and long training times allow for the weights to update too quickly in counterproductive directions.

Other high-level observations regarding the hyperparameters include the sensitivity to the DMC hyperparameters, τ and q. In general, large values of τ produced more productive acceleration measurements for the PINN-GM than smaller values. This suggests that longer time correlation between unmodelled accelerations can generate more stable training.

The PINN-GM that achieved the lowest error is plotted in the middle row of Figure 5. The error decreases most prominently in the YZ plane, where the trajectory intersects the region of high-error (originally $>$30%) and reduces it down to approximately 10%. In the XY and XZ planes, some regions show a decrease in error — namely around the semi-minor axis of the body – whereas regions at higher altitudes show a slight increase in error. This increase in error is undesirable though expected. Given that the in-situ training data is produced at low altitudes, the PINN-GM will give preference to these samples accurately at the expense of mismodeling regions for which there is little to no data. This problem is exacerbated when the entire mission duration is used to update the model. Future work will investigate how to mitigate this behavior.

5. Conclusion

This paper introduces the Physics-Informed Neural Network Gravity Model (PINN-GM) into a Kalman filter for small-body gravity estimation. While the preliminary results show that there is some marginal improvements to the on-board gravity model, there are a number of sensitivities of which the researcher must be attuned. Specifically, the PINN-GM can produce considerably different field representations given the same measurement when provided different values for the network learning rate, training epochs, and DMC characterization. This sensitivity is undesirable, forcing engineers and scientists to carefully tune the network learning parameters before deployment. Despite this, there remains a number of encouraging outlets for future research. For one, the hyperparameter study shown in Figure 6 demonstrates that there exists a wide range of possible parameters in which the network can achieve better performance than the baseline model. With further optimization and testing it is likely that these results can be further exploited.

In addition, higher-order forms of dynamical model compensation have been shown to produce more reliable estimates of high-order gravitational perturbations [23]. By incorporating such models into the filter, it is possible that the training data used to update the on-board PINN-GM will be of higher quality such that the gradient descent is less noisy and less likely to introduce unwanted error into the system.

Another area for future research is improving the filter implementation. This filter implementation is not numerically desirable for training neural networks. By framing the problem in a dimensionalized (as opposed to non-dimensionalized) manner, the numerics of the algorithm are more likely to be unstable or inaccurate. In small-body settings, this choice can be detrimental to the accuracy of the filter solutions. The unmodelled accelerations in question are orders-of-magnitude smaller than the measurements and remaining state variables. When performing the corresponding matrix inversions for state updates, naturally there will be a loss of precision given such large differences in scales. One potential solution is to use a non-dimensionalized form of the dynamics to ensure greater numerical stability. The other remedy is to use alternative filter types, such as the square-root information filter (SRIF). Using a SRIF cuts the required numerical precision in half making it particularly useful for deep space navigation. A non-dimensionalized SRIF can therefore produce more precise acceleration predictions / training data for the on-board PINN-GM.

Finally, in addition to continuing to improve the training pipeline, future work will investigate a more representative mission scenario. This paper focused on a largely equatorial trajectory which makes observability of the state and acceleration along the z-axis particularly challenging. Moreover, this scenario only simulates gravitational effects on the spacecraft. Future work will investigate how these new filters perform in a more representative setting with additional perturbations like solar radiation pressure, thermal radiation pressure, and N-body effects.

REFERENCES

[1] S. i. Watanabe, Y. Tsuda, M. Yoshikawa, S. Tanaka, T. Saiki, and S. Nakazawa, "Hayabusa2 Mission Overview," *Space Science Reviews*, Vol. 208, No. 1-4, 2017, pp. 3–16, 10.1007/s11214-017-0377-1.

