ESA GNC & ICATT Conference 2026
28 Sep. – 2 Oct. 2026, Sevilla, Spain
https://atpi.eventsair.com/gncicatt/

Neural Correction of Center-of-Brightness for Robust Centroid Estimation in RAMSES' Farinella Mission

Umberto Mondini[1,*], Pietro Califano[1], Paolo Panicucci[1], Paolo Lunghi[1], and Francesco Topputo[1]

[1] *Politecnico di Milano, Via La Masa 34, Milan, 20156, Italy*
[*] *Corresponding author: umberto.mondini@polimi.it*

ABSTRACT

Within RAMSES, the Farinella CubeSat will approach and characterize near-Earth asteroid (99942) Apophis, requiring robust onboard centroid estimation for autonomous line-of-sight determination during proximity operations. The center of brightness is simple but degrades at high phase angles, where the visible surface is partly unilluminated, causing a systematic bias relative to the center of figure. Closed-form analytical corrections partly compensate but lose robustness for non-spherical shapes, extreme illumination, and unavailable or unreliable range information. Parametric fitting remains sensitive to a-priori shape uncertainty, while image-based convolutional networks exceed CubeSat resources and pose black-box verification and validation challenges. We propose NeuralCOB, a compact multilayer perceptron with 10,562 trainable parameters. Rather than the full image, it uses hand-crafted geometric and photometric features, namely phase and Sun-direction angles, bounding box, and brightness moments, mapped to an enhanced centroid estimate. It requires no externally supplied range. Across synthetic datasets spanning Sun–spacecraft–target geometries and Apophis shape realizations, NeuralCOB consistently outperforms the analytical centroid baseline, especially at high phase angles and for strongly non-spherical shapes where analytical assumptions fail. Explainability analyses show which inputs dominate the learned correction, providing evidence for flight qualification.
NeuralCOB will fly as a navigation experiment and as backup solution onboard Farinella.

1 Introduction

In April 2029 the near-Earth asteroid (99942) Apophis will fly past the Earth at an exceptionally small distance. The event is a rare chance to observe a potentially hazardous object from close range. The European Space Agency (ESA) has responded with the Rapid Apophis Mission for Space Safety (RAMSES) [1]. Within it, the Farinella CubeSat will approach Apophis and characterize it during proximity operations [2]. In this phase, light-time delay and limited ground contact rule out human-in-the-loop control, so navigation must run onboard. Autonomy of this kind rests on a reliable estimate of the target centroid in the image. From that centroid, the line-of-sight (LoS) direction to the asteroid is derived and passed to the navigation filter [3, 4].

The cheapest way to obtain that centroid is the center of brightness (COB). Its negligible computational cost makes it well suited to CubeSat hardware. The center of figure (COF) is here defined as the image projection of the target center of mass and the COB coincides with it only when the visible disk is fully and uniformly lit. The issue is that proximity operations rarely meet that condition. Indeed, at the high phase angles that dominate the mission trajectories, the lit portion collapses into a sunward crescent. The COB then drifts systematically towards the Sun, and the LoS measurement inherits a bias that degrades close-range navigation accuracy.

Compensating this offset is a long-standing problem, and each of the established methods gives up something that Farinella cannot afford. Analytical photocenter corrections predict the sunward shift in closed form from the

© *Politecnico di Milano*

phase angle and the apparent radius [5]. They are compact and transparent, and they underpin the analytical COB (ACOB) baseline flown on Farinella (Section 2.2). Their accuracy, however, depends on the assumption of a spherical body, a fixed reflectance law and an external range measurement. Parametric fitting functions replace that closed-form expression with correction curves fitted to the expected observation conditions [6, 7]. Yet they still fix the basis functions beforehand and calibrate only their weights, a limitation that constrained the recalibration options on the Milani mission [8].

A second family avoids the bias at its source. Limb-based schemes fit a conic to the illuminated edge and recover the body center analytically [9, 10]. They assume an ellipsoidal silhouette, which a small body at close range typically does not present. They also become ill-conditioned over the short lit arcs typical of high phase angles, which is where the bias is largest. Correlation against rendered references, as in the correlation with a Lambertian sphere (CLAMB) scheme adopted for Hera, tolerates illumination changes and needs no external range [11]. It still relies on a spherical Lambertian model, and its repeated two-dimensional FFTs carry a cost that is significant on a flight processor.

At the opposite extreme, convolutional neural networks (CNNs) acting on the full image drop the shape assumption entirely [12]. The price is a memory footprint and computational cost beyond CubeSat budgets [13]. In addition, these methods are opaque, which complicates the verification and validation of autonomous flight software [14]. For this reason, ESA is increasingly pushing onboard AI applications toward greater transparency [15]. Section 5 applies explainability techniques of this kind to NeuralCOB.

None of the approaches reviewed above combines the three properties the Farinella experiment requires: a correction that adapts to shape and illumination without external range or target models, a footprint suited to onboard execution, and transparency compatible with flight qualification.

This paper proposes *NeuralCOB*, a lightweight learned correction algorithm designed to fill that gap. It maps the biased COB onto an enhanced centroid that approximates the COF. Rather than consuming the full image, it operates on a handful of hand-crafted features. These encode both the illumination geometry and the observed brightness distribution, and thus blend physical prior knowledge with photometric descriptors such as the moments of brightness. A compact multilayer perceptron (MLP), small enough for the Farinella processor, turns the features into the centroiding correction.

Learning the mapping function instead of prescribing it is a deliberate choice. A classical parametric algorithm tunes the coefficients of predetermined basis functions. In a neural network, by contrast, the hidden units are themselves shaped by the data during training, so the functional form of the COB-to-COF map is inferred rather than imposed. This matters because the shape and reflectance of Apophis remain largely unknown at the time the image processing algorithm is developed. Should a refined shape model become available before the encounter, the network can be fine-tuned on updated data rather than redesigned.

A second difference concerns the range information. NeuralCOB draws the apparent-scale information it needs from the image itself, so it never requires an externally supplied range. This keeps it usable when range estimates are unreliable or absent.

2 Problem Formulation

2.1 Centroiding Measurement

The camera is modeled using the pinhole camera model with intrinsic matrix $\mathbf{K} \in \mathbb{R}^{3\times3}$, and returns a grayscale image $I : \Omega_{\text{img}} \to [0, I_{\max}]$ of the target, with Ω_{img} the detector pixel grid. A pixel $\mathbf{p} = (u, v)^\top$, where u is measured rightward and v downward from the upper-left detector corner, back-projects to a unit line-of-sight direction in the camera frame

$$\text{LoS}(\mathbf{p}) = \frac{\mathbf{K}^{-1}\tilde{\mathbf{p}}}{\left\|\mathbf{K}^{-1}\tilde{\mathbf{p}}\right\|}, \qquad \tilde{\mathbf{p}} = (u, v, 1)^\top. \tag{1}$$

Segmentation then separates the target from the background. The illuminated region is the set of pixels brighter than a detection threshold τ, that is $\Omega = \{\mathbf{p} \in \Omega_{\text{img}} : I(\mathbf{p}) \geq \tau\}$. The threshold is set adaptively by Otsu's method, which maximizes the between-class variance of the intensity histogram [16].
Over this region, the brightness-weighted raw moments are computed

$$M_{pq} = \sum_{\mathbf{p}\in\Omega} I(\mathbf{p})\, u^p v^q, \qquad p, q \in \{0, 1, 2, \dots\}, \tag{2}$$

and the COB follows as the intensity-weighted mean of the illuminated pixels

$$\mathbf{COB} = \frac{1}{M_{00}} (M_{10},\, M_{01})^{\top}. \tag{3}$$

The navigation filter, however, needs the direction to the target center of mass, that is the reference direction LoS(**COF**). Any centroiding scheme is then judged by the pointing error it induces,

$$\varepsilon(\mathbf{p}) = \arccos\big(\mathrm{LoS}(\mathbf{p})^{\top}\,\mathrm{LoS}(\mathbf{COF})\big). \tag{4}$$

The COB and the COF coincide only when the apparent disk is fully and symmetrically illuminated. The quantity governing this condition is the phase angle α, measured at the target between the Sun and the spacecraft. As α grows, the terminator sweeps across the visible disk. The illuminated region thins into a crescent leaning towards the Sun, and the brightness distribution loses the symmetry that tied the COB to the COF. To first order, the residual bias is a scalar displacement along the image-plane unit direction of the Sun $\hat{\mathbf{u}}_{\odot}$,

$$\boldsymbol{\beta} \;=\; \mathbf{COB} - \mathbf{COF} \;\approx\; \delta(\alpha, r_{\mathrm{app}}; \Theta)\, \hat{\mathbf{u}}_{\odot}. \tag{5}$$

The magnitude δ vanishes as $\alpha \to 0$ and grows monotonically with α. It also scales with the apparent radius r_{app} of the target in pixels, and depends on its shape and reflectance parameters Θ. The bias $\boldsymbol{\beta}$ propagates directly into Eq. (4), with nothing to attenuate it. The bias therefore corrupts the LoS most severely during close proximity operations, where the range is short, r_{app} is large and α is high.

