Uniqueness and nonuniqueness in reconstruction of nonconvex bodies from photometry

Mikko Kaasalainen

June 28, 2019

Abstract

The main text is the same as the corresponding section in Viikinkoski et al. 2017, A&A 607, A117, included here for introduction and clarity. The essential part of this separate study is the appendix that presents a complete proof and discussion of fundamental uniqueness and nonuniqueness properties of the nonconvex photometric inverse problem. We also include a related mathematical joke on how to make a binary target look like a single one and yet describe the surface accurately.

Information content analysis, including uniqueness proofs of the reconstructions of bodies based on various projectionlike data sources as well as the weighting of those sources when used simultaneously, has been presented in a number of works (see, e.g., Kaasalainen & Lamberg 2006 and references therein, Kaasalainen 2011, Viikinkoski & Kaasalainen 2014, Nortunen et al. 2017). We revisit here the reconstruction of nonconvex shapes from photometry only. The numerical results discussed, e.g., in Kaasalainen et al. (2001), Ďurech & Kaasalainen (2003), and Kaasalainen & Ďurech (2007) give a good practical overview of the inverse problem. The problem, however, is a mathematical one, so any nonuniqueness or uniqueness proofs are invaluable for understanding the actual characteristics of the case. We present here fundamental results that require a somewhat special setup, but they give insight to the general problem, and are among the very rare proofs that can be given about the problem in the first place.

In the following, the term *illuminated projection area* denotes the total area of the projections of the visible and illuminated (hereafter VI) parts of a body in the viewing direction (from which the illumination direction can differ). *Brightness data* is the generalization of this, where the surface elements contributing to the illuminated projection area are each weighted by a scattering function depending on the local viewing and illumination conditions (the scattering of illuminated projection areas is called geometric). Brightness data are also called generalized projections (Kaasalainen & Lamberg 2006). In two dimensions, the body is a planar curve and the projection area is the sum of the widths of the VI parts of the curve seen from the viewing direction.

Tangent-covered bodies or TCBs are bodies for which each point on the surface has at least one tangent that does not intersect any other part of the body (but can be tangent to them). The *tangent hull* of a body is the set of surface points for which the above criterion is true, augmented by their tangents to form a closed, connected surface of a TCB. By a *concavity*, we mean a part of the surface of a body that is not part of its tangent hull. As discussed

in Kaasalainen (2011), tangent-covered bodies (and the tangent hull of a body, also called its profile hull) are reconstructable from their disc-resolved silhouette or profile curves. TCBs are thus the set of all convex bodies and all nonconvex bodies without concavities. By definition, bodies with concavities are not TCBs, a TCB is identical to its tangent hull, in three dimensions convex bodies are a subset of TCBs, and in two dimensions the sets of convex bodies and TCBs are identical.

Two-dimensional nonconvex bodies cannot be uniquely determined from their brightness data. By definition, the tangent hull of a 2D body is its convex hull, so any nonconvexities of a 2D body are concavities. Each concavity is covered by a line that is part of the convex hull, and all parts of a line have the same visibility and illumination. Thus the effect of the concavity on brightness data can be replaced by smaller concavities, covered by the same line, that are isomorphic to the original concavity whose length along the line equals their combined lengths. It is also possible to have concavities of different shapes along the same line that together produce a shadow effect that can be attributed to one concavity of still another shape (see below). Therefore the brightness data of any nonconvex 2D body can be reproduced by infinitely many other nonconvex versions. We call this *scale ambiguity*.

Concavities of three-dimensional bodies cannot be uniquely determined from the brightness data of the bodies. For simplicity, we assume here any concavity to be contained in a plane that is part of the convex hull of the body. This is just to avoid lengthy discussions of special shadowing conditions that are not material to the argument. Then the 3D case is a direct generalization of the scale ambiguity of the 2D case: the effect of the concavity on brightness data can be replaced by smaller isomorphic (or possibly other) concavities in the plane whose combined surface area equals that of the concavity. We assume that the plane is suitably larger than the part of it occupied by the concavity so that the smaller concavities can be arranged within the plane. Note that this arrangement is also nonunique: even if one uses a size constraint for the concavities, their locations in the plane cannot be deduced from brightness data.

An interesting corollary of this nonuniqueness is that even a large-scale concavity is actually indistinguishable from a locally rugged surface: in other words, a concavity within a plane can be replaced by the same plane with scattering properties caused by small-scale roughness (cf. the discussion in Kaasalainen et al. 2004).

