



Essentials of Strong Gravitational Lensing

Prasenjit Saha¹ · Dominique Sluse² · Jenny Wagner³ · Liliya L.R. Williams⁴

Received: 14 April 2023 / Accepted: 30 December 2023 / Published online: 5 February 2024
© The Author(s) 2024

Abstract

Of order one in 10^3 quasars and high-redshift galaxies appears in the sky as multiple images as a result of gravitational lensing by unrelated galaxies and clusters that happen to be in the foreground. While the basic phenomenon is a straightforward consequence of general relativity, there are many non-obvious consequences that make multiple-image lensing systems (aka strong gravitational lenses) remarkable astrophysical probes in several different ways. This article is an introduction to the essential concepts and terminology in this area, emphasizing physical insight. The key construct is the Fermat potential or arrival-time surface: from it the standard lens equation, and the notions of image parities, magnification, critical curves, caustics, and degeneracies all follow. The advantages and limitations of the usual simplifying assumptions (geometrical optics, small angles, weak fields, thin lenses) are noted, and to the extent possible briefly, it is explained how to go beyond these. Some less well-known ideas are discussed at length: arguments using wavefronts show that much of the theory carries over unchanged to the regime of strong gravitational fields; saddle-point contours explain how even the most complicated image configurations are made up of just two ingredients. Orders of magnitude, and the question of why strong lensing is most common for objects at cosmological distance, are also discussed. The challenges of lens modeling, and diverse strategies developed to overcome them, are discussed in general terms, without many technical details.

Keywords Gravitational lensing · Strong gravitational lensing

✉ P. Saha
psaha@physik.uzh.ch

D. Sluse
dsluse@uliege.be

J. Wagner
thegravitygrinch@gmail.com

L.L.R. Williams
llrw@umn.edu


¹ Physik-Institut, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

² STAR Institute, Quartier Agora Allée du six Août, 19c B-4000 Liège, Belgium

³ Bahamas Advanced Study Institute and Conferences, 4A Ocean Heights, Hill View Circle, Stella Maris, Long Island, Bahamas

⁴ School of Physics and Astronomy, University of Minnesota, 116 Church Street, Minneapolis, MN 55455, USA

Fig. 1 The archetype of strong gravitational lensing: Q0957+561. The two bright bluish objects with diffraction spikes are lensed images of a quasar at redshift $z = 1.41$. The elliptical galaxy near the lower quasar image, and the other galaxies in the field are part of a cluster at $z = 0.36$, which together form the gravitational lens. The field shown is $30'' \times 30''$ with North up and East to the left



1 The General Picture

Some time in centuries past, pieces of glass in shapes resembling lentil seeds (*Lens culinaris*) came to be known as lenses. The name remains with us: the optical elements developed for lighthouses in the 19th century are still called lenses, even though they look more like giant pineapples. A lens, then, is anything that transmits light with some non-trivial optical path or light travel time, causing light deflection. If an interesting optical path is created by a gravitational field, meaning the spacetime metric, we have a gravitational lens.

Light deflection in accordance with Einstein gravity was measured in 1919, and theorists have written about gravitational lensing for even longer,¹ but it was in the Einstein centennial year of 1979, with the discovery of the “double quasar” Q0957+561 by Walsh et al. (1979), that gravitational lensing became part of astronomy. A beautiful image of this system is shown in Fig. 1.

In the decades since 1979 more than a thousand multiply-imaging gravitational-lens systems comparable to Q0957+561 have been discovered, and the discovery rate is accelerating. This article will introduce the essential theory needed to study such systems, which are also called “strong-lensing” systems. The latter term contrasts the phenomenon of multiple image with weak-lensing, where there is light deflection and image distortion but not multiple images — such as in Giblin et al. (2021) on weak lensing by large-scale structure, or Crosta et al. (2017) on weak lensing in the solar system. Note, however, that strong lensing still occurs in weak gravitational fields — we will not consider light propagation near event horizons (as in the well-known images from Akiyama et al. 2022). Some of the theory is also applicable to gravitational microlensing by stars in or near the Milky way (e.g., Mróz et al. 2020), but mainly we will be concerned with lensing by galaxies and clusters at cosmological distances.

¹For the fascinating scientific history, often at odds with textbook narratives, see Valls-Gabaud (2006) and Kennefick (2012).

1.1 Light Deflection

The practice of strong gravitational lensing involves some simplifying assumptions.

First, geometrical optics is assumed for light sources and the gravitational-lensing processes. Wave optics is essential at the telescope, but we do not consider interference between multiple images. Appendix C briefly discusses how the standard formalism can be extended to wave optics, but for the rest of the present article, light means rays.

Second, gravitational fields are assumed to be weak. As a consequence, light rays are deflected from straight lines only by very small angles. Furthermore, the contributions of many masses can be simply added (or integrated over) without the need to solve Einstein's field equations. Thus, lensing is treated as a first-order perturbation from light travelling along straight lines.

A third assumption is that light deflection occurs over regions much smaller than the distances between light source, lensing mass, and observer. This assumption is expressed in the notions of a *lens plane* and a *source plane*. It is not meant that sources and lenses are flat, but rather that the gradual deflection of light through a gravitational field is integrated along the line of sight through the comparatively small active region, and approximated as a discrete deflection. Distances transverse to the line of sight, however, are treated with respect, and not unceremoniously integrated out.

The well-known expression

$$\frac{4GM_{\odot}}{c^2 R_{\odot}} = 1.75'' \simeq 1 \times 10^{-5} \text{ rad} \quad (1)$$

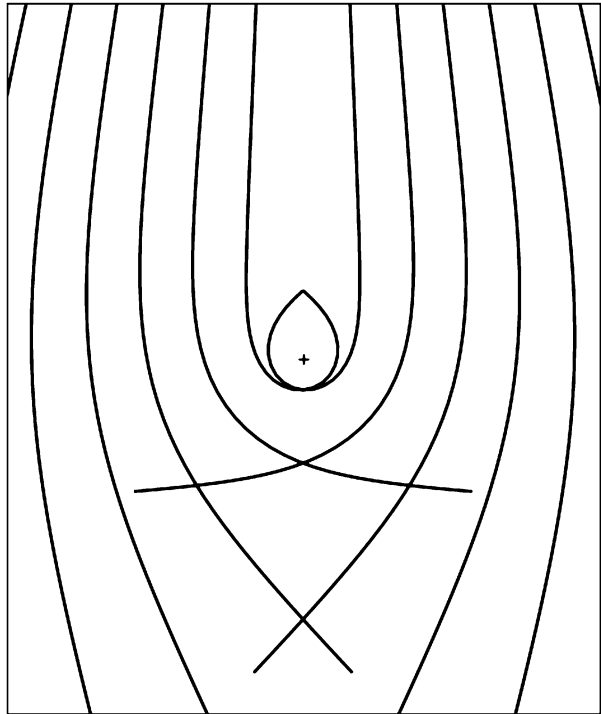
for gravitational light deflection at the rim of the Sun exemplifies all three assumptions: (i) geometrical optics is understood, (ii) the metric perturbations are of the same order as the deflection in radians, and thus small, and (iii) the Sun is approximated as a flat deflector in the sky. Light deflections by galaxies are coincidentally of the same order as that by the Sun, while clusters of galaxies produce deflections an order of magnitude larger. Thus, the light deflections that will concern us in practice are $\lesssim 1 \times 10^{-4}$ rad and entirely in the regime of weak gravitational fields.

Introductions to gravitational lensing (e.g., Meneghetti 2021) conventionally start by specializing to the regime of geometrical optics, weak fields, and a planar lens. Some of the essential concepts, however, are also valid for strong fields. So let us start a little outside the usual comfort zone, with an example of lensing by a strong gravitational field.

Figure 2 shows light deflection through one of the simplest curved spacetimes. The metric (see Appendix A for details) is similar to a Schwarzschild metric, but has no event horizon (which, though fascinating, would take too long to discuss here). Light rays originate from a point source high above the figure, and are gradually diverging as they approach the strong field in the middle of the figure, and then the lens pulls the rays together. There is no focal point, though there is a sort of focal line along which rays converge symmetrically from paths on opposite sides of the lens. Two such ray intersections on the focal line are shown in the figure. Observers located there will observe a ring of light, provided the lens is circularly symmetric. There are also asymmetric ray intersections (two of which are indicated) in the lower part of the figure, not on the focal line. Observers located there will see the source light coming from two directions in their sky.

The rays near the middle of Fig. 2 undergo very large deflections, even more than 180° . Large deflections occur in strong gravitational fields near neutron stars and black holes, leading to the exotic multiple imaging now observed near the M87 black hole (Broderick et al. 2022). It does not, however, occur in the systems we will consider.

Fig. 2 A representation of the deflection of light rays. The source is far above, outside the figure. There is a strong gravitational field near the center of the figure (marked with a +), which produces large deflections (even $> 180^\circ$). Farther from the center, the field is weaker and the deflections smaller, though still much larger than the real systems we are considering



Lensing theory conventionally uses observer-centric coordinates. To a given observer, a source appears to be located along a direction (say θ), whereas its true location is along some different direction (say β). The directions θ and β are conventionally expressed as two-dimensional angles on the sky, measured in radians or arc-seconds. The mapping relating θ and β is called the lens equation, and is conventionally expressed as

$$\beta = \theta - \alpha(\theta) \quad (2)$$

with $\alpha(\theta)$ having the interpretation of an apparent deflection angle. Adding angles in this way, one is tacitly assuming that the sky has been mapped to a plane. For large deflections we would need to specify the sky projection, but if all angles are small, at any time, only a very small patch of sky is relevant, and angles can be simply added.

The convention of using observer-centric coordinates seems peculiar to gravitational lensing. In optics, it is usual to put the origin at the lens. Even within astronomy, the study of planetary atmospheres as lenses for occulted stars (see e.g., Elliot and Olkin 1996) uses lens-centric coordinates. Observer-centric coordinates are, however, a good idea, because some useful facts get built into the notation. As we can see from the form of the lens equation (2), the mapping from image plane θ to source plane β is single valued, but the inverse can be multiple valued. Thus each image corresponds to a unique source, whereas each source can have multiple images. Ray-tracing using the lens equation (2) using your favorite deflection field $\alpha(\theta)$ is easy and can be entertaining.² Each pixel in the image plane is given the bright-

²Try visiting

<https://phdenzel.github.io/streaming-lens/>

on a device with a camera.

ness of the corresponding pixel in the source plane. (Since β is a single-valued function of θ , but not the other way round, ray tracing is usually done backwards, from the observer's sky to the source.) Note that ray tracing in this way automatically preserves surface brightness (photons per steradian). A patch of uniform brightness on the source stays the same uniform brightness when lensed, but it can change in size and shape.

The conventional minus sign for the deflection in Eq. (2) means that $\alpha(\theta)$ points away from the deflector rather towards it. This is an example of a well-known principle in astrophysics,³ which is that you first think of how a rational person would do it, and then you *change the sign*.

1.2 Wavefronts and Fermat's Principle

The lens equation (2) would actually be too general for lenses, if we let $\alpha(\theta)$ be arbitrary. This is because you cannot send light any way you like using just lenses — for that you would need optical fibers. Imagine a cable consisting of a bundle of twisting optical fibers, and a signal input simultaneously into all the fibers. The overall direction of travel of the light signal along the cable will be different from the direction of travel within individual fibers. Lenses cannot make light travel like this. In lensing, the surface or front described by a bundle of light rays as light travels is always normal to the rays themselves. This front is known as the wavefront. (Despite the name, the wavefront is well-defined in geometrical optics. If the source is small but not effectively a point, the wavefront will be wiggly at the wavelength scale, and will not produce interference.)

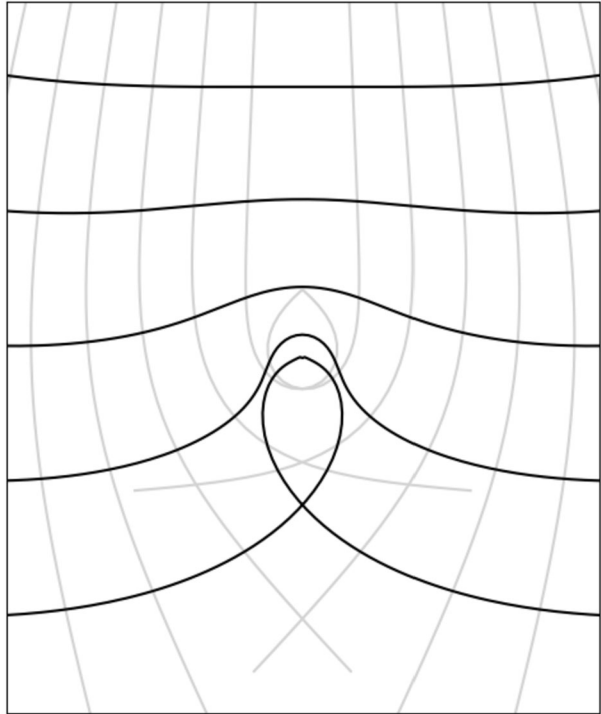
Figure 3 is a representation of the wavefront corresponding to the rays in Fig. 2. Here, we imagine the source emitting an omni-directional light flash, and the wavefront shows how far that light flash has travelled in a given time in different directions. Light rays are simply the local normals to the wavefront. Notice that the lens delays the wavefront — drastically in the strong-field region, less in other places. Most interesting for strong gravitational lensing, some observers will be crossed by the wavefront more than once, at different times and from different directions.

Let us now modify the preceding to the scenario shown in Fig. 4. Let us pivot the source (located above and outside the figure) by a small angle β about the lens. The incoming wavefront is then inclined by β . After the wavefront has passed the lens and been delayed by it, we freeze the wavefront. Then we imagine a second backwards wavefront, which emanates from an observer at the bottom of the figure. Any part of this second wavefront can be labelled by the angle θ on the observer's sky. It is like the observer's sky rising up to meet the frozen wavefront. Let $t(\theta)$ be the time at which a point θ on the rising-sky wavefront meets the original frozen wavefront. Let us now concentrate on any places where the two wavefronts meet tangentially. These correspond to $\nabla t(\theta) = 0$. At these places the normals to the two wavefronts are the same. Since normals to the wavefront are light rays, $\nabla t(\theta) = 0$ corresponds to light rays travelling out from the source with a wavefront and then down the sky to the observer. Moreover, $t(\theta)$ at these sky locations is the light travel time, apart from an additive constant. This is Fermat's principle in gravitational lensing.

The two-wavefront picture of lensing (due to Nityananda 1990) is remarkable in that, from the basic postulate that light rays are normal to a wavefront, it leads inexorably to Fermat's principle and the notion of an abstract surface $t(\theta)$, even with large deflections and strong fields. We assumed a stationary lens, but in fact Fermat's principle in terms of an arrival time can be extended to arbitrary time-varying gravitational fields in general relativity

³Attributed to J. Binney.

Fig. 3 Wavefronts corresponding to the light rays in Fig. 2. (The source is above outside the figure.) Light gray curves are taken over from that figure. Each of the black curves is the endpoint of all the rays at a particular time. The small loop is the result of very large deflections and would only occur in a strong gravitational field



(Nityananda and Samuel 1992). In the following we will specialize to small deflections, but it is worth emphasizing that the notion of the arrival-time surface $t(\theta)$ and its consequences have much wider applicability.