2.2 The Analytical COB Baseline

ACOB, the baseline adopted onboard Farinella, models the bias in closed form. The target is approximated as a Lommel–Seeliger sphere, for which the sunward photocenter offset is taken linear in phase angle and apparent radius [5]. Since the crescent leans sunward, the offset is applied in the anti-sunward direction to recover the COF estimate

$$\widehat{\mathbf{COF}}_{\mathrm{ACOB}} \;=\; \mathbf{COB} \;-\; \hat{\delta}_{\mathrm{ACOB}}\, \hat{\mathbf{u}}_{\odot} \;=\; \mathbf{COB} \;-\; \kappa\, \alpha\, r_{\mathrm{app}}\, \hat{\mathbf{u}}_{\odot}, \tag{6}$$

with α in degrees, r_{app} in pixels, and κ a calibrated reflectance coefficient ($\kappa \approx 6.2 \times 10^{-3}$ deg^{-1}). We consider two ACOB variants. They share Eq. (6) and differ only in how the apparent radius is obtained

$$r_{\mathrm{app}}^{\mathrm{range}} = \frac{1}{\mathrm{IFOV}}\, \arctan\frac{R}{\rho}, \qquad r_{\mathrm{app}}^{\mathrm{blob}} = \tfrac{1}{2}\, a_{\mathrm{major}}. \tag{7}$$

The *range-based* variant relies on metric knowledge of the geometry: the a-priori mean radius R of the target, the spacecraft–target range ρ, and the average instantaneous field of view IFOV, that is, the angular size of a pixel. The *blob-based* variant instead reads the apparent radius directly from the segmented image, as half the major-axis length a_{major} of the segmented blob. It therefore requires neither range nor a-priori shape information.

Both variants inherit the limits of the spherical, linear-in-α model beneath them. The range-based form is only as accurate as the range solution and the assumed radius R, and it degrades when the true shape departs from the a-priori sphere. The blob-based form removes that dependency, but ties the correction to a_{major}, which itself shrinks and stops tracking the true apparent size as the crescent thins at high α. These shortcomings motivate the learned correction developed in Section 3.

3 NeuralCOB: Feature Design and Network Architecture

NeuralCOB replaces the closed-form correction of Eq. (6) with a learned one, which can represent the shape-dependent part of the bias that a spherical model cannot capture. As in ACOB, the correction is applied as a residual on the COB,

$$\widehat{\mathbf{COF}}_{\mathrm{NCOB}} \;=\; \mathbf{COB} \;+\; \boldsymbol{f}_{\theta}(\mathbf{x}), \qquad \boldsymbol{f}_{\theta}(\mathbf{x}) \in \mathbb{R}^2, \tag{8}$$

where $\boldsymbol{f}_{\theta}$ is a small multilayer perceptron (MLP) with parameters θ, evaluated on a normalized feature vector $\mathbf{x}$. Regressing the residual instead of the COF keeps the regression label small and centered on zero, and preserves the additive structure of Eq. (6). NeuralCOB is therefore a drop-in replacement for the ACOB correction term.

3.1 Input Features

The vector $\mathbf{x}$ is composed of two complementary families of hand-crafted quantities, which are summarized in Table 1. The first family describes the illumination geometry, and is obtained from the Sun sensor and the known attitude: it consists of the phase angle α and the image-plane Sun unit direction $\hat{\mathbf{u}}_\odot$. These are the two quantities on which the ACOB bias model is built, and they allow the network to modulate the correction with the illumination conditions.

The second family describes the observed brightness distribution, and is extracted from the segmented blob of Section 2.1. It consists of four quantities: the COB, which acts as the residual anchor of Eq. (8); the bounding box (bbox); the offset of the COB from the bbox center; and the normalized central moments μ_{pq} of Eq. (2) up to third order. The second moments provide the root-mean-square (RMS) radius $\sqrt{\mathrm{tr}}$, with $\mathrm{tr} = \mu_{20} + \mu_{02}$, which represents the apparent size of the brightness distribution relative to the blob extent, and the anisotropy pair $(\mu_{20} - \mu_{02})/\mathrm{tr}$ and $2\mu_{11}/\mathrm{tr}$, which encode the elongation and the tilt of the brightness distribution. The third moments provide the per-axis skewness, that is the photometric signature of the asymmetry of the illuminated crescent.

Every feature is brought to order unity by a deterministic rule fixed a priori, and never by dataset statistics. Most of these rules are geometric bounds. Pixel coordinates are normalized by the sensor size, the phase angle by 180 deg, and the RMS radius by half the bbox diagonal. For the COB–bbox offset, the normalization is half the bbox width and height, $0.5 \times \mathrm{bbox}_{wh}$. Since the COB always lies inside the bbox, the result is a scale-free fraction in $[-1, 1]$. The skewness is the exception, as it admits no natural geometric bound: it is divided by the fixed constant 0.6, which brings its typical magnitude to order unity. The use of fixed constants instead of the train-set mean and variance prevents data leakage and preserves the physical meaning of each input.

Feature	Role	Normalization
Phase angle α	Sun–target–spacecraft angle	$\alpha/180$ deg
Sun direction $\hat{\mathbf{u}}_\odot$	image-plane Sun unit vector	unit vector
Weighted COB	brightness centroid (anchor)	/ sensor size
Bounding box	region extent $[u_{\mathrm{bb}}, v_{\mathrm{bb}}, w_{\mathrm{bb}}, h_{\mathrm{bb}}]$	/ sensor size
COB–bbox offset	COB minus bbox center	/ $0.5 \times \mathrm{bbox}_{wh}$
RMS radius $\sqrt{\mathrm{tr}}$	brightness spread (apparent size)	/ $0.5 \times \mathrm{bbox}_{diag}$
Anisotropy	elongation and tilt of the blob	$O(1)$
Skewness (x, y)	third normalized moments	/ 0.6

Table 1 *NeuralCOB input features and their deterministic normalization.*

3.2 Image Characterization

All the features of Table 1 are produced by a single image-characterization routine, designed to run within the flight processor budget. The segmentation is based on the adaptive Otsu threshold τ, which is clamped to a fixed interval $[\tau_{\min}, \tau_{\max}]$ so that an all-dark or a saturated image cannot drive the threshold to a degenerate value. At high phase angles the illuminated crescent can break into two disconnected blobs. The routine therefore retains the two largest regions and merges their raw moments, while minimum-area and brightness gates reject spurious noise blobs. The descriptors extracted from a representative rendered image are shown in Figure 1.

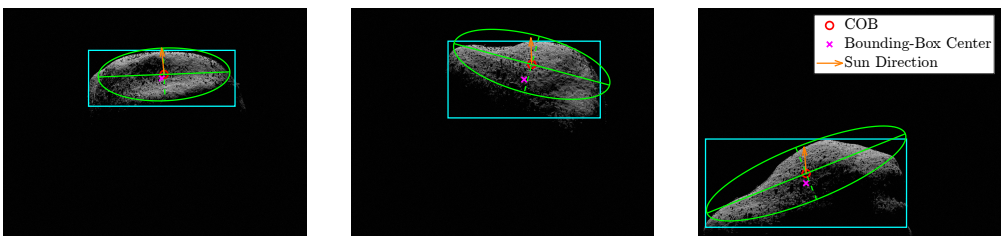


Figure 1 *Image characterization on a representative rendered frame.*

The efficiency of the routine follows from the way each region is analyzed. The routine does not label the connected components and integrate over them in a second pass. It instead grows each region from an above-threshold seed:

a rectangular scan expands the bounding frontier outward until all four of its edges fall below τ, accumulating the brightness-weighted moments M_{pq} of Eq. (2) as it goes. The reciprocal $1/M_{00}$ is computed once and cached, so that the moment normalizations reduce to multiplications. The central moments then provide the RMS radius, the anisotropy, and the ellipse major axis a_{major}. Since every buffer is statically sized, the routine requires no dynamic memory allocation, is compatible with code generation, and has a bounded worst-case execution time.

