SmallSat swarm gravimetry: Revealing the interior structure of asteroids and comets

William Gordon Ledbetter¹(✉), Rohan Sood¹, James Keane², and Jeffrey Stuart²

1. *Astrodynamics and Space Research Laboratory, The University of Alabama, Tuscaloosa, AL 35487, USA*

2. *Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA*

ABSTRACT

A growing interest in small body exploration has motivated research into the rapid characterization of near-Earth objects to meet economic or scientific objectives. Specifically, knowledge of the internal density structure can aid with target selection and enables an understanding of prehistoric planetary formation to be developed. To this end, multi-layer extensions to the polyhedral gravity model are suggested, and an inversion technique is implemented to present their effectiveness. On-orbit gravity gradiometry is simulated and employed in stochastic and deterministic algorithms, with results that imply robustness in both cases.

KEYWORDS

gravimetry
asteroid exploration
astrodynamics
stochastic optimization

Research Article

Received: 26 May 2020

Accepted: 26 October 2020

© The Author(s) 2020

1 Introduction

Exploratory missions to small, unknown celestial bodies are limited in their ability to adequately characterize and investigate the resource potential of a body of interest. Multiple near-Earth objects (NEOs) have been identified as potential targets for human exploration with the goal of harnessing available resources. However, limited knowledge of the physical characteristics of NEOs, such as their shape, density, gravity field, and composition, poses a challenge to any manned exploration. Additionally, both asteroids and comets are of significant importance to answer fundamental questions regarding the origin and evolution of our solar system. Driven by the need for scientific exploration and considering our economic interest, the proposed strategy utilizes a swarm of SmallSats to address specific strategic knowledge gaps (SKGs), enabling resource investigation and acting as a precursor to human exploration. The algorithm proposes a multi-vehicle approach to significantly enhance the quality of data and assist in close-proximity guidance of spacecraft. Specifically, a swarm of small probes, deployed from the primary spacecraft, performs flybys of the celestial object, enabling high-fidelity, *in-situ* gravimetry measurements by the chief. These measurements are then used to construct a gravity model, which can be used

for model-predictive control [1], or, as is the topic of this paper, obtaining information about the internal structure of the body.

Previous studies considered the use of gravimetry with passive craft [2], and existing methods on polyhedral density estimation require an existing gravity model of a body prior to any analysis [3, 4]. Taking inspiration from the GOCE mission [5] and work by Park *et al.* [6], in this investigation, sensor-equipped probes derived from the work of Carroll and Faber [7] enable direct gravity measurements. Additionally, the probes are tracked by the chief as they move in the vicinity of a central body.

Carroll and Faber proposed using a bias-free accelerometer at a fixed distance from the center of mass of the spacecraft to measure the tidal acceleration on the spacecraft. Taking into account that the tidal acceleration is a function of the gravity gradient (GG) acting on the spacecraft, GG was chosen as the measurement type for the recovery.

Our procedure requires only three sets of data to obtain a density map: a shape model, probe positions, and the corresponding gravity gradient measurements at those positions. This study focuses on the development and initial testing of a noise-free recovery algorithm, and also undertakes a stochastic analysis of positional uncertainty.

To simulate this problem, a dynamical model was constructed that consists of a swarm of probes, which is ejected from the chief, and which flies by the central body while collecting gravity gradient data. We then execute the gravity inversion algorithm, as detailed below, to obtain densities for each “node” of the shape model of the central body.

2 Dynamical model

Modern research in asteroid gravimetry utilizes the polyhedral gravity model developed by Werner and Scheeres [8]. The model is attractive for two primary reasons: ease of implementation and benefits over other models. First, the model takes advantage of a three-dimensional (3D) shape reconstruction, data that can be approximated from Earth via a light curve [9] or radar [10] observations. The existence of reliable methods for shape model construction enables the implementation of the polyhedral model in multiple scenarios [11]. Shape models can also be used in conjunction with imaging for navigation algorithms [12, 13], further motivating the shape model construction as a primary mission objective. Second, polyhedral models guarantee convergence of the gravity field in close proximity to irregular bodies. The infinite sum of the spherical harmonics diverges within a sphere circumscribing the body, creating a significant invalid zone for irregularly shaped objects. Although methods exist for mitigating this issue [14], the polyhedral formulation presents itself as an appropriate vehicle for investigating physical characteristics such as the shape and density of an asteroid.

2.1 Heterogeneous polyhedral gravity model

The goal of this study is to use gravimetry techniques to reconstruct the internal density maps. These reconstructions are often known as inversions because they progress in the direction opposite to that in which the equation is typically used. This practice necessitates the initial expression of the gravity of a polyhedron with some arbitrary internal structure as though the density were already known. The formulation begins with the commonly used polyhedral gravity model, as detailed below.

Werner and Scheeres’ formulation for a homogeneous

polyhedral model is given in Eq. (1):

$$U = G\rho \sum_{e \in \text{edges}} \bar{\mathbf{r}}_e \cdot \mathbf{E}_e \cdot \bar{\mathbf{r}}_e \cdot L_e - G\rho \sum_{f \in \text{facets}} \bar{\mathbf{r}}_f \cdot \mathbf{F}_f \cdot \bar{\mathbf{r}}_f \cdot \omega_f \quad (1)$$

where G is the gravitational constant, ρ is the average density of the body, $\bar{\mathbf{r}}_e$ is a vector from the point in space to a point on edge e , $\mathbf{E}_e$ is a matrix describing the geometry of edge e , and L_e is the potential of a 1D “wire” expressed in terms of the distances to the edge endpoints. In the second term, $\bar{\mathbf{r}}_f$ is a vector from the query point in space to a point on facet f , $\mathbf{F}_f$ is a matrix describing the geometry and orientation of facet f , and ω_f is the signed solid angle subtended by facet f , from the perspective of the point in space. Equations (2)–(5) detail the calculation of parameters used in Eq. (1):

$$\mathbf{E}_e = \hat{n}_A \hat{n}_{12}^A + \hat{n}_B \hat{n}_{21}^B \quad (2)$$

$$\mathbf{F}_f = \hat{n}_f \hat{n}_f \quad (3)$$

$$L_e = \log \frac{r_{e,1} + r_{e,2} + e_{12}}{r_{e,1} + r_{e,2} - e_{12}} \quad (4)$$

$$\omega_f = 2 \arctan \left[\frac{\bar{\mathbf{r}}_1 \cdot (\bar{\mathbf{r}}_2 \times \bar{\mathbf{r}}_3)}{r_1 r_2 r_3 + r_1 \bar{\mathbf{r}}_2 \cdot \bar{\mathbf{r}}_3 + r_2 \bar{\mathbf{r}}_3 \cdot \bar{\mathbf{r}}_1 + r_3 \bar{\mathbf{r}}_1 \cdot \bar{\mathbf{r}}_2} \right] \quad (5)$$

In Eqs. (2) and (3), $\hat{n}_f$ is equivalent to $\hat{n}_A$ and $\hat{n}_B$, all of which describe a unit outward normal vector for face f , A , or B , respectively. In Eq. (2), $\hat{n}_{12}^A$ is a unit vector that is in the plane of face A , is orthogonal to the edge between faces A and B , and points out of face A . Likewise, the unit vector $\hat{n}_{21}^B$ is in the plane of face B , orthogonal to the edge between B and A , and points out of face B . With the safe assumption that the shape model is static in the chosen reference frame, these values can be pre-computed and need not be recalculated for every evaluation of Eq. (1).