1.3 Geometrical and Gravitational Time Delays

Let us now turn Fig. 4 into equations, under the following simplifying assumptions.


- All angles are taken to be small, in the sense of $\sin \theta \approx \theta$ and $\cos \theta \approx 1 - \frac{1}{2}\theta^2$.
- Fig. 4 has the source far away from the lens. We now assume that the source is at infinity. (We will relax this assumption later.) As before, on the observer’s sky the source is located at location β .
- The gravitational field is assumed localized to the neighborhood of the lens. The backward wavefront is therefore spherical.

Under these assumptions, the frozen wavefront just after passing the lens is a slightly warped plane inclined by β , while the dashed backward wavefront is $\propto \theta^2$.

Without any lensing mass, the frozen wavefront is simply a plane, and the time taken for backward wavefront to reach it follows from geometry as

$$t_{\text{geom}}(\theta) = \frac{D}{c} \left(\frac{1}{2}|\theta|^2 - \theta \cdot \beta \right) + \text{const} , \tag{3}$$

Fig. 4 Modified from Fig. 3. The source (outside the figure) has been moved slightly to the left. The wavefront proceeds as the dotted curves and is then frozen at the solid black curve. Then a second backward wavefront (dashed) emerges from the observer on the ground. Each of the dashed curves is the endpoint at a particular time of rays emanating from the observer, but those rays are not shown



where D is the shortest distance from the observer to the frozen wavefront. Since β is constant, we can equivalently write

$$t_{\text{geom}}(\theta) = \frac{D}{2c} |\theta - \beta|^2 + \text{const} \tag{4}$$

Adding a lensing mass, the frozen wavefront gets warped, and the time taken by the backward wavefront to reach it changes by say $t_{\text{grav}}(\theta)$. The light travel time, up to a constant, would then be

$$t(\theta) = t_{\text{grav}}(\theta) + t_{\text{geom}}(\theta) . \tag{5}$$

The form of $t_{\text{grav}}(\theta)$ is not obvious, but let us put in the ansatz

$$t_{\text{grav}}(\theta) = -\frac{G}{c^3} \sum_i M_i \ln |\theta - \theta_i| \tag{6}$$

pending verification that it leads to the correct deflection angle. Note here, that, since the angles are all small, the logarithm will always be negative, and because of the minus sign the contribution of the mass term will always be positive. Light rays will reach the observer and images will appear where the gradient of $\nabla t(\theta) = 0$. Working out the gradient we have

$$\beta = \theta - \frac{4G}{c^2 D} \sum_i M_i \frac{\theta - \theta_i}{|\theta - \theta_i|^2} \tag{7}$$

which, as we can see, is the lens equation (2) for a sum of Sun-like deflectors at distance D . As throughout this section, the source is much further away, a restriction we will lift later. This verifies our ansatz (6) for the gravitational time delay.

The combined geometric and gravitational time delay

$$t(\boldsymbol{\theta}) = \frac{D}{2c} |\boldsymbol{\theta} - \boldsymbol{\beta}|^2 - \frac{G}{c^3} \sum_i M_i \ln |\boldsymbol{\theta} - \boldsymbol{\theta}_i|, \quad (8)$$

also variously known as the arrival time or the Fermat potential, is a central concept in gravitational lensing. As we have seen, the lens equation is the zero-gradient condition on the time delay. The second and higher derivatives are also important, as detailed later.

Fermat's principle in gravitational lensing was introduced using different approaches by Schneider (1985) and Blandford and Narayan (1986). Expressions like (8) do, however, appear in earlier work (Cooke and Kantowski 1975; Borgeest 1983) albeit without the interpretation as Fermat's principle. Wavefront arguments were introduced to gravitational lensing by Refsdal (1964). The gravitational time delay (6) is known as the Shapiro delay after its deduction by Shapiro (1964) as a property of null geodesics from general relativity. That time delays also emerge simply as an integral of the deflection angle was shown independently (and actually slightly earlier) by Refsdal (1964).

2 The Cosmological Context

2.1 Why Cosmological Distances?

The lensing galaxy in Q0957+561 is at redshift $z = 0.36$, which is a typical distance for strong lenses. Why don't we see strong lenses nearby?

To see why, let us briefly consider the Sun. The Sun is assuredly a gravitational lens, but we do not observe it as a strong gravitational lens, because the light deflection at its rim (Eq. (1)) is much smaller than its angular size on our sky. If one observed from a much greater distance D , the apparent size of the Sun would become smaller than the deflection angle, and sources directly behind the Sun would be lensed into images on either side. We can express this fact as

$$\frac{4GM}{c^2\xi} = \frac{\xi}{D} \quad (9)$$

together with the condition that ξ (the physical radial distance of the apparent images on the observer's sky) is larger than the solar radius.

What if the Sun were transparent, or we were looking for strongly-lensed neutrinos? Then the requirement that ξ is more than the solar radius would be lifted, and the condition (9) would apply as

$$\frac{4GM(<\xi)}{\xi^2} = \frac{c^2}{D} \quad (10)$$

with the mass enclosed within ξ . Note the latter form — it says that strong lensing can occur if the Newtonian acceleration at the lens is at least c^2/D . (This condition is usually expressed as a critical 2D density M/ξ^2 , but let us stay with threshold acceleration for now.) From $D = 1$ au the threshold is too high even for the solar interior, but for $D > 25$ au

a transparent Sun will indeed strongly lens neutrinos (Demkov and Puchkov 2000). For $D > 550$ au the optical solar gravity lens would be observable, and is indeed the subject of some space-mission concepts (Turyshev and Toth 2020). But for now, strong lensing needs much larger distances and lower acceleration thresholds.

Let us consider some example cases of the threshold acceleration (10). The values⁴

$$M = M_{\odot} \quad \xi \simeq 2 \text{ c s} \quad D \simeq 400 \text{ au} \quad (11)$$

correspond roughly to optical strong lensing by the Sun (the solar gravity lens). If we scale ξ up by 10^3 and D up by 10^6 we get

$$M = M_{\odot} \quad \xi \simeq 4 \text{ au} \quad D \simeq 2 \text{ kpc} \quad (12)$$

which is in the regime of microlensing in the Milky Way. However, the angular scale $\xi/D = 1 \times 10^{-8} \text{ rad} \simeq 2 \text{ mas}$ which is challenging to image even with optical interferometry (cf. Cassan et al. 2022). If, instead, we scale all three quantities by 10^{12} we have

$$M = 10^{12} M_{\odot} \quad R \simeq 20 \text{ kpc} \quad D \simeq 2 \text{ Gpc} , \quad (13)$$

which is typical of massive galaxies at cosmological distances. Moreover the angular scale $\xi/D = 1 \times 10^{-5} \text{ rad} \simeq 2''$ is comparatively easy to resolve.

Thus, we see that strong lensing is very difficult to observe nearby, but once you can observe galaxies at cosmological distance with arcsecond resolution, strong lensing becomes more frequent.

2.2 Distances in Cosmology

With lenses being at cosmological distances, we need to account for the universe expanding as the light is travelling through it. This is quite straightforward to do, using the usual machinery of modern cosmology, which is covered in many textbooks, and need not be repeated here. Let us nonetheless draw attention to the main concept relevant to lensing, namely the different ways of expressing distance.

The observable associated with distance is the redshift z . While an observed z may contain a small kinematic contribution (of order 10^{-3}), cosmological redshifts are essentially wavelengths expanding with the universe. Thus a redshift of z refers to an epoch when the scale factor of the universe was $1 + z$ smaller than now. This applies to any cosmological model that is homogeneous and isotropic on large scales (that is, satisfies the cosmological principle).

In the current standard cosmological model, the universe is spatially flat and its present total energy density is

$$\rho_{\text{cr}} = \frac{3H_0^2}{8\pi G} , \quad (14)$$

in which H_0 is the current expansion rate, that is, the Hubble constant.⁵ The density at redshift z is ρ_{cr} times

$$E(z) = \Omega_m (1+z)^3 + \Omega_r (1+z)^4 + \Omega_{\Lambda} , \quad (15)$$

⁴Here it is useful to remember that $1 \text{ au} \simeq 500 \text{ c s}$ (light-seconds), and $1 \text{ pc} \simeq 10^8 \text{ c s}$ (see Appendix B for more precise values).

⁵Distances, times, and angle-size relations needed for studying lenses involve several unit conversions for H_0 , which sets the scales. Appendix B gives the relevant conversions.

in which Ω_m is the current matter density fraction, Ω_r is the current relativistic energy-density fraction, and Ω_Λ the “dark energy” density fraction.

The lookback time to some redshift z_d is given by

$$t_d = \frac{1}{H_0} \int_{z_d}^1 \frac{dz}{(1+z)E(z)} \tag{16}$$

and ct_d is by definition the distance travelled by light from the source to the observer. The comoving distance between two points along the line of sight at redshifts z_d and z_s is given by

$$\chi_{ds} = \frac{c}{H_0} \int_{z_d}^{z_s} \frac{dz}{E(z)} \tag{17}$$

As the name indicates, the comoving distance scales up distances according to the present scale of the universe.

Gravitationally bound structures like galaxies and clusters do not expand with the universe. But the light coming from them knows⁶ nothing of such subtleties, and shows the angular sizes of galaxies and clusters as if they *had* been expanding with the universe. To account for this, angular-size or angular-diameter distances are defined. The angular-diameter distance to the lens (or deflector) and source are

$$D_d = \frac{\chi_d}{1+z_d} \quad D_s = \frac{\chi_s}{1+z_s}, \tag{18}$$

respectively, with the subscripts having the obvious meanings. Also important is the angular-diameter distance

$$D_{ds} = \frac{\chi_s - \chi_d}{1+z_s} \tag{19}$$

from the lens to the source. The latter needs clarification. An observer at the lens would measure the angular distance as $(1+z_d)/(1+z_s)$ times the comoving distance. But that observer’s comoving distances are all smaller by $(1+z_d)$ than ours. Hence the expression (19).

In the older literature (notably in the well-known book by Schneider et al. 1992) cosmologies with total density different from ρ_{cr} are also considered, as are more general considerations of the angular-diameter distances. The expressions given here are restricted to cosmologies that are spatially flat (or equivalently, have mean density equal to ρ_{cr}), but this is a standard assumption made in nearly all of the recent literature.

Yet another quantity which one occasionally meets in lensing formalism is the conformal time. A conformal time interval is given by $(1+z)^{-1} dt$. Using the conformal time, one can trace light rays as if in a static universe. In the previous section, we used some arguments involving running light rays backwards, and the arguments remain valid in an expanding universe if conformal time is understood.

2.3 The Lens and Source Planes

With angular-diameter distances in hand, we can now extend the lens equation and the time delay to the expanding universe, and relax the assumption that sources are at infinity.

⁶This is known as anthropomorphism.

Fig. 5 A diagram (reproduced from Bartelmann and Schneider 2001) of the various lengths and angles used in the lensing formalism. The dashed line is a reference direction, sometimes called the optical axis. The angle $\hat{\alpha}$ is the light deflection by the lens. Note that D_d, D_{ds}, D_s are angular-diameter distances, and $D_d + D_{ds} > D_s$, whereas ξ and η are physical distances




Figure 5 shows the various quantities involved. Through the small-angle approximation, the regions where the source and lens are located, are approximated as planes. The lens plane and source plane have redshifts z_d and z_s respectively. The lens (deflector) is at angular-diameter distance D_d from the observer, so that an angle θ on the observer’s sky corresponds to a physical displacement ξ on the lens plane. Analogously, the source is at angular-diameter distance D_s from the observer, and an angle β corresponds to a physical displacement η on the source plane. Note again that the angular-diameter distance D_{ds} is not equal to $D_s - D_d$ but larger, because of the way angular-diameter distances are defined.

Projecting distances on the source plane, we have

$$D_s \beta = D_s \theta - D_{ds} \hat{\alpha}(\theta) \tag{20}$$

and hence

$$\beta = \theta - \frac{D_{ds}}{D_s} \hat{\alpha}(\theta) \tag{21}$$

which is our lens equation (2) generalized to an expanding universe. Introducing the apparent (or scaled) deflection angle

$$\alpha \equiv \frac{D_{ds}}{D_s} \hat{\alpha} \tag{22}$$

makes the lens equation look exactly as before.

Assembling the lens out of Sun-like deflectors (as in Eq. (7)), we have

$$\beta = \theta - \frac{4G}{c^2} \frac{D_{ds}}{D_d D_s} \sum_i M_i \frac{\theta - \theta_i}{|\theta - \theta_i|^2}, \tag{23}$$

where we have used $\xi = D_d \theta$. As is easy to verify, this equation is the zero-gradient condition of the time delay

$$\frac{t(\theta, \beta)}{1 + z_d} = \frac{D_d D_s}{2c D_{ds}} |\theta - \beta|^2 - \frac{4G}{c^3} \sum_i M_i \ln |\theta - \theta_i|. \tag{24}$$

We did not derive the $1 + z_d$ factor in the time delay, but one can argue that this factor must be there, because the gravitational time delay occurs at the lens plane and gets redshifted.

One issue we have swept under the carpet here, is that, in Eq. (20), we should have parallel transported the lens plane to the source plane, instead of simply projecting. When one extends to multiple lens planes, as we will do below, disregarding parallel transport produces spurious image rotation at leading order (cf. Grimm and Yoo 2018). Such artifacts are harmless for the applications considered here, but it is good to be aware of them.

2.4 Multiple Lens or Source Planes

Although our equations so far have assumed one source plane and one lens plane, it is not difficult to generalize to multiple planes.

The lensing of a source plane is independent of any other source planes. Hence, for a single lens plane, each source plane β_s simply has its own lens equation

$$\beta_s = \theta - \frac{D_{ds}}{D_s} \hat{\alpha}_d(\theta). \tag{25}$$

The form of Eq. (25) tells us that a deflection at one z_d is equivalent to a scaled deflection at another z_d . This is useful in that minor lensing masses along the line of sight can be replaced by effective masses on the main lens plane.

It is also possible to generalize the original lens equation (21) explicitly to multiple planes. To do so, let us first rewrite Eq. (21)

$$\frac{\xi}{D_s} = \frac{\eta}{D_d} - \frac{D_{ds}}{D_s} \hat{\alpha}. \tag{26}$$

Then let us introduce new subscripts: 0 for the observer plane, 1 for the lens plane, and 2 for the source plane. Let us further introduce $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ for comoving positions on these respective planes. This gives

$$\xi = \frac{\mathbf{x}_2 - \mathbf{x}_1}{1 + z_2} \quad \eta = \frac{\mathbf{x}_1 - \mathbf{x}_0}{1 + z_1}. \tag{27}$$

Meanwhile we can also write the angular-diameter distances in terms of comoving distances

$$D_d = \frac{\chi_1 - \chi_0}{1 + z_1} \quad D_{ds} = \frac{\chi_2 - \chi_1}{1 + z_2} \quad D_s = \frac{\chi_2 - \chi_0}{1 + z_2}. \quad (28)$$

Here \mathbf{x}_0 and χ_0 are both zero, as they go from the observer to the same observer, but they will play a useful role presently.