This front-end is not specific to NeuralCOB. It also provides the COB of Eq. (3) and the major axis a_{major} used by the blob-based ACOB of Eq. (7). The two methods therefore operate on the same segmentation and moment estimates, so that the benchmark of Section 6 isolates the correction stages, rather than differences in the upstream image processing.

3.3 Network and Training

The map $\boldsymbol{f}_\theta$ processes the 16-dimensional feature vector of Table 1 through two hidden layers of 128 and 64 units, each followed by a $tanh$ activation, and a linear head that produces the two components of the correction. Being smooth and bounded, the $tanh$ activation provides well-behaved gradients and maps cleanly onto the operations available on the flight processor. With 10,562 trainable parameters, the network occupies less than 43 KB in single precision. It is implemented in PyTorch [17] and exported to the Open Neural Network Exchange (ONNX) format [18] for onboard inference. At this size it fits within the Farinella processor budget. The exported computational graph is shown in Figure 2.

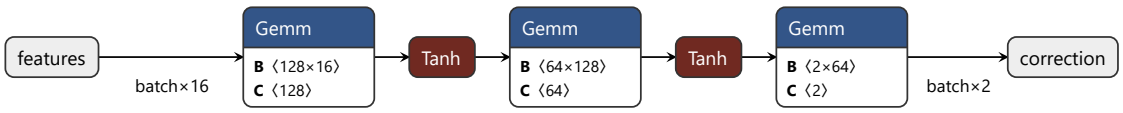


Figure 2 *Exported ONNX computational graph of the NeuralCOB network.*

The network is trained by minimizing a Huber loss between the reconstructed estimate $\widehat{\mathbf{COF}}_{\text{NCOB}}$ and the ground-truth COF, with a transition $\delta_H = 40$ px. Being quadratic below the transition and linear above it, this loss prevents occasional large errors from dominating the gradient [19]. The optimization is performed with AdamW [20] and a decoupled weight decay of 10^{-2}, over 1000 epochs with batches of 64. The learning rate is initialized at 10^{-3} and decreased to 10^{-6} by a cosine-annealing schedule [21]. Only the checkpoint with the best validation score is retained. Finally, the correction is made robust to the pointing error of the Sun sensor selected for Farinella. The two illumination features are perturbed with simulated sun-sensor noise, while the photometric features and the COB anchor are left intact. The perturbation is redrawn at every training step, so that the network sees a fresh realization of the illumination knowledge error each time a sample is presented. The corresponding noise model is described in Section 4.2.

4 Synthetic Dataset Generation

Supervised training of the correction requires a large set of images in which the COF is known exactly. Apophis has not yet been observed from close range, and even for bodies that have been visited the true centroid is not directly observable, but only reconstructed approximately from ephemerides. The dataset is therefore synthetic. Rendering offers a clear advantage: every scene is built from a small set of geometric quantities, so the COF is an input to the construction rather than a measured quantity, and is known exactly. The same quantities can then be swept deliberately to span the a-priori uncertainty in shape and illumination. The remainder of this section describes the generation pipeline.

4.1 Rendering and Camera Model

The images are generated with the physically based Cycles path tracer [22], configured for photorealistic output. The view transform is set to *Raw*, which disables tone mapping and gamma encoding, so that the stored pixel values are linear in radiance and consistent with the response of the flight detector. The target is illuminated by a single directional light source representing the Sun, against a black background representing empty space. The surface reflectance is Lambertian, with a base color given by a procedural noise texture, so that the albedo is spatially non-uniform.

The camera of Farinella is HORUS, whose design derives from the Milani NavCam [7]. Its intrinsic matrix $\mathbf{K}$ is that of the pinhole model of Section 2.1, with a 2048×1536 detector, 2.2 μm pixels, and a 32 deg horizontal field of view. Optical and detector effects are reproduced in post-processing, by applying a Gaussian point-spread blur ($\sigma \approx 2$ px) and adding Gaussian read noise of variance 5×10^{-4} in normalized intensity.

4.2 Scene Definition and A-priori Knowledge Error

A scene is defined by the true COF pixel location, the spacecraft–target range ρ, the phase angle α, the Sun azimuth ψ in the image plane, and the target attitude. The Sun is placed at 1 AU. Since the COF is an input, it provides the regression label directly, while the pair (α, ψ) determines the illumination geometry that drives the bias of Eq. (5).

The two dataset families differ in how these quantities are obtained. In the *random-sampling* family they are drawn independently and uniformly for each pose. The azimuth ψ spans the full circle $[0, 360)$ deg, and the COF is constrained to the box $u \in [151, 1898]$ px and $v \in [151, 1386]$ px, which corresponds to a margin of about 150 px from each detector edge. The only exception is the near-edge dataset of Section 4.4, in which the COF is deliberately placed outside this box. The range and the phase angle are not sampled over a common interval: their bounds are selected for each dataset to target a specific regime, and are reported in Table 2. In the *trajectory* family, the same quantities are obtained from the propagated spacecraft state and the Apophis ephemeris. Only the camera attitude is randomized around the nominal target-pointing direction, with a Gaussian boresight error of $1\sigma = 3$ deg and a Gaussian roll error of $1\sigma = 5$ deg, so that the center of figure is not systematically located at the image center.

These quantities define the true scene that is rendered. The a-priori geometry provided to the correctors is instead degraded by a realistic knowledge error, modeled as independent Gaussian perturbations of the range, the phase angle, and the Sun direction. The range error has $1\sigma = 20$ m, which is representative of the navigation-filter solution covariance, while the two illumination terms have $1\sigma = 3.67$ deg, the accuracy of the Sun sensor selected for Farinella. The range is used only by the range-based ACOB variant, whereas the blob-based variant and NeuralCOB receive the illumination terms alone. The rendered image and the COF label remain exact, so that the network is trained and evaluated against ground truth while its inputs carry the knowledge errors expected in flight.

4.3 Shape Realizations of Apophis

The a-priori shape of Apophis is uncertain, and a correction intended to outperform the spherical ACOB model must be exposed to this uncertainty during training. The base body is a shape reconstruction of Apophis [23], for which a mean radius of 170 m is assumed. Two mechanisms are used to generate departures from it.

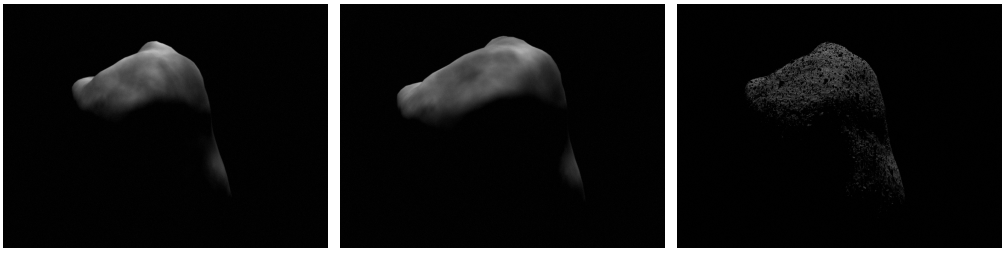


Figure 3 *Representative Apophis shape realizations used in the synthetic dataset. From left to right: the nominal shape model currently available, a sample from the tri-axial shape-variation set, and a realization produced by the asteroid-generation pipeline.*

The first mechanism is a parametric tri-axial variation. The body is described as a tri-axial ellipsoid with ordered semi-axes $a \geq b \geq c$, and the base mesh is rescaled anisotropically so that its effective diameter and elongation match sampled target values,

$$D = 2\,(a\,b\,c)^{1/3}, \qquad \eta_{ac} = \frac{a}{c}, \qquad \eta_{bc} = \frac{b}{c}, \tag{9}$$

where D is the mean diameter and η_{ac}, η_{bc} are the axis ratios. The three parameters are drawn from Gaussian distributions, reported as mean $\pm 1\sigma$: $D = 340 \pm 48$ m, $\eta_{ac} = 1.64 \pm 0.11$, and $\eta_{bc} = 1.14 \pm 0.10$. The longest extent of the body is mapped onto a and the shortest onto c, so that the rescaling preserves its natural elongation. The variation is applied in one of two modes. It is either redrawn at each pose, as in the random-sampling set

`dataset_006`, or arranged as a systematic pose × shape grid, as in the shape-sweep family `shape_mod_###`. All the other datasets, including training, validation, the held-out test trajectory, and the near-edge dataset, use the fixed base mesh.