A simple modification of Eq. (1) via the distributive property enables the quantification of heterogeneous density, ρ_i , where each facet and edge has a unique density.

$$U = G \sum_{e \in \text{edges}} \bar{\mathbf{r}}_e \cdot \mathbf{E}_e \cdot \bar{\mathbf{r}}_e \cdot L_e \rho_e - G \sum_{f \in \text{facets}} \bar{\mathbf{r}}_f \cdot \mathbf{F}_f \cdot \bar{\mathbf{r}}_f \cdot \omega_f \rho_f \quad (6)$$

Equation (6) can then be converted into a vector operation:

$$U = \bar{u} \cdot \bar{\rho} = [\bar{u}_e, \bar{u}_f] \cdot [\bar{\rho}_e, \bar{\rho}_f] \quad (7)$$

The vector $\bar{u}$ encapsulates both the asteroid shape and the satellite position via Eqs. (8) and (9), which are extracted from Eq. (6):

$$\bar{u}_e = G[L_{e1}(\bar{r}_{e1} \cdot \mathbf{E}_{e1} \cdot \bar{r}_{e1}), L_{e2}(\bar{r}_{e2} \cdot \mathbf{E}_{e2} \cdot \bar{r}_{e2}), \dots] \quad (8)$$

$$\bar{u}_f = -G[\omega_{f1}(\bar{r}_{f1} \cdot \mathbf{F}_{f1} \cdot \bar{r}_{f1}), \omega_{f2}(\bar{r}_{f2} \cdot \mathbf{F}_{f2} \cdot \bar{r}_{f2}), \dots] \quad (9)$$

Each ρ_e and ρ_f are evaluated at the associated node, which is the midpoint of the respective facet or edge. Our notation convention is as follows: $\bar{u}$ is a vector of all u_e and u_f , whereas $\bar{u}_e$ is a vector of all u_e and $\bar{u}_f$ is a vector of all u_f . The density vectors $\bar{\rho}$, $\bar{\rho}_e$, and $\bar{\rho}_f$ are paired with their corresponding $\bar{u}_i$. This convention is consistent for forthcoming derivations of the acceleration and gravity gradient, with the dimension of shape-based quantities increasing by one for each derivative ($\bar{u}_{(1 \times n)}$, $\mathbf{a}_{(3 \times n)}$, $\mathbf{g}_{(3 \times 3 \times n)}$).

Thus far, the dynamical model is capable of calculating the gravitational potential of a polyhedron with non-uniform density by assuming that each facet- and edge-defined tetrahedron of the shape has a unique density. Each tetrahedron consists of one facet of the shape model and the origin of the coordinate system. This formulation assumes variability only in the latitude and longitude dimensions. The density of a facet represents the density of the volume of the tetrahedron, from facet to origin. In the derivation of Werner and Scheeres' equation, each edge must be counted twice, once for each facet it borders. In their formulation of the constant-density gravity field, common terms cancel and combine such that each edge is considered only once. For a heterogeneous body, the edge is still considered once, but with a density that is the average of that of the two adjacent facets. As a result, each tetrahedron is homogeneous from the center of the body to the assumed surface.

Radial variability is achieved by layering multiple polyhedra. In previous work, non-concentric homogeneous layering was utilized to represent various morphologies [4].

However, they assumed complete heterogeneity, thus simplifying the layering assumptions without sacrificing the scope of the expressible internal structures. The construction of radial variability begins by assuming that each internal layer is a concentric, scaled copy of the outermost polyhedron. This assumption preserves the $\mathbf{E}_e$ and $\mathbf{F}_f$ matrices from the reference shape. Therefore, only the relative position vectors $\bar{r}_e$ and $\bar{r}_f$, and the scalars L_e and ω_f vary with each subsequent layer. The mathematical procedure necessary for layering is explained in Fig. 1 and Eq. (10). In Eqs. (10)–(12), the size of each color wheel is analogous to $\bar{u}_i$ and the orientation of the color pattern is analogous to $\bar{\rho}_i$. In this study, the outermost shape and density pattern are assigned $i = 1$, with an increasing value of i for the inner layers.

$$U = (\bar{u}_1 - \bar{u}_2)\bar{\rho}_1 + (\bar{u}_2 - \bar{u}_3)\bar{\rho}_2 + \bar{u}_3\bar{\rho}_3 = \bar{u} \cdot \bar{\rho} \quad (10)$$

$$\bar{u} = [(\bar{u}_1 - \bar{u}_2), (\bar{u}_2 - \bar{u}_3), \bar{u}_3] \quad (11)$$

$$\bar{\rho} = [\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3] \quad (12)$$

Acceleration is analytically obtained by taking the gradient of Eq. (6), shown in Eq. (13), and can be reformulated as represented by Eq. (14):

$$\bar{\mathbf{A}} = -G \sum_{e \in \text{edges}} \bar{r}_e \cdot \mathbf{E}_e \cdot L_e \rho_e + G \sum_{f \in \text{facets}} \bar{r}_f \cdot \mathbf{F}_f \cdot \omega_f \rho_f \quad (13)$$

$$\bar{\mathbf{A}} = \mathbf{a} \cdot \bar{\rho} = [\mathbf{a}_e, \mathbf{a}_f] \cdot [\bar{\rho}_e, \bar{\rho}_f] \quad (14)$$

Here, $\bar{u}$ is a $n \times 1$ row vector, whereas $\mathbf{a}$ is a $3 \times n$ matrix. As a result, $\bar{\mathbf{A}}$ is a 3-element vector. The first step of the acceleration calculation entails building the $\mathbf{a}$ matrix according to Eqs. (15) and (16):

$$\bar{\mathbf{a}}_e = -GL_e(\bar{r}_e \cdot \mathbf{E}_e) \quad (15)$$

$$\bar{\mathbf{a}}_f = G\omega_f(\bar{r}_f \cdot \mathbf{F}_f) \quad (16)$$

To maximize the computational efficiency, each $\bar{\mathbf{a}}_e$ and $\bar{\mathbf{a}}_f$ vector is calculated in parallel and stored in the $\mathbf{a}$ matrix, as is the case for u_e and u_f above. Considering a multi-layered body, the $\mathbf{a}$ matrices for each layer need

Fig. 1 Visualization of polyhedral layering. The wheel size is analogous to the asteroid shape $\bar{u}_i$, and the color pattern is analogous to the density $\bar{\rho}_i$. By assuming concentric layering, the angular coordinates of vertices are consistent across layers.

to be subtracted, as shown in Fig. 1 and Eq. (17):

$$\begin{aligned}\bar{\mathbf{A}} &= (\mathbf{a}_1 - \mathbf{a}_2)\bar{\rho}_1 + (\mathbf{a}_2 - \mathbf{a}_3)\bar{\rho}_2 + \cdots \\ &\quad + (\mathbf{a}_{n-1} - \mathbf{a}_n)\bar{\rho}_{n-1} + \mathbf{a}_n\bar{\rho}_n\end{aligned}\quad (17)$$

$$\mathbf{a} = [(\mathbf{a}_1 - \mathbf{a}_2), (\mathbf{a}_2 - \mathbf{a}_3), \dots, (\mathbf{a}_{n-1} - \mathbf{a}_n), \mathbf{a}_n] \quad (18)$$

The acceleration ($\bar{\mathbf{A}}$) is the gradient of potential, whereas the gravity gradient ($\mathbf{GG}$) is the gradient of acceleration, and can be represented by Eq. (19) [8]:

$$\mathbf{GG} = -G \sum_{e \in \text{edges}} \mathbf{E}_e \cdot L_e \rho_e + G \sum_{f \in \text{facets}} \mathbf{F}_f \cdot \omega_f \rho_f \quad (19)$$

Again, the density is extracted via vectorization, resulting in the $\bar{\mathbf{g}}\bar{\mathbf{g}}$ term, which is a vector of 3×3 matrices.

