PREDICTING ASTEROID PHYSICAL PROPERTIES FROM PHOTOMETRIC DATA USING MACHINE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Accurately determining the physical parameters of asteroids, such as shape, rotation period, and rotation axis orientation, is crucial to understanding their structure and evolution, as well as to planning space missions. Since direct observation of the shapes of these objects is usually impossible, these parameters are determined indirectly through photometric observations. This paper describes a method to support this process using machine learning techniques. The final model developed – a two-branch convolutional neural network combining light-curve analysis with geometric metadata -— achieves a coefficient of determination (R^2) of about 0.40 on an independent test set, a result significantly above the random baseline.

1 INTRODUCTION

Photometric light curves are one of the primary sources of information about the physical properties of asteroids. The observed brightness variations, measured by ground-based telescopes and space observatories, depend mainly on an object's rotation and shape, the illumination and viewing geometry, and the light-scattering properties of its surface. Their analysis makes it possible to estimate an asteroid's rotation period, spin-axis orientation, and approximate shape (Kaasalainen & Torppa, 2001; Kaasalainen et al., 2001; Ďurech et al., 2020). These parameters are important for understanding the formation and evolution of small Solar System bodies and for characterizing asteroid populations.

The need for efficient photometric analysis is increasing with the development of large-scale sky surveys. In particular, the Legacy Survey of Space and Time (LSST), conducted by the Vera C. Rubin Observatory, is expected to provide multi-epoch photometric observations of millions of small Solar System bodies (Ivezić et al., 2019). This volume of data creates an opportunity to study the physical properties of asteroids at the population level, but it also requires methods capable of processing large numbers of sparsely and irregularly sampled observations automatically and efficiently.

Asteroid physical parameters are traditionally estimated through photometric inversion, in which synthetic light curves are fitted to observations. Existing methods differ, among other aspects, in how they represent an asteroid's shape. Simple representations such as triaxial ellipsoids enable efficient preliminary estimation of rotational parameters and overall axis ratios (Muinonen et al., 2015). Their symmetry, however, prevents them from capturing asymmetric features present in the light curves of irregular bodies. More flexible representations based on convex or non-convex surface meshes can describe asteroid shapes in greater detail, but generally introduce more free parameters and require richer observational data and greater computational effort (Kaasalainen & Torppa, 2001; Kaasalainen et al., 2001; Bartczak & Dudziński, 2018).

The Cellinoid shape model provides an intermediate representation between a triaxial ellipsoid and a high-dimensional surface model (Lu et al., 2014; Lu & Ip, 2015). It consists of eight ellipsoidal octants whose adjacent components share the corresponding semi-axes. The resulting model introduces only three additional shape parameters relative to a triaxial ellipsoid, while allowing global shape asymmetry to be represented. Previous studies have shown that the Cellinoid model can reproduce asymmetric light curves more accurately than a standard ellipsoid and can be applied to sparse photometric measurements (Lu & Ip, 2015; Lu et al., 2016).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Despite their physical interpretability, conventional inversion methods require an iterative search over a multidimensional parameter space. This optimization must be performed separately for each object, may converge to a local optimum, and can become computationally expensive as the number of observations and model parameters increases (Bartczak & Dudziński, 2018; Tang et al., 2025). These properties limit the number of asteroids for which physical parameters can be estimated efficiently and hinder homogeneous statistical analyses of large asteroid populations.

Machine learning offers a potential way to reduce this computational cost by replacing object-specific iterative optimization with direct inference after training. Deep learning has previously been applied to the classification of space objects from photometric light curves (Linares et al., 2020; Allworth et al., 2021). More recently, Tang et al. (2025) proposed a neural method that maps photometric observations and the corresponding observing geometries to three-dimensional point-cloud representations of asteroid shapes. Their approach reconstructs convex hulls, predicts potential concave regions, and achieves millisecond-scale inference. This result demonstrates the potential of learned mappings for accelerating asteroid shape inversion.