The second mechanism, used for the highest-fidelity family (`asteroid_gen_###`), is a custom procedural asteroid-generation pipeline. It enriches a given shape model with multi-scale terrain and regolith detail, while preserving its global shape. The mesh is loop-subdivided and displaced with layered fractional Brownian motion (fBm) [24] and ridged noise, which introduce relief at macroscopic and microscopic scales. The mesh is then UV-unwrapped into a UDIM atlas. A procedural field of boulders is scattered over it, and albedo and bump maps are baked. The boulders are modeled as irregular three-dimensional Voronoi cells, optionally connected into clusters, and are distributed by a global planner that reserves non-overlapping surface regions according to coverage and size class. The resulting textured asset is rendered along a trajectory. Since these models violate both the ellipsoidal assumption of the tri-axial variation and the uniform-albedo spherical assumption of ACOB, they provide the most out-of-distribution appearance of the ensemble, and are reserved for the robustness assessment of Section 6. Figure 3 compares the three realizations at a common geometry.

4.4 Dataset Families, Naming and Split

The ensemble is composed of several datasets, which differ in their generation mechanism and in the regime they target, and which are reported in Table 2. Each name carries its family: `dataset` for independently sampled (i.i.d.) poses, and `trajectory` for temporally correlated poses taken from a propagated Sun-Stabilized Terminator Orbit (SSTO) [2] arc. The range and phase-angle intervals given in the table are the uniform sampling bounds for the random-sampling family, and the values obtained from the propagated arc for the trajectory family.

Dataset	Poses	ρ, km	α, deg	Shape	Role
`dataset_004`	10k	1.00–2.00	0–120	base mesh	Train, broad coverage
`dataset_005`	5k	1.00–2.00	60–120	base mesh	Train, high phase angle
`trajectory_004`	500	0.94–1.21	89–102	base mesh	Train, SSTO 1.0 km
`trajectory_005`	500	1.39–1.54	96–102	base mesh	Train, SSTO 1.4 km
`dataset_003`	3k	1.00–2.00	0–120	base mesh	Validation
`trajectory_002`	3k	1.13–1.26	95–98	base mesh	Test, SSTO 1.2 km
`dataset_007`	1k	0.80–2.00	30–120	base mesh	Robustness, near-edge COF
`dataset_006`	5k	1.00–2.00	60–120	tri-axial, per pose	Robustness, shape var.
`shape_mod_002`	100 × 30	1.00–2.00	60–120	tri-axial, 30-shape grid	Robustness, shape sweep
`asteroid_gen_002`	3k	1.13–1.26	95–98	procedural	Robustness, terrain & boulders

Table 2 *Catalog of the datasets used in this work.*

The split is designed to assess generalization rather than memorization. The training set combines two random-sampling datasets with two propagated trajectories. The random-sampling datasets provide broad and uniform coverage of the observation geometry, and one of them is dedicated to high phase angles. A disjoint random-sampling dataset is used as validation set for model selection. The test set is a held-out propagated trajectory, `trajectory_002`, which is never seen during training or validation and therefore provides an independent estimate of the flight-like performance.

Although the three SSTO arcs belong to the same family, they do not originate from a single propagation. Each is propagated from its own initial condition, targets its own osculating circular radius, and covers its own epoch window: `trajectory_005` at 1.4 km, from 2029-03-12 to 2029-03-22; `trajectory_002` at 1.2 km, from 2029-03-20 to 2029-03-31; and `trajectory_004` at 1.0 km, from 2029-04-20 to 2029-05-01. Their orbit planes, phasing, and body-fixed illumination histories are therefore distinct, and the test arc is not an interpolation of the training ones. The sampling step also differs, being 30 min on the training arcs and 5 min on the test arc, so that no pose is shared and the body is observed at different rotational phases.

The extreme-regime, shape-sweep, and procedural families are excluded from training entirely. They are used only for the robustness benchmark of Section 6, in which NeuralCOB is evaluated under conditions well outside its training distribution.

5 Explainability Analysis

Qualifying a neural network for flight requires more than a measurement of its accuracy. The rule that the network has learned must itself be made explicit, and checked against the physics it is expected to reproduce. Two properties of NeuralCOB make such evidence affordable to obtain: the input is the low-dimensional, deterministically normalized and bounded feature vector of Table 1, and the map $\boldsymbol{f}_\theta$ has only 10,562 parameters.

The analysis is carried out on the random-sampling validation set of 3,000 poses (`dataset_003` in Table 2), whose uniform and independent sampling of the observation geometry spans a wider input support than any single trajectory. Three complementary tools are applied. Per-pose attribution is obtained with exact interventional Shapley additive explanations (SHAP), dataset-level response shapes with first-order accumulated local effects (ALE), and pairwise interactions with second-order ALE. All the attributions are expressed in pixels, that is in the units of the correction itself.

5.1 Attribution Methods

SHAP quantifies, for a single pose, the contribution of each feature to its correction relative to a baseline. The Shapley value ϕ_i of a feature i is its marginal contribution to the model output, averaged over every ordering in which the inputs can be added [25, 26]. The contribution of a subset of features is taken here as an *interventional* expectation: the features in the subset are held at their values in the pose under analysis, while the remaining ones are drawn from a background set. With only $d = 16$ inputs, all $2^d = 65{,}536$ coalitions are enumerated exactly, so no sampling approximation enters the attributions. They satisfy the additivity property $\boldsymbol{f}_\theta(\mathbf{x}) = \mathbb{E}[\boldsymbol{f}_\theta] + \sum_i \phi_i(\mathbf{x})$ to a residual of about 10^{-13} px, which verifies the implementation. The values are computed per correction component, and averaged as $\overline{|\phi_i|}$ to obtain a global ranking.

First-order ALE describes how the correction changes when a single feature is varied. Partial-dependence importance is biased when the inputs are correlated, which is the case here. ALE instead takes local differences of the model within bins and accumulates them, which leaves the resulting curve $f_{\mathrm{ALE},j}$ unbiased under such correlation [27]. The curve is estimated over 150 quantile bins using all 3,000 poses, with the conditional derivative of the model replaced by a finite difference between the bin edges, and is centered to zero mean. Its mean absolute value $\overline{|f_{\mathrm{ALE},j}|}$ measures the main effect of the feature.

Second-order ALE establishes whether two features act together beyond their separate effects. The pure interaction of a pair (j, k) is the part of the response that the two main effects leave unexplained. On each cell of a two-dimensional quantile grid, it is estimated from the mixed second difference of the model at the four corners and accumulated in both directions. Subtracting the two first-order effects and centering the surface leaves a flat, identically zero surface whenever the pair is purely additive [27]. All the $\binom{16}{2} = 120$ feature pairs are evaluated on a 20×20 quantile grid, and ranked by the mean absolute surface value.

5.2 Learned Correction Structure

The three tools agree on a physically coherent picture, in which each correction component is driven mainly by the same-axis inputs (Figure 4). SHAP and ALE rank the same inputs highest, and the inputs where the two rankings diverge are those acting through interactions rather than through a main effect.

Illumination direction. For each correction component, the most important input is the same-axis projected Sun direction, according to both SHAP and ALE ($\overline{|\phi|} \approx 58$ and 71 px for the x and y corrections). Its first-order ALE curve, on the left of Figure 5, is almost antisymmetric, and sweeps the same-axis correction from about $+105$ px at $\hat{u}_\odot = -1$ to about -100 px at $\hat{u}_\odot = +1$, while the cross-axis response remains flat. This matches the expected physics: the COB is pulled towards the illuminated limb, so the correction that recovers the COF points *opposite* to the Sun, along the illumination axis. Thus, without this structure being explicitly imposed, the network recovers the geometry of the phase-angle bias, $\boldsymbol{\beta} \approx \delta\, \hat{\mathbf{u}}_\odot$ from Eq. (5), and learns its negative as the correction.