$$\mathbf{GG} = \bar{\mathbf{g}}\bar{\mathbf{g}} \cdot \bar{\rho} = [\bar{\mathbf{g}}\bar{\mathbf{g}}_e, \bar{\mathbf{g}}\bar{\mathbf{g}}_f] \cdot [\bar{\rho}_e, \bar{\rho}_f] \quad (20)$$

$$\bar{\mathbf{g}}\bar{\mathbf{g}}_e = -GL_e \mathbf{E}_e \quad (21)$$

$$\bar{\mathbf{g}}\bar{\mathbf{g}}_f = G\omega_f \mathbf{F}_f \quad (22)$$

Each 3×3 $\mathbf{GG}$ matrix is then organized into a 6-element vector such that $\bar{\mathbf{g}}\bar{\mathbf{g}}'$ is a $6 \times n$ matrix. Below-diagonal components are omitted because the $\mathbf{GG}$ matrix is symmetric:

$$\begin{aligned}&\begin{bmatrix} GG_{xx} & GG_{xy} & GG_{xz} \\ GG_{yx} & GG_{yy} & GG_{yz} \\ GG_{zx} & GG_{zy} & GG_{zz} \end{bmatrix} \\ &\rightarrow [GG_{xx} \ GG_{xy} \ GG_{yy} \ GG_{xz} \ GG_{yz} \ GG_{zz}]^T \quad (23) \\ &\bar{\mathbf{g}}\bar{\mathbf{g}} \rightarrow \bar{\mathbf{g}}\bar{\mathbf{g}}' \quad (24)\end{aligned}$$

The procedure for layering is identical to the potential and acceleration:

$$\bar{\mathbf{g}}\bar{\mathbf{g}}' = [(\bar{\mathbf{g}}\bar{\mathbf{g}}'_1 - \bar{\mathbf{g}}\bar{\mathbf{g}}'_2), (\bar{\mathbf{g}}\bar{\mathbf{g}}'_2 - \bar{\mathbf{g}}\bar{\mathbf{g}}'_3), \dots, \bar{\mathbf{g}}\bar{\mathbf{g}}'_n] \quad (25)$$

Because the gravitational potential must satisfy Laplace's equation, only five of the six terms are independent. Nonetheless, in consideration of future stochastic analyses, all six components are simulated.

3 Procedure

3.1 Theory: mathematical recovery

This section presents the development of a method that takes a known shape model and gravity measurements as inputs, and outputs a density vector that, when paired with the shape model, reproduces the gravity measurements. This relationship is expressed via matrix multiplication:

$$\bar{\mathbf{X}} = \mathbf{x}\bar{\rho} \quad (26)$$

where $\bar{\mathbf{X}}$ is a gravity measurement, $\bar{\rho}$ is the density vector of the body, and $\mathbf{x}$ is the matrix function that relates the two vectors. In this case, the $\mathbf{x}$ matrix is a function of shape and position, as explained in Eqs. (7)–(9). The recovery procedure was chosen to process the gravity gradient as measurements to mimic the GOCE mission [5].

3.2 Observability and correlation

In Ref. [6], an analysis is presented regarding finite-cube and finite-sphere gravity models wherein the observability of internal density was estimated from on-orbit data. The method utilized a batch-least-squares filter that was applied to the range and range-rate measurements; no direct gravity measurements were considered. In recent years, innovations in sensor technology have opened the possibility of measuring the $\mathbf{GG}$ directly from small-scale spacecraft [7]. The same mathematical technique used by Park *et al.* [6] can now be applied to predict the effectiveness of this new gravimetry technique.

Assume that a gravity gradient measurement is taken at a point in space, and the measurement can be defined by a polyhedral gravity model. The equation that would predict the measurement is

$$\bar{\mathbf{GG}} = \bar{\mathbf{g}}\bar{\mathbf{g}}' \cdot \bar{\rho} = [\bar{\mathbf{g}}\bar{\mathbf{g}}'_e, \bar{\mathbf{g}}\bar{\mathbf{g}}'_f] \cdot [\bar{\rho}_e, \bar{\rho}_f] \quad (27)$$

Recall that $\bar{\mathbf{GG}}$ is the vectorized gravity gradient matrix, $\bar{\mathbf{g}}\bar{\mathbf{g}}'$ is a $6 \times n$ matrix of unwrapped $\bar{\mathbf{g}}\bar{\mathbf{g}}_{(e/f)}$ matrices for each edge and facet, and $\bar{\rho}$ is the density vector describing the body. Each matrix in the edge and facet vector is given in Eqs. (21) and (22), and is restructured according to Eq. (23). The recovery problem assumes that the position and gravity gradient can be measured. The position is used to calculate $\bar{\mathbf{g}}\bar{\mathbf{g}}'$, the gravity measurements are substituted for $\bar{\mathbf{GG}}$, and the goal of the problem is to find the density vector, $\bar{\rho}$ that most closely satisfies Eq. (27).

For a given position, the relationship between the density and gravity gradient is linear. Subsequently, partial derivatives with respect to the density are found simply to be

$$\frac{\partial \bar{\mathbf{GG}}}{\partial \bar{\rho}} = \bar{\mathbf{g}}\bar{\mathbf{g}}' \quad (28)$$

Then, following Park *et al.*, define

$$\mathbf{\Lambda} = \bar{\mathbf{g}}\bar{\mathbf{g}}'^T \bar{\mathbf{g}}\bar{\mathbf{g}}' \quad (29)$$

to obtain the covariance matrix:

$$\mathbf{P} = \mathbf{A}^{-1} \quad (30)$$

The covariance matrix given in Eq. (30) provides the correlations and standard deviations of elements in the density vector for a given set of measurement data. Analysis of these values informs predictions about the accuracy of a recovery attempt. For example, a high standard deviation indicates high uncertainty about the value of a density node, and a high correlation with other nodes implies that the gravitational effects caused by that node are difficult to distinguish from the effects of other nodes.

3.3 Adam optimizer

The problem of high-resolution polyhedral density reconstruction bears a certain rudimentary resemblance to the problem of training a neural network: a large number of parameters must be tuned to recreate a set of noisy data. Methods that have delivered the best performance in high-dimensional neural network training can be classified as stochastic gradient methods. One such algorithm, Adam, has quickly become one of the most popular methods since its introduction in 2014 [15]. By keeping the running average of the second moment of the gradient, Adam is able to scale the step size of each iterative cycle for empirically improved convergence. Certain situations result in slower convergence [16], and prior work has sought to modify the base algorithm to accommodate such cases [17]. In this investigation, the original formulation based on adaptive moment estimation is implemented. The core loop, Algorithm 1, is shown for reference. Note that the loop control parameter, $i_{\max}$, limits the procedure to a fixed number of iterative cycles. This allows the user to terminate the computation and investigate the initial behavior of the algorithm before investing significant time into a convergence-evaluated test case.

The stochastic gradient function, ∇f , is minimized by using the running averages of the mean, $\bar{m}$, and the uncentered variance, $\bar{v}$, to calculate an appropriate step size. Kingma and Ba draw an analogy between the ratio $\bar{m}_i/\sqrt{\bar{v}_i}$ and the signal-to-noise ratio (SNR) [15]. A smaller SNR (larger $\bar{v}$) indicates greater uncertainty in the direction of the true gradient, and the step size is appropriately scaled in line 10 of Algorithm 1.