A subset of tangent-covered bodies can be determined from their illuminated projection areas at least with a simple scale constraint. Here we show that the set of bodies essentially reconstructable from their illuminated projection areas is larger than the set of convex surfaces. Essentially means here the use of a natural constraint. The uniqueness proof for the brightness data of convex bodies is discussed in Kaasalainen & Lamberg (2006) and references therein.

Proving anything about the integrals over a nonconvex body dependent on directions is notoriously difficult for the simple reason that usually such integrals are not analytically calculable, requiring numerical ray-tracing or the finding of the roots of equations containing high-order functions. Therefore we must resort to a number of special assumptions, considering first the case in two dimensions. We can construct a proof there since in 2D the location of the shadow boundary point is essentially the same as its projection area (length in 2D). In 3D, the shadow boundary is resolved in 2D (see Kaasalainen 2011), and the area requires the computation of an integral (as would brightness data in 2D).

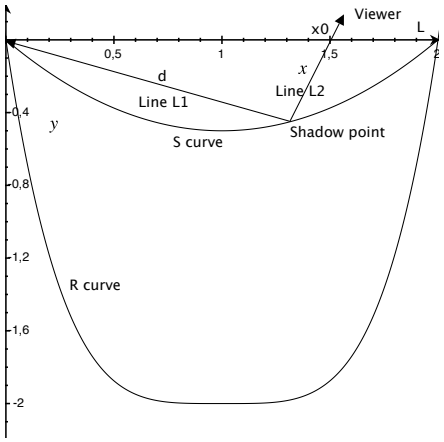

Figure 1: Sketch of the geometry of the nonconvex shape in the xy-plane.

Our reconstruction consists of three parts: *1.* the uniqueness of a concavity up to scaling in 2D, *2.* the separation of the uniqueness of this shape and that of the rest of the 2D curve, and *3.* the transformation of the 2D curve into a 3D TCB and the scale constraints.

1. For simplicity, we consider an isolated concave section $\mathcal{S}$ of a curve $\mathcal{C}$ in the xy-plane, with end points on the x-axis, one at the origin and the other at $(L, 0)$, and the rest of $\mathcal{S}$ is below the x-axis. We are interested in the shadow caused by the end point at $(0, 0)$. Let the illumination direction be ϕ and the viewing direction θ, with $\phi - \theta = \alpha$. When $\alpha \neq 0$, $\mathcal{S}$ can be uniquely reconstructed by finding the point (x_s, y_s) on $\mathcal{S}$ separating the shadow and the illuminated part for successive values of ϕ and θ. We require $\mathcal{S}$ to be starlike w.r.t. the origin: this enables simple shape parametrization and yields only one shadow and one illuminated section on the concavity. Further, we require VI parts of $\mathcal{S}$ to cover the whole of the line between (x_s, y_s) and $(L, 0)$ projected in θ, and that there are θ covering the whole of the motion of (x_s, y_s) from $(0, 0)$ to $(L, 0)$ as θ increases. A simple sufficient but by no means necessary arrangement for this is to let $\mathcal{S}$ be an inverted convex curve with the angles between the y-axis and the tangents of $\mathcal{S}$ both $\alpha/2$ at $(0, 0)$ and $(L, 0)$.

The shadow point (x_s, y_s) is the intersection point of two lines L_1 and L_2. L_1 is the shadow line through the origin in ϕ, and L_2 is the line in θ such that its distance from the corresponding line through $(L, 0)$ is the observed illuminated projection length l. Thus

$$x_s = -d\cos\phi = x_0 - s\cos\theta, \quad y_s = -d\sin\phi = -s\sin\theta, \tag{1}$$

where d and s are the length parameters of the lines L_1 and L_2, respectively, and L_2 passes through x_0 for which

$$l = (L - x_0)\sin\theta. \tag{2}$$

Combining these, one obtains the unique solution for d, and thus x_s, y_s at given θ, α, L, and

observed l. Parametrizing l with the usual polar angle φ, obtained directly from $\varphi = \phi - \pi$, we have

$$d(\varphi)\sin\alpha = L\sin\theta - l(\theta), \quad \theta = \varphi - \alpha + \pi, \quad 0 \le \theta \le \pi - \alpha, \tag{3}$$

so

$$[x(\varphi), y(\varphi)]_{\mathcal{S}} = [d(\varphi)\cos\varphi, d(\varphi)\sin\varphi], \tag{4}$$

and $d(\varphi) = 0$ for some values of θ when $l(\theta) = L\sin\theta$; $d(\varphi) > 0$ for some interval of φ ending at $\varphi = 2\pi$ when $d = L$ and $l = 0$.