Phase angle gates the illumination correction. SHAP ranks the phase angle third ($\approx$ 20–21 px), and yet its first-order ALE curve is nearly flat, of a few pixels only. The apparent contradiction is resolved by the second-order analysis: among all the 120 pairs, the phase × same-axis-Sun pair is by far the strongest interaction (Figure 6, $\approx$ 20 and 21 px, more than twice the next one), with a sign-changing surface. The phase angle therefore acts almost entirely by modulating the magnitude of the Sun-direction correction, rather than through a standalone offset. This is the data-driven counterpart of the ACOB law of Eq. (6), in which the bias grows with the phase angle. The network has thus recovered not only the direction of the analytical bias model, but also its dependence on the

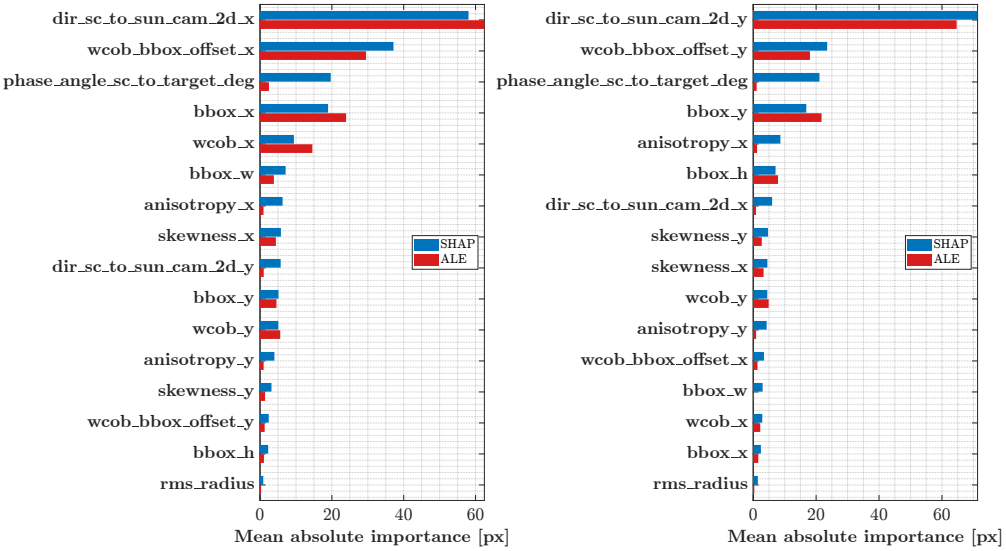


Figure 4 *Global importance of every input for the two correction components (x left, y right): paired mean absolute SHAP (blue) and first-order ALE (red) bars, in pixels, on the* 3,000*-pose validation set, sorted by importance.*

illumination geometry.

Data-driven geometric refinement. Beyond the illumination, the network exploits geometry that the closed-form ACOB cannot represent. The COB–bbox offset is the second most important input, and enters as a smooth and monotone same-axis correction, shown on the right of Figure 5: from about +70 to −75 px for x, and ±46 px for y. The relative position of the brightness centroid within the blob is therefore read directly as a recentering signal. The bounding-box position and size add a nonlinear location and scale calibration. The higher-order shape descriptors, anisotropy and skewness, contribute little on average but become significant in the interactions and in the tails of their distributions. The correction is therefore shape-dependent.

Taken together, the three tools show that the learned correction is dominated by an interpretable and physically grounded illumination rule, which recovers the analytical bias structure of Section 2, and is augmented by data-driven geometric refinements.

6 Results and Benchmarking

6.1 Evaluation Protocol

All the algorithms are evaluated on the same rendered poses and receive the a-priori geometry degraded by the knowledge error of Section 4.2. Four estimators span the information available onboard. The uncorrected COB of Eq. (3) is the reference against which any improvement is measured, and it sizes the phase-angle bias that the correction must remove. The two ACOB variants of Eq. (7) follow, the blob-based one requiring neither range nor a-priori shape, the range-based one requiring both. The metric range granted to the latter is itself uncertain onboard, since it arrives either as the noisy estimate of the navigation filter, or from a lidar, which is available only when the target lies within its design range.

Scoring proceeds through the pinhole model of Eq. (1), which back-projects the estimated centroid to a unit line-of-sight, and then through the angular pointing error ε of Eq. (4). It is evaluated at $\mathbf{p} = \widehat{\mathbf{COF}}$ against LoS(**COF**). The error distribution is heavy-tailed over every test set, so the mean $\bar{\varepsilon}$ and the median ε_{50} are reported together with the upper quantiles ε_{95} and ε_{99}, collected for every corrector and every set in Table 4.

That angle is a magnitude: it measures how large the error is and not where it points, so a given value is consistent with any split between a persistent offset and a random dispersion. The correction acts along one specific direction, the sunward offset of Eq. (5), so the signed residual of the estimate is reported as well,

$$\mathbf{r} = \widehat{\mathbf{COF}} - \mathbf{COF}, \qquad r_{\parallel} = \mathbf{r} \cdot \hat{\mathbf{u}}_{\odot}, \qquad r_{\perp} = \mathbf{r} \cdot \hat{\mathbf{u}}_{\odot}^{\perp}, \tag{10}$$

resolved in pixels along the Sun image-plane direction and its in-plane normal. A positive $r_{\parallel}$ means that the

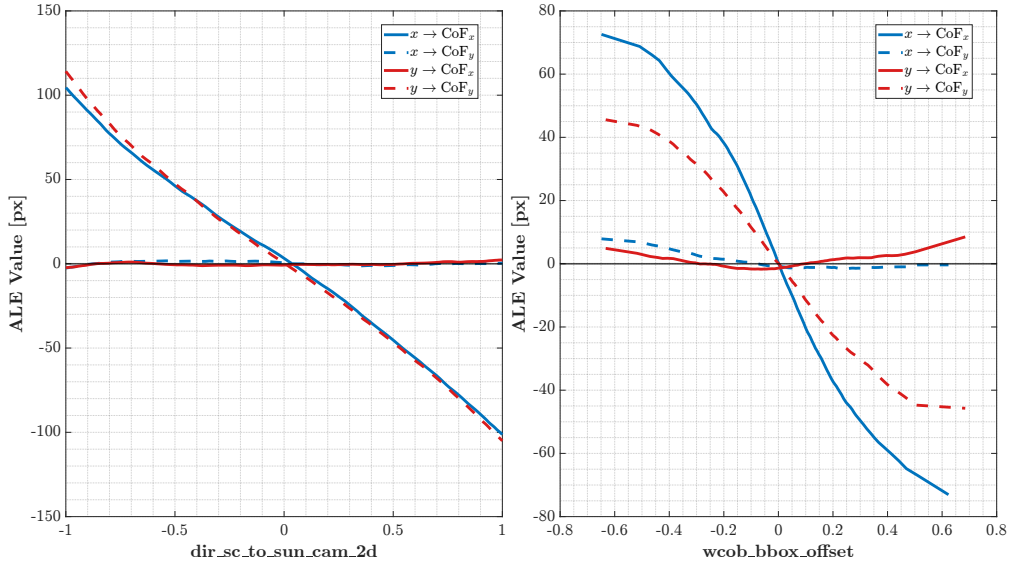


Figure 5 *First-order ALE of the correction against the projected Sun direction (left) and the COB–bbox offset (right), on normalized feature axes. Color marks the input component (x blue, y red) and line style the output component (x solid, y dashed).*

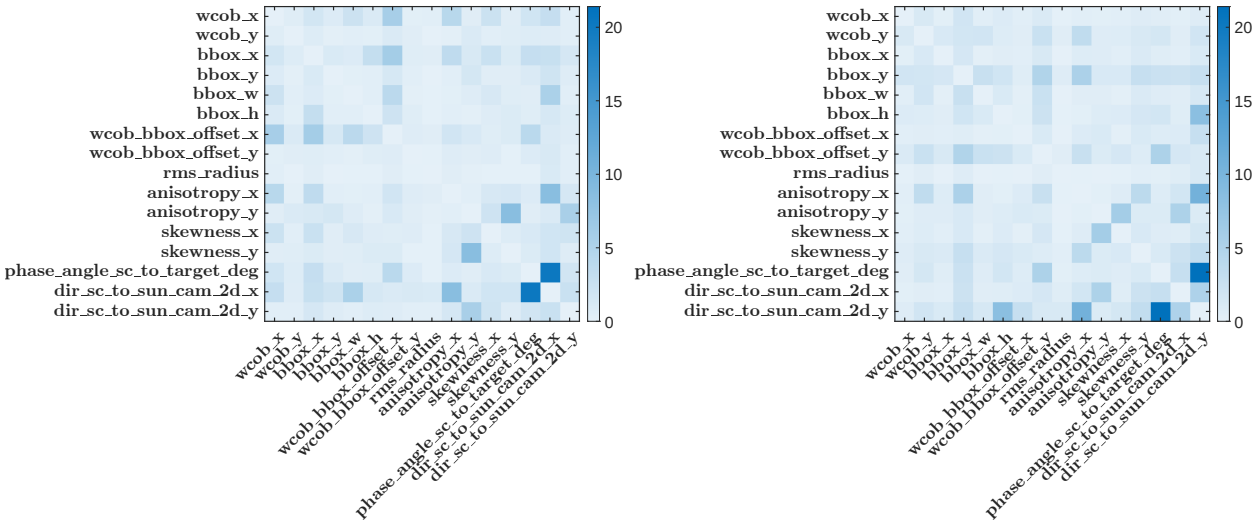


Figure 6 *Second-order ALE interaction strength over all the feature pairs, for the two correction components.*

estimate still lies sunward of the true COF. What a navigation filter accumulates is not the offset of a single frame, but the part of the residual shared by a whole sequence of frames. Here the *bias* is defined as the ensemble mean $\mathbb{E}[r_{\parallel}]$ over a test set, and $\sigma_{\parallel}, \sigma_{\perp}$ as the random dispersion about it.