The optimization problem in this study is formulated as a least-squares minimization. Because of its broad applicability, the framework for solving such problems is

Algorithm 1 Adam algorithm (reproduced for reference)

Input: α : Step size

Input: $\beta_1, \beta_2 \in [0, 1)$: Exponential decay rates

Input: $i_{\max}$: Maximum iterations

Input: $\nabla f(\bar{x})$: Gradient function

Input: $\bar{x}_0$: Initial guess

```

1:  $i \leftarrow 0$ 
2:  $\bar{m}_0 \leftarrow \bar{0}$ 
3:  $\bar{v}_0 \leftarrow \bar{0}$ 
4: while  $i < i_{\max}$  do
5:    $i \leftarrow i + 1$ 
6:    $\bar{g}_i \leftarrow \nabla f_i(\bar{x}_{i-1})$ 
7:    $\bar{m}_i \leftarrow \beta_1 \bar{m}_{i-1} + (1 - \beta_1) \bar{g}_i$ 
8:    $\bar{v}_i \leftarrow \beta_2 \bar{v}_{i-1} + (1 - \beta_2) \bar{g}_i^{[2]}$  (square bracket
    exponent indicates element-wise power)
9:    $\alpha_i \leftarrow \alpha \cdot \sqrt{1 - \beta_2^i} / (1 - \beta_1^i)$  (correct for
    initialization bias)
10:   $\bar{x}_i \leftarrow \bar{x}_{i-1} - \alpha_i \bar{m}_i / (\sqrt{\bar{v}_i} + \varepsilon)$  (element-wise
    square root and division)
11: end while
12: return  $\bar{x}_i$ 

```

written in the same C++ class as the main optimizer, and is set as the default mode. The objective of the least-squares optimization is

$$\min(\bar{\mathbf{y}} - \mathbf{A}\bar{\mathbf{x}})^2 \quad (31)$$

where $\bar{\mathbf{x}}$ is the design vector, $\bar{\mathbf{y}}$ is the data vector, and $\mathbf{A}$ relates the two, such that, for the optimal case, $\bar{\mathbf{y}} = \mathbf{A}\bar{\mathbf{x}}$. Adam operates using only the gradient, which is

$$\bar{\mathbf{g}}(\bar{\mathbf{x}}) = 2\mathbf{A}^T \mathbf{A}\bar{\mathbf{x}} - 2\mathbf{A}^T \bar{\mathbf{y}} \quad (32)$$

Batching is often utilized when training deep neural networks. In its most common form, batching consists of choosing a subset of data points to calculate the gradient function, performing a few iterative cycles and then randomly re-selecting those data points. This technique can prevent the optimizer from being trapped in the non-convex regions of the solution space. For the least-squares formulation, the data points are re-sampled to select rows of the $\mathbf{A}$ matrix and the corresponding elements in the $\bar{\mathbf{y}}$ vector, and substituting these truncated values in Eq. (32). Two parameters that are often used to control batching are the batch size and batch frequency. Each specification is fairly intuitive: the batch size refers to the number of points extracted from matrix $\mathbf{A}$ and vector $\bar{\mathbf{y}}$ in each sample, and the batch frequency determines how often the data are resampled.

3.4 Implementation

3.4.1 Gravity model

A strategic implementation process was structured around a sequence of events similar to that of an actual mission. Based on an object-oriented programming paradigm, a C++ class that contains the necessary procedures for gravity calculations and recovery was developed to assist with investigations. First, the target is defined by constructing and importing a polyhedral model. A polyhedron is fully defined by two sets of data: vertices and facets. Moreover, two derivative datasets are used to provide information that is relevant for Eq. (1): edges and facets-of-edge. Vertices are simply the indexed X – Y – Z coordinates of points on the surface of the body, with dimension $numVert \times 3$ in three-dimensional space. Each facet is a set of three vertex indices that form a counterclockwise triangle when viewed from outside the body, with dimension $numFacet \times 3$. The edge file is a list of all pairs of points that are connected in the facet definition, with dimensions of $numEdge \times 2$. Each edge borders two facets, and the facets-of-edge file defines both the index of the adjacent facets and the associated clockwise (CW) or counter-clockwise (CCW) direction of the edge relative to each facet. This file also has the dimension $numEdge \times 2$.

Considering the two-dimensional (2D) shape in Fig. 2, the coordinates of points 1–4 are defined in the vertex file in Table 1.

Table 1 Vertex file stores the coordinates of each vertex

	X	Y
1	1	1
2	2	3
3	6	2
4	7	4

The corresponding facet file is provided in Table 2.

Table 2 Facet file stores the indices of the vertices (V_1 , V_2 , and V_3) composing each facet

	V_1	V_2	V_3
1	1	3	2
2	2	3	4

The edge file (Table 3) defines all connected vertex pairs.

Table 3 Edge file stores the indices of the vertices (V_1 and V_2) comprising each edge

	V_1	V_2
1	1	2
2	1	3
3	2	3
4	2	4
5	3	4

Finally, the facets-of-edge file (Table 4) defines which two facets border each edge, along with their respective orientations (with x representing a facet not shown in Fig. 2).

Fig. 2 2D polyhedron example. Planar points $P1$ – $P4$ are used to define five edges and two facets. Tables 1–4 are used to define this shape in the C++ implementation.

Table 4 Facets-of-edge file stores the two facets adjacent to every edge, and their CCW or CW relationship

	CCW	CW
1	x	1
2	1	x
3	2	1
4	x	2
5	2	x

Each text file is imported into a matrix and used to initialize an object of the C++ class. The number of layers is simply passed to the object as an integer.

After importing the shape model, a density map is defined for each layer. This task is accomplished by using spherical harmonics to create deviations from the average density. A set of Stokes coefficients (C_l^m , S_l^m) is passed to the object for each layer of density, enabling different patterns of variation for each layer. Within each layer, the density is only a function of the latitude and longitude. The radial variation must be defined by selecting different patterns for each layer. The use of the spherical-harmonic-defined density enables the same density map to be tested on different shape models. In addition, this method enforces reasonable continuity and maintenance of the average. An option to explicitly define the density of each node via the vector input also exists.

Subsequently, the object undergoes an initialization process that calculates and stores certain parameters that are required for later calculations. These parameters include a point in each facet (PIF), a point on each edge (POE) as well as $\mathbf{F}_f$, $\mathbf{E}_e$, μ , and, if the spherical harmonic density is used, the explicit density of each node. POE is the midpoint of each edge, and PIF is the centroid of each facet. The matrices $\mathbf{E}_e$ and $\mathbf{F}_f$ are evaluated by employing Eqs. (2) and (3), respectively. The explicit density of each node is calculated using the spherical coordinates of that node via PIF and POE. The latitude (φ) and longitude (λ) of each PIF and POE are input into Eq. (33) along with the appropriate set of density coefficients.

$$\rho_i(\varphi, \lambda) = \rho_{i,\text{avg}} \sum_{l=0}^{\infty} \sum_{m=0}^1 P_l^m(\sin \varphi) \cdot [C_{l,i}^m \cos(m\lambda) + S_{l,i}^m \sin(m\lambda)] \quad (33)$$

The parameters $\rho_{i,\text{avg}}$, $C_{l,i}^m$, and $S_{l,i}^m$ are defined by the user to construct the body they wish to recover, and may be different for each layer. In this context, spherical harmonics are used as a scalar function on a sphere.

Recall the assumption that each interior layer is a scaled copy of the outermost; thus, the λ and φ angles are consistent for all layers. The discretization of λ and φ is a function of the shape model and may not be uniformly distributed.

3.4.2 Matrix inversion recovery

The $\mathbf{GG}$ and position data must now be recovered into a density vector. This procedure requires access to the shape definition of the central body. The core of the recovery is the solution to the linear algebra equation:

$$\mathbf{A}\bar{x} = \bar{b}$$

where $\bar{x}$ is density, and $\bar{b}$ is the $\mathbf{GG}$ measurement vector. The number of unknowns is the product of the number of layers with the sum of edges and facets.