In this work, we investigate a complementary approach in which a neural network directly estimates a compact, physically interpretable Cellinoid representation from photometric observations. Unlike the point-cloud reconstruction method of Tang et al. (2025), our objective is not to predict a dense set of surface points, but to infer a low-dimensional set of parameters describing the asteroid's asymmetric shape and spin state. This compact output representation is intended to combine efficient inference with the ability to generate a physical surface model and its corresponding synthetic light curves. It is therefore particularly suitable for population-scale processing of sparse survey photometry.

Our main contributions are as follows:

- We introduce a two-branch convolutional neural network that combines photometric light curves with observing-geometry metadata to estimate the normalized parameters of a Cellinoid shape model.
- We employ a compact and physically interpretable Cellinoid representation that captures global asteroid-shape asymmetry using substantially fewer output variables than a dense mesh or point cloud.
- We construct a large-scale synthetic dataset containing Cellinoid models and their corresponding light curves generated under diverse illumination and viewing geometries.
- We evaluate the method using both a conventional train–test split and an independent test set containing previously unseen asteroid shapes. The proposed model achieves a coefficient of determination of approximately 0.40 on the independent test set, outperforming the random baseline.

2 Related Work

2.1 Photometric inversion of asteroid properties

Classical photometric inversion estimates an asteroid's shape and spin state by minimizing the discrepancy between observed and synthetic light curves. The optimization framework developed by Kaasalainen & Torppa (2001) and Kaasalainen et al. (2001) established a practical approach for recovering convex shape models, rotation periods, pole orientations, and scattering parameters from observations acquired under different illumination and viewing geometries. These techniques have subsequently been applied to both dense light curves and sparse survey photometry (Ďurech et al., 2020). The reliability of the recovered solution depends on the number and quality of observations, the coverage of rotational phase and observing geometry, the noise level, and the adopted scattering model.

2.2 Asteroid shape representations

Shape representations used in photometric inversion offer different trade-offs between expressiveness, identifiability, and computational cost. Triaxial ellipsoids provide a low-dimensional approximation and allow efficient initial estimation of rotation periods, pole orientations, and axis ratios

(Muinonen et al., 2015). Because they are symmetric, however, they cannot represent the global asymmetry of irregular asteroids.

Convex surface models provide greater flexibility while retaining useful stability properties in light-curve inversion (Kaasalainen & Torppa, 2001; Kaasalainen et al., 2001). Recovering concavities from disk-integrated photometry alone is substantially more difficult and may require high-quality measurements at informative phase angles or complementary observations. Non-convex methods such as SAGE use evolutionary optimization to estimate detailed shape models together with rotational parameters (Bartczak & Dudziński, 2018). This greater expressiveness is accompanied by a larger search space and higher computational cost.

The Cellinoid model lies between ellipsoidal and high-dimensional surface representations. Its shape is controlled by six semi-axes defining eight connected ellipsoidal octants (Lu et al., 2014). Consequently, it captures global asymmetry using only three more shape parameters than a triaxial ellipsoid. Lu & Ip (2015) applied this representation to multiple light curves observed under different geometries, while Lu et al. (2016) demonstrated its use with sparse photometric data. These properties make the Cellinoid a suitable candidate for low-dimensional, population-scale inference.

2.3 Learning from photometric light curves

Machine-learning methods have been used to extract object properties from photometric time series. Linares et al. (2020) trained convolutional neural networks to classify resident space objects from light-curve measurements. Allworth et al. (2021) studied shape classification of space debris and showed that transfer from simulated to observational light curves can improve performance on real data. These methods support the feasibility of learning informative representations from photometric measurements, but they address classification of artificial space objects rather than continuous estimation of asteroid shape and spin parameters.

The work most closely related to ours is that of Tang et al. (2025), who directly mapped brightness measurements and the corresponding illumination and viewing geometries to a three-dimensional point cloud. Their architecture uses convolutional and Transformer-based components to reconstruct an asteroid's convex hull and identify potential concave regions. The method was evaluated on synthetic data and on observations of two real asteroids, demonstrating fast inference and the feasibility of neural asteroid-shape reconstruction.