This result shows formally the scale ambiguity: d is unique only if there is one concavity with d, since otherwise one may define

$$\sum_i d_i \sin\alpha = \sum_i L_i \sin\theta - l(\theta), \quad \sum_i L_i = L, \quad \sum_i d_i = d, \tag{5}$$

so scaled-down arrangements of adjacent concavities with similar or various $d_i(\varphi)$ will produce the same $l(\theta)$.

2. Recall that any convex curve can be defined by its curvature function $C(\psi)$, $0 \le \psi < 2\pi$. The length element is $ds = C(\psi)\,d\psi$, so the points $[x(\psi), y(\psi)]$ on the curve are given by

$$x(\psi) = -\int_0^{\psi} C(\psi')\sin\psi'\,d\psi', \quad y(\psi) = \int_0^{\psi} C(\psi')\cos\psi'\,d\psi'. \tag{6}$$

The observed projection length of the VI part of the curve at θ is

$$l(\theta) = \int_{-\pi/2+\alpha}^{\pi/2} C(\psi+\theta)\cos\psi\,d\psi. \tag{7}$$

Expanding $C(\psi) = \sum_n a_n \cos n\psi + b_n \sin n\psi$, $n \ge 0$, $a_1 = b_1 = 0$, and similarly $l(\theta) = \sum_n c_n \cos n\theta + d_n \sin n\theta$, we have (Ostro & Connelly 1984, Kaasalainen 2016)

$$c_n = a_n I_n^c(\alpha) + b_n I_n^s(\alpha), \quad d_n = b_n I_n^c(\alpha) - a_n I_n^s(\alpha), \tag{8}$$

where

$$I_n^c(\alpha) = \int_{-\pi/2+\alpha}^{\pi/2} \cos n\psi\cos\psi\,d\psi, \quad I_n^s(\alpha) = \int_{-\pi/2+\alpha}^{\pi/2} \sin n\psi\cos\psi\,d\psi, \tag{9}$$

so the solution of the inverse problem is

$$a_n = \frac{I_n^c c_n - I_n^s d_n}{(I_n^c)^2 + (I_n^s)^2}, \quad b_n = \frac{I_n^s c_n + I_n^c d_n}{(I_n^c)^2 + (I_n^s)^2}. \tag{10}$$

The denominator $(I_n^c)^2 + (I_n^s)^2$ is nonzero at $0 < \alpha < \pi$, so the convex curve is uniquely obtained from the $l(\theta)$-data at any α.

When the integral defining $l(\theta)$ mixes parts of $\mathcal{S}$ and other parts of $\mathcal{C}$, there is no simple way to parametrize $\mathcal{C}$ so as to allow analytical computations of $l(\theta)$. Thus we separate the observations of $\mathcal{S}$ from the rest of the curve $\mathcal{C}$ denoted by $\mathcal{R} = \mathcal{C}\setminus\mathcal{S}$. We assume that $\mathcal{R}$ is convex and without linear sections. Then, denoting the angle interior to $\mathcal{C}$ between the x-axis and the tangent of $\mathcal{R}$ at the origin and $(L, 0)$ by, respectively, γ_0 and γ_L, we must have $\gamma_0 \le \alpha, \gamma_L \le \alpha$. With these

assumptions, either parts of $\mathcal{S}$ or parts of $\mathcal{R}$ are VI at any given θ. The data from $\mathcal{S}$ cannot be mimicked by a convex curve since they are like those of a line for some range of θ.

Given a full range of observations $l(\theta)$ of $\mathcal{C}$, one can proceed as follows. First find the values of θ at which $l(\theta)$ vanishes at the given α. Assuming $\mathcal{R}$ not to contain straight lines ending at sharp corners elsewhere, there are two intervals (or points) in θ with $l(\theta) = 0$, one ending at $\theta = 0$ and one starting at $\theta = \alpha$ (the lengths of the intervals are $\alpha - \gamma_0$ and $\alpha - \gamma_L$). Between these, $l(\theta)$ is sinusoidal for an interval of θ starting at the end of one interval of $l = 0$: $l = L \sin\theta$. This reveals the existence of the corners at the origin and $(L, 0)$. It also fixes the xy-frame in our convention, and gives the size of $\mathcal{S}$: $L = l/\sin\theta$ from the sinusoidal interval. After the shadow effect appears at some θ, the shape profile of $\mathcal{S}$ is uniquely obtained from $l(\theta)$-data at $0 \leq \theta \leq \pi - \alpha$. Observations of $l(\theta)$ in the range $\pi - \alpha \leq \theta \leq 2\pi$ are sufficient to determine $C(\psi)$ of $\mathcal{R}$ (or any convex curve) in the interval $3\pi/2 - \alpha \leq \psi \leq 3\pi/2 + \alpha$ since $C(\psi)$ in this interval does not affect $l(\theta)$ at other values of θ. For this range of θ, these data are exactly the same as those from $\mathcal{C}$ with $\mathcal{S}$ a straight line from the origin to $(L, 0)$ would be.