Substituting these various expressions into our rewritten lens equation (26) and simplifying the result, we obtain

$$\frac{\mathbf{x}_2 - \mathbf{x}_1}{\chi_2 - \chi_1} = \frac{\mathbf{x}_1 - \mathbf{x}_0}{\chi_1 - \chi_0} - \hat{\boldsymbol{\alpha}}_1. \quad (29)$$

This equation relates comoving quantities at three consecutive planes. Since the planes 0, 1, 2 were not special in the derivation, we can generalize to

$$\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\chi_{n+1} - \chi_n} = \frac{\mathbf{x}_n - \mathbf{x}_{n-1}}{\chi_n - \chi_{n-1}} - \hat{\boldsymbol{\alpha}}_n. \quad (30)$$

This is equivalent to Eq. (19) from Petkova et al. (2014). To use it, we need to know the comoving distances χ_n to all the planes, the deflections $\hat{\boldsymbol{\alpha}}_n$ at those planes, and the starting values $\mathbf{x}_0 = 0$ (the observer) and \mathbf{x}_1 (the scaled location on the observer's sky). We can then compute \mathbf{x}_n for arbitrarily many lens planes.

So far, the best-known multi-plane system is the ‘‘Jackpot lens’’ J0946+1006 (Gavazzi et al. 2008; Smith and Collett 2021), which has three sources at different redshifts, one of which is also a second lens. Upcoming surveys will surely find more such systems, and the topic is likely to become important.

3 The Standard Formalism of Gravitational Lensing

Several introductions to strong lensing have been published over the years, laying out the basic formalism, though the emphasis and amount of detail differs. Schneider et al. (1992) was the first book on lensing, and was very influential. The historical introduction is quite interesting, and the quantitative portions of the book are very comprehensive and often at an advanced level; however, the portions of the book that deal with observational state of the field have been superseded many times over in the last three decades. Bartelmann and Schneider (2001) is an extension of the book to weak lensing, but can also be used as an introduction to lensing theory in its own right. Blandford and Narayan (1986) give a short but well-explained article, providing a gentle introduction among other things to catastrophe theory as relevant to gravitational lensing (Petters et al. 2001, is a deep dive into that area). They also discuss the similarities with atmospheric mirages, which is further developed in Refsdal and Surdej (1994) who propose an original introduction to the field and explain how to design an optical lens simulator with a piece of glass. (Kochanek 2006) provides an interesting perspective specifically on strong gravitational lensing. Meneghetti (2021) is the most recent review, and includes Jupyter notebooks. The above list is non-exhaustive; in addition to dedicated lensing reviews many text-books today contain chapters on lensing.

Common to all these sources, and indeed to the literature on strong gravitational lensing generally, is a set of concepts and cryptic-sounding terms that do not appear elsewhere in astrophysics. With the physical basis and cosmological context in hand, we now explain what these are.

3.1 Introducing Scaled Quantities

Let us replace the sum over discrete masses in Eq. (24) with an integral over a projected mass distribution

$$\sum_i M_i \ln |\boldsymbol{\theta} - \boldsymbol{\theta}_i| \rightarrow D_d^2 \int \Sigma(\boldsymbol{\theta}') \ln |\boldsymbol{\theta} - \boldsymbol{\theta}'| d^2\boldsymbol{\theta}' . \tag{31}$$

Here $\Sigma(\boldsymbol{\theta})$ is a function of the angular position, but physically, it is a projected density in kg m^{-2} on the lens plane. Hence the D_d^2 factor.

Next, we introduce two scales

$$\begin{aligned} T_{\text{ds}} &\equiv (1 + z_d) \frac{D_d D_s}{c D_{\text{ds}}} , \\ \Sigma_{\text{cr}} &\equiv \frac{c^2}{4\pi G} \frac{D_s}{D_d D_{\text{ds}}} \end{aligned} \tag{32}$$

and present their interpretations: As defined, cT_{ds} is a combination of lens and source distances, sometimes called the time-delay distance. Like other cosmological distances, $T_{\text{ds}} \propto H_0^{-1}$. In the limit $D_{\text{ds}} \rightarrow D_s$, which is to say sources at infinity, $cT_{\text{ds}} \rightarrow D_d$. Thus cT_{ds} is a replacement for distance D which we used in Sect. 1.3 when introducing the general picture.

The quantity Σ_{cr} is known as the critical density for strong lensing, because it turns out to be the threshold for appearance of multiple images. (It is *not* the cosmological critical density, which has different dimensions.) For typical lens and source redshifts Σ_{cr} comes to a few kg m^{-2} — about the surface density of window glass (see Appendix B). The proverbial astute reader will have noticed that $G\Sigma_{\text{cr}}$ is an acceleration. Indeed the threshold acceleration we mentioned earlier (Eq. (10)) is $4\pi G\Sigma_{\text{cr}}$.

Using the scale parameters T_{ds} and Σ_{cr} we can replace the time delay and projected mass distribution with the dimensionless quantities $\tau(\boldsymbol{\theta})$ and $\kappa(\boldsymbol{\theta})$

$$\begin{aligned} t(\boldsymbol{\theta}, \boldsymbol{\beta}) &= T_{\text{ds}} \tau(\boldsymbol{\theta}, \boldsymbol{\beta}) , \\ \Sigma(\boldsymbol{\theta}) &= \Sigma_{\text{cr}} \kappa(\boldsymbol{\theta}) \end{aligned} \tag{33}$$

and express the arrival time in the following, commonly used form

$$\begin{aligned} \tau(\boldsymbol{\theta}, \boldsymbol{\beta}) &= \frac{1}{2} |\boldsymbol{\theta} - \boldsymbol{\beta}|^2 - \psi(\boldsymbol{\theta}) , \\ \psi(\boldsymbol{\theta}) &= \frac{1}{\pi} \int \kappa(\boldsymbol{\theta}') \ln |\boldsymbol{\theta} - \boldsymbol{\theta}'| d^2\boldsymbol{\theta}' . \end{aligned} \tag{34}$$

Fermat’s principle, $\nabla\tau(\boldsymbol{\theta}) = 0$, gives the lens equation (2) in the form

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \nabla\psi(\boldsymbol{\theta}) . \tag{35}$$

The last term on the right is

$$\boldsymbol{\alpha} = \frac{1}{\pi} \int \frac{\boldsymbol{\theta} - \boldsymbol{\theta}'}{|\boldsymbol{\theta} - \boldsymbol{\theta}'|^2} \kappa(\boldsymbol{\theta}') d^2\boldsymbol{\theta}' . \tag{36}$$

Thus $\psi(\boldsymbol{\theta})$ behaves as a scaled version of a gravitational potential. It is known as the lens potential and is the solution of the 2D Poisson equation

$$\nabla^2 \psi(\boldsymbol{\theta}) = 2\kappa(\boldsymbol{\theta}). \quad (37)$$

Since the logarithmic kernel in Eq. (34) is very broad, the potential will be much smoother than the density.

The scaled density $\kappa(\boldsymbol{\theta})$ is known as the convergence. We will see the reason for this name in Sect. 3.4. There is no single standard name for $\tau(\boldsymbol{\theta}, \boldsymbol{\beta})$, but Fermat potential, (scaled) time delay, or arrival time are all used.

3.2 The Einstein Radius

The quantities κ , ψ , and $\boldsymbol{\alpha}$ behave analogously to the density, classical gravitational potential, and force. In particular, a circular mass distribution behaves like a point mass when observed from outside. This fact has useful consequences, and in particular it helps understand the way κ is defined.

Consider a point mass at the origin, that is

$$\kappa(\boldsymbol{\theta}) = \pi\theta_E^2 \delta(\boldsymbol{\theta}), \quad (38)$$

in which θ_E is a constant whose meaning will be clear in a moment. The lens equation for this mass is

$$\boldsymbol{\beta} = \left(1 - \theta_E^2/|\boldsymbol{\theta}|^2\right) \boldsymbol{\theta}. \quad (39)$$

A source at $\boldsymbol{\beta} = 0$ will result in a ring image at $|\boldsymbol{\theta}| = \theta_E$ — the well-known Einstein ring. Its radius θ_E is known as the Einstein radius.

Because of the aforementioned property of circular mass distributions, the Einstein ring will be the same for any circular distribution of mass within θ_E provided the integrated κ equals $\pi\theta_E^2$ — in particular a uniform distribution with $\kappa = 1$ inside the disk.

This property lets us define a notional θ_E even when a mass distribution is not circular: as the radius of a region whose mean enclosed κ is unity.

The mass M corresponding to a given Einstein radius will be $\pi\theta_E^2 \Sigma_{cr}$ and thus

$$\theta_E^2 = \frac{4GM}{c^2 D_d} \frac{D_{ds}}{D_s}. \quad (40)$$

For a given mass, θ_E decreases with distance. However, the corresponding physical Einstein radius

$$R_E = D_d \theta_E \quad (41)$$

increases with increasing distance.

One further interpretation of θ_E follows from rewriting Eq. (40) as

$$\theta_E = \frac{4GM}{c^2 \xi} \frac{D_{ds}}{D_s}. \quad (42)$$

Recalling that GM/ξ is the squared circular speed v_c^2 at radius ξ we see that θ_E is related to kinematics. A circular speed of 300 km s^{-1} corresponds to an Einstein radius of up to $4v_c^2/c^2$ which is roughly an arcsecond.




Fig. 6 Contour maps of arrival-time surfaces $\tau(\theta)$ image positions marked. Minima are in light-blue, saddle points in orange-red, and maxima in green. Contours through saddle points are highlighted. The potential $\psi(\theta)$ is the same in all three panels, but the source positions are different

3.3 The Three Kinds of Images

Let us consider the Fermat potential or arrival time from Eq. (34), that is

$$\tau(\theta, \beta) = \frac{1}{2}|\theta|^2 - \psi(\theta) - \theta \cdot \beta \tag{43}$$

plus the constant $\frac{1}{2}|\beta|^2$, which can be discarded. It is interesting and useful to visualize $\tau(\theta, \beta)$ as a surface: the time-delay or arrival-time surface.

Without a lens, the arrival-time surface looks like a parabolic well with a minimum at $\theta = \beta$. A lensing mass pushes the surface up at nearby locations. If the lensing mass profile is circular and directly in front of the source, the minimum of the parabola will be replaced by a hill with a circular valley around it. However, unless the lens is perfectly circular and the source is perfectly centered behind it, the valley will not be perfectly symmetric, but will have a bowl on one side. Elongated mass distributions can lead to two bowls, and for more complicated mass distributions, further minima and maxima may appear. Each bowl and hill corresponds to an image location, because, by definition, $\nabla\tau(\theta) = 0$ at minima and maxima. In addition to maxima and minima there are saddle points which also have $\nabla\tau(\theta) = 0$ and hence correspond to images.

The relationship between maxima, minima, and saddle points is conveniently illustrated by contour maps of the arrival-time surface $\tau(\theta)$, such as in Fig. 6. The highlighted contours are self-crossing, and the self-intersections are saddle points. The saddle-point contours separate regions of nested simple closed curves, and each such region is either a bowl going down to a minimum, or a hill rising to a maximum. As we can see from the figure, there are two kinds of saddle-point contours. One kind looks like an infinity sign, resembling a Bernoulli lemniscate (which is given in polar coordinates by $r^2 = \cos(2\phi)$). This kind occurs between two bowls, or between two hills. The other kind of saddle-point contours resembles a limaçon trisectrix ($r = 1 - 2 \cos \phi$ in polar coordinates), which corresponds to a hill within a bowl.

The lens in Fig. 6 is an elliptical mass distribution (a commonly-used form called a PIEMD — see Appendix A). Each panel shows $\tau(\theta, \beta)$ for a different β . As is evident from the expression (43) we have a function of θ only plus a tilt due to the $\theta \cdot \beta$ term. In the right-hand panel of Fig. 6 we have a hill in the middle and two bowls above and below it. Turning to the middle panel we see the result of moving the source a little to the left, which tilts the arrival-time surface to the left; the two bowls move to the left. The left panel shows

how moving the source further to the left tilts the surface more, and makes the bowls merge into one. A still larger tilt would make the hill disappear too, leaving a single bowl.

More images can form, through nesting of the two cases illustrated in Fig. 6.

1. A minimum is split by a limaçon-like contour into a minimum and a saddle point opposite, or
2. a minimum is split by a lemniscate-like contour into two minima with a saddle point in between.

Less common is

1. the inverted form of the second case, where a maximum is split into two maxima with a saddle point in between.

The inverted form of the first case can occur in ordinary topography, but does not occur in lensing. A circular mass distribution produces case 1; case 2 (whether nested or by itself) is characteristic of elongated mass distributions. For case 3 one needs multiple, very distinct mass peaks. Images always appear and disappear in pairs: a saddle point and either a maximum or minimum. Thus, the number of images is always odd.

3.4 Magnification

Although the number of images is always odd, the number of observable images may not be. This is because different images are variously magnified, and some images can be demagnified so much that they cannot be detected anymore. Magnification and demagnification are implicit in the arrival-time surface, as we will see now.

The contour lines in Fig. 6 hint at magnification. If one looks at this figure from afar (and tries to ignore the heavy saddle-point contours as much as possible) the regions where contour lines are close together look dark, while the regions with widely-spaced contours appear to be bright. Now, consider a solution of the lens equation (2) and take a differential around it as

$$\Delta\boldsymbol{\beta} = (1 - \nabla\nabla\psi(\boldsymbol{\theta})) \Delta\boldsymbol{\theta} = \nabla\nabla\tau(\boldsymbol{\theta}) \Delta\boldsymbol{\theta}. \quad (44)$$

If the Hessian $\nabla\nabla\tau(\boldsymbol{\theta})$ is large (small), $\Delta\boldsymbol{\theta}$ will cause us to move off-source quickly (slowly), and meanwhile the contours will be close together (widely spaced). Thus, the apparent brightness between the contour lines gives some qualitative indication of magnification.

The Hessian

$$\mathbf{M}^{-1} \equiv \nabla\nabla\tau(\boldsymbol{\theta}, \boldsymbol{\beta}) = 1 - \nabla\nabla\psi(\boldsymbol{\theta}) \quad (45)$$

is an important quantity. Formally inverting Eq. (44) we can write

$$\Delta\boldsymbol{\theta} = \mathbf{M}^{-1} \Delta\boldsymbol{\beta}, \quad (46)$$

to express the lensing result of small finite sources. The inverse Hessian (or inverse curvature) of the arrival time is a tensor magnification. It is called the magnification matrix.

The inverse magnification matrix is conventionally separated into a trace and a traceless part. From the Poisson equation (37) it follows that the trace will be $1 - \kappa$. The traceless part is called the shear and denoted by γ . Thus we have

$$\mathbf{M}^{-1} = (1 - \kappa) \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \gamma \begin{pmatrix} \cos 2\varphi & \sin 2\varphi \\ \sin 2\varphi & -\cos 2\varphi \end{pmatrix}, \quad (47)$$

with κ, γ, φ all depending on θ . The meaning of the terms convergence and shear now becomes evident: κ represents an isotropic scaling, while γ represents a shearing oriented about position angle φ . The notation

$$\begin{pmatrix} 1 - \kappa - \gamma_1 & \gamma_2 \\ \gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix} \quad (48)$$

for \mathbf{M}^{-1} is also commonly used. Either way, the matrix can be diagonalized (taking $\gamma_1 \rightarrow \gamma, \gamma_2 \rightarrow 0$) by a suitable rotation of the coordinate system. The resulting diagonal elements, which are the eigenvalues, are $1 - \kappa \pm \gamma$. For minima both eigenvalues are positive; for maxima both are negative; at saddle points the eigenvalues have opposite sign. The determinant of \mathbf{M}

$$\mu \equiv [(1 - \kappa)^2 + \gamma^2]^{-1} \quad (49)$$

is a scalar magnification. It is positive for minima and maxima, and negative for saddle points, the absolute value representing the total size change of a small source. Clearly, if κ or γ become very large, the corresponding image will be demagnified and may be unobservable. At the centers of galaxies and clusters, κ does become extremely large (though γ remains modest). This produces an extreme peak in the arrival time, resulting in a demagnified image (a maximum) at that location. These demagnified maxima are often observed at the centres of lensing clusters. Lensing galaxies, however, nearly always demagnify the central maxima to unobservability.