First, the $\mathbf{A}$ matrix is constructed. Recall Eq. (27), where density extraction results in a vector of matrices $\bar{g}\bar{g}$. By unwrapping each matrix in $\bar{g}\bar{g}$ (Eqs. (21) and (22)), as shown by Eq. (23), an n -dimensional vector of symmetric 3×3 matrices becomes one $6 \times n$ matrix. The $\mathbf{GG}$ equation now resembles the acceleration formula with a height of 6 instead of 3:

$$\overline{\mathbf{GG}} = \bar{g}\bar{g}' \cdot \bar{\rho} \quad (34)$$

Note that the $\bar{g}\bar{g}'$ matrix is only a function of the position. Attempting to invert Eq. (34) with a single sample point is an underconstrained problem, given that the number of density nodes, even for a polyhedron of medium resolution, is much greater than 6. However, by accumulating multiple gravity samples into a larger matrix, the problem can be sufficiently constrained. For s samples, we have $\overline{\mathbf{GG}}_{(6s \times 1)}$ and $\bar{g}\bar{g}'_{(6s \times n)}$. Equation (35) outlines the fundamental requirements for the necessary number of sample points and the number of probes required to obtain those points:

$$s = p \cdot k \geq n/6 \quad (35)$$

where s is the total number of samples, p is the number of probes, k is the number of samples taken on each probe, and n is the number of density nodes of the body. As is apparent from Eq. (35), as the number of probes increases, a surplus of sample points is generated. Although the option to use all the points does exist, making the $\bar{g}\bar{g}'$ matrix very tall, full inclusion may significantly increase the computational time necessary to obtain the solution. Therefore, a parameter termed the overconstraint factor

(OCF), which defines the height-to-width ratio of the matrix, is introduced. If p is the number of density coefficients, and c is the overconstraint factor, the number of sample points, N , in the inversion is defined by

$$N = cp \quad (36)$$

The sample points are sorted by increasing the orbital radius, and are sequentially pulled into the inversion matrix until the overconstraint factor is met. Therefore, the most distant points may be unused. All data points in the matrix are assigned uniform weights. The linear algebra problem is solved by executing a QR factorization algorithm.

Further analysis is then performed on the calculated density solution. One immediate test, performed as a rough heuristic, was used to evaluate the validity of the density distribution. The C++ object is duplicated to preserve all settings, and the output of the density algorithm is set as its density map. Then, the potential of the original and duplicate objects are evaluated at the same point, and each is displayed on the screen for user interpretation. If the inversion is successful, the outputs will be identical. Although the test is not comprehensive, it provides a reasonable starting point for investigating unexpected outputs.

4 Results

The two shape models, an octahedron and a disturbed spheroid, shown in Fig. 3, were used to validate the algo-

rithm. Each body represents a different level of resolution and, therefore, a different level of computational complexity. Additionally, recall that the addition of layers causes a multiplicative increase in the dimensionality of the problem. In the following analysis, the bodies are used in the order of increasing fidelity. This choice allows the performance of the algorithm to be observed on relatively fast tests before committing significant resources for more intensive analyses. Both bodies are assumed to be non-rotating.

A range of potential scenarios were simulated by varying certain parameters listed in Table 5. The number of probes is varied to investigate the effects of ground track coverage and can directly affect the total number of sample points available in the inversion. The layers of all polyhedra are evenly spaced in the radial distance, and by considering more layers, the recovery procedure may return a more accurate representation of the central body. Additionally, the number of layers affects the dimensionality of the inversion problem. The overconstraint factor affects the number of samples that are accumulated in

Table 5 Range of parameter variation

Property	Range
Probe	1–25
Layer	1–3
Overconstraint factor	1–8
Ejection cone half-angle	22°–39°
Batch size	8–8000
Noise magnitude	0–10 m

(a) Octahedral body: 20 nodes per layer, 10 km semimajor axis, $\mu \approx 2.37 \times 10^{-4} \text{ km}^3/(\text{kg}\cdot\text{s}^2)$

(b) Disturbed spheroid body: 560 nodes per layer, 15 km semimajor axis, $\mu \approx 2.35 \times 10^{-3} \text{ km}^3/(\text{kg}\cdot\text{s}^2)$

Fig. 3 Polyhedral bodies used for validation.

the information matrix.

The inertial velocity of the probes is then calculated such that the vectors are evenly spaced around a cone of which the centerline connects the chief to the central body. The half-angle of the ejection cone is conservatively chosen to prevent the probes from impacting the central body, and is proportional to the maximum radius of the body. The resultant minimum altitudes range from 2 to 8 km, which is reasonable, because the probes are not recovered in this analysis.

Positional noise is assumed to follow a zero-mean Gaussian distribution, the magnitude of which is controlled by a user-defined standard deviation. All stochastic scenarios in the Results section have a positional standard deviation of 10 m. In reality, the positional noise depends on the navigation solution of the mission and can follow non-standard distribution curves. For the purposes of this study, the Gaussian distribution is a reasonable test case because the use thereof is consistently able to cause the matrix inversion technique to fail.

The gravity gradient measurements are assumed to originate from a system such as the tidal acceleration gravity gradiometer proposed by Carroll and Faber [7]. Unlike existing inversion techniques, this approach relies on on-board sensors as well as remote position determination. Additionally, even though we consider the position uncertainty directly, we currently assume perfect GG measurements. Recalling the matrix formulation of the inversion, $\overline{GG} = gg \cdot \bar{\rho}$, the uncertainty in the position perturbs the matrix gg , whereas the uncertainty in the gravity measurement affects the vector $\overline{GG}$.

The values in Table 6 are held constant for all simulations. The average density was selected to correspond

to the estimated density of asteroid Eros [18]. Recall that spherical harmonic coefficients are used to define density variation. To obtain a valid density map, the first 6×6 coefficients from the 50×50 degrees and order GEM-T3 Earth gravity model are utilized. In their normal context, these coefficients are used to define small gravitational perturbations; therefore, this set of coefficients also results in bounded variation from the mean density value. For the following tests, the sample rate of the probes was arbitrarily selected to be 0.1 Hz. This parameter directly affects the number of sample points available for a given simulation duration. After manual testing, the ejection velocity was selected to approximate the escape velocity. This choice, in conjunction with the cone half-angle, prevents the probes from impacting the body for the range of explored parameters.

Table 6 Constant parameters: the mean density [18], and the density Stokes coefficients [19, 20]

Property	Value	Unit
Mean density	2.67×10^{12}	kg/km ³
Density Stokes' coeffs.	GEM-T3 Earth model	N/A
Sample rate	0.1	Hz
Chief position	(50, 0, 0)	km
Ejection velocity	$1.4\sqrt{\mu/r_{\text{chief}}}$	km/s

A typical simulation for a multiple-probe flyby of a sample asteroid is illustrated in Fig. 4. In this case, five probes are ejected from a relative position of (50, 0, 0) km from the asteroid, and are propagated forward for 6 h. The integration scheme is an Adams–Bashforth 10th order method, and the asteroid is assumed to be static. In reality, an asteroid undergoes non-negligible rotation, but the static assumption helps to verify the algorithm

Fig. 4 Sample simulation using an Eros shape model. Five probes are ejected in a conical pattern and fly by the heterogeneous model of Eros.

and to isolate recovery behavior. The extension to the rotating case simply consists of a rotation transformation on the input data. The initial velocities for the probes are calculated to be evenly spaced around the body, resulting in more complete ground track coverage. The green and red trajectories are “behind” the body from the given perspective. Owing to the inherent dynamical nonlinearities, the probes do not maintain perfect spacing, as evidenced by their non-singular reconvergence on the right side of Fig. 4. However, they maintain sufficient spacing to achieve comprehensive ground track coverage, as was intended.

Before noise is applied to the data, the theoretical observability is analyzed using Park’s method, as detailed in Section 3.2. The standard deviations and the correlation matrix are saved in full precision and visualized later using applicable tools (MATLAB). Noise is sampled from a Gaussian distribution with zero mean, where the level of noise is defined by a standard deviation. A set is sampled from the defined distribution and is added to nominal trajectories.