This concludes the determination of $\mathcal{C}$ up to the scale ambiguity on $\mathcal{S}$: one obtains $\mathcal{R}$ and the shape $\mathcal{S}$, but there is no way of telling from $l(\theta)$ whether the data are due to $\mathcal{S}$ or smaller concavities on the line section L_0 between the origin and $(L, 0)$.

3. Next we define a class of 3D bodies by extending the 2D curve $\mathcal{C}$ to apply to some interval $z_0 \leq z \leq z_1$, and covering the cylindrical shape of height $z_1 - z_0$ with planes at z_0 and z_1. This surface is a TCB, and we call the shape class cylindrical TCBs. Concavities in 2D are changed into saddle surfaces rather than concavities in 3D, with much richer information possibilities. Viewing directions along the z-axis then reveal the true area A between the real concavity curve and L_0: if A_P is the observed projected area of the body and A_C the area contained inside the curve formed of $\mathcal{R}$ and L_0, $A = A_C - A_P$. This information strongly constrains the possible scale ambiguities. Now we must have $\sum_i A_i = A$ in addition to $\sum_i L_i = L$ for each concavity i, and the combined area of many concavities is smaller than the area of one due to the quadratic scaling of the areas A_i w.r.t. the lengths L_i. Also, viewing and illumination directions other than those along the z-axis or in the xy-plane offer additional information. In these other observing geometries, scale ambiguity is at least partly removed for the saddle-surface versions of the concavities, but it is difficult to show the exact properties analytically.

If A equals the area A_S by S (on the whole of L_0) or is close to it, the assumption of one concavity is well justified. Indeed, if $A = A_S$, only one occurrence of the concavity shape $\mathcal{S}$ is possible if only this shape is considered. A simple way to enforce the assumption is to attribute almost all area A to one scaled concavity $\mathcal{S}$ and fill the remaining length on L_0 by negligibly small-scale roughness.

The above construction applies also to other arrangements than that of $\mathcal{S}$ and $\mathcal{R}$. For example, the profiles of the concave sides of a "nonconvex Reuleaux triangle" are solvable at least for $\alpha \geq \pi/6$. Viewing directions tilted from the z-axis can be used to isolate the area information for each concave section of the surface by placing them in the opposite azimuthal direction.

Not all tangent-covered bodies can be uniquely determined from their brightness data. Even the use of all observing geometries and size constraints cannot resolve the case where sections on L_0 are without concavities (just straight lines). Then it is impossible to say where along the line the sections and concavities are. Another example of TCB ambiguity: consider

the surface of a wedge-shaped 3D tangent-covered body (or a convex body with a wedgelike part), with two intersecting planes forming the wedge. A new tangent-covered surface can be formed by making a tangent-covered hole that is contained between the two planes. The hole need not be cylindrical as long as it is tangent-covered so that the new body is still a TCB. If the hole is suitably smaller than the wedge, all brightness data of the body can be reproduced by having a collection of smaller isomorphic holes within the wedge arranged such that their combined area equals that of the original concavity. This is possible because the wedge shape allows an infinite number of isomorphic holes to be created (two parallel planes instead of a wedge would not allow this). A corollary of this is that a hole can be replaced by arbitrarily small-scale perforation.

Even when formally removable, ambiguities lead to instabilities near limit conditions. For example, if the straight line on which two scale-ambiguous concavities are adjacent is bent between the concavities, the above uniqueness results are obtained for both separately. However, this requires α to be at least as large as $\pi - \beta$ where β is the bending angle. A small β is a typical depiction for many nearby nonconvex features on asteroid surfaces, so reconstructions are unstable even at high α.

The reconstruction of a nonconvex body from its brightness data is fundamentally nonunique. Scale and location ambiguities are inherent to the inverse problem. Nevertheless, the reconstructable class of bodies from brightness data is, in a certain sense and with constraints, larger than the set of convex shapes especially due to the projection area information from TCBs. This corroborates the numerical success in simulations such as those in Ďurech & Kaasalainen (2003). However, the reconstruction of nonconvex bodies from photometry has neither the fundamental uniqueness properties of the convex case nor the Minkowski stability that pertains to the global shape and applies to both data and model errors. These aspects are illustrated by, e.g., the case of the asteroid Eros. Eros, with its sizable nonconvex TCB-like feature, can be roughly approximated by the simple cylindrical model when viewed from the direction of its rotation axis. One might thus expect its photometry to yield a unique nonconvex solution. Even so, as discussed in Kaasalainen & Ďurech (2007), the convex model fits the data as well as a nonconvex one and, above all, better than the real shape with usual scattering models. There are various nonconvex shapes that fit the data equally well. This underlines the ambiguous properties and the instabilities of the photometry of nonconvex bodies and the need for large α.