A curious property of magnification is that there is always at least one image with absolute magnification more than 1. This appears to say that lensing can increase the total amount of light. However, this unphysical property is really an artifact of the small-angle approximation (Wucknitz 2008).

3.5 Critical Curves and Caustics

As we have seen, there are three kinds of images: minima, maxima, and saddle points of the arrival-time surface $\tau(\theta)$. As with any topographic surface, the three kinds can be identified according to the two eigenvalues of the curvature matrix $\nabla\nabla\tau(\theta)$: minima have both eigenvalues positive; maxima have both eigenvalues negative, whereas saddle points have one of each sign.

Since $\tau(\theta)$ has only a quadratic dependence on β , the curvature matrix $\nabla\nabla\tau(\theta)$ is independent of β . That is to say, although image location depends on β , the image type—if an image is present—depends only on θ . It follows that image types must be territorial: some regions can have only minima, and so on. If there are regions, there must be boundaries between the regions. The boundaries between regions of different image types are known as critical curves. One can find them by computing where the eigenvalues of the curvature matrix change sign. In general, both eigenvalues will not change sign together, so that minima and maxima do not border each other and there is always a buffer of saddle points between them.

One can map critical curves to the source plane via the lens equation, obtaining curves on the source plane, which are called caustics. Critical curves are abstract curves on the sky, they are no locations in real space. Caustics, on the other hand, are real locations on the source plane. If a point source lies on a caustic, it will light up a point on the critical curve. There may be additional images at other positions, but the image on the critical curve

is special because it has formally infinite magnification, since the curvature of the arrival-time surface is zero. This does not create a physical paradox, because (a) physically realistic sources have a finite size, and the singularity can be integrated over, and (b) wave optics would intervene to remove the singularity. If a moving source crosses from one side of a caustic to the other, one of two things happens, depending on the sense of crossing: Either two images are created at the critical curve and leave it in opposite directions, or the converse happens and two images merge and disappear on the critical curve. Positive-parity images always move in the same general direction as the source, while saddle points move the opposite way, so there is no ambiguity about which of these outcomes to expect.

In summary, critical curves are curves of extreme magnification on the sky, whereas caustics are borders on the source plane at which the number of lensed images changes. In ordinary optics, caustics usually refer to bright curves on a screen resulting from lensing of a small bright source; this just means that source and observer are swapped compared to our considerations here. In another difference of terminology, rainbows are sometimes referred to as caustics, but in the gravitational lensing convention they are critical curves.

The consequences of critical curves and caustics are enough to fill a book (Petters et al. 2001) and still leave room for new research. Fortunately, the most common cases are fairly straightforward, and in the following we will concentrate on those.

3.6 Common Image Configurations

The critical curves and caustics of a lens depend only on its potential. The images and saddle-point contours move as the source moves. Figure 7 revisits the lens used for Fig. 6 and illustrates example image configurations. The first panel shows the mass distribution and the corresponding lens potential, the second panel shows the critical curves and caustics. The rest of the figure shows image positions, types, and magnifications for six placements of the source. The lower row has the same source positions as Fig. 6. The outer critical curve separates minima and saddle points, while the inner critical curve separates saddle points and maxima. The outer critical curve corresponds to the inner caustic, and vice versa. The inner caustic has the qualitative form of an astroid ($x^{2/3} + y^{2/3} = 1$ in cartesian coordinates) but is also often called a diamond caustic.

A drastic modification is shown in Fig. 8. The elliptical PIEMD is replaced with a rectangular tile shaped like a candy bar, with a uniform $\kappa = 4$ (for details see Appendix A again). The size is chosen so that the outer critical curve has the same area as before. We see that the potential contours are still ovals, because of the logarithmic kernel in Eq. (34). The caustics are larger, and somewhat differently shaped. The critical curves are shaped very differently. Yet, if we place the source at similar locations with respect to the diamond caustic, the resulting image configurations look similar.

Figure 9 shows six lensed quasars with image configurations similar to those in Figs. 7 and 8. These have been rotated (and some cases also flipped) from the usual North-up/East-left orientation, for better comparison with the panels labeled 1–6 in Figs. 7 and 8. Let us consider each of these examples in turn.

1. Q0957+561 is a typical double (not counting the presumed demagnified image coinciding with the lensing galaxy). Doubles arise when the source is outside the diamond caustic. One image lies outside both of the critical curves, while the other lies between the two critical curves. The former is a minimum, while the latter is a saddle point in the arrival time. The minimum will naturally have a shorter arrival time. The two images form an obtuse angle with respect to the lensing galaxy, which happens when the source location is not aligned with the long or the short axes of the lensing mass. We label it an “inclined double” (cf. Saha and Williams 2003).

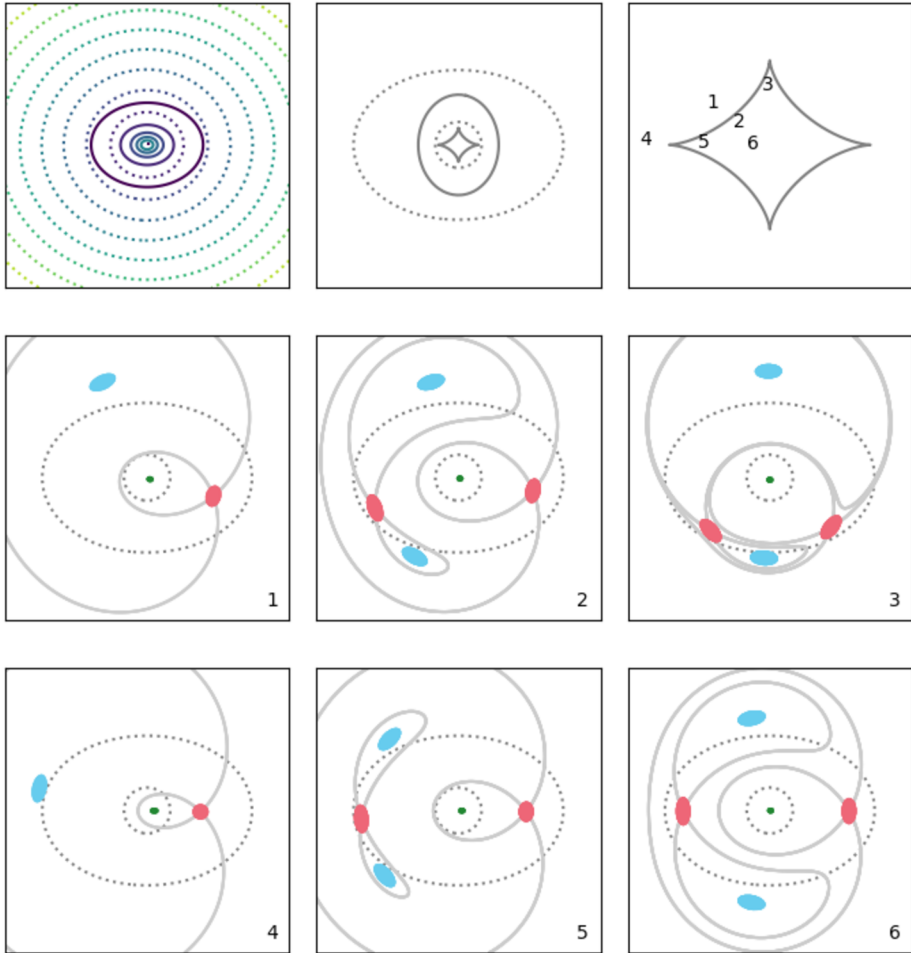


Fig. 7 Lens and image properties for a PIEMD. The upper left panel shows contours of the surface density $\kappa(\theta)$ (solid curves) and the $\psi(\theta)$ (dotted curves). The upper middle panel shows the critical curves (dotted) and caustic curves (solid). The upper right panel is a detail of the ‘diamond’ caustic with numbered source positions. The numbers correspond to panels in the middle and lower rows. The numbered panels show critical curves (dotted), saddle-point contours (solid) and images (as before, light blue: minima, orange red: saddle points, green: maxima). The contrast between magnification and demagnification, and the stretching of the lensed images, is actually much more than appears here

2. PG1115+080 is an “inclined quad”, which is the type of configuration that arises if the source lies inside the diamond caustic, and is not near any of the cusps. Compared to Q0957+561, two additional images, one a minimum and one a saddle point, form on either side of the outer critical curve. This arrangement is sometimes called a “fold type” because the source lies near a fold caustic.
3. RXJ0911+0551, with three images close together near the lensing galaxy and one image further away on the opposite side, indicates a source near the cusp aligned with the short axis of the lensing mass. It is a “short-axis quad”. Three nearly-merging images are characteristic of a source near a cusp.

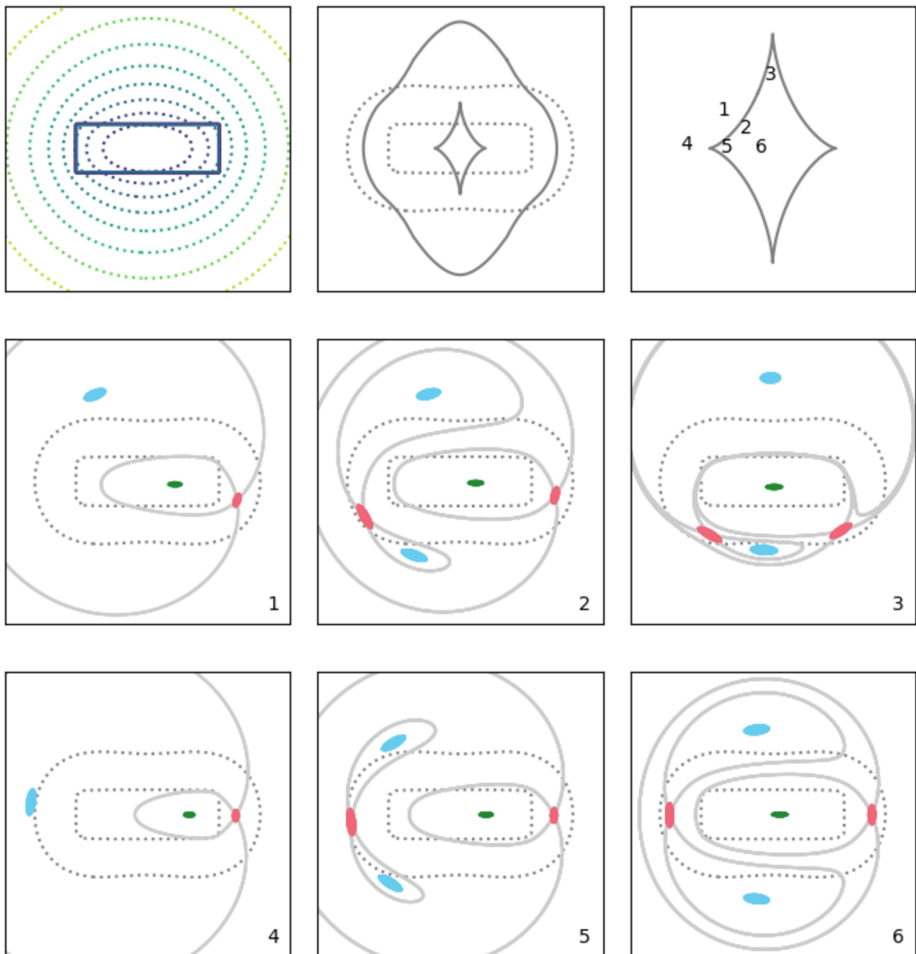


Fig. 8 As Fig. 7 but for a candybar lens

4. HE1104-1805, with two images forming nearly a straight line with the lensing galaxy, is an “axial double”. This configuration arises when the source location is aligned with either the long or short axis of the lensing mass.
5. RXJ1131-1231 is a “long-axis quad”. It differs from a short-axis quad in that the three nearly-merging images are further from the lensing galaxy than the fourth image. In this case, the source is near a cusp aligned with the long axis of the lensing mass.
6. Q2237+030 with its cross-like arrangement, is a “core quad” and must have its source near the center of the diamond caustic.

With some practice, it is possible to stare at a lens candidate and sketch conjectural saddle-point contours, thereby identifying the type of each image (minimum, maximum, or saddle point) and predicting the time ordering. Figure 10 shows an example.





Fig. 9 A collage of HST images of lensed quasars oriented so that the lensed-image positions correspond to the example configurations in Figs. 7 and 8. The systems (from top left in reading order) are Q0957+561, PG1115+080, RXJ0911+0551, HE1104-1805, RXJ1131-1231, and Q2237+030. The horizontal white tick marks $1''$ angular scale. Images were observed in wavelengths corresponding to I -, J - or H -band filter

Fig. 10 The lens candidate SW05 J1434+5228 with a conjectural identification of the image structure (taken from Küng et al. 2018) which is similar to those on panel 2 of Figs. 7 and 8, as well as PG1115+080 in Fig. 9. This type of markup can be used as input for lens-modeling, as discussed later in Sect. 4.7



3.7 Degeneracies

The various examples in Figs. 7–10 indicate two contrasting things about image configurations. On the one hand, some basic properties of the lensing mass, the general placement of the source, and the likely time-ordering of the images can be read off directly from the arrangement of images. On the other hand, the lensing mass distribution can be altered quite

drastically with little change in the images. In fact, as we will see presently, it is possible to modify a lens while keeping lensing observables exactly fixed. This is the problem of lensing degeneracies.

Let us go back to Fig. 3, and run the light backwards. That is to say, let us consider a single observer above (at the origin of the rays), and imagine sources scattered over the figure. If the sources are densely distributed and their locations are known, the observer can expect to reconstruct the entire wavefront and infer the gravitational field that caused it.

The origin of lensing degeneracies is in that lensing gives information about the gravitational field only along light paths to the observer from lensed sources. If the sources are sparse, the observer's information will be incomplete. Some clusters lenses go a long way towards providing a dense sampling of sources over a range of redshifts, and the situation will improve further with infrared telescopes in space. But for most lenses, one is working with a single source plane, of which only a small part emits light.