Our method shares the goal of replacing object-wise iterative optimization with direct neural inference, but differs in its output representation and intended use. Instead of reconstructing a high-dimensional point cloud, our method predicts a compact Cellinoid shape representation. The model is defined by six semi-axes (a_1, a_2, b_1, b_2, c_1, and c_2); however, because a_1 is fixed to unity to ensure scale invariance, the network estimates the remaining five normalized semi-axes. This representation provides a compact and physically interpretable description from which both the object's surface and its synthetic light curves can be generated. We further investigate the model's performance under sparse photometric sampling and assess its suitability for population-scale inference.

3 Dataset

3.1 Synthetic Data Generation and Shape Model

Because photometric inversion is highly ambiguous and analytical mappings from 1D light curves to 3D shapes do not exist, we employ a data-driven strategy. However, the number of real asteroids with confirmed, high-resolution ground-truth shapes is insufficient to train deep neural networks. Therefore, we constructed a large-scale synthetic dataset by generating random Cellinoid shapes and rendering their light curves under various observational geometries.

The Cellinoid is parameterized by six semi-axes $(a_1, a_2, b_1, b_2, c_1, c_2)$. To ensure scale invariance, we normalize the geometries by fixing $a_1 = 1.0$. The remaining five parameters are drawn uniformly from the range $[0.30, 0.99]$. The geometry of a single observational series is defined by the rotation axis orientation and the unit direction vectors pointing toward the Sun and the observer. The corresponding solar phase angles range from 5° to 45° (mean 29°).

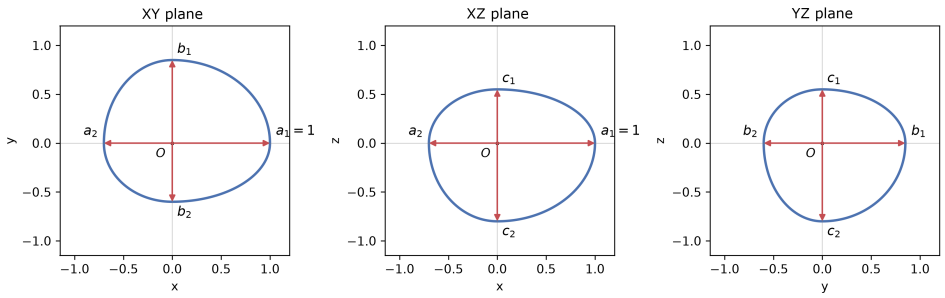


Figure 1: A cellinoid shape model depicted across its three primary cross-sections.

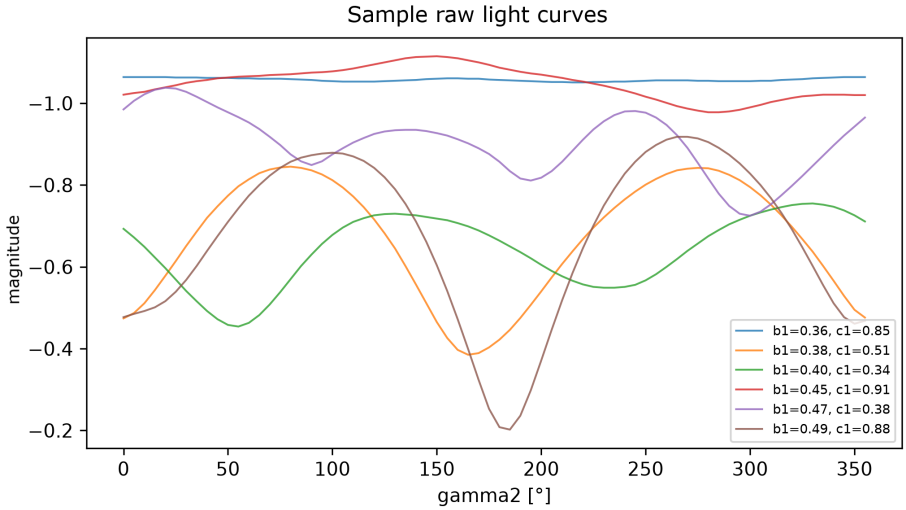


Figure 2: Sample synthetic light curves generated for different shape configurations.