[2] D. S. Lauretta, S. S. Balram-Knutson, E. Beshore, W. V. Boynton, C. Drouet d'Aubigny, D. N. DellaGiustina, H. L. Enos, D. R. Golish, C. W. Hergenrother, E. S. Howell, C. A. Bennett, E. T. Morton, M. C. Nolan, B. Rizk, H. L. Roper, A. E. Bartels, B. J. Bos, J. P. Dworkin, D. E. Highsmith, D. A. Lorenz, L. F. Lim, R. Mink, M. C. Moreau, J. A. Nuth, D. C. Reuter, A. A. Simon, E. B. Bierhaus, B. H. Bryan, R. Ballouz, O. S. Barnouin, R. P. Binzel, W. F. Bottke, V. E. Hamilton, K. J. Walsh, S. R. Chesley, P. R. Christensen, B. E. Clark, H. C. Connolly, M. K. Crombie, M. G. Daly, J. P. Emery, T. J. McCoy, J. W. McMahon, D. J. Scheeres, S. Messenger, K. Nakamura-Messenger, K. Righter, and S. A. Sandford, "OSIRIS-REx: Sample Return from Asteroid (101955) Bennu," *Space Science Reviews*, Vol. 212, No. 1-2, 2017, pp. 925–984, 10.1007/s11214-017-0405-1.

[3] C. T. Russell, F. Capaccioni, A. Coradini, M. C. De Sanctis, W. C. Feldman, R. Jaumann, H. U. Keller, T. B. McCord, L. A. McFadden, S. Mottola, C. M. Pieters, T. H. Prettyman, C. A. Raymond, M. V. Sykes, D. E. Smith, and M. T. Zuber, "Dawn Mission to Vesta and Ceres," *Earth, Moon and Planets*, Vol. 101, No. 1-2, 2007, pp. 65–91, 10.1007/s11038-007-9151-9.

[4] D. Y. Oh, S. Collins, T. Drain, W. Hart, T. Imken, K. Larson, D. Marsh, D. Muthulingam, J. S. Snyder, D. Trofimov, L. T. Elkins-Tanton, I. Johnson, P. Lord, and Z. Pirkl, "Development of the Psyche Mission for NASA's Discovery Program," *36th International Electric Propulsion Conference*, Vienna, Austria, 2019.

[5] D. J. Scheeres, J. W. McMahon, E. B. Bierhaus, J. Wood, L. A. Benner, C. Hartzell, P. Hayne, J. Hopkins, R. Jedicke, L. LeCorre, S. Naidu, P. Pravec, and M. Ravine, "Janus: A NASA SIMPLEx Mission to Explore Two NEO Binary Asteroids," *Proceedings of the International Astronautical Congress, IAC*, Vol. 2020-Octob, 2020.

[6] J. K. Miller, A. S. Konopliv, P. G. Antreasian, J. J. Bordi, S. Chesley, C. E. Helfrich, W. M. Owen, T. C. Wang, B. G. Williams, D. K. Yeomans, and D. J. Scheeres, "Determination of Shape, Gravity, and Rotational State of Asteroid 433 Eros," *Icarus*, Vol. 155, No. 1, 2002, pp. 3–17, 10.1006/icar.2001.6753.

[7] A. S. French, *Precise Orbit Determination And Gravity Field Estimation During Small Body Missions*. PhD thesis, University of Colorado Boulder, 2020.

[8] D. J. Scheeres, B. Khushalani, and R. A. Werner, "Estimating Asteroid Density Distributions from Shape and Gravity Information," *Planetary and Space Science*, Vol. 48, No. 10, 2000, pp. 965–971, 10.1016/s0032-0633(00)00064-7.