A more formal treatment appears in Wagner (2018) and Wagner (2019). The earlier paper characterizes the most general class of degeneracies in the strong lensing formalism for a single source and lens plane without making any assumption about the mass distribution within the lens. The second paper extends the analysis from a single lens plane to additional lenses along the line of sight, connecting the degeneracies in the lensing mass distribution in a single lens plane to degeneracies in the overall matter distribution along the entire line of sight (see especially Fig. 3 in Wagner 2019). These two works explain that one physical cause of these degeneracies is rooted in the fact that lensing only probes the integrated mass density along the entire line of sight from the observer to the source. Consequently, the distribution of the not directly observable dark matter can be redistributed in many ways. At the same time, as shown in Wagner et al. (2018) and Wagner (2019), the observable properties of multiply-imaged galaxies behind strong gravitational lenses uniquely constrain *local* properties of the lens as the shear per mass density at the locations of the multiple images and ratios of mass densities between the different multiple images. This is the maximum information that can be extracted from static multiple images without using any additional information about the overall mass density distribution. All other properties of lenses, especially in the areas far from any multiple image, are determined by the assumptions contained in the lens models.

Among these general degeneracies, there are some simple but still important ones, which we now discuss.

The first is a monopole degeneracy, which we have actually already seen. In Sect. 3.2 we noted that a circular mass distribution behaves from outside like a point mass. From inside it has no deflection. These two properties are analogous to the well-known property of spherical shells in Newtonian gravity, and can be derived using similar methods. It follows that if we replace a mass ring with another mass ring, both rings having the same center and the same mass but different radii, there will be no change in the deflection, except in the annular region between the two rings. The same holds for derivatives of the deflection, including shear and magnification. Thus, images not between the two rings will be unaffected. As a result, we are free to choose any circular or annular region not containing any images, and radially redistribute the mass in that region, with zero effect on images. Moreover, we can repeat this operation in several round regions, as long as they contain no images. At distances outside all the images, we can freely add or remove circular rings of mass.

Our next degeneracy is also very simple, though its consequences are significant. Consider the original lensing equation (2) and multiply both sides by the same factor. This will of course preserve any solution of the equation, but will scale $t(\theta)$. That is, the observed image positions will stay the same, but the time delays between different images will get

scaled, and so will the implied source sizes. The corresponding transformation of the lensing mass is a little more subtle. To derive it, we first recall Eqs. (36)–(37), and from them we first derive

$$\nabla \cdot \boldsymbol{\alpha} = 2 \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{\text{cr}}} \quad (50)$$

We also have $\nabla \cdot \boldsymbol{\theta} = 2$. Combining these relations, it is clear that multiplying both

$$\boldsymbol{\beta} \text{ and } 1 - \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{\text{cr}}} \quad (51)$$

by a constant factor amounts to multiplying both sides of the lens equation by the same factor. This degeneracy goes by various names. One is the mass-sheet degeneracy or mass-sheet transformation (or MST), because its limiting case (multiplying the lens equation by zero) makes $\Sigma(\boldsymbol{\theta}) = \Sigma_{\text{cr}}$ everywhere. Another name is the steepness degeneracy, since making the mass distribution less or more like a uniform sheet makes the mass profile steeper or shallower. Yet another name is the magnification transformation, because rescaling $\boldsymbol{\beta}$ with $\boldsymbol{\theta}$ fixed implies changing the magnification.

As expected from the wavefront picture, having more lensed sources reduces the effect of lensing degeneracies. Having more images reduces the scope of the monopole degeneracy, because there is less area available to redistribute mass. The mass-sheet transformation (51) is prevented from being applied globally if sources at different redshifts are present (because Σ_{cr} is not the same for all images), but an extension that stitches together different scalings in different regions of $\boldsymbol{\theta}$ is possible (Liesenborgs et al. 2012; Schneider 2014). There are also more complicated redistributions of mass, coupled with transformations in $\boldsymbol{\beta}$, often constructed numerically in particular cases. These are collectively referred to as shape degeneracies (Saha and Williams 2006), or source-position transformations (Schneider and Sluse 2014).

4 Lens Modeling

4.1 General Considerations

Lens modeling is the process of reconstructing the lensing mass distribution and the unlensed source brightness distribution from the observables. The principal observable is the lensed brightness distribution. Other observables — time delays, or complementary data like stellar kinematics — may be available, depending on the system. Lens and source redshifts enter as parameters, and so does a cosmological model with its parameters to convert these redshifts into distances.

There are several different lens-modeling strategies in use and many codes implementing them are available. Some results, such as measurements of enclosed mass, are robust across different modeling strategies. Other results may depend strongly on the assumptions that go into the lens models. For a quick visual impression, one can compare the mass maps and magnification maps of the lensing cluster Abell 2744⁷ obtained by five different groups using the same observations.

⁷<https://archive.stsci.edu/prepds/frontier/lensmodels/>.

There are many differences of details in how different modeling treat the data. But there is a deeper reason why different modeling methods using exactly the same data can disagree about some of the results, which applies even if the data are noiseless. The problem is lensing degeneracies, which we discussed earlier. Given a gravitational lens and a light source, computing the lensing observables is a well-posed problem. However, the inverse problem of inferring the lensing mass and the unlensed source does not generally have a unique solution, and additional information has to be provided.

The issue is that we only have information about $t(\boldsymbol{\theta})$ where the images are. At these points, we know that $\nabla t(\boldsymbol{\theta}) = 0$ and the values of $\nabla \nabla t(\boldsymbol{\theta})$. We may also know the value of $t(\boldsymbol{\theta})$ (up to a constant), if time delays are measured. Roughly speaking, time delays measure the potential difference between the image locations, the image locations themselves are probes of the gradient of the potential, while fluxes inform about the tidal field, but these pieces of information are available only at image locations. Elsewhere the information from lensing must be supplemented with other data where available, and somehow interpolated where necessary.

In the following we will discuss the essential ideas behind different modeling strategies in the literature.

4.2 Spherical Lens Models

Let us begin with the simplest.

Earlier, we have already briefly considered a point-mass lens (see Sect. 3.2 and Eq. (39)). However, while a point-mass lens is useful for modeling microlensing by individual stars within our Galaxy, it is not an adequate approximation for strong lensing by galaxies and clusters. The light paths corresponding to observed multiple images come through the dark halos of lensing galaxies, or even through regions with significant stellar mass. Thus, the first requirement of a realistic model for a galaxy or cluster lens is an extended mass distribution.

An elegant modification of a point mass, that does have an extended mass distribution, is the Plummer model from stellar dynamics,⁸ originally proposed a century ago as a model for globular clusters. The Newtonian gravitational potential of a Plummer sphere is

$$\Phi(r) = \frac{GM}{(r^2 + r_c^2)^{1/2}} \quad , \quad (52)$$

with r_c being interpreted as a core radius, which softens the singularity of a point mass. The corresponding density distribution is

$$\rho(r) = \frac{3r_c^2}{4\pi} \frac{M}{(r^2 + r_c^2)^{5/2}} \quad . \quad (53)$$

Projecting on the sky from distance D_d and writing the projected r as $D_d\theta$ and r_c as $D_d\theta_c$ we have

$$\Sigma(\theta) = \frac{M}{\pi D_d^2} \frac{1}{(\theta^2 + \theta_c^2)^2} \quad (54)$$

⁸This subsection quotes a number of results and models from stellar dynamics, all of which are explained in more detail in Binney and Tremaine (2008).

Dividing by the critical projected density gives

$$\kappa(\theta) = \frac{\theta_E^2 \theta_c^2}{(\theta^2 + \theta_c^2)^2}, \quad \theta_E^2 = \frac{4GM}{c^2 D_d} \frac{D_{ds}}{D_s}. \tag{55}$$

The Einstein radius is the same as in Eq. (40) for a point mass. It is easy to verify that the mean κ within a disk of radius θ_E is unity. The lens potential and deflection angle read

$$\begin{aligned} \psi(\theta) &= \frac{1}{2} \theta_E^2 \ln(\theta^2 + \theta_c^2), \\ \alpha(\theta) &= \frac{\theta_E^2 \theta}{\theta^2 + \theta_c^2}. \end{aligned} \tag{56}$$

As expected, the $\theta_c \rightarrow 0$ limit for all the expressions is the point mass.

The Plummer lens is useful as a component in multi-component lens models (which we will discuss below in Sect. 4.6), but by itself, it is not plausible as a galaxy halo. Galaxy halos have circular velocities which are roughly constant. Now, the circular velocity is related to the Newtonian potential by

$$v_c^2 = r \frac{\partial \Phi}{\partial r} \tag{57}$$

and from Eq. (52) we can see that $v_c(r)$ for a Plummer sphere is far from being constant.

A better first approximation to galaxy halos is the isothermal sphere and its variants. An isothermal sphere in stellar dynamics is a collisionless self-gravitating fluid of stars with a Maxwellian velocity distribution. The density is given by

$$G\rho(r) = \frac{\sigma^2}{2\pi r^2}, \tag{58}$$

in which σ is the stellar velocity dispersion which is constant for a given system. The circular velocity is easily derived from the density (58) and is given by

$$v_c^2 = 2\sigma^2 \tag{59}$$

for all r . This is the “flat rotation curve” property. The constituent stars need not be on circular orbits. The velocity distribution could be isotropic or even predominantly radial, but

$$\langle v^2 \rangle = 2\sigma^2 \tag{60}$$

will apply in all cases due to the virial theorem. The observable stellar velocity dispersion will be the average along one direction (the line of sight). Thus

$$\langle v_{\text{los}}^2 \rangle \simeq \frac{2}{3} \sigma^2, \tag{61}$$

with equality applying for an isotropic velocity distribution.

To derive the lensing properties of an isothermal sphere, we first project the density (58) in 3D to

$$G\Sigma(\theta) = \frac{\sigma^2}{2D_d\theta} \tag{62}$$

on the sky and express it in terms of a new parameter θ_E as

$$\kappa(\theta) = \frac{\theta_E}{2\theta}, \quad \theta_E = 4\pi \frac{D_{ds}}{D_s} \frac{\sigma^2}{c^2}. \tag{63}$$

As with the Plummer lens, it is easy to verify that the mean κ within a disk of radius θ_E is unity. Hence θ_E really is the Einstein radius. The lens potential and deflection are extraordinarily simple:

$$\begin{aligned} \psi(\theta) &= \theta_E \theta, \\ \alpha(\theta) &= \theta_E \frac{\theta}{\theta}. \end{aligned} \tag{64}$$

The deflection angle is simply θ_E times the radial unit vector. The deflection angle always having the same magnitude θ_E is the lensing analog of the circular speed being the same everywhere.

The simple isothermal lens described so far is often called the SIS, short for “singular isothermal sphere”. This refers to the density singularity at the center. The singularity can be removed by putting in a small core, similar to the core in the Plummer sphere, yielding a non-singular isothermal sphere, NSIS for short. In models, however, the singularity is often retained. Redistributing the central density cusp to a circular core has no effect on lensing behavior further out. The only practical effect of a central singular density is to make the central image unobservable (cf. Sect. 3.4 above). Thus, if no central image is observable, modelling the configuration with a central singularity is convenient.

Another singular property of the isothermal sphere is that the total mass is formally infinite. This follows from the density expression (58) and is the reason for M not appearing in the expressions for this model. There are variants of the isothermal sphere (called King models) which have a finite mass. This model class was used in the very first lens models (Young et al. 1980). It is, however, more common to use the SIS as is, tacitly assuming that the density is truncated outside the region of interest.

There are many more circularly symmetric lens models in use. We will not discuss more examples here, as the Plummer and SIS models illustrate the essentials. The relation between kinematics and lensing is especially worth noting. We have seen earlier (Eq. (42)) that for a point lens, θ_E is proportional to the orbital speed squared. We now see that an expression like (63) relating θ_E and the velocity dispersion is a consequence of the virial theorem. A relation of this type can always be set up for any lens in dynamical equilibrium, but the precise coefficient depends on details of the mass distribution.

4.3 Elliptical Lens Models

If a lensing mass is circular, any and all images must be in the same or opposite direction as β , because there is no other direction in the system. The arrival-time surface will be qualitatively like the left panel of Fig. 6. There will be a minimum on the same side as the source, a maximum near the lens center, and a saddle point on the opposite side. If the lens mass has a sharp peak at the center (as in the SIS), the maximum will be at the center and demagnified to vanishing brightness. Observed multiple-image systems never have all images (including demagnified images at the lens center) in a straight line. Thus, observed lenses are manifestly non-circular.

Generalising a circular lens to an elliptical one is easy, one only has to add a term

$$\gamma_1(\theta_x^2 - \theta_y^2) + 2\gamma_2\theta_x\theta_y \tag{65}$$

to the lens potential ψ of a circular lens. The additional term is called external shear (or XS). By construction, external shear is traceless and hence implies no change in κ , that is, in the mass distribution. Instead, it implies mass *outside* the region of interest. External shear is the lensing analog of a tidal field in orbital dynamics.

Additionally, one can modify the SIS or any other circular lens to make it elliptical. Thus, there is the SIEP (singular isothermal elliptical potential) where the potential of the SIS has elliptical isocontours. Then there is the PIEMD (pseudo-isothermal elliptical mass distribution) where the mass distribution κ of the SIS is both given a core and made elliptical, leading to a quite complicated potential, given below in Appendix A.2. The PIEMD was used for Figs. 6 and 7 earlier. These lens models are discussed in detail in Kassiola and Kovner (1993) and Kormann et al. (1994). A large catalog of commonly-used mass models is given in Keeton (2001).

The disadvantage of ad hoc modifications of the SIS is that they break the connection with dynamics. For circular lenses there are spherical stellar-dynamical models, which make predictions for observable kinematics in a lensing galaxy. For the SIEP or the PIEMD there are no known stellar dynamical models. One then has to either build an ellipsoidal galaxy model numerically, as done by Barnabè and Koopmans (2007) for example (and more recently Shajib 2019, has shown that an arbitrary elliptical lens can be expressed as a sum of stellar-dynamical components) or fall back on spherical models for the kinematics.

How well are real gravitational lenses approximated by elliptical lens models? To answer this question we turn to several studies, which concentrate on quasar quads, because of the precise astrometry provided by quasars. Here there are two contrasting strategies that have been developed: one is to model the observations with elliptical vs more complicated lens models; the other other is to derive and use some diagnostic for departures from ellipticity, without necessarily fitting a model. Let us consider the latter approach first.

One diagnostic for non-elliptic lenses starts from the case of an SIEP, which has

$$\psi(\theta_r, \phi) \propto r(1 + q \cos(\phi - \phi_0)). \quad (66)$$

This potential exhibits a remarkable relation between the position angles ϕ_i of the images about the lens center. For this lens Kassiola and Kovner (1995) derive what they call a configuration invariant, which can be written as

$$\cos(\phi_1 + \phi_2 - \phi_3) + \cos(\phi_2 + \phi_3 - \phi_1) + \cos(\phi_3 + \phi_1 - \phi_2) = 0 \quad (67)$$

taking the remaining position angle as $\phi_0 = 0$. This expression can be interpreted as a surface in the positional-angle space (ϕ_1, ϕ_2, ϕ_3) . Woldesenbet and Williams (2012, 2015) found that about half of the quasars lie close to a polynomial surface, which they call the fundamental surface of quads (FSQ), and image positions from the SIEP satisfy it precisely. Falor and Schechter (2022) argue that the FSQ must be an approximation to Eq. (67).