3.2 Data Structure and Preprocessing

The complete dataset comprises over 10 million rows. Each fundamental unit is an observational series representing a single light curve at a fixed Sun-asteroid-observer geometry (detailed in Table 1).

While the raw data stores all pairwise rotation angle combinations (yielding $72 \times 71 = 5112$ rows per series to support differential photometry), our input pipeline extracts a standardized sequential series. By fixing a reference rotation angle, we retrieve the apparent magnitudes sampled at uniform 5° intervals over a full 360° rotation cycle. This transforms the tabular data into a continuous 1D tensor of dimension 72×1, which is fed directly into the convolutional network.

Frequency analysis of the generated light curves confirms their physical validity: the averaged Fourier spectrum is dominated by the second harmonic, directly reflecting the characteristic double-dip feature of elongated, rotating bodies.

4 Methodology

4.1 Input Representation and Network Architecture

To handle cyclical variables correctly, the object's rotation axis angles were encoded as sine and cosine pairs to prevent boundary discontinuities at $360^\circ/0^\circ$. The model thus receives two structurally distinct inputs: the ordered 72×1 light-curve signal, and a 9-element vector containing the encoded scalar geometric metadata.

Table 1: Input and Target Features of the Synthetic Dataset

PARAMETER	DESCRIPTION
`a2,b1,b2,c1,c2`	Target Cellinoid shape parameters ($a_1 = 1.0$)
`alpha, beta`	Ecliptic longitude/latitude of the rotation axis
`x_sun,y_sun,z_sun`	Sun-target unit vector components
`x_obs,y_obs,z_obs`	Observer-target unit vector components
`gamma1,gamma2`	Target rotation angles
`mag1,mag2`	Resulting apparent magnitudes
`phase`	Solar phase angle

To process these disparate data types, we designed a dual-branch architecture (Fig. 3). The first branch acts as a feature extractor for the light curve. It utilizes two 1D convolutional layers (64 and 128 filters, kernel sizes 5 and 3) with ReLU activations and Batch Normalization. A `GlobalAveragePooling1D` layer terminates this branch, reducing the feature maps to a 128-element vector. This global pooling mechanism averages the temporal signal, making the network highly invariant to the asteroid's initial rotational phase.

Concurrently, a dense branch processes the 9-element geometric metadata vector using a single 64-neuron fully connected layer equipped with Batch Normalization and Dropout. The outputs of both branches are concatenated into a 192-feature vector and passed through a deep classification block comprising two fully connected layers (256 and 128 neurons) with Dropout (0.2–0.4). The final linear layer outputs the five geometric shape parameters. The total parameter count is 111,237.