A convex surface actually represents merely a nonconvex case in which regularisation suppressing local nonconvex features has been given infinite weight. This, however, does generally not deteriorate the fit as discussed in Ďurech & Kaasalainen (2003); so far, the only asteroid requiring a nonconvex shape to explain its photometry is Eger (Ďurech et al. 2012). All others can be explained down to the noise level by convex shapes, which means that, from the point of view of regularisation theory, there is no optimal regularisation weight so it is best to use full weight to avoid the inevitable instabilities and nonuniqueness at lower weights. Statistically, no result between the two extreme weights can be shown to be the best one, so the safest result is a convex shape because of its strong uniqueness and stability properties. From the Bayesian point of view, the problem is the lack of proper statistics to cover the systematic data and model errors (dominating over the data noise) and the difficulty of finding shape sampling covering the whole of the shape space (single or a few shape supports cannot do this properly in the sense of Markov chain Monte Carlo; see Viikinkoski et al. 2015).

References

J. Ďurech, M. Kaasalainen, *A&A* **404**, 709 (2003)

J. Ďurech et al., *A&A* **547**, A10 (2012)

M. Kaasalainen et al., *Icarus*, **153**, 37 (2001)

M. Kaasalainen et al., *Icarus*, **167**, 78 (2004)

M. Kaasalainen, L. Lamberg, *Inverse Problems* **22**, 749 (2006)

M. Kaasalainen, J. Ďurech, *Proc. of IAU, Symp. 236*, **2**, 151 (2007)

M. Kaasalainen, *Inverse Problems and Imaging* **5**, 37 (2011)

M. Kaasalainen, *XVI Special Courses of the Observatorio Nacional*, AIP conf. proc. 1732, 020003 (2016)

H. Nortunen et al., *A&A*, **601**, A139 (2017)

S. Ostro, D. Connelly, *Icarus*, **57**, 443 (1984)

M. Viikinkoski, M. Kaasalainen, *Inverse Problems and Imaging* **8**, 885 (2014)

M. Viikinkoski et al., *A&A* **576**, A8 (2015)

Appendix: full shape determination

In the above, we used a fixed z-height or thickness $h = z_1 - z_0$ of the disc $\mathcal{C} \times [z_0, z_1]$ for simplicity, and derived the constraints on the size of the concavity from views along the z-axis. We fix $z_0 = 0$ here (without loss of generality) so the disc is now $\mathcal{C} \times [0, h]$. We can also prove that it is possible to derive h uniquely and thus indeed define a class of nonconvex bodies that can be fully determined from brightness data. By a reconstructable shape class we mean a category of surfaces defined by a number of non-trivial parameters that extensively alter the shapes, and that the values of these parameters can be uniquely determined from sufficient data.[1]

If the disc height h is known, then the length L of the nonconvex feature follows directly from the observed area (brightness value) Lh obtained at $\alpha = 0$ and $\theta = \pi/2$. Then the illuminated area data from the direction of the z-axis automatically show whether the $\mathcal{S}$-curve covers the largest possible area; i.e., is a single curve of length L, or if there are multiple $\mathcal{S}$-curves. But if h is not known, L cannot be known either, in which case these observations do not settle the case of the possibly multiple $\mathcal{S}$-curves. There are infinitely many combinations of h and L that reproduce the data and will allow $\mathcal{S}$-multiplicities. Thus, to determine the size and height of the disc as well as the $\mathcal{S}$-multiplicity constraints, one needs illumination directions other than those in the xy-plane or along the z-axis.

To derive the disc height h we use the same self-shadowing phenomenon of nonconvex bodies (distinct from the occluding shadow of convex bodies) that provided the information needed

[1] We note that the three-dimensional shape of this proof is constrained also because we only know from the convex proof which illumination and viewing coverage is sufficient for uniquely reconstructing a given part of a convex surface. If we knew which coverage is necessary, we might sharpen the result. Again, from numerical inversion results we know that the coverage needed is much smaller than in the original proof.

for determining the curve $\mathcal{S}$. We choose $\phi = \pi$ and $\alpha = \pi/2$ so $\theta = \pi/2$. In this case, the concave surface $\mathcal{S}_h = \mathcal{S} \times [0, h]$ of height h and defined by the curve $\mathcal{S}$ is entirely in shadow. By introducing a latitude ϑ for the illumination source we can illuminate the concave part in another way. The illumination source is given in three dimensions by the angles (ϑ, ϕ), where ϑ is the polar angle measured from the z-axis. $\vartheta = 0$ means that the illumination is parallel to the z-axis, and $\vartheta = \pi/2$ keeps it in the xy-plane. Because of symmetry, we choose $0 \leq \vartheta \leq \pi/2$. We now have illumination at (ϑ, π) and start to decrease ϑ from $\pi/2$ so that illuminated parts start to appear on the surface $\mathcal{S}_h$. We also assume that both L and h are nonzero.