Another diagnostic for deviations from elliptical symmetry comes from Witt (1996), who considered a general, purely elliptical potential

$$\psi(\theta_x^2 + \theta_y^2/q^2) \quad (68)$$

and showed that all images, as well as the source, all lie on a rectangular hyperbola. It is not difficult to construct Witt's hyperbola. We start with a basic rectangular hyperbola $\theta_x^2 - \theta_y^2 = 1$, which has two branches. There are four simple linear transformations one can apply to the curve: (i) rotate, (ii) shift along θ_x , (iii) shift along θ_y , and (iv) rescale. Together,

they transform the curve to

$$\theta_x^2 - \theta_y^2 + c_1\theta_x\theta_y + c_2\theta_x + c_3\theta_y = c_4, \quad (69)$$

which is a general rectangular hyperbola. Writing down this equation for any four points (θ_x^k, θ_y^k) gives four linear equations for the coefficients, and solving these equations gives a rectangular hyperbola passing through all four points. An important predictive property of Witt's hyperbola is that it applies to any central image as well, and since for a centrally peaked potential the central image is at the lens center, the center of the lensing galaxy will lie on the hyperbola.

Witt's hyperbola itself is independent of a specific choice of an elliptical lens model, but can be used to find lens models via a further construction (Wynne and Schechter 2018). This is to construct an ellipse (Wynne's ellipse) aligned with Witt's hyperbola that passes through the four images of the quad. The four images serve to set the (i) ellipticity, (ii) shift along θ_x , (iii) shift along θ_y , and (iv) scale of Wynne's ellipse. The parameters of the Witt-Wynne construction (see Fig. 11 for an example) translate into parameters for two distinct lens models: SIS+XS (Wynne and Schechter 2018) and SIEP with aligned XS (Luhtarv et al. 2021).

The other approach, that of fitting elliptical and beyond-elliptical models to the data (not only the lensed quasar positions but also extended images) has been pursued by several authors. An early example is the study of PG 1115+080 by Yoo et al. (2005), who concluded that the lens is indeed an ellipse. Sluse et al. (2012) carried out a comparable study of diverse quasar quads, while also identifying perturbers in the form of nearby galaxies. There are also several works (Gomer and Williams 2018; Van de Vyvere et al. 2022a,b) which introduce non-elliptical features (twists, substructures) into simulated lenses, to test how much they affect perturbed models and see how they affect modeling.


While these studies are restricted to quadruply-imaged quasars, they support a more general conclusion that simple elliptical lenses are a good first-order approximation for most lenses, and the counter-examples of large departures from elliptical lenses are usually flagged by the presence of additional lensing galaxies that perturb the main lens. Accurately fitting observations thus requires more features to be added. Schmidt et al. (2023) and Etherington et al. (2022) use elliptical lenses, including components for additional lensing galaxies where necessary, to model 30 and 59 lenses, respectively. In contrast, Collett et al. (2018) and Shajib et al. (2023) each study a single galaxy lens in more detail, using 3D galaxy models and spatially resolved kinematic data. Meanwhile, machine-learning approaches are also being developed (e.g., Gomer et al. 2023) towards more efficient implementations of these models.

4.4 Bayesian Terminology

Lens models, especially if they are rather elaborate with many parameters, are often formulated in Bayesian terms, the key words being *likelihood*, *prior*, and *posterior*, all three being probability distributions.

To introduce these terms, let us write ϖ for the model parameters and \mathcal{D} for the data. The ϖ are to be understood in a very general sense, they can consist not only of numbers but also of flags specifying the type of model, such as SIEP or PIEMD and so on, and include the location and other properties of the unlensed source. The \mathcal{D} consist of the images of the lensed system and of other available observations (e.g. time-delays). The likelihood is $p(\mathcal{D} | \varpi)$, meaning the conditional probability of \mathcal{D} being as measured for given ϖ . Basically, the likelihood is the probability of the noise difference between \mathcal{D} and hypothetical

Fig. 11 The Witt-Wynne construction for the quad PS J0147+4630 (from Luhtaru et al. 2021). The image configuration is similar to those on panel 5 of Figs. 7 and 8. The ellipse is similar though not identical to the critical-shear ellipse introduced by Wucknitz (2002) and its short axis is aligned with the long axis of the potential



noiseless data from ϖ . The prior is $p(\varpi)$, and it expresses our knowledge of the parameters before having any data. The posterior

$$p(\varpi | \mathcal{D}) \propto p(\mathcal{D} | \varpi) p(\varpi) \tag{70}$$

is then given by Bayes’ theorem in probability, and expresses our knowledge of the parameters given the data.

The likelihood is usually conceptually clear, though it may involve a lot of computation. Typically it has a Gaussian form $p(\mathcal{D} | \varpi) \propto \exp(-\frac{1}{2}\chi^2)$. The prior $p(\varpi)$, on the other hand, may be the subject of debate. If the likelihood factor in Eq. (70) has a sharp peak in ϖ , then the posterior will have the same sharp peak, and the prior will have little influence. If, however, the likelihood does not have a single sharp peak, the choice of prior will strongly influence the posterior.

Actually, the situation with ϖ is more complicated. Because of lensing degeneracies, the likelihood can be insensitive to some parameters. Let us divide the parameters into ‘sharp’ and ‘blunt’ $\varpi = (\varpi_s, \varpi_b)$ such that the likelihood $p(\mathcal{D} | \varpi_s, \varpi_b)$ is sharply peaked in ϖ_s and blunt in ϖ_b . We can then marginalize out the blunt parameters

$$p(\varpi_s | \mathcal{D}) = \int p(\varpi | \mathcal{D}) d\varpi_b \tag{71}$$

to extract the sharp (that is, well-constrained) parameters. As a simple example, consider a system with a clean Einstein ring and no other data. In this example ϖ_s as the enclosed mass and ϖ_b as the radial profile are a possible parameter separation. Suppose, however, that we use a different parametrization that mixes sharp and blunt — say we take $\varpi = (\varpi_1, \varpi_2)$ where ϖ_1 is a power-law slope and ϖ_2 is a normalization — and marginalize out the normalization. We would then get an apparent constraint on the power-law slope, that is really coming from the prior. In this simple example, the problem is easy to diagnose and fix, but the situation gets worse with increasing the number of parameters (distributed between ϖ_s and ϖ_b), and hence the hyper-volume of the parameter space.

4.5 Source Reconstruction

When considering image formation earlier, we assumed that lensed images are small enough to have an unambiguous position. This is clearly valid for lensed quasars and supernovae, and reasonable when speaking of small components in a lensed galaxy, but it is not valid for arcs formed from the lensing of entire galaxies. In such cases one is observing the superposition of a continuous family of lensed sources. In fact, the recent modeling papers we mentioned above have to reconstruct the sources, as well as modeling the lenses.

It is possible to reconstruct a lensed source independently of any lens model, provided there is enough detail in the lensed images. Tessore (2017) and Birrer (2021) develop formalism for doing so, but to explain the general idea, let us consider an example, namely the spectacular five-image system in the lensing cluster CL 0024+1654. Wagner et al. (2018) draw attention to six small components in the source galaxy, all of which appear unambiguously in all five lensed images. They then reconstruct the source in a model-independent way, and compare the results with those obtained by modeling the lens first. Such systems are not common, but nevertheless it is helpful to discuss why model-independent source reconstruction works. Let us think of the six components on the source as being one reference β_0 and five $\Delta\beta$ vectors. These $\Delta\beta$ are magnified and sheared into five sets of five observable $\Delta\theta$ vectors each by five different magnification matrices. The original $\Delta\beta$ vectors are unknown (5×2 numbers, because we are in 2D). The five magnification matrices are also unknown (3×5 numbers, because magnification matrices are symmetric and have 3 independent components). This yields a total of 25 unknowns. Meanwhile, the $\Delta\theta$ vectors provide $5 \times 5 \times 2 = 50$ constraints. The latter suffice to solve for all $\Delta\beta$ vectors, except for an overall scale, which is degenerate. The reference position β_0 of the source is not solved for, either, as it also degenerate. Another example, this time for a galaxy lens, appears in the study of B0128+437 by Wagner and Williams (2020), using four images with three components each.

More commonly, however, the source and lens are reconstructed together. Consider an extended source with brightness distribution $s(\beta)$. When lensed, this source will produce an extended image system

$$I(\theta) = \int \delta[\theta - \alpha(\theta) - \beta] s(\beta) d^2\beta, \quad (72)$$

where $\delta[\dots]$ is the delta function. The lens modeling problem is now to recover $\alpha(\theta)$ and $s(\beta)$ from observations of $I(\theta)$ and any other available information.

Let us suppose that a lens model is available. This may be a trial lens model produced in one step of some larger iterative process. But let us say the lens is given, and in particular the deflection field $\alpha(\theta)$ is given. Equation (72) then takes the form

$$I(\theta) = \int L(\theta, \beta) s(\beta) d^2\beta, \quad (73)$$

where $L(\theta, \beta)$ is given. In practice, there will be a further convolution by a point spread function (psf) or other instrumental response function, but if the psf is known, it can be absorbed inside $L(\theta, \beta)$. Under these assumptions, the source reconstruction becomes a linear image-reconstruction problem.

For optical imaging, noise is generally Gaussian. We can then write the likelihood as $\exp(-\frac{1}{2}\chi^2)$, where

$$\chi^2 = \int \sigma^{-2}(\theta) (I^{\text{data}}(\theta) - I(\theta))^2 d^2\theta, \quad (74)$$

with $\sigma(\boldsymbol{\theta})$ being the noise dispersion. We are writing an integral here for notational convenience, but it is actually a sum over pixels. There are several possible ways to go about solving for $s(\boldsymbol{\alpha})$. One can assume a parametric form for $s(\boldsymbol{\beta})$ and fit for the parameters by minimizing χ^2 . Alternatively, one can adopt a basis set for $s(\boldsymbol{\beta})$ and fit for the coefficients in that basis set by minimizing χ^2 . Birrer et al. (2015) present a suitable basis set for doing so (shapelets), while Galan et al. (2021) consider wavelets. A third option is to put $s(\boldsymbol{\beta})$ on a grid. In this approach, a simple minimization of χ^2 is actually not recommended, because it tends to amplify the effect of noise, leading to artifacts like negative-brightness pixels or checkerboard patterns. To suppress such artifacts, a penalty function is added to χ^2 and the sum is minimized. The choice of penalty function — or equivalently, the prior on the source light distribution — itself has a large literature in image reconstruction (for example, in interferometry — see Thiébaud and Young 2017) and is known as regularization. Warren and Dye (2003) in an influential paper recommend

$$|\nabla^2 s(\boldsymbol{\beta})|^2 \quad (75)$$

as the regularization function, which favors smoother $s(\boldsymbol{\beta})$. It can be given a small or large weight, relative to χ^2 , for gentle or aggressive smoothing, as desired. Further developments of the idea appear in several later works (Suyu et al. 2006; Vegetti and Koopmans 2009; Vernardos et al. 202). An advantage of a quadratic regularization function like (75) is that the minimization procedure requires solving only linear equations. But it is not a requirement, and source reconstruction using more complicated regularization functions is also possible, notably in the early work by Kochanek et al. (1989) and later in Brewer and Lewis (2006).

4.6 Many-Component Lens Models

Whereas for galaxy lenses a single elliptical component is often used, while multiple components are becoming common, in cluster-lens models multiple components are a must from the start. This is partly because clusters are further from a long-term stable state than galaxies and tend to have more complicated mass distributions, but mainly because cluster lenses present many more lensed images and thus more constraints. An illustrative example are the models of Abell 2744 by Bergamini et al. (2023), consisting of components for individual galaxies, with a larger-scale component for the diffuse matter in the cluster not identified with any individual galaxy.

The various modeling approaches we have discussed so far introduce a lot of parameters to be fitted, but the system remains overdetermined, meaning that there are more data points than parameters. A different category of multi-component methods reconstructs the lens in a basis set such as a grid, where the number of components exceeds the number of data points. By itself, such a free-form reconstruction is underdetermined and non-unique. To solve that problem, regularization can be applied, much as in the source reconstruction discussed above. Using a regularization, one can search, for instance, the smoothest model that fits the data, or the model the most similar to the light distribution. An early example of this technique in lensing is the reconstruction of Abell 370 by Abdelsalam et al. (1998). The mass map of Abell 2744 by Finney et al. (2018) is typical of more recent work on these lines. Hybrids with basic elliptical components and an additional substructure component which is free-form with regularization are also possible (e.g. Vernardos et al. 202; Galan et al. 2022).

4.7 Model Ensembles

What if we specify neither a simple parametric form nor a regularization? In Bayesian terms, this means there are more parameters in ϖ than there are data points in \mathcal{D} , while $p(\varpi)$ imposes some requirements on ϖ but has no hierarchy of favored and less-favored models. In practical terms, there is no single optimal model, only an ensemble of models. The ensemble provides confidence regions on any desired quantity, which can be easily extracted by marginalization. Conceptually, lens-model ensembles are like Monte-Carlo Markov chains widely used for likelihood or posterior-probability sampling.

Model ensembles may contain extreme models, but that does not mean the posterior $p(\varpi | \mathcal{D})$ is uniform in all properties between the extremes. To give an analogy, consider an ensemble of random walks with a given number and length of steps, and no preference between different random walks. The ensemble will span every random walk from zero net displacement to a straight line (i.e., the maximum displacement possible), but the distribution of displacements will not be uniform between these extremes.

A very general implementation of these ideas for the reconstruction of cluster lenses is the GRALE code (introduced by Liesenborgs et al. 2006, 2007, with many later enhancements). In GRALE each model consists of a large number of Plummer lenses (Eq. (56)) and there is an ensemble or population of models that evolves according to a genetic algorithm to fit the lensing data. The prior is implicitly defined by the algorithm rather than explicit, but is minimal: the mass distribution must be non-negative. The likelihood is not Gaussian, however it does incorporate the data on image locations and shapes, and also the absence of lensed images where they are not observed. The minimal prior means that the model ensemble tends to include a diversity of models degenerate with respect to the observables. Hence the inferred uncertainties are generally larger than in other techniques. The reconstruction of Abell 2744 by Sebesta et al. (2019) is illustrative.

Since galaxy lenses provide much less data than cluster lenses, a more informative prior $p(\kappa)$ is necessary. An explicit but still very weak prior (amounting to requiring the lens to be centrally concentrated) has been used in several modeling codes (Williams and Saha 2000; Saha and Williams 2004; Coles et al. 2014; Küng 2018). Here we mention two contrasting applications. Denzel et al. (2021) used the model-ensemble method to model eight time-delay quasars simultaneous, with H_0 as a shared parameter to be inferred. In (Küng et al. 2018) citizen scientists modeled tens of galaxy lenses by marking up each lens candidate and using a model-ensemble code to quickly explore the range of models consistent with the markup. Figure 10 shows an example of input markup from that work.

Finally, we mention yet another recent strategy, which amounts to finding plausible matches in galaxy-formation simulations instead of trying to reconstruct the lens. That is, the prior $p(\kappa)$ can be a large ensemble of simulated galaxies, which is then filtered to a much smaller ensemble that is consistent with the available data. Any such technique will be computationally expensive, perhaps unrealistically so for clusters, but is starting to become feasible for galaxy lenses (Denzel et al. 2021). What is attractive is the prospect of bypassing lens modeling altogether, and using lensing observations to directly address questions about galaxy structure and formation.