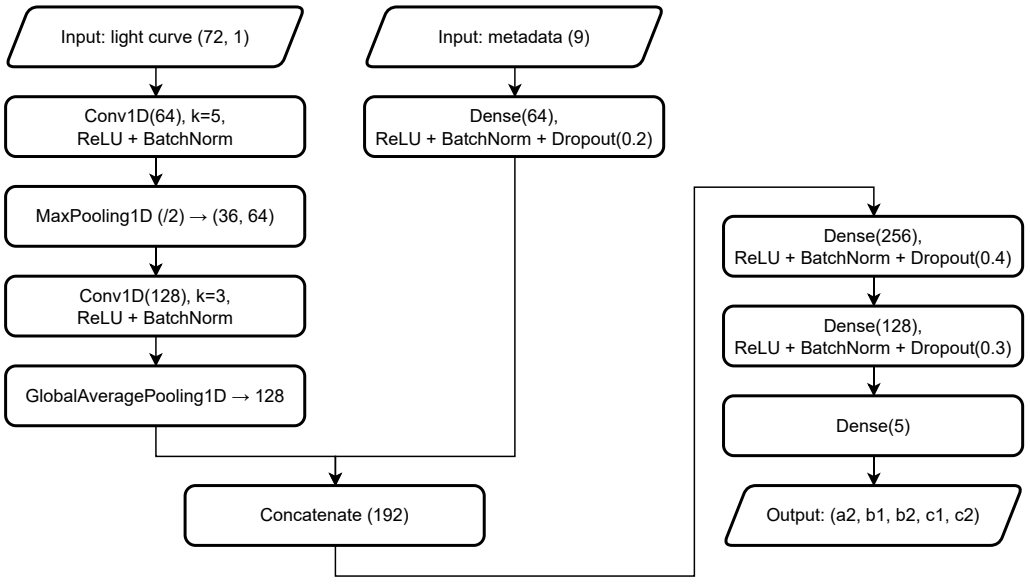


Figure 3: Dual-branch convolutional network architecture.

4.2 Training Setup

The model was optimized using the Adam optimizer with an initial learning rate of 10^{-3}, minimizing the Mean Squared Error (MSE). Training was scheduled for 150 epochs with a batch size of 32, using a 90:10 train-validation split. We employed Early Stopping (patience of 15 epochs) and learning rate decay (factor of 0.5 upon 5 epochs of stagnation, minimum 10^{-5}).

5 EXPERIMENTS AND RESULTS

5.1 EVALUATION PROTOCOLS

We evaluate the framework using two independent protocols: 1. **Generalization Benchmark:** An independent test set comprising 11 previously unseen asteroid shapes. Each shape was simulated across $\sim$95 random observational geometries, yielding 1,046 light curves. This tests the model's ability to infer the properties of entirely novel objects. 2. **Standard Benchmark:** A conventional 80:20 split of the main synthetic dataset, providing a baseline metric for large-sample prediction error.

Performance is assessed using Mean Squared Error (MSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2), calculated both globally and per-parameter.

5.2 SINGLE-LIGHTCURVE EVALUATION

On individual light curves from the independent test set, the model achieves a global $R^2 \approx 0.40$ and an MAE of roughly 0.08. This error margin indicates that the network successfully captures the overall proportions and general aspect ratios of the objects, even if fine-grained exactness remains challenging. The global R^2 score encompasses two sources of variance: the inherent prediction error of the neural network, and geometry-induced projection biases (where certain viewing angles mask shape details).

5.3 MULTI-GEOMETRY AGGREGATION

In actual astronomical workflows, true shape invariants are derived by combining observations of a target across multiple distinct epochs and geometries. To emulate this, we adopted a post-hoc multi-geometry aggregation scheme.

Independent predictions were generated for all available light curves corresponding to a single object ($\sim$95 geometries per shape). These individual output vectors were then averaged into a single consensus shape estimate. Because the CNN learns to extract underlying 3D invariants, individual predictions across diverse viewing angles cluster around the true parameters. This simple consensus averaging acts as a powerful noise-reduction filter, effectively canceling out single-view geometric projection biases. Consequently, evaluating the aggregated shape estimates yields a dramatic performance increase, pushing the R^2 score up to approximately 0.85 without requiring complex multi-sequence architectures during training.

5.4 PER-PARAMETER ANALYSIS

Decomposing the metrics reveals varying reconstruction accuracies among the five Cellinoid parameters (Fig. 4). The highest accuracy is observed for the equatorial dimension b_1 ($R^2 \approx 0.48$), whereas a_2 proves the most difficult to reconstruct ($R^2 \approx 0.32$).

This discrepancy aligns with the photometric properties of the data. Equatorial dimensions (b_1, b_2) strongly dictate the peak-to-peak amplitude of the light curve—a distinct feature easily captured by convolutions. Conversely, a_2, c_1, and c_2 exhibit weak correlations with light-curve morphology.

The low retrieval accuracy for c_1 and c_2 is fundamentally a physical constraint: these parameters define the object's extent along its spin axis (the Z-axis). For typical viewing geometries, the visible cross-section along the spin axis remains relatively constant during rotation, primarily influencing the absolute mean baseline of the signal rather than its dynamic shape. Because our input light curves are normalized, this absolute magnitude baseline is removed, effectively erasing the primary photometric signature of the Z-axis dimensions. When faced with this loss of signal, the network defaults to predicting values near the dataset mean, minimizing MSE but suppressing the R^2 score for these specific parameters.