For Adam to operate, the problem must be cast as a least-squares matrix formulation. Thus, the objective is

$$\min(\overline{\mathbf{G}\mathbf{G}} - \mathbf{g}\mathbf{g}' \cdot \bar{\rho})^2 \quad (37)$$

where $\overline{\mathbf{G}\mathbf{G}}$ is the measurement vector, $\bar{\rho}$ is the density vector, and $\mathbf{g}\mathbf{g}'$ is the system matrix defined from Eq. (27). The noise associated with the position is considered when calculating $\mathbf{g}\mathbf{g}'$, and the noise in the gravity measurements is included in $\overline{\mathbf{G}\mathbf{G}}$. Adam also requires an initial guess to start the procedure. In most cases, the average body density is given as the initial guess. This assumption is reasonable because the average density of small bodies can often be estimated from the Earth.

The optimization is initialized using the default values for exponential decay $\beta_1 = 0.9$ and $\beta_2 = 0.999$ [15]. The step size, $\alpha = 20$, is empirically evaluated, and is problem-specific (Kingma and Ba suggested an α of 0.001 for their application). Performance is quantified by calculating the L2-norm of the difference between the optimized density and the true density at each iterative cycle.

4.1 Simulation 1: octahedron, 1 layer

4.1.1 Noise-free case

In case one, a single probe flyby of a one-layer octahedron was performed. In this scenario, the reconstructed

densities were accurate to machine precision for all 20 nodes in the body. Although this preliminary investigation instills confidence in the concept and validity of the implementation, higher fidelity modeling is required to test the robustness of the strategy. Additionally, visual representations of the density are difficult to interpret for such a low-resolution shape.

The theoretical correlations for this case are given in Fig. 5. For each axis, node ID refers to the index of that node in the density vector, $\bar{\rho}$. Note that all diagonal components were set to zero to ensure that the color scale is more intuitive. The most evident difference is the emergence of a visual distinction between edges and facets. Figure 5(b) shows consistent low-magnitude negative correlations of edges with facets, and relatively small positive correlations between facets. The outliers within the facets, such as (13, 18), are likely caused by the symmetry of the body.

The average standard deviation for the 10-probe case is 5.57×10^5 g/cm³, an improvement of one order of magnitude over the average of the one-probe case of 4.64×10^6 g/cm³. However, because the average density of the body is 2.67 g/cm³, both deviations suggest that a recovery would be imprecise.

4.1.2 Adding positional noise

Preliminary investigation suggested that optimization with noise in the measurements may be nontrivial. However, adding up to 10 m in positional uncertainty resulted in convergence profiles that are nearly indistinguishable from the noise-free case. Zooming in to the level of Fig. 6 is required to discern the difference. Note the scaling of the y -axis. It is observed that in the 10 cm and 1 m cases, the optimization outperformed that of the clean case. However, this result may be misleading: if the optimization is continued, the noisy cases will overshoot the target and settle with a higher final error.

4.1.3 Parameter tuning

The concept of batching, as discussed in the section on the Adam Optimizer, is implemented in all simulations, and two variables are introduced that control the process: the batch size and batch frequency. In the previous analyses, a large number of samples were available for use in the optimization, but because the batch size was constant throughout, each iteration only utilized 800 samples. An increase in the batch size enables more samples to be used in each iteration. The effects of varying the batch

Fig. 5 Noise-free correlation comparison for case 1. The edges are top-left, and the facets are bottom-right, marked by bold lines. Symmetry is an intrinsic property of correlation matrices. Values on the diagonal are always equal to 1 by definition in all correlation matrices, but for intuitive coloration, these are plotted as zero.

Fig. 6 Case 1: final iterations of noise comparison. The effect of adding up to 10 m of error on the accuracy of the Adam algorithm after 50,000 iterative cycles is nearly negligible.

size are illustrated in Fig. 7. The drastic improvement in the convergence speed from a larger batch size can be explained by Kingma and Ba's SNR analogy: a more comprehensive sample of the data results in a consistent gradient vector.

If the gradient is more stable, the algorithm takes larger steps. However, certain trade-offs exist with respect to the runtime. Increasing the batch size is the same as increasing the height of the $\mathbf{A}$ matrix in Eq. (32); as such, the linear algebra operations may be computationally expensive.

Considering the standard deviation predictions from above, the result for the 8000 batch size is particularly

surprising. Despite the standard deviation being on the order of 10^5 g/cm³, the densities are recovered to an accuracy of approximately 6×10^{-4} g/cm³, a disparity of approximately 9 orders of magnitude.

4.2 Simulation 2: disturbed spheroid, 1 layer

4.2.1 Noise-free case

To further continue the analysis and increase the complexity, a "disturbed spheroid" replaces the octahedron as the central body. The spheroid has 560 density nodes per layer, 28 times the resolution of the octahedron. The simulations start with the same setup as the octahedron

Fig. 7 Case 1: batch size comparison. The inclusion of more data points per iteration results in more rapid convergence at the cost of robustness against non-convexity.

—one layer and one probe. The increased resolution and geometric irregularities drastically impact the reconstruction accuracy. A 3D representation of the orbit of the probe around the central body is represented in Fig. 8(a).

Figure 8(b) shows a lat/lon map of the true density for one layer, with the ground track of the probe superimposed in red. In the 3D graph, probes fly according to the arrows, and in lat/lon plots, probes start at the center and fly outwards. Recall that the ejection velocity is approximately equal to the escape velocity. Therefore, for the short duration of these simulations, orbital capture is not a concern. The minimum orbital radius for all disturbed spheroid simulations is approximately 23.7 km. Figures 8(c) and 8(d) have the same X (latitude) and Y (longitude) axes as Fig. 8(b), but with different Z -axis data (density or error). Figure 8(c) is the recovered density, and Fig. 8(d) shows the error between the recovered and actual density. This collection of graphs forms the standard layout for subsequent analyses as well. The color conventions for the density plots are as follows: red points are midpoints of edges, black points are midpoints of facets, blue points are vertices, yellow surface is for higher z -value, and blue surface is for lower z -value. Some points are unfortunately obscured by the 3D topography in the MATLAB plot; however, the colored interpolation surface is an important analysis tool. Specific points are given to demonstrate the resolution of the body.

With these details in mind, the oscillating pattern of the true density generated by the spherical harmonic definition is apparent in Fig. 8(b). These variations are not immediately evident in the reconstruction, although,

because Fig. 8(c) is dominated by an incongruous spike around $(-40, 10)$. The precise location of the spike does not seem to correspond with any particular feature of the polyhedron, but it is notable that it occurs on the hemisphere with no ground track coverage. Because of the apparent outlier in the reconstruction, the percent error plot, Fig. 8(d), is shown on a common logarithmic scale. A top-down visual inspection, Fig. 8(d), reveals an interesting phenomenon with respect to the ground track of the probe. The lowest errors occur directly below the closest pass of the probe. Inspecting the error magnitude via the color bar in Fig. 8(d) reveals that the error is too large for these data to be reliable. Three parameters may affect the error: the number of probes, data sample rate, and overconstraint factor. First, we investigated the effects of improved ground-track coverage by increasing the number of probes.

The correlation analysis can also provide insight into the poor accuracy of this recovery attempt and predict any improvements that may arise by increasing the number of probes. Analysis of theoretical observability using the same technique as case 1 produced a less intuitive output. The standard deviation and correlation coefficient for many nodes was returned as NaN, implying that these areas cannot be observed at all. These NaN nodes are illustrated as dark lines in Fig. 9(a). This unobservability is not surprising, considering the results in Fig. 8, but the complete absence of any apparent patterns was unexpected.

However, for the 10-probe case in Fig. 9(b), the body is fully observable with correlations that generally approx-

Fig. 8 Simulation: disturbed spheroid, 1 probe, 1 layer. Color conventions: red points = edges, black points = facets, blue points = vertices, yellow surface = higher z -value, and blue surface = lower z -value.

imate zero, predicting superior recovery performance. The average standard deviation for the 10-probe case is 6.46×10^{10} g/cm³, which is significantly larger than the mean of 2.67 g/cm³; however, previous results suggest that recovery may still be possible.