4.8 Model-Comparison Studies

It is acceptable for different lens models (i.e. models that differ in lens/source reconstruction) to disagree on some derived quantities as long as the quantities important for a specific science goal are accurately recovered. Strong lensing very accurately recover the Einstein

radius and the mass enclosed within that radius, but other quantities (e.g. the ellipticity of the mass distribution) may be less accurately recovered. Some of these uncertainties on recovered parameters arise because of the degeneracies mentioned earlier, this is what is commonly referred to as “known unknowns”, but other differences arise because of “unknown unknowns”. The latter are harder to identify and require model comparisons.

We have already mentioned Abell 2744 several times, in connection with different modeling approaches. The system is one of six lensing clusters (the Frontier Field clusters) which have been likewise modeled by several groups. There have further been what one may call meta-analyses (Priewe et al. 2017; Remolina González et al. 2018; Raney et al. 2020) which study and interpret the effect of different choices made by different methods. These comparisons demonstrate a global robustness of models in regions well constrained by the data. They also highlight the complementarity of the parametric and free-form models.

Another kind of model-comparison study are lens modeling challenges, where different groups are invited to reconstruct a lens from mock data. Meneghetti et al. (2017) and (Ding et al. 2021) are two such blind tests of the methods of several groups, covering cluster lenses and galaxy lenses respectively. Such studies are a useful indicator of robustness and uncertainties in the current state of the art.

A third kind of model comparison consists in blindly forecasting some observables with various models. One spectacular case of this exercise happened with the observations of the lensed supernova “SN Refsdal”. After the observations of the first four lensed images of the supernova, several groups gathered to predict the reappearance position, brightness of the 5th image (located $\sim 8''$ from the first quadruplet) about hundreds of days after the previous image (Kelly et al. 2023a,b). A fraction of the models successfully reproduced the data, enabling the use of this first multiply imaged lensed supernova for H_0 inference. The success of such a model comparison is very encouraging as numerous supernovae may be discovered in upcoming “all-sky” surveys (see the chapter on transients).

The comparison of models is a useful but difficult exercise because differences between models may arise not only from fundamental degeneracies and prior choices as discussed above, but also from subtle bugs that do not teach much about lensing but reveal a need to fix a code. Importantly, it is also hampered by the difficulty of comparing codes on the same grounds. The comparison of deflection and magnification maps offers a simple mean of comparison but extra work is required to derive differences and biases on specific parameters, such as e.g. the slope of the mean slope of the density profile, or the position angle of the main axis of the mass distribution. An additional hurdle is the difference of conventions in defining some of these quantities. The COOLEST framework (Galan et al. 2023), which provides an infrastructure to gather and compare results of lens models in a painless way, may help and trigger more model comparisons on large ensembles of lenses and ensure a good control of systematic uncertainties.

4.9 Reflections on Lens Modeling

As we remarked earlier (in Sect. 4.1) strong-lensing observations measure the projected gravitational potential its derivatives at the image locations, and the task of lens modeling is, in effect, to interpolate the values elsewhere. Doing so requires additional assumptions, and preferably complementary data as well.

A useful analogy is the density inside the Earth. We can measure gravitational acceleration at the surface or from orbit, and this gives the total mass very accurately, and also provides information on multipoles, but it does not give the radial density profile. That requires theory on the structure of planets, and further data like seismography.

Likewise lens modeling depends on our knowledge of and uncertainties about the structure and dynamics of galaxies and clusters. Are lensing galaxies all isothermal ellipsoids to the few-percent level? Perhaps not. Are dark halos so full of ongoing mergers that all equilibrium models are useless? Perhaps not, either. But between these extremes, there is a diversity of views among researchers in lensing.

The plurality of methods and data used for modeling lens systems may look difficult to unite at first sight. However, the reconstruction of the lensing mass density generally results from a trade-off between model rigidity and flexibility chosen in the light of the data and scientific questions posed, a situation that machine learning scientists may frame in terms of bias and variance arising for under-/over-fitting of the data. But beware, the choice of a model over another is not uniquely governed by those considerations. Inductive biases, namely biases introduced by the choice of a method, also exist. They are inherent to scientific research and evolve hands in hand with scientific results. A diversity of views and methods is often a good sign that our understanding of the world progresses. Strong lensing is no exception to this research walk.

5 Conclusions and Outlook

The introduction to strong gravitational lensing presented here starts by inviting the reader to think how gravitational fields will affect geometrical wavefronts propagating through the Universe. This leads to the main theoretical construct, the Fermat potential or arrival-time surface. The first and second derivatives of the arrival time surface are the lens equation and the distortion (or magnification) matrix, respectively, which are the main players in the quantitative analysis of the problem. The equations are the same as in other approaches that use only geometrical optics to describe light propagation, but many properties of lensing follow from geometrical arguments even without equations. In particular, visualising the arrival-time surface for a specific lensing configuration allows the researcher to gain a qualitative understanding of the problem.

The paper also introduced the reader to the other two essential components of strong lensing: critical curves and caustics. With these theoretical concepts under their belt, the reader is ready to learn about lensing as a probe in astrophysics. As its fundamentals are well understood, the primary application of lensing in astrophysics is to build mass models of the lens, and to recover the structure of the background source. Most of the mass in galaxy and cluster lenses is dark, so lensing is very handy here. Distant, high-redshift sources tend to be faint and small; cosmic lenses make sources larger, allowing for better spatial resolution. We also consider some simple lens models, and briefly describe more sophisticated ones used in the literature.

This short introduction to the essentials of strong lensing should provide the reader with the basic principles and physical intuition as a starting point to dive deeper into this growing field. The applications of lensing are ubiquitous: hardly any aspect of modern astrophysics has been untouched by it. The other papers in this collection cover most, if not all the current, as well as emerging applications of lensing: Search for strong gravitational lensing through serendipitous and deliberate discoveries, Strong lensing by galaxies, Strong lensing by galaxy clusters, Strong lensing and microlensing of quasars, Microcaustics near macrocaustics, Gravitational lensing as a cosmic telescope, Probing the nature of dark matter on (sub-)galactic scales with strong lensing, Time-Delay Cosmography and the estimation of the Hubble Constant and other cosmological parameters, Strong lensing and microlensing

of supernovae, Strong lensing of Gamma-Ray Bursts, Fast Radio Bursts, and Gravitational waves, and Diffraction in strong gravitational lensing.

We hope the reader enjoys this and other papers in this series, and finds strong gravitational lensing as enthralling as the authors do!

Appendix A: Example Lens Models Used

A.1 A Softened Schwarzschild Metric

Consider the metric

$$ds^2 = -e^{2\Phi(r)} dt^2 + e^{-2\Phi(r)} dr^2 + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \quad (76)$$

If we put

$$2\Phi(r) = \ln \left(1 - \frac{2GM}{c^2 r} \right) \quad (77)$$

we would have a Schwarzschild metric. For the example in Figs. 2–4 we instead we put

$$\Phi(r) = -\frac{GM}{c^2} \frac{1}{\sqrt{r_c^2 + r^2}} \quad (78)$$

where r_c is a small softening length.

This metric has no especially interesting physical interpretation. It is just used as an example of a strong gravitational field without the additional complication of an event horizon.

A.2 PIEMD

The so-called pseudo-isothermal elliptical mass distribution, used as an example in Figs. 6 and 7 is defined via two elliptical radii

$$\begin{aligned} R_e^2 &= q^2(s^2 + \theta_x^2) + \theta_y^2 \\ R_f^2 &= (1 - q^2)\theta_x^2 + (R_e + s)^2 \end{aligned} \quad (79)$$

Then the density is

$$\kappa = \frac{1}{2R_e} \quad (80)$$

the deflection is

$$\begin{aligned} \alpha_x &= \frac{1}{\sqrt{1 - q^2}} \arctan \left(\frac{\sqrt{1 - q^2} \theta_x}{R_e + s} \right) \\ \alpha_y &= \frac{1}{\sqrt{1 - q^2}} \operatorname{arctanh} \left(\frac{\sqrt{1 - q^2} \theta_y}{R_e + q^2 s} \right) \end{aligned} \quad (81)$$

while the potential and its second derivatives are

$$\begin{aligned}\psi &= \theta_x \alpha_x + \theta_y \alpha_y - s \ln R_f \\ \psi_{xx} &= \frac{(qs)^2 + R_e s + \theta_y^2}{R_e R_f} \\ \psi_{yy} &= \frac{s^2 + R_e s + \theta_x^2}{R_e R_f} \\ \psi_{xy} &= -\frac{\theta_x \theta_y}{R_e R_f}\end{aligned}\quad (82)$$

A.3 A Candy-Bar Lens

The candy-bar lens, used as an example in Fig. 8, is not by itself a serious lens model, but can be a useful component of one. Square mass tiles are used in many free-form models. Another application is in microlensing, where a rectangular mass mass tile is removed from a smooth mass distribution and replaced with point masses. To compute the potential, we need the integral from Eq. (34) for constant κ , that is to say

$$\int \ln(\theta_x^2 + \theta_y^2) d\theta_x d\theta_y \quad (83)$$

over a rectangular region. The result and its derivatives can be expressed in terms of arctans and logs. Expressions for a square are given Saha and Williams (1997) and for a rectangle in Zheng et al. (2022).

Appendix B: Constants and Unit Conversions

Astronomical distances are commonly given in many different units. When contrasting scales are involved, as they often are in gravitational lensing, it is rather convenient to convert everything to light-seconds. Thus, while a parsec is defined as the distance from which an astronomical unit (au) subtends an angle of 1 arcsec, with the au being $1.495978707 \times 10^{11}$ m, the values

$$\begin{aligned}1 \text{ au} &= c \times 499.0 \text{ s} \\ 1 \text{ pc} &= c \times 1.029 \times 10^8 \text{ s}\end{aligned}\quad (84)$$

are easier to work with, and more digits can be added if necessary. Similarly, the solar mass can be expressed as

$$\frac{GM_\odot}{c^3} = 4.925 \times 10^{-6} \text{ s} \quad (85)$$

The Hubble constant is commonly given in the form

$$H_0 = h_{70} \times 70 \text{ km s}^{-1} \text{ Mpc}^{-1} \quad (86)$$

The following conversions may be useful.

$$\begin{aligned}
 H_0^{-1} &= h_{70}^{-1} \times 14.0 \text{ Gyr} \\
 cH_0^{-1} &= h_{70}^{-1} \times 4.28 \text{ Gpc} \\
 &= h_{70}^{-1} \times 20.8 \text{ kpc arcsec}^{-1}
 \end{aligned}
 \tag{87}$$

The critical density (from Eq. (32)) comes to

$$\frac{cH_0}{4\pi G} = h_{70} \times 0.811 \text{ kg m}^{-2}
 \tag{88}$$

times a dimensionless factor which is typically $\sim 1/z_d$.

The cosmological total density ρ_{cr} (from Eq. (14)) can be expressed in several ways, such as

$$\frac{3H_0^2}{8\pi G} \frac{c^2}{e} = h_{70}^2 \times 5.15 \text{ GeV m}^{-3}
 \tag{89}$$

Appendix C: Lensed Waves

If a source is large enough, that the distance to the observer from different parts of the source differ by more than a wavelength, radiation from it is incoherent at the observer. This is the case for nearly all astronomical sources. For these, geometrical optics applies. In particular, the arrival-time surface, whether in its scaled form $\tau(\boldsymbol{\theta}, \boldsymbol{\beta})$ or in its dimensional form $t(\boldsymbol{\theta}, \boldsymbol{\beta})$ (see Eqs. (32)–(34)), is defined over the whole sky, but only its zero-gradient points correspond to images.

For sources that are small enough and far enough, however, the whole $t(\boldsymbol{\theta})$ surface interferes. This applies to gravitational waves (see e.g., Dai et al. 2020) and may apply to fast radio bursts (Leung et al. 2022). At frequency ω the resulting electric field of the light (or gravitational wave strain) is given by

$$\int \exp(i\omega t(\boldsymbol{\theta})) d^2\boldsymbol{\theta}
 \tag{90}$$

This is sometimes called the Fresnel-Kirchoff diffraction integral. Of particular interest is the quotient

$$\frac{\int \exp(i\omega t_L(\boldsymbol{\theta})) d^2\boldsymbol{\theta}}{\int \exp(i\omega t_{NL}(\boldsymbol{\theta})) d^2\boldsymbol{\theta}}
 \tag{91}$$

of the lens and no-lens cases.

With no lensing

$$t_{NL}(\boldsymbol{\theta}) = \frac{1}{2} T_{\text{ds}} \left((\theta_x - \beta_x)^2 + (\theta_y - \beta_y)^2 \right)
 \tag{92}$$

with the scale T_{ds} being as defined in from Eq. (32). With a lens present, things are of course more complicated. For one regime, however, namely where the wave period is much shorter than gravitational time delays (that is, $\omega t(\boldsymbol{\theta}) \gg 1$), the integrals can be usefully estimated. The idea is to Taylor-expand $t(\boldsymbol{\theta})$ near a stationary point (say $\boldsymbol{\theta}^*$) where of course

an ordinary image would be located. Near a stationary point there is no linear dependence on $\theta - \theta^*$ and the quadratic terms are given by the Hessian (47) matrix. We can assume, without loss of generality, that the Hessian is diagonal. Let us introduce

$$a, b = 1 - \kappa \pm \gamma \quad \mu = 1/(ab) \quad (93)$$

and then μ will be the scalar magnification. In terms of these we have

$$t_L(\theta) = t(\theta^*) + \frac{1}{2} T_{ds} (a(\theta_x - \theta_x^*)^2 + b(\theta_y - \theta_y^*)^2) \quad (94)$$

Now both the integrals in (91) have the form

$$\int_{-\infty}^{\infty} \exp\left(\frac{\pm i u^2}{2\sigma^2}\right) du = \frac{e^{\pm i\pi/4}}{\sqrt{2\pi} \sigma} \quad (95)$$

Note changing the sign in integrand changes the phase of the integral by $\pi/2$. We need to pay attention to the signs of ω , a , and b . Bearing this mind and simplifying, we get

$$|\mu|^{1/2} \exp(i\omega t(\theta^*) - i\Delta) \quad (96)$$

where

$$\Delta = \frac{1}{2}\pi \operatorname{sgn} \omega \times \begin{cases} 0 & \text{for minima} \\ 1 & \text{for saddle points} \\ 2 & \text{for maxima} \end{cases} \quad (97)$$

The Δ term here is known as the Morse phase.

Acknowledgements We would like to thank the International Space Science Institute in Bern, Switzerland (ISSI) for their hospitality, and the conveners for organizing the stimulating workshop on “Strong Gravitational Lensing”. We thank our co-participants at the workshop for many discussions and suggestions, and the reviewers for their comments. We especially thank P. L. Schechter and P. Schneider for sharing many insights.

Funding Open access funding provided by University of Zurich.