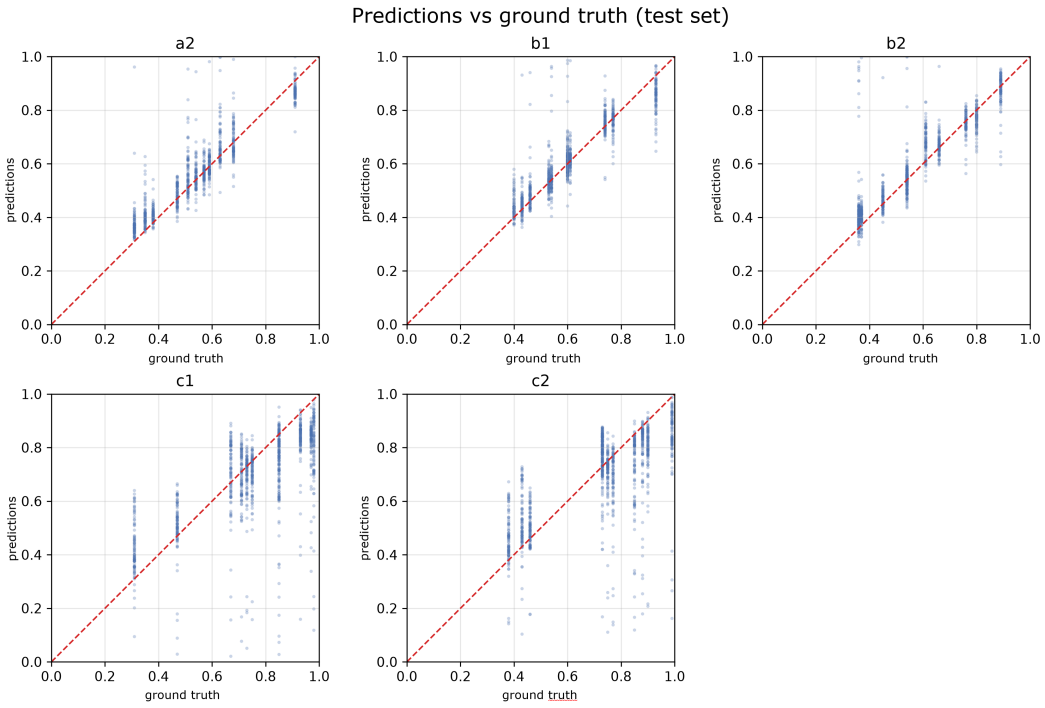


Figure 4: Predicted versus ground-truth values for the five shape parameters on the independent test set. The dashed line ($y = x$) represents an ideal fit.

5.5 Impact of Data Volume and Architecture Capacity

To systematically investigate whether the model's performance is limited by the volume of training data or by the fundamental observability of the physical parameters, we conducted scaling and architectural comparisons.

In the first experiment, the final CNN architecture was evaluated under two data regimes: a partial-data setting ($\sim$58,000 samples) and a full-data setting ($\sim$673,000 samples). As shown in Table 2, increasing the dataset size by a factor of 11 yielded only marginal improvements. The global R^2 increased by less than 0.03. This early plateauing, despite a 3.5-fold increase in training time, indicates that providing vastly more single-geometry examples does not unlock hidden information regarding the Z-axis dimensions (c_1, c_2), but merely fine-tunes the estimation of observable equatorial features.

Table 2: Impact of training dataset size on the final Dual-Branch CNN performance (single-lightcurve predictions).

METRIC	PARTIAL DATA	FULL DATA
Training Samples	58,731	673,607
MSE	0.0195	0.0190
RMSE	0.1396	0.1380
MAE	0.0837	0.0780
Global R^2	0.374	0.400

To ensure these saturation effects are specific to the problem's physical constraints and not a symptom of limited architectural capacity, we conducted a second experiment evaluating four models trained from scratch: a trivial mean baseline, a multi-layer perceptron (MLP) and a Random Forest operating on hand-crafted spectral features (Fourier harmonics), and our proposed Dual-Branch

CNN operating on raw light curves. Models were evaluated using an internal validation set (15% of shapes) and the external 11-shape test set.