6.2 Nominal Performance

Two nominal-shape sets bracket in-distribution and held-out performance. The random-sampling validation set `dataset_003` covers the observation geometry uniformly, with the COF over the full detector, the range over 1.0–2.0 km, the phase angle over 0–120 deg, and the Sun azimuth over the whole circle, so it probes the envelope of proximity-operations geometries rather than any single arc. On it, NeuralCOB is the most accurate corrector at every quantile of Table 4, cutting the median error of the uncorrected COB by two thirds.

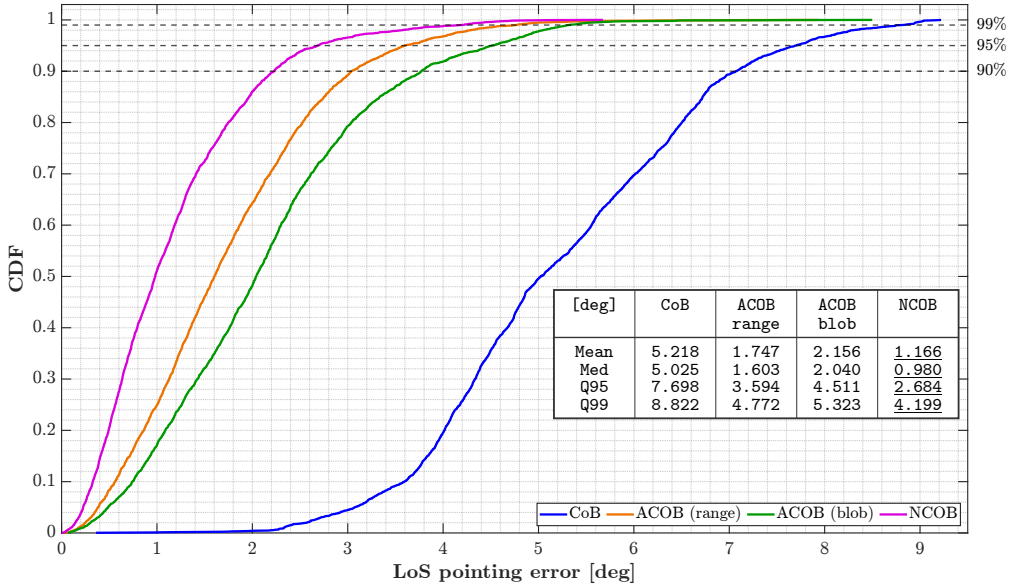


Figure 7 *Cumulative distribution of the LoS pointing error ε on the held-out trajectory* `trajectory_002`.

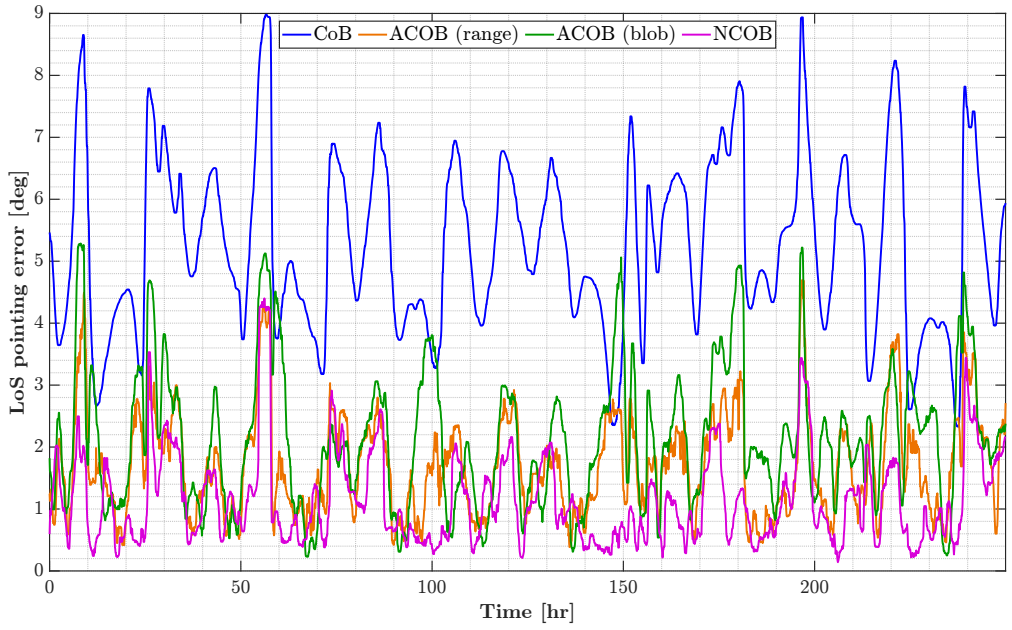


Figure 8 *LoS pointing error along the held-out trajectory* `trajectory_002`, *smoothed over a* 60*-minute window.*

The held-out trajectory `trajectory_002` is a propagated open-loop SSTO arc of 3,000 temporally correlated poses sampled every 5 minutes, never seen in training or validation. The target is held near boresight with a randomized attitude control error, the range stays within 1.13–1.26 km, and the phase angle is concentrated around 96 deg. The Sun azimuth is nearly fixed. This reproduces the actual mission conditions, in which the solar panels are held perpendicular to the Sun, up to a roll control error with 5 deg standard deviation. This set therefore exercises a single sustained high-phase regime, in which the bias of Section 2.1 is both large and persistent.

The uncorrected bias is far larger here, the COB median roughly doubling with respect to the random set, because the arc sits at a sustained phase angle of about 96 deg where the crescent is strongly asymmetric. Figure 7 shows that NeuralCOB holds its median a factor of five below the COB, and 39% and 52% below the range- and blob-based ACOB. The margin over the analytical baselines is thus far wider than on the random set, precisely

in the regime that the correction targets. The range-based ACOB now outperforms the blob-based one at every statistic. The range is well determined along the arc, whereas the thinning crescent distorts the major axis that the blob variant reads as apparent size.

6.3 Systematic and Random Error Components

Because the poses are ordered in time, this set also exposes the temporal structure of the error. Figure 8 traces the pointing error along the arc, smoothed over a 60-minute window for clarity. The illumination geometry is nearly constant, and yet the COB error oscillates quasi-periodically between about 3 and 9 deg as the target rotates and presents a changing illuminated cross-section. This is the shape-driven part of the bias that the spherical model cannot capture. NeuralCOB stays lowest and most stable, within a sub-degree to 2 deg band over most of the arc. The blob-based ACOB tracks well on average but suffers transient excursions of up to about 5 deg when the thinning crescent distorts its apparent-size cue. The range-based ACOB is smoother, and yet sits consistently above NeuralCOB, its spherical correction unable to follow the rotation-driven modulation that the learned correction recovers from the photometric moments.

A navigation filter models its measurements as zero-mean, so a persistent bias is absorbed into the state estimate instead of being averaged away. Decomposing the signed residual of Eq. (10) separates that component from the random one. On the random-sampling set the Sun azimuth is uniform over the full circle, so the COB residual averages to $[-0.4,\ -1.9]$ px in image axes and appears unbiased, whereas its sunward component amounts to 158.1 px. Only the projection onto $\hat{\mathbf{u}}_{\odot}$ exposes the effect that the correction targets.