Keeping all other variables the same, the number of probes used in the recovery algorithm is increased to ten. With the increase in ground coverage, the simulation generates extremely promising results. The oscillating pattern is clearly visible in the reconstruction, Fig. 10(a), and a side-view visual comparison between the truth and result reveals the accurate reproduction of small graphical nuances. The error plot, Fig. 10(c), confirms the visual intuition, with large central zones vanishing completely.

This graphical artifact may be misleading because an error still exists even though it was not captured owing to truncation.

4.2.2 Adding positional noise

Given the promising performance of the matrix inversion technique in the noise-free case, the Adam optimizer is now employed in the same scenario with positional uncertainty. A batch size of 8000 was selected (based on the results of Fig. 7), and 10 m is used as the standard deviation of the positional uncertainty. In Fig. 11, the error is consistently reduced from the initial guess and is unaffected by the number of probes in the simulation. This optimization proceeded significantly slower than the octahedral case, with 50,000 iterative cycles resulting in

(a) Case 2: 1 probe correlation matrix. Dark regions indicate unobservable regions. The absence of any patterns indicates that accurate recovery is unlikely

(b) Case 2: 10 probe correlation matrix. The near-zero correlations when using 10 probes indicate that each node is distinguishable from the others

Fig. 9 Noise-free correlation comparison for case 2. Taken with Fig. 8, the improvement of the 10-probe case constitutes strong evidence of a relationship between the ground track coverage and recovery accuracy.

an approximately $3 \times 10^{-5} \text{ g/cm}^3$ improvement in the density estimate.

Because these analyses do not converge, a different simulation is necessary to understand the convergence behavior. Instead of using the average density as the initial guess, the correct answer (true density) is passed to the optimizer. The accuracy and convergence behavior are interpreted from Fig. 12 based on the extent to which the algorithm is able to successfully maintain the solution, given the noisy data.

It is evident that the number of probes has a stronger effect on the tail end of the optimization than at the outset. With the data from 20 probes, the Adam optimizer is able to maintain the solution to within $1 \times 10^{-6} \text{ g/cm}^3$, whereas the 5-probe simulation diverges much more rapidly. Noteworthy is that for any batch size less than the full dataset, the algorithm is not guaranteed to converge in the deterministic sense. Each iteration is based on a different random subset of the available data, which will cause the stochastic directional changes seen in Fig. 12. Because of these considerations, a limit in the number of iterative cycles is suggested as the stopping criterion instead of the step size tolerance.

4.3 Simulation 3: disturbed spheroid, 3 layers

4.3.1 Noise-free case

Next, multi-layer reconstruction of the central body was considered. Beginning with ten probes, a 3-layer ver-

sion of the disturbed spheroid was set as the central body. Density maps of the internal layers follow the same pattern as the outermost layer, but with a reduced magnitude. The amplitude reduction is defined by

$$[C_l, S_l] = [C_1, S_1]/5^{l-1} \quad (38)$$

where C and S are the Stokes coefficients, and l is the layer, with layer $l = 1$ being the outermost layer. Recall that the first 6×6 GEM-T3 Earth gravity model coefficients are used to define the outer-layer densities for these simulations. Inspecting the error as a function of layers, the outermost layer was found to have reduced in accuracy by two orders of magnitude relative to the single-layer body. The middle and innermost layers show a further increase in error beyond the acceptable range. The addition of five more probes incrementally improved the reconstruction for the outer and middle layers; however, the central layer is largely unaffected even in a scenario with 20 probes.

Recall that the overconstraint factor (OCF) controls the aspect ratio of the inversion matrix, and therefore affects the number of sample points considered in the analysis. By increasing the OCF from 4 to 8, the number of utilized sample points doubles. This modification is expected to have the same effect as increasing Adam's batch size parameter. For both the fifteen- and twenty-probe cases, increasing the overconstraint factor to 8 results in an error reduction of approximately an order

Fig. 10 Simulation: disturbed spheroid, 10 probes, 1 layer. First successful recovery of complex, irregular bodies.

of magnitude in layer 3. A full comparison of the average percent error is given in Table 7, which demonstrates the layer-3 average error reduction from 444% to 5.28%.

4.3.2 Adding positional noise

In the stochastic analysis, the correlation matrix again predicts large areas that are not observable. For the areas that returned a number, the average standard deviation was 2.18×10^{14} g/cm³; however, the results of case 1 suggest that a reasonably accurate answer can still be found. The NaN nodes are illustrated by dark lines in Fig. 13. Each layer of the polyhedron is represented by a third of the x - and y -axes, with the outermost layer shown on the top left. The concentration of dark lines in the bottom right implies that the inner layers are less

Table 7 Average magnitude of percent error for 3-layer disturbed spheroid. The combination of additional probes with a greater number of considered data points, via the OCF, results in a two-order-of-magnitude improvement in the central layer recovery accuracy

Probes	Layer	OCF: 4	OCF: 8
15	1	0.3009%	0.00094%
	2	16.6801%	0.2164%
	3	443.9877%	85.1575%
20	1	0.1195%	0.00057%
	2	4.2104%	0.0525%
	3	296.6921%	24.0605%
25	1	0.0310%	0.00022%
	2	1.9908%	0.0133%
	3	242.6755%	5.2795%

Fig. 11 1-Layer disturbed spheroid: effects of probe number on Adam optimization performance. No strong improvement is observed by adding more measurement data.

Fig. 12 1-Layer disturbed spheroid: convergence behavior. The availability of more comprehensive measurements allows the algorithm to maintain the solution more steadily.

discernible than the outer layer.

The results for this 3-layer case are consistent with the behavior of the 1-layer case. Error reduction is consistent, if slow, and the algorithm produces an improvement of approximately $2.5 \times 10^{-5} \text{ g/cm}^3$ over 50,000 iterative cycles. It is also apparent that the number of probes is not a major factor in the initial stages of recovery. The asymptotic shape of Fig. 14 has become a common theme for these initial plots.

The behavior close to the solution is also in good correspondence with previous results. The addition of more probes enables the solution to be more effectively maintained, whereas simulations with fewer probes tend

to diverge faster. In Fig. 15, unlike previous cases, returns beyond those of the 15-probe simulations appear to be diminishing.

5 Conclusions

This paper proposes a new technique for small body gravimetry and analysis. Based on recent developments in sensor technology, a method capable of recovering the internal density distribution of a polyhedral body using only the position and gravity gradient data was developed. A multi-layer gravity model was developed with reference to previous work, and an expression for the gravity gradient was derived as a linear function

Fig. 13 Case 3: correlations of 3-layer disturbed spheroid with 10 probes (dark blue = unobservable). The concentration of dark, unobservable regions in the bottom right coincides with the aforementioned results and Park *et al.*'s conclusion [6].

Fig. 14 3-Layer disturbed spheroid: effects of probe number on Adam optimization performance. Even for a more complex body, the number of probes has a minimal effect on the initial phases of optimization.

of the density. This formulation was then used as a basis for recovering the gravity model from simulated measurements. Both a GPU-based QR factorization inversion algorithm and a stochastic method from the field of deep learning were employed to solve the matrix problem.

The most significant conclusion from the research is the apparent disagreement between the theoretical observ-

ability of a body and the empirical performance of the Adam algorithm. Even with significant positional noise applied to the dataset, Adam was able to consistently move from a decent initial guess to a more accurate solution. In future, we intend investigating the disconnection between the predicted performance and test results, and aim to attempt to develop techniques and heuristics for more effective accuracy prediction.

Fig. 15 3-Layer disturbed spheroid: convergence behavior. The availability of more data points again helps to maintain the solution, but diminishing returns are observed beyond 15 probes.

Acknowledgements

The researcher would like to thank the University of Alabama Graduate Council for financially supporting this work. Additionally, Dr. Kieran A. Carroll, CTO of Gedex, is thanked for providing guidance at the 49th LPSC when this concept was still in development. A portion of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (No. 80NM0018D0004).