Table 3: Comparison of different methods trained on the full dataset, averaged over three random seeds. Results are reported for single-curve (R^2) and aggregated per-shape predictions (R^2/shape) on the external test set.

Method	R^2 **(single-curve)**	R^2**/shape**	**Train Time [s]**	**Inf. Time [ms]**
Mean Baseline	-0.214	-0.212	0.005	0.00
MLP + FFT	-0.449	0.127	29	0.07
Random Forest + FFT	0.312	0.477	475	5.40
Dual-Branch CNN	0.290	0.835	$14,075$	0.12

The results (Table 3) validate the shift from spectral feature engineering to end-to-end deep learning. The earlier MLP iteration saturated almost immediately ($R^2/\text{shape} = 0.127$). While the Random Forest performed respectably on single curves, it failed to benefit significantly from multi-geometry aggregation ($R^2/\text{shape} = 0.477$).

In contrast, the proposed CNN strongly dominates the per-shape metric ($R^2/\text{shape} = 0.835$). While its single-curve R^2 shows variance depending on the random seed (averaging 0.290 ± 0.06, with the previously reported 0.40 representing a highly favorable run), the aggregated per-shape metric remains remarkably stable across runs (± 0.001). This confirms that the deep features extracted by the CNN encode geometrically consistent shape information that reliably converges to the true physical object when averaged across multiple views. Finally, despite longer training times, the CNN maintains near-instantaneous inference (0.12 ms per light curve), making it highly suitable for large-scale survey pipelines.

6 Conclusions

In this study, we proposed a dual-branch Convolutional Neural Network designed to invert asteroid light curves into physical shape parameters based on a convex Cellinoid model. Our primary conclusions are:

1. **Feasibility and Aggregation:** The proposed architecture successfully extracts geometric features from single, normalized light curves. While single-curve predictions yield a baseline $R^2 \approx 0.40$, aggregating these predictions across multiple observation geometries for a single object dramatically elevates performance to $R^2 \approx 0.85$. This proves the viability of deep learning as a rapid heuristic tool for initial shape estimation.
2. **Physical Observability:** Prediction accuracy is strictly linked to physical observability. Equatorial dimensions (b_1, b_2) are well-constrained by the light-curve amplitude and are reconstructed efficiently. Conversely, polar dimensions (c_1, c_2) exhibit weak correlations due to projection effects and the loss of absolute brightness baselines during normalization.
3. **Performance Bottlenecks:** The current ceiling in prediction quality is no longer dictated by the volume of training data, but by the inherent unobservability of Z-axis dimensions in isolated, normalized light curves, alongside the finite capacity of single-curve architectures.

While the size of the independent test set ($N = 11$) limits broad statistical claims, the substantial performance jump achieved via multi-geometry aggregation serves as a strong proof-of-concept. Future work will focus on validating this approach using real-world photometry and developing architectures capable of natively ingesting and jointly processing multiple light curves. Additionally, we plan to investigate novel normalization strategies to better constrain Z-axis dimensions and extend the network to reconstruct complex, non-convex asteroid shapes.

AI use statement

In this work, we have not used generative AI tools for generating synthetic data sets, cleaning and reformatting datasets, designing research methodology or experiments, implementing methods, or

interpreting results. Helping develop theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in writing proofs, proposing or refining hypotheses, assisting with translation, and supporting qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools for drafting parts of a research paper, suggesting a structure for a research paper, summarizing or analyzing existing literature, identifying relevant literature, editing a research paper to improve readability, proposing keywords for a research paper, creating or editing software code, suggesting experimental parameters, and brainstorming potential modeling approaches. We have reviewed all AI-assisted work. Specifically, all LLM-generated literature references and summaries were manually cross-checked against original publications, AI-generated code snippets and suggested experimental parameters were thoroughly tested and validated by the authors, and all AI-assisted text, structures, and keywords were evaluated and edited to ensure domain consistency and scientific precision. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Reproducibility statement

To ensure the reproducibility of our results, all synthetic data generation parameters, light-curve sampling intervals, and preprocessing steps are detailed in Section 3. The complete dual-branch network architecture, including layer specifications, loss functions, and optimization hyperparameters, is fully described in Section 4 and Section 5.