	`dataset_003`, random i.i.d.				`trajectory_002`, held-out			
Corrector	$\mathbb{E}[r_{\parallel}]$	$\sigma_{\parallel}$	$\mathbb{E}[r_{\perp}]$	$\sigma_{\perp}$	$\mathbb{E}[r_{\parallel}]$	$\sigma_{\parallel}$	$\mathbb{E}[r_{\perp}]$	$\sigma_{\perp}$
COB	158.1 ± 4.2	117.7	-1.3 ± 2.5	71.0	319.8 ± 3.3	92.0	-3.6 ± 2.3	63.8
ACOB (blob)	8.4 ± 3.0	82.8	-1.1 ± 2.6	72.1	$\mathbf{20.4 \pm 4.9}$	137.1	-3.5 ± 2.4	66.6
ACOB (range)	7.3 ± 2.5	70.7	-1.1 ± 2.6	72.2	32.4 ± 3.7	102.3	-3.6 ± 2.4	66.4
NeuralCOB	$\mathbf{6.0 \pm 2.3}$	**62.9**	$\mathbf{-0.8 \pm 2.0}$	**57.1**	24.1 ± 2.7	**74.9**	$\mathbf{-2.2 \pm 1.5}$	**41.8**

Table 3 *Signed residual in the Sun frame on the two nominal sets, in px. The $\pm$ denotes the 95% confidence interval on the mean.*

Three results follow from Table 3. First, every corrector removes the great majority of the systematic offset, between 94 and 96% on the random set and between 90 and 94% on the trajectory. Splitting the along-sun mean-square error as $\mathbb{E}[r_{\parallel}^2] = \mathbb{E}[r_{\parallel}]^2 + \sigma_{\parallel}^2$, the share carried by the systematic term falls from 64% to 1% on the random set, and from 92% to below 10% on the trajectory.

Second, the residual bias alone does not separate the correctors, which sit within a few pixels of one another with overlapping confidence intervals. The dispersion is what separates them. On both sets NeuralCOB has the smallest $\sigma_{\parallel}$ and $\sigma_{\perp}$. On the trajectory it is the only corrector to reduce either quantity below the uncorrected COB: both analytical variants raise the along-sun dispersion above the uncorrected 92.0 px, to 102.3 and 137.1 px. The trajectory holds α near 96 deg and the range near 1.2 km, so $\hat{\delta}_{\mathrm{ACOB}}$ is nearly constant along the arc and removes the mean offset but none of the variance. What remains is the shape- and rotation-driven variation shown in Figure 8, for which Eq. (6) has no correction term.

Third, the cross-sun component reveals what only the learned correction can do. ACOB acts purely along $\hat{\mathbf{u}}_{\odot}$, so by construction it leaves $r_{\perp}$ untouched, and both variants indeed reproduce the cross-sun mean of the COB to within 0.25 px. NeuralCOB alone reduces $\sigma_{\perp}$, by 20% on the random set and by 35% on the trajectory. It is therefore the only corrector that acts on the direction the analytical model cannot reach, which is where the shape-driven part of the bias lies. The practical consequence for the navigation filter is a reduction of the persistent sunward offset from about 5 deg to about 0.4 deg, with a smaller random spread. That offset is held in a nearly fixed image-plane direction along the arc, and is therefore hard to distinguish from a real state error.

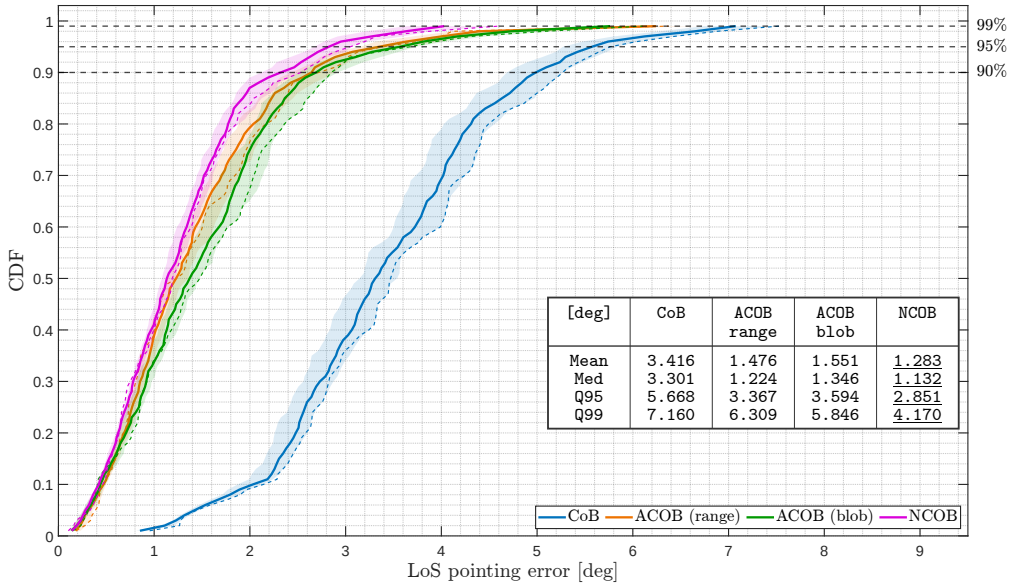
6.4 Robustness

Four held-out sets then stress the correction outside its training conditions. All the statistics are in Table 4.
Near-edge COF. The set `dataset_007` places the target against the detector boundary, concentrating its 1,000 poses in a thin frame along all four image edges, with the range over 0.8–2.0 km and the phase angle over 30–120 deg. The illuminated blob is then partly clipped by the field of view, which corrupts the segmentation, the

brightness moments, and the apparent-size cues on which every corrector relies. This is the hardest set of the benchmark, and the absolute errors are the largest recorded. Even there, NeuralCOB is the best method at every quantile. The two ACOB variants nearly coincide, the truncated blob and the closed-form model failing in the same way, whereas NeuralCOB, which reads the full photometric signature, improves on the better of them by 17% in median and 19% at ε_{95}.

Shape variation at high phase angle. The set `dataset_006` restricts the phase angle to 60–120 deg while covering COF location, range and Sun azimuth broadly, and varies the body shape at every pose. Illumination and shape are therefore stressed jointly. Every corrector degrades with respect to the nominal random set, and yet the ordering is unchanged. The analytical variants have no term for a body that changes from pose to pose, whereas the learned correction reads that change out of the brightness moments.

Shape sweep. The controlled grid `shape_mod_002` renders the same 100 moderate-to-high-phase geometries under 30 shape realizations, drawn independently from the a-priori Gaussians of Section 4.3. Holding the geometry fixed while varying only the body isolates the sensitivity to shape, and directly stresses the spherical assumption of ACOB. In Figure 9, the shaded band spans the 30 shapes. NeuralCOB is the most accurate corrector at every quantile, and also the most consistent: its band is the narrowest of the four. Its accuracy therefore barely moves as the body changes, whereas the COB spread is the widest. Its advantage is largest in the tail, exactly where the elongated realizations break the spherical assumption, with a reduction of 29 to 34% at ε_{99} over the two analytical variants. It is also there that the range-based ACOB, ahead of the blob-based variant at every other statistic, falls behind it: a radius fixed a priori to the mean body is worst for the most elongated realizations, which is the failure mode that this sweep was built to expose.



[deg]	CoB	ACOB range	ACOB blob	NCOB
Mean	3.416	1.476	1.551	1.283
Med	3.301	1.224	1.346	1.132
Q95	5.668	3.367	3.594	2.851
Q99	7.160	6.309	5.846	4.170

Figure 9 *Cumulative distribution of the LoS pointing error ε on the shape sweep `shape_mod_002`, with a band spanning the* 30 *shape realizations.*

Procedurally generated asteroid. The set `asteroid_gen_002` renders the procedurally enriched body of Section 4.3, with terrain, boulders and albedo maps, at the same 3,000 geometries as the held-out trajectory. The matched geometry isolates the effect of the altered shape and appearance, which violate the ellipsoidal and uniform-albedo assumptions of the analytical model. Despite the unfamiliar appearance, NeuralCOB degrades only marginally and remains the best corrector at every quantile, with a median lead of 45% and 33% over the blob- and range-based ACOB that persists into the tail. Along the arc, the COB keeps its rotation-driven oscillation, now roughened by the surface features, while NeuralCOB holds the lowest and most stable trace. The terrain and the albedo add high-frequency fluctuations, without destabilizing the learned correction.

6.5 Benchmark Summary

Across these stresses, that is, edge framing, high-phase shape variation, strong elongation, and realistic terrain and albedo, NeuralCOB has the lowest mean, median and upper quantiles of every corrector on every set, shown in Table 4. The correction therefore degrades gracefully outside its training distribution, which is the kind of generalization evidence that onboard deployment requires.