[9] D. J. Scheeres, A. S. French, P. Tricarico, S. R. Chesley, Y. Takahashi, D. Farnocchia, J. W. McMahon, D. N. Brack, A. B. Davis, R. L. Ballouz, E. R. Jawin, B. Rozitis, J. P. Emery, A. J. Ryan, R. S. Park, B. P. Rush, N. Mastrodemos, B. M. Kennedy, J. Bellerose, D. P. Lubey, D. Velez, A. T. Vaughan, J. M. Leonard, J. Geeraert, B. Page, P. Antreasian, E. Mazarico, K. Getzandanner, D. Rowlands, M. C. Moreau, J. Small, D. E. Highsmith, S. Goossens, E. E. Palmer, J. R. Weirich, R. W. Gaskell, O. S. Barnouin, M. G. Daly, J. A. Seabrook, M. M. Al Asad, L. C. Philpott, C. L. Johnson, C. M. Hartzell, V. E. Hamilton, P. Michel, K. J. Walsh, M. C. Nolan, and D. S. Lauretta, "Heterogeneous Mass Distribution of the Rubble-Pile Asteroid (101955) Bennu," *Science Advances*, Vol. 6, No. 41, 2020, 10.1126/sciadv.abc3350.

[10] R. Werner and D. Scheeres, "Exterior Gravitation of a Polyhedron Derived and Compared with Harmonic and

Mascon Gravitation Representations of Asteroid 4769 Castalia," *Celestial Mechanics and Dynamical Astronomy*, Vol. 65, No. 3, 1997, pp. 313–344, 10.1007/BF00053511.

[11] Y. Takahashi and D. Scheeres, "Morphology Driven Density Distribution Estimation for Small Bodies," *Icarus*, Vol. 233, May 2014, pp. 179–193, 10.1016/j.icarus.2014.02.004.

[12] J. Martin and H. Schaub, "Physics-Informed Neural Networks for Gravity Field Modeling of the Earth and Moon," *Celestial Mechanics and Dynamical Astronomy*, Vol. 134, Apr. 2022, 10.1007/s10569-022-10069-5.

[13] W. M. Kaula, *Theory of Satellite Geodesy: Applications of Satellites to Geodesy*. Waltham, Mass.: Blaisdell Publishing Co, 1966.

[14] G. Romain and B. Jean-Pierre, "Ellipsoidal Harmonic Expansions of the Gravitational Potential: Theory and Application," *Celestial Mechanics and Dynamical Astronomy*, Vol. 79, No. 4, 2001, pp. 235–275, 10.1023/A:1017555515763.

[15] J. Martin and H. Schaub, "Physics-Informed Neural Networks for Gravity Field Modeling of Small Bodies," *Celestial Mechanics and Dynamical Astronomy*, Sept. 2022, p. 28.

[16] E. Hewitt and R. E. Hewitt, "The Gibbs-Wilbraham Phenomenon: An Episode in Fourier Analysis," *Archive for History of Exact Sciences*, Vol. 21, No. 2, 1979, pp. 129–160, 10.1007/BF00330404.

[17] C. Jekeli, "Potential Theory and the Static Gravity Field of the Earth," *Treatise on Geophysics: Second Edition*, Vol. 3, 2015, pp. 9–35, 10.1016/B978-0-444-53802-4.00056-7.

[18] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations," *Journal of Computational Physics*, Vol. 378, Feb. 2019, pp. 686–707, 10.1016/j.jcp.2018.10.045.

[19] A. G. Baydin, B. A. Pearlmutter, and J. M. Siskind, "Automatic Differentiation in Machine Learning: A Survey," *Journal of Machine Learning Research*, Vol. 18, 2018, pp. 1–43.

[20] L. Bottou, *Stochastic Gradient Descent Tricks*, Vol. 7700. Springer, Jan. 2012, 10.1007/978-3-642-35289-8-25.

[21] G. Welch, "An Introduction to the Kalman Filter," tech. rep., 1997.

[22] B. Schutz, B. Tapley, and G. H. Born, *Statistical Orbit Determination*. Elsevier, 2004.

[23] J. M. Leonard, F. G. Nievinski, and G. H. Born, "Gravity Error Compensation Using Second-Order Gauss-Markov Processes," *Journal of Spacecraft and Rockets*, Vol. 50, Jan. 2013, pp. 217–229, 10.2514/1.A32262.