Since the curve $\mathcal{S}$ is fully known from earlier reconstruction (up to the unknown scale length L), the projected area (up to the unknown height h) of the VI part V_+ of $\mathcal{S}_h$ in the direction of the viewer, $P[V_+(\mathcal{S}_h)]$, is readily computed at least by numerical ray-tracing (the shape of the shadow boundary depends on the shape of $\mathcal{S}$). Note that an analytical expression of this is not necessary for the following proof, it is sufficient to know that $P[V_+(\mathcal{S}_h)]$ can be computed to any desired accuracy.

We assume, as above, that the shape of $\mathcal{S}$ is such that there is only one boundary curve between shadow and light. Also, we assume for simplicity that $\mathcal{S}$ is bounded by a rectangle, one side of which is L_0. These constraints are not strictly necessary and are straightforward to generalise.

First we assume that there is only one curve $\mathcal{S}$ covering L_0. Define a computable function

$$f(\vartheta) = P[\vartheta, V_+(\mathcal{S}_\infty)], \tag{11}$$

where the ϑ-argument on the r.h.s. gives the illumination angle for the VI part. Thus $f(\vartheta)$ gives the projected area when $h \to \infty$. We define $f(\vartheta)$ only for $\vartheta > 0$ since $f(\vartheta) \to \infty$ as $\vartheta \to 0$. Then, with the viewer in the xy-plane, the shadow boundary curve starts in the xz-plane of projection from the point (x_0, h), where

$$x_0 = \inf\{x | \eta_1(x) \leq 0\}, \tag{12}$$

where $\eta \in \mathbb{R}^2$ is the normal of the curve $\mathcal{S}$ in the xy-plane, and ends at the point (L, z_*), where

$$z_* = h - L \cot \vartheta. \tag{13}$$

Choosing unit length for L in $f(\vartheta)$ (with the above viewing direction), the observed projected area is $P_o(\vartheta) = L^2 f(\vartheta)$, so we obtain

$$L = \sqrt{P_o(\vartheta)/f(\vartheta)} \tag{14}$$

for some span of ϑ.

When h is finite and ϑ decreases from $\pi/2$, at some ϑ the shapes of the functions $f(\vartheta)$ and $P_o(\vartheta)$ start to differ because the illuminated part of $\mathcal{S}_h$ no longer increases as expected for $f(\vartheta)$ since the shadow falls beyond the surface: $\mathcal{S}_h$ necessarily cannot accommodate all the illumination. By our earlier assumed constraints on the shape of $\mathcal{S}$, this missing illuminated area must be nonzero. This means that we have detected at which ϑ_h the whole of h is spanned; i.e., $z_* = 0$, and thus $h/L = \cot \vartheta_h$, which yields h (Fig. 2):

$$h = \sqrt{\frac{P_o(\vartheta_h)}{f(\vartheta_h)}} \cot \vartheta_h. \tag{15}$$

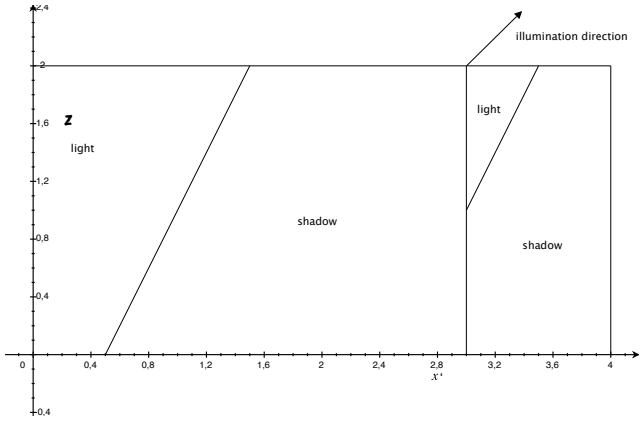

Figure 2: A schematic of the shadows as seen by the viewer in the $x'z$-plane ($x' = -x + c$, c a constant, x as in Fig. 1) at $\vartheta = \pi/4$, $\phi = \pi$, and $\theta = \pi/2$. Now the nonconvex feature consists of two $\mathcal{S}$-curve concavities in the xy-plane: $l_1 = 3$, $l_2 = 1$. The top points of the boundaries of the shadow projections in the $x'z$-plane are the $(-x_0 + c, h)$ of Eq. (12), and the boundaries tend towards the values z_* of Eq. (13) (with L replaced by l_1 and l_2) at the left-hand vertical boundaries of the concavities, one reaching it and the other not.