References

- [1] Liao-Mcpherson, D., Dunham, W. D., Kolmanovsky, I. Model predictive control strategies for constrained soft landing on an asteroid. In: Proceedings of the AIAA/AAS Astrodynamics Specialist Conference, **2016**: 2016-5507.
- [2] Atchison, J., Mitch, R., Rivkin, A. Swarm flyby gravimetry. Technical report. The Johns Hopkins University Applied Physics Laboratory, **2015**.
- [3] Scheeres, D. J., Khushalani, B., Werner, R. A. Estimating asteroid density distributions from shape and gravity information. *Planetary and Space Science*, **2000**, 48(10): 965–971.
- [4] Takahashi, Y., Scheeres, D. J. Morphology driven density distribution estimation for small bodies. *Icarus*, **2014**, 233: 179–193.
- [5] Ditmar, P., Klees, R., Kostenko, F. Fast and accurate computation of spherical harmonic coefficients from satellite gravity gradiometry data. *Journal of Geodesy*, **2003**, 76(11–12): 690–705.
- [6] Park, R. S., Werner, R. A., Bhaskaran, S. Estimating small-body gravity field from shape model and navigation data. *Journal of Guidance, Control, and Dynamics*, **2010**, 33(1): 212–221.
- [7] Carroll, K., Faber, D. Tidal acceleration gravity gradiometry for measuring asteroid gravity field from orbit. In: Proceedings of the 69th International Astronautical Congress, **2018**: IAC-18-A3.4B.3.
- [8] Werner, R. A., Scheeres, D. J. Exterior gravitation of a polyhedron derived and compared with harmonic and mascon gravitation representations of asteroid 4769 Castalia. *Celestial Mechanics and Dynamical Astronomy*, **1996**, 65(3): 313–344.
- [9] Kaasalainen, M. Optimization methods for asteroid lightcurve inversion: I. shape determination. *Icarus*, **2001**, 153(1): 24–36.
- [10] Shepard, M. K., Richardson, J., Taylor, P. A., Rodriguez-Ford, L. A., Conrad, A., de Pater, I., Adamkovics, M., de Kleer, K., Males, J. R., Morzinski, K. M. *et al.* Radar observations and shape model of asteroid 16 Psyche. *Icarus*, **2017**, 281: 388–403.
- [11] Bercovici, B., McMahon, J. W. Improved shape determination for autonomous state estimation. In: Proceedings of the 31st International Symposium on Space Technology and Science, **2017**.
- [12] Cocard, C., Kubota, T. Autonomous navigation near asteroid based on visual SLAM. In: Proceedings of the 23rd International Symposium on Space Flight Dynamics, **2012**.
- [13] Bercovici, B., McMahon, J. W. Point-cloud processing using modified Rodrigues parameters for relative navigation. *Journal of Guidance, Control, and Dynamics*, **2017**, 40(12): 3167–3179.

- [14] Takahashi, Y., Scheeres, D. J. Small body surface gravity fields via spherical harmonic expansions. *Celestial Mechanics and Dynamical Astronomy*, **2014**, 119(2): 169–206.
- [15] Kingma, D. P., Ba, J. Adam: A method for stochastic optimization. *arXiv preprint*, **2014**, arXiv:1412.6980.
- [16] Reddi, S. J., Kale, S., Kumar, S. On the convergence of Adam and beyond. In: Proceedings of the International Conference on Learning Representations, **2018**.
- [17] Dozat, T. Incorporating Nesterov momentum into Adam. **2015**.
- [18] Yeomans, D. K. Radio science results during the NEAR-shoemaker spacecraft rendezvous with Eros. *Science*, **2000**, 289(5487): 2085–2088.
- [19] Cazenave, A. *Geoid, Topography and Distribution of Landforms*. Washington, D. C.: American Geophysical Union, **1995**: 32–39.
- [20] Lerch, F. J., Nerem, R. S., Putney, B. H., Felsentreger, T. L., Sanchez, B. V., Klosko, S. M., Patel, G. B., Williamson, R. G., Chinn, D. S., Chan, J. C. *et al.* Geopotential models of the Earth from satellite tracking, altimeter and surface gravity observations: GEM-T3 and GEM-T3S. NASA Technical Memorandum 104555. NASA, Goddard Space Flight Center Greenbelt, MD, **1992**.

William Gordon Ledbetter is in his fourth year of Ph.D. study after receiving his bachelor of science degree in aerospace engineering from The University of Alabama, summa cum laude. He has been interested in space for as long as he can remember, and learned how to code by reading the TI-83 Plus manual. As the

capstone of his senior year he participated in design projects performing mission architecture analysis for proposed missions to Mars. William is currently exploring autonomous mission design for swarm exploration of asteroids, including landing and rendezvous. He is also investigating embedded systems for rapid, *in-situ* multi-agent optimal control. William received the 2017–2018 Graduate Council Fellowship and was named the 2019 Aerospace Engineering graduate student of the year at UA by the Engineering Council of Birmingham. William was selected as the 2019–2020 and 2020–2021 Alabama Space Grant Consortium Fellow. Beginning in the summer of 2020, William has an ongoing internship at Sandia National Labo-

ratories. E-mail: wgledbetter@crimson.ua.edu.

Rohan Sood is an assistant professor of aerospace engineering at the University of Alabama. He leads the Astrodynamics and Space Research Laboratory where his team investigates cost-effective, innovative, resilient trajectories and spacecraft behavior. Some of the research topics are asteroid exploration, solar sailing, multi-

body dynamics, spacecraft attitude control, and applications of artificial intelligence in astrodynamics. Dr. Sood received his bachelor of science degree in aerospace engineering from the University at Buffalo. He received his Ph.D. degree in aeronautics and astronautics with concentration in astrodynamics and space applications from Purdue University under the mentorship of Dr. Kathleen C. Howell and Dr. Henry (Jay) Melosh. Dr. Sood was a member of the GRAIL Science Team. He has since collaborated with NASA MSFC, ARC, and JPL on upcoming missions and to explore innovative mission concepts. Dr. Sood is a pianist, an artist (oil on canvas, charcoal), and enjoys participating in competitive sports. E-mail: rsood@eng.ua.edu.

James Tuttle Keane hails from Cedar Rapids, Iowa. He received his bachelor degree in astronomy and geology from the University of Maryland, College Park, and his Ph.D. degree in planetary science from the University of Arizona, under the mentorship of Dr. Isamu Matsuyama. After a postdoctoral position

at the California Institute of Technology, he started at JPL in 2020. Dr. Keane is a planetary scientist, studying the interactions between orbital dynamics, rotational dynamics, and geologic processes on rocky and icy worlds across the solar system. He uses a combination of theoretical methods, coupled with the analysis of spacecraft-derived datasets to investigate the dynamics, structure, origin, and evolution of solar system bodies. He has extensive experience with NASA missions, including GRAIL, New Horizons, and the proposed Io Volcano Observer. In addition to science, Dr. Keane is an accomplished science illustrator and communicator, specializing in pen and pencil illustrations of planetary science concepts, new results, and exciting missions. E-mail: james.t.keane@jpl.nasa.gov.

Jeffrey Stuart received his B.S., M.S., and Ph.D. degrees in aeronautics and astronautics from Purdue University in 2008, 2011, and 2014, respectively. He is currently technical staff in Mission Design & Navigation at the Jet Propulsion Laboratory, California Institute of Technology. He is the MDNav Lead for the SunRISE

formation flying mission and works on a variety of projects from early formulation to flight operations. His interests include automated trajectory design, advanced navigation techniques, formation flying, combinatorial optimization, and visual analytic methods. E-mail: Jeffrey.r.stuart@jpl.nasa.gov.

Open Access This article is licensed under a Creative Com-

mons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.