Declarations

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Abdelsalam HM, Saha P, Williams LLR (1998) Mon Not R Astron Soc 294:734. <https://doi.org/10.1046/j.1365-8711.1998.01356.x>

- Akiyama K, Alberdi A, Alef W et al (2022) *Astrophys J Lett* 930:L17. <https://doi.org/10.3847/2041-8213/ac6756>
- Barnabè M, Koopmans LVE (2007) *Astrophys J* 666:726. <https://doi.org/10.1086/520495>
- Bartelmann M, Schneider P (2001) *Phys Rep* 340:291. [https://doi.org/10.1016/S0370-1573\(00\)00082-X](https://doi.org/10.1016/S0370-1573(00)00082-X)
- Bergamini P, Acebron A, Grillo C et al (2023) *Astrophys J* 952:84. <https://doi.org/10.3847/1538-4357/acd643>
- Binney J, Tremaine S (2008) *Galactic dynamics*, 2nd edn. Princeton University Press, Princeton
- Birrer S (2021) *Astrophys J* 919:38. <https://doi.org/10.3847/1538-4357/ac1108>
- Birrer S, Amara A, Refregier A (2015) *Astrophys J* 813:102. <https://doi.org/10.1088/0004-637X/813/2/102>
- Blandford R, Narayan R (1986) *Astrophys J* 310:568. <https://doi.org/10.1086/164709>
- Borgeest U (1983) *Astron Astrophys* 128:162
- Brewer BJ, Lewis GF (2006) *Astrophys J* 637:608. <https://doi.org/10.1086/498409>
- Broderick AE, Pesce DW, Gold R et al (2022) *Astrophys J* 935:61. <https://doi.org/10.3847/1538-4357/ac7c1d>
- Cassan A, Ranc C, Abil O et al (2022) *Nat Astron* 6:121. <https://doi.org/10.1038/s41550-021-01514-w>
- Coles JP, Read JI, Saha P (2014) *Mon Not R Astron Soc* 445:2181. <https://doi.org/10.1093/mnras/stu1781>
- Collett TE, Oldham LJ, Smith RJ et al (2018) *Science* 360:1342. <https://doi.org/10.1126/science.aao2469>
- Cooke JH, Kantowski R (1975) *Astrophys J Lett* 195:L11. <https://doi.org/10.1086/181697>
- Crosta M, Geralico A, Lattanzi MG, Vecchiato A (2017) *Phys Rev D* 96:104030. <https://doi.org/10.1103/PhysRevD.96.104030>
- Dai L, Zackay B, Venumadhav T, Roulet J, Zalzarriaga M (2020). ArXiv:e-prints. [arXiv:2007.12709](https://arxiv.org/abs/2007.12709)
- Demkov YN, Puchkov AM (2000) *Phys Rev D* 61:083001. <https://doi.org/10.1103/PhysRevD.61.083001>
- Denzel P, Coles JP, Saha P, Williams LLR (2021) *Mon Not R Astron Soc* 501:784. <https://doi.org/10.1093/mnras/staa3603>
- Denzel P, Mukherjee S, Saha P (2021) *Mon Not R Astron Soc* 506:1815. <https://doi.org/10.1093/mnras/stab1716>
- Ding X, Treu T, Birrer S et al (2021) *Mon Not R Astron Soc* 503:1096. <https://doi.org/10.1093/mnras/stab484>
- Elliot JL, Olkin CB (1996) *Annu Rev Earth Planet Sci* 24:89. <https://doi.org/10.1146/annurev.earth.24.1.89>
- Etherington A, Nightingale JW, Massey R et al (2022) *Mon Not R Astron Soc* 517:3275. <https://doi.org/10.1093/mnras/stac2639>
- Falor C, Schechter PL (2022) *Astron J* 164:120. <https://doi.org/10.3847/1538-3881/ac80bc>
- Finney EQ, Bradač M, Huang K-H et al (2018) *Astrophys J* 859:58. <https://doi.org/10.3847/1538-4357/aabf97>
- Galan A, Peel A, Joseph R, Courbin F, Starck JL (2021) *Astron Astrophys* 647:A176. <https://doi.org/10.1051/0004-6361/202039363>
- Galan A, Van de Vyvere L, Gomer M, Vernardos G, Sluse D (2023) *J Open Sour Softw* 8:5567. <https://doi.org/10.21105/joss.05567>
- Galan A, Vernardos G, Peel A, Courbin F, Starck JL (2022) *Astron Astrophys* 668:A155. <https://doi.org/10.1051/0004-6361/202244464>
- Gavazzi R, Treu T, Koopmans LVE et al (2008) *Astrophys J* 677:1046. <https://doi.org/10.1086/529541>
- Giblin B, Heymans C, Asgari M et al (2021) *Astron Astrophys* 645:A105. <https://doi.org/10.1051/0004-6361/202038850>
- Gomer MR, Williams LLR (2018) *Mon Not R Astron Soc* 475:1987. <https://doi.org/10.1093/mnras/stx3294>
- Gomer MR, Ertl S, Biggio L et al (2023) *Astron Astrophys* 679:A59. <https://doi.org/10.1051/0004-6361/202347507>
- Grimm N, Yoo J (2018) *J Cosmol Astropart Phys* 2018:067. <https://doi.org/10.1088/1475-7516/2018/07/067>
- Kassiola A, Kovner I (1993) *Astrophys J* 417:450. <https://doi.org/10.1086/173325>
- Kassiola A, Kovner I (1995) *Mon Not R Astron Soc* 272:363. <https://doi.org/10.1093/mnras/272.2.363>
- Keeton CR (2001). ArXiv e-prints. [arXiv:astro-ph/0102341](https://arxiv.org/abs/astro-ph/0102341)
- Kelly PL, Rodney S, Treu T et al (2023a) *Astrophys J* 948:93. <https://doi.org/10.3847/1538-4357/ac4ccb>
- Kelly PL, Rodney S, Treu T et al (2023b) *Science* 380:abh1322. <https://doi.org/10.1126/science.abh1322>
- Kennefick D (2012) In: Lehner C, Renn J, Schemmel M (eds) *Einstein and the changing worldviews of physics*. Birkhäuser, Boston, pp 201–232. https://doi.org/10.1007/978-0-8176-4940-1_9
- Kochanek CS (2006) In: Meylan G, Jetzer P, North P, Schneider P, Kochanek CS, Wambsgans J (eds) *Gravitational lensing: strong, weak and micro*. Saas-fee advanced courses, vol. 33. Springer, Berlin, Heidelberg, pp 91–268. https://doi.org/10.1007/978-3-540-30310-7_2
- Kochanek CS, Blandford RD, Lawrence CR, Narayan R (1989) *Mon Not R Astron Soc* 238:43. <https://doi.org/10.1093/mnras/238.1.43>
- Kormann R, Schneider P, Bartelmann M (1994) *Astron Astrophys* 284:285
- Küing R (2018) *Astron Comput* 23:115. <https://doi.org/10.1016/j.ascom.2018.02.007>
- Küing R, Saha P, Ferreras I et al (2018) *Mon Not R Astron Soc* 474:3700. <https://doi.org/10.1093/mnras/stx3012>

- Leung C, Kader Z, Masui KW et al (2022) *Phys Rev D* 106:043017. <https://doi.org/10.1103/PhysRevD.106.043017>
- Liesenborgs J, De Rijcke S (2012) *Mon Not R Astron Soc* 425:1772. <https://doi.org/10.1111/j.1365-2966.2012.21751.x>
- Liesenborgs J, De Rijcke S, Dejonghe H (2006) *Mon Not R Astron Soc* 367:1209. <https://doi.org/10.1111/j.1365-2966.2006.10040.x>
- Liesenborgs J, de Rijcke S, Dejonghe H, Bekaert P (2007) *Mon Not R Astron Soc* 380:1729. <https://doi.org/10.1111/j.1365-2966.2007.12236.x>
- Luhtaru R, Schechter PL, de Soto KM (2021) *Astrophys J* 915:4. <https://doi.org/10.3847/1538-4357/abfda1>
- Meneghetti M (2021) Introduction to gravitational lensing; with Python examples. Lecture Notes in Physics, vol 956. Springer, Cham. <https://doi.org/10.1007/978-3-030-73582-1>
- Meneghetti M, Natarajan P, Coe D et al (2017) *Mon Not R Astron Soc* 472:3177. <https://doi.org/10.1093/mnras/stx2064>
- Mrólz P, Udalski A, Szymański MK et al (2020) *Astrophys J Supp* 249:16. <https://doi.org/10.3847/1538-4365/ab9366>
- Nityananda R (1990) *Curr Sci* 59:1044. <http://www.jstor.org/stable/24094893>
- Nityananda R, Samuel J (1992) *Phys Rev D* 45:3862. <https://doi.org/10.1103/PhysRevD.45.3862>
- Petkova M, Metcalf RB, Giocoli C (2014) *Mon Not R Astron Soc* 445:1954. <https://doi.org/10.1093/mnras/stu1860>
- Petters AO, Levine H, Wambsgans J (2001) Singularity theory and gravitational lensing. Progress in mathematical physics, vol 21. Birkhäuser, Boston. <https://doi.org/10.1007/978-1-4612-0145-8>
- Priewe J, Williams LLR, Liesenborgs J, Coe D, Rodney SA (2017) *Mon Not R Astron Soc* 465:1030. <https://doi.org/10.1093/mnras/stw2785>
- Raney CA, Keeton CR, Brennan S, Fan H (2020) *Mon Not R Astron Soc* 494:4771. <https://doi.org/10.1093/mnras/staa921>
- Refsdal S (1964) *Mon Not R Astron Soc* 128:295. <https://doi.org/10.1093/mnras/128.4.295>
- Refsdal S (1964) *Mon Not R Astron Soc* 128:307. <https://doi.org/10.1093/mnras/128.4.307>
- Refsdal S, Surdej J (1994) *Rep Prog Phys* 57:117. <https://doi.org/10.1088/0034-4885/57/2/001>
- Remolina González JD, Sharon K, Mahler G (2018) *Astrophys J* 863:60. <https://doi.org/10.3847/1538-4357/aacf8e>
- Saha P, Williams LLR (1997) *Mon Not R Astron Soc* 292:148. <https://doi.org/10.1093/mnras/292.1.148>
- Saha P, Williams LLR (2003) *Astron J* 125:2769. <https://doi.org/10.1086/375204>
- Saha P, Williams LLR (2004) *Astron J* 127:2604. <https://doi.org/10.1086/383544>
- Saha P, Williams LLR (2006) *Astrophys J* 653:936. <https://doi.org/10.1086/508798>
- Schmidt T, Treu T, Birrer S et al (2023) *Mon Not R Astron Soc* 518:1260. <https://doi.org/10.1093/mnras/stac2235>
- Schneider P (1985) *Astron Astrophys* 143:413
- Schneider P (2014) *Astron Astrophys* 568:L2. <https://doi.org/10.1051/0004-6361/201424450>
- Schneider P, Ehlers J, Falco EE (1992) Gravitational lenses. Astronomy and astrophysics library. Springer, New York. <https://doi.org/10.1007/978-3-662-03758-4>
- Schneider P, Sluse D (2014) *Astron Astrophys* 564:A103. <https://doi.org/10.1051/0004-6361/201322106>
- Sebesta K, Williams LLR, Liesenborgs J, Medezinski E, Okabe N (2019) *Mon Not R Astron Soc* 488:3251. <https://doi.org/10.1093/mnras/stz1950>
- Shajib AJ (2019) *Mon Not R Astron Soc* 488:1387. <https://doi.org/10.1093/mnras/stz1796>
- Shajib AJ, Mozumdar P, Chen GCF et al (2023) *Astron Astrophys* 673:A9. <https://doi.org/10.1051/0004-6361/202345878>
- Shapiro II (1964) *Phys Rev Lett* 13:789. <https://doi.org/10.1103/PhysRevLett.13.789>
- Sluse D, Chantry V, Magain P, Courbin F, Meylan G (2012) *Astron Astrophys* 538:A99. <https://doi.org/10.1051/0004-6361/201015844>
- Smith RJ, Collett TE (2021) *Mon Not R Astron Soc* 505:2136. <https://doi.org/10.1093/mnras/stab1399>
- Suyu SH, Marshall PJ, Hobson MP, Blandford RD (2006) *Mon Not R Astron Soc* 371:983. <https://doi.org/10.1111/j.1365-2966.2006.10733.x>
- Tessore N (2017) *Astron Astrophys* 597:L1. <https://doi.org/10.1051/0004-6361/201629947>
- Thiébaud É, Young J (2017) *J Opt Soc Am A* 34:904. <https://doi.org/10.1364/JOSAA.34.000904>
- Turyshv SG, Toth VT (2020) *Phys Rev D* 102:024038. <https://doi.org/10.1103/PhysRevD.102.024038>
- Valls-Gabaud D (2006) In: Alimi J-M, Füzfa A (eds) Albert Einstein century international conference. AIP Conf Proc, vol. 861, pp 1163–1163. <https://doi.org/10.1063/1.2399715>
- Van de Vyvere L, Gomer MR, Sluse D et al (2022a) *Astron Astrophys* 659:A127. <https://doi.org/10.1051/0004-6361/202141551>
- Van de Vyvere L, Sluse D, Gomer MR, Mukherjee S (2022b) *Astron Astrophys* 663:A179. <https://doi.org/10.1051/0004-6361/202243382>

- Vegetti S, Koopmans LVE (2009) *Mon Not R Astron Soc* 392:945. <https://doi.org/10.1111/j.1365-2966.2008.14005.x>
- Vernardos G, Koopmans LVE (2022) *Mon Not R Astron Soc* 516:1347. <https://doi.org/10.1093/mnras/stac1924>
- Wagner J (2018) *Astron Astrophys* 620:A86. <https://doi.org/10.1051/0004-6361/201834218>
- Wagner J (2019) *Mon Not R Astron Soc* 487:4492. <https://doi.org/10.1093/mnras/stz1587>
- Wagner J (2019) *Universe* 5:177. <https://doi.org/10.3390/universe5070177>
- Wagner J, Liesenborgs J, Tessore N (2018) *Astron Astrophys* 612:A17. <https://doi.org/10.1051/0004-6361/201731932>
- Wagner J, Williams LLR (2020) *Astron Astrophys* 635:A86. <https://doi.org/10.1051/0004-6361/201936628>
- Walsh D, Carswell RF, Weymann RJ (1979) *Nature* 279:381. <https://doi.org/10.1038/279381a0>
- Warren SJ, Dye S (2003) *Astrophys J* 590:673. <https://doi.org/10.1086/375132>
- Williams LLR, Saha P (2000) *Astron J* 119:439. <https://doi.org/10.1086/301234>
- Witt HJ (1996) *Astrophys J Lett* 472:L1. <https://doi.org/10.1086/310358>
- Woldesenbet AG, Williams LLR (2012) *Mon Not R Astron Soc* 420:2944. <https://doi.org/10.1111/j.1365-2966.2011.20110.x>
- Woldesenbet AG, Williams LLR (2015) *Mon Not R Astron Soc* 454:862. <https://doi.org/10.1093/mnras/stv1960>
- Wucknitz O (2002) *Mon Not R Astron Soc* 332:951. <https://doi.org/10.1046/j.1365-8711.2002.05426.x>
- Wucknitz O (2008) *Mon Not R Astron Soc* 386:230. <https://doi.org/10.1111/j.1365-2966.2008.13017.x>
- Wynne RA, Schechter PL (2018). *ArXiv e-prints*. [arXiv:1808.06151](https://arxiv.org/abs/1808.06151)
- Yoo J, Kochanek CS, Falco EE, McLeod BA (2005) *Astrophys J* 626:51. <https://doi.org/10.1086/429959>
- Young P, Gunn JE, Kristian J, Oke JB, Westphal JA (1980) *Astrophys J* 241:507. <https://doi.org/10.1086/158365>
- Zheng W, Chen X, Li G, Chen H-Z (2022) *Astrophys J* 931:114. <https://doi.org/10.3847/1538-4357/ac68ea>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.