Next, define

$$f_h(\vartheta) = P[\vartheta, V_+(\mathcal{S}_h)]. \tag{16}$$

Thus, $f_h(\vartheta) = f(\vartheta)$ for $\vartheta \geq \vartheta_h$; elsewhere, $f_h(\vartheta) < f(\vartheta)$.

If the observations show that

$$P_o(\vartheta) = L^2 f_h(\vartheta) \qquad \forall \vartheta, \quad 0 < \vartheta \leq \pi/2, \tag{17}$$

we know that either there is only one curve $\mathcal{S}$ or there is a degeneracy and our deduced L is incorrect: there are two or more curves of exactly the same length. If the length L is divided into N equal-sized curves of shape $\mathcal{S}$, the h' deduced as above, based on the assumption of a single curve, will be $\sqrt{N}$ times larger than the actual height h to account for the combined projected area of the N edge sections. Similarly, the deduced length L' will be $\sqrt{N}(L/N) = L/\sqrt{N}$, so the circumference of the $\mathcal{C}$-shaped boundary curve of the disc is underestimated by $\sqrt{N}$. Thus, if the expected and observed projected areas in the direction of the z-axis are, respectively, P'_z and P_z, we find that if $p \equiv P_z/P'_z > 1$, there is an N-fold degeneracy of

$$N = p + (1 - p)c, \tag{18}$$

where c is the expected ratio of the missing projected area of the nonconvex feature to the projected area of the disc if the feature were replaced by a straight line.

It remains to check the case of L being split into $\mathcal{S}$-shaped curves of unequal sizes. We omit degeneracies there for simplicity, and the case of differently shaped curves obviously has generic ambiguities. Let ϑ' be the first ϑ at which the observed $P_o(\vartheta)$ differs from $L'^2 f(\vartheta)$ as ϑ decreases

from $\pi/2$, and h' and L' are the height and length deduced based on ϑ' as above. If there are smaller versions of $\mathcal{S}$, illumination through them requires $\vartheta < \vartheta'$ to meet the lower rim of the disc, in which case not as much illuminated area is missed as in $f_{h'}$. Thus, if

$$L'^2 f_{h'}(\vartheta) < P_o(\vartheta) < L'^2 f(\vartheta), \qquad \vartheta < \vartheta', \tag{19}$$

we then know that there are multiple curves of shape $\mathcal{S}$ in various sizes.

Here the information from illumination at various ϑ yields both h and the sizes of the curves (although their ordering along the line L_0 cannot be deduced). Let each curve i have a projected length l_i along L_0 so that the total length is $L = \sum_i l_i$ (Fig. 2), and we observe the ϑ_i when the corresponding illuminated portion on $\mathcal{S}_h$ meets the lower rim of the disc. This is seen from the behaviour of P_o in the same way as above. The ϑ_i occur in the order of decreasing l_i as ϑ decreases. Then

$$l_i/l_j = \tan\vartheta_i / \tan\vartheta_j. \tag{20}$$

By observing all the ϑ_i when scanning $0 < \vartheta < \pi/2$ we find all the possible curves (in Fig. 2, $\vartheta_2 < \vartheta < \vartheta_1$), and scale them to correct lengths by matching the sum of the projected areas of all illuminated portions with any observed $P_o(\vartheta)$, $\vartheta > \vartheta_1$. This also yields h.

Our proof has thus shown that it is possible to define a nontrivial class of nonconvex shapes that can be uniquely reconstructed from brightness observations. One can also verify if the object does not belong to this class, or if it does not have a strictly unique single $\mathcal{S}$-curve. The proof makes use of the self-shadowing of nonconvex bodies, the related high illumination phase angles, and the constraining nature of the "missing profile area" information of tangent-covered bodies. The proof is not strictly constructive in parts (in that we do not construct the Fourier series but prove that the corresponding information is in the data), but the numerical inversion computations corroborate it as discussed earlier. Such computations also emphasise the unstable and fundamentally nonunique nature of the general nonconvex inversion, especially due to the model errors.

The main points of the proof are that:

1. There are a number of fundamental unavoidable scale and location ambiguities in nonconvex inversion, as well as strong instabilities in the vicinities of these (regions with low global curvature or close to concave; low phase angles).