References

James Allworth, Lloyd Windrim, James Bennett, and Mitch Bryson. A transfer learning approach to space debris classification using observational light curve data. *Acta Astronautica*, 181:301–315, 2021. doi: 10.1016/j.actaastro.2021.01.048. URL https://doi.org/10.1016/j.actaastro.2021.01.048.

Przemysław Bartczak and Grzegorz Dudziński. Shaping asteroid models using genetic evolution (SAGE). *Monthly Notices of the Royal Astronomical Society*, 473(4):5050–5065, 2018. doi: 10.1093/mnras/stx2535. URL https://doi.org/10.1093/mnras/stx2535.

Josef Ďurech, Josef Hanuš, and Radim Vančo. Inversion of asteroid photometry from Gaia DR2 and the Lowell Observatory photometric database. *Astronomy & Astrophysics*, 643:A59, 2020. doi: 10.1051/0004-6361/202037729. URL https://doi.org/10.1051/0004-6361/202037729.

Željko Ivezić et al. LSST: From science drivers to reference design and anticipated data products. *The Astrophysical Journal*, 873(2):111, 2019. doi: 10.3847/1538-4357/ab042c. URL https://doi.org/10.3847/1538-4357/ab042c.

Mikko Kaasalainen and Johanna Torppa. Optimization methods for asteroid lightcurve inversion. i. shape determination. *Icarus*, 153(1):24–36, 2001. doi: 10.1006/icar.2001.6673. URL https://doi.org/10.1006/icar.2001.6673.

Mikko Kaasalainen, Johanna Torppa, and Karri Muinonen. Optimization methods for asteroid lightcurve inversion. ii. the complete inverse problem. *Icarus*, 153(1):37–51, 2001. doi: 10.1006/icar.2001.6674. URL https://doi.org/10.1006/icar.2001.6674.

Richard Linares, Roberto Furfaro, and Vishnu Reddy. Space objects classification via light-curve measurements using deep convolutional neural networks. *The Journal of the Astronautical Sciences*, 67(3):1063–1091, 2020. doi: 10.1007/s40295-019-00208-w. URL https://doi.org/10.1007/s40295-019-00208-w.

Xiao-Ping Lu and Wing-Huen Ip. Cellinoid shape model for multiple light curves. *Planetary and Space Science*, 108:31–40, 2015. doi: 10.1016/j.pss.2014.12.019. URL https://doi.org/10.1016/j.pss.2014.12.019.

Xiao-Ping Lu, Alberto Cellino, Daniel Hestroffer, and Wing-Huen Ip. Cellinoid shape model for Hipparcos data. *Icarus*, 267:24–33, 2016. doi: 10.1016/j.icarus.2015.12.006. URL https://doi.org/10.1016/j.icarus.2015.12.006.

Xiaoping Lu, Haibin Zhao, and Zhong You. Cellinoid shape model for asteroids. *Earth, Moon, and Planets*, 112(1–4):73–87, 2014. doi: 10.1007/s11038-014-9439-5. URL https://doi.org/10.1007/s11038-014-9439-5.

Karri Muinonen, Olli Wilkman, Alberto Cellino, Xiaobin Wang, and Yibo Wang. Asteroid lightcurve inversion with Lommel–Seeliger ellipsoids. *Planetary and Space Science*, 118:227–241, 2015. doi: 10.1016/j.pss.2015.09.005. URL https://doi.org/10.1016/j.pss.2015.09.005.

YiJun Tang, ChenChen Ying, ChengZhe Xia, XiaoMing Zhang, and XiaoJun Jiang. Asteroid shape inversion with light curves using deep learning. *Astronomy & Astrophysics*, 696:A55, 2025. doi: 10.1051/0004-6361/202452058. URL https://doi.org/10.1051/0004-6361/202452058.