	$\bar{\varepsilon}$				ε_{50}				ε_{95}				ε_{99}			
Evaluation set	COB	A_b	A_r	NC	COB	A_b	A_r	NC	COB	A_b	A_r	NC	COB	A_b	A_r	NC
`dataset_003`	2.79	1.35	1.25	**1.06**	2.60	1.07	0.99	**0.86**	6.04	3.38	3.07	**2.64**	7.81	5.18	4.85	**4.07**
`trajectory_002`	5.22	2.16	1.75	**1.17**	5.03	2.04	1.60	**0.98**	7.70	4.51	3.59	**2.68**	8.82	5.32	4.77	**4.20**
`dataset_007`	4.10	3.11	3.05	**2.57**	3.80	2.73	2.69	**2.22**	8.31	7.01	6.99	**5.68**	10.77	9.69	10.00	**8.58**
`dataset_006`	3.95	1.77	1.66	**1.43**	3.62	1.42	1.32	**1.18**	7.17	4.37	4.19	**3.54**	8.99	6.71	6.26	**5.46**
`shape_mod_002`	3.42	1.55	1.48	**1.28**	3.30	1.35	1.22	**1.13**	5.67	3.59	3.37	**2.85**	7.16	5.85	6.31	**4.17**
`asteroid_gen_002`	5.21	2.29	1.87	**1.35**	4.94	2.10	1.74	**1.16**	7.83	4.63	3.85	**3.18**	9.02	6.04	4.93	**4.46**

Table 4 *Benchmark statistics: mean $\bar{\varepsilon}$, median ε_{50} and upper quantiles ε_{95}, ε_{99} of the LoS pointing error (deg), for every corrector on every evaluation set. A_b and A_r denote the blob- and range-based ACOB, NC denotes NeuralCOB.*

7 Conclusions

This paper has presented NeuralCOB, a compact and feature-driven neural correction that maps the biased center of brightness onto an enhanced centroid approximating the center of the target asteroid, so that the line-of-sight can be determined onboard during the proximity operations of Farinella at Apophis. The full image is never processed. A small set of hand-crafted illumination-geometry and photometric features is fed to a multilayer perceptron of 10,562 parameters, which captures the nonlinear and shape-dependent COB-to-COF bias that the closed-form ACOB model cannot represent.

Over a large synthetic dataset, NeuralCOB is the most accurate corrector at every quantile, on the nominal sets and equally on the held-out sets that break the analytical assumptions. The median pointing error falls to about 1 deg on all but the near-edge set, up to a factor of five below the uncorrected COB, and below the range-based ACOB even though the latter is granted a metric range and an a-priori radius that NeuralCOB does not need. What distinguishes it is not the residual bias, which every corrector removes to a comparable degree, but the random dispersion: on the flight-like trajectory it is the only corrector to lower that dispersion below the uncorrected COB, and the only one to improve the component orthogonal to the Sun direction, where the shape-driven part of the bias lies.

The explainability analysis of Section 5 shows that the model learns the physical illumination geometry of the bias together with its phase-angle modulation, the kind of interpretability that a safety-critical deployment demands. A markedly different target would still require retraining on its own shape distribution, leaving the feature design and the network architecture unchanged.

On the strength of these results, NeuralCOB will fly aboard Farinella as a navigation experiment, providing the first in-flight demonstration of a learning-based navigation approach at a largely unknown asteroid, and supporting future missions with a higher degree of autonomy.

Acknowledgements

The work described in this paper is part of Farinella Phases A/B/C/D study, which has been funded by the European Space Agency under ESA contract No. 4000147437/25/NL/RK. The authors would like to acknowledge Tyvak International and the whole Farinella Team, including ESA's experts, for their valuable support.

References

[1] M. Kueppers, P. Martino, I. Carnelli, and P. Michel. Ramses – ESA's study for a small mission to Apophis. In *AAS/Division for Planetary Sciences Meeting Abstracts # 55*, volume 55, page 201.08, 2023.

[2] C. Giordano, L. F. Civati, C. Buonagura, A. Beshaj, I. Fodde, P. Panicucci, F. Piccolo, E. Pilo, P. Califano, F. Ferrari, et al.

Trajectory options for RAMSES' Farinella CubeSat around (99942) Apophis. *Advances in Space Research*, 78:5131–5150, 2026. doi: 10.1016/j.asr.2026.06.067.

[3] W. M. Owen. Methods of optical navigation. In *21st AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 11-215, New Orleans, LA, 2011.

[4] J. A. Christian. Optical navigation using planet's centroid and apparent diameter in image. *Journal of Guidance, Control, and Dynamics*, 38(2):192–204, 2015. doi: 10.2514/1.G000872.

[5] M. Kaasalainen and P. Tanga. Photocentre offset in ultraprecise astrometry: Implications for barycentre determination and asteroid modelling. *Astronomy & Astrophysics*, 416:367–374, 2004. doi: 10.1051/0004-6361:20031711.

[6] M. Pugliatti, V. Franzese, A. Rizza, F. Piccolo, C. Bottiglieri, C. Giordano, F. Ferrari, and F. Topputo. Design of the on-board image processing of the Milani mission. In *Proceedings of the 44th Annual American Astronautical Society Guidance, Navigation, and Control Conference*, pages 1303–1323, 2024. doi: 10.1007/978-3-031-51928-4_72.

[7] F. Ferrari et al. The scientific operations of Milani NavCam. *Space Science Reviews*, 2026. In press.

[8] M. Pugliatti, C. Giordano, and F. Topputo. The image processing of Milani: Challenges after DART impact. In *12th International Conference on Guidance, Navigation & Control Systems (GNC) and 9th International Conference on Astrodynamics Tools and Techniques (ICATT)*, pages 1–15, 2023. doi: 10.5270/esa-gnc-icatt-2023-174.

[9] J. A. Christian. A tutorial on horizon-based optical navigation and attitude determination with space imaging systems. *IEEE Access*, 9:19819–19853, 2021. doi: 10.1109/ACCESS.2021.3051914.

[10] P. Panicucci, F. Piccolo, A. Rizza, G. Merisio, F. Topputo, and R. Walker. Vision-based navigation for the LUMIO CubeSat mission. In *46th AAS Guidance, Navigation and Control Conference*, AAS Paper 24-152, 2024.

[11] F. A. Stancu, M. Avilés Rodrigálvarez, A. Pellacani, Á Palomino Aguado, A. Alcalde Barahona, F. Pace, P. Bajanaru, V. M. Moreno Villa, and J. Gil-Fernández. HERA centroiding image processing algorithm based on the normalised correlation with a Lambertian sphere. In *ICINCO (2)*, pages 109–116, 2024. doi: 10.5220/0012899200003822.

[12] M. Pugliatti, V. Franzese, and F. Topputo. Data-driven image processing for onboard optical navigation around a binary asteroid. *Journal of Spacecraft and Rockets*, 59(3):943–959, 2022. doi: 10.2514/1.A35213.

[13] K. Cosmas and K. Asami. Utilization of FPGA for onboard inference of landmark localization in CNN-based spacecraft pose estimation. *Aerospace*, 7(11):159, 2020. doi: 10.3390/aerospace7110159.

[14] R. C. Cardoso, G. Kourtis, L. A. Dennis, C. Dixon, M. Farrell, M. Fisher, and M. Webster. A review of verification and validation for space autonomous systems. *Current Robotics Reports*, 2:273–283, 2021. doi: 10.1007/s43154-021-00058-1.

[15] European Space Agency, KP Labs, and Warsaw University of Technology. Catalogue of explainability techniques for AI applications in space operations. Technical Report ESA Grant No. 4000144194/23/D/BL, European Space Agency (ESA), 2025. URL `https://assurance-ai.space-codev.org/`.

[16] N. Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979. doi: 10.1109/TSMC.1979.4310076.

[17] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019.

[18] ONNX Community. ONNX: Open neural network exchange. `https://onnx.ai/`, 2026. Accessed 7 August 2026.

[19] P. J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964. doi: 10.1214/aoms/1177703732.

[20] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[21] I. Loshchilov and F. Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.

[22] Blender Foundation. Cycles: Open source production rendering. `https://www.cycles-renderer.org/`, 2026. Accessed 23 July 2026.

[23] P. Pravec, P. Scheirich, J. Ďurech, J. Pollock, P. Kušnirák, K. Hornoch, A. Galád, D. Vokrouhlický, A. W. Harris, et al. The tumbling spin state of (99942) Apophis. *Icarus*, 233:48–60, 2014. doi: 10.1016/j.icarus.2014.01.026.

[24] B. B. Mandelbrot and J. W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10 (4):422–437, 1968. doi: 10.1137/1010093.

[25] L. S. Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, 1953.

[26] S. M. Lundberg and S. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.

[27] D. W. Apley and J. Zhu. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):1059–1086, 2020. doi: 10.1111/rssb.12377.