2. Tangent-covered bodies and suitable lighting geometries offer unique information on nonconvex features and alleviate scale ambiguities.

It is also possible to show uniqueness results for a concave feature on a bent version of the above construction (two planar sections forming a wedge instead of a planar region), approximating a more general saddle-like feature on a curved surface. This, however, pertains to a very restricted shape class and does not illustrate the case of general shapes in $\mathbb{R}^3$. Also, it does not offer more insights to the above points already detected in the more clearly defined shape class represented by $\mathcal{C} \times [0, h]$.

An often asked question is: can't we just unravel some nonconvex features by regularisation or some other trick? To which the answer is, unfortunately, simply no if the convex version fits the data as well as the others. This is because regularisation can only reveal probable detail if the

continuous removing of the apparent detail (complexity) in the result starts to worsen the data fit at some point. The morphological complexity of a nonconvex shape is, by any measure, always larger than that of its convex hull (which is not the same as the corresponding convex result but approximating it well enough for the purpose of the argument if not for fit comparison). Thus it follows that, if a convex shape fits the data better or as well as the more complex ones, there is no way to assign better probability to any of the latter on purely morphological basis.

If we could use physical and geological arguments for structural rather than morphological complexity to create a quantitative regularisation function, based on physical likelihood, the situation might be different. This, however, essentially requires us to already know what any target should basically look like both inside and on the surface, which is rather against the whole point of the exercise in the first place. The case of an outright unfeasible convex solution is a different matter, if we decide we just have to use a more likely shape representation (e.g., a bi-lobated shape) even if we cannot give it a quantitative likelihood. This is a discrete or categorisation problem of model selection rather than a continuous problem of regularisation. Here a useful way to see what actually causes the modelled lightcurve is to plot the rotating plane-of-sky animation of the asteroid together with the evolving model lightcurve and its fit to the observed one. Then the changes in the projected area and the shadows become apparent. This is particularly helpful in appraising nonconvex solutions qualitatively when one tries to identify the features most essential for the shapes of the observed lightcurves.

To wrap up, we note an interesting mathematical curiosity, or a setup for pulling a prank on an overconfident observer. Consider an arbitrary convex curve $\mathcal{C}$ in $\mathbb{R}^2$, and its version $\mathcal{C} \times [0, h]$ in $\mathbb{R}^3$, with the origin at the centroid of the area enclosed by $\mathcal{C}$. Let the illumination phase angle α be $-\pi/2$ for the viewing angles $0 \leq \theta < \pi/2$ and $\pi \leq \theta < 3\pi/2$, and $\alpha = \pi/2$ for $\pi/2 \leq \theta < \pi$ and $3\pi/2 \leq \theta < 2\pi$. By $\alpha < 0$ we mean that the integration limits in Eq. (7) are from $-\pi/2$ to $\pi/2 + \alpha$; i.e., the light source is trailing instead of leading on the unit circle S^1. Now the observed lightcurve $l(\theta)$ is piecewise defined and discontinuous at intervals of $\pi/2$, so the Fourier analysis of the inverse problem is not applicable as above. However, we see from Eq. (7) that information on all the Fourier components of the curvature $C(\psi)$ is retained in $l(\theta)$ (the effect of $C(\psi)$ on $l(\theta)$ is not well sampled only at and near the special points $\psi = \pi/2$ and $\psi = 3\pi/2$). Thus the observer can expect to obtain a good approximation of $\mathcal{C}$ (at least with suitable regularisation, as corroborated by numerical experiments) by solving a set of linear equations by some truncation or discretisation corresponding to Eq. (7) for $l(\theta)$ observed at a large number of θ covering all of S^1. The height h is obtained by observing the area enclosed by $\mathcal{C}$ by viewing $\mathcal{C} \times [0, h]$ from the direction of the z-axis. Thus the observer has a unique accurate solution for $\mathcal{C} \times [0, h]$ from the observations.

Suppose now that $\mathcal{C}$ is actually split in two parts along the x-axis such that we have two closed convex curves $\mathcal{C}_1$ and $\mathcal{C}_2$ with straight lines parallel to the x-axis as parts. Now $\mathcal{C}_1$ and $\mathcal{C}_2$ can be moved arbitrarily far away from each other along the y-axis while the observations remain invariant. No information on the gap reaches the observer who nevertheless has a unique single-body reconstruction of $\mathcal{C}$ that perfectly explains the data. This scenario is readily extended to convex bodies in $\mathbb{R}^3$ with partial observing geometries analogous to those that are typically sufficient for reconstructing the single body but in this case not for detecting the split body.