optica

Optical triangulations of curved spaces: supplementary material

DIMITRIS GEORGANTZIS GARCIA[1,2], GREGORY J. CHAPLAIN[1,3], JAKUB BĚLÍN[1], TOMÁŠ TYC[4], CHRISTOPH ENGLERT[1], AND JOHANNES COURTIAL[1,*]

[1] School of Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, UK
[2] Current address: Sheffield University Management School, University of Sheffield, Conduit Road, Sheffield, S10 1FL, UK
[3] Current address: Department of Mathematics, Imperial College London, London SW7 2AZ, UK
[4] Institute of Theoretical Physics and Astrophysics, Masaryk University,, Kotlářská 2, 61137 Brno, Czech Republic
[*] Corresponding author: johannes.courtial@glasgow.ac.uk

Published 3 February 2020

This document provides supplementary information to "Optical triangulations of curved spaces," https://doi.org/10.1364/OPTICA.378357. It provides details on the optical cancelling of a wedge of space using different ap-proaches, specifically using negative refraction, using absolute instruments, using combinations of skew lenses, and using light-field transfer. It also provides additional information on the ray-tracing simulations shown in the main document.

1. CANCELLING A WEDGE OF SPACE USING NEGATIVE REFRACTION

Consider some flat 2-manifold M, a piece of paper for instance. If we pick two arbitrary points $V, P \in M$, it is clear that going around V from P back to itself corresponds to travelling an angle of 2π around V. Now, let us remove a wedge of space with angle ε at the vertex V from M and identify its edges to one another. Going around V from P back to itself now corresponds to an angle of $2\pi - \varepsilon$ instead of 2π, see Fig. 1(c) in the main text. The result of this construction is a cone, a manifold that is flat everywhere except for its vertex V; at V itself it is not flat but has *deficit angle* $\varepsilon_V = \varepsilon$.

Fig. S1 shows how to optically cancel a wedge of space of angle ε. To see the principle, consider a wedge of space with refractive index $n = -1$ and with apex angle $\varepsilon/2$ embedded in vacuum ($n = 1$), see Fig. S1(a). It is easy to see from Snell's law that a ray impinging point P at the interface 1 between the $n = 1$ and $n = -1$ media is refracted on interface 1 such that, after having been refracted also at interface 2, passes through the point P', the mirror image of P with respect to plane 2. Moreover, the direction of the outgoing ray is obtained by rotating the incoming ray direction by the angle of ε around the wedge edge (see ray A in Fig. S1(a)). This way, the space between the half-planes 1 and 1' appears to be "cancelled". This can be interpreted such that the wedge of angle $\varepsilon/2$ of negative space ($n = -1$) cancels out an equal amount of adjacent positive space ($n = +1$). This way, we are effectively removing a wedge of angle ε from the 3D space in consideration, creating a 3D ε-cone.

However, the method we have just described has the following deficiency. Light-rays striking interface 1 such that their projections to a plane perpendicular to the wedge edge make an angle $\beta \leq \varepsilon/2$ with it (such as ray B in Fig. S1(a)) have a resulting direction that is either parallel to or divergent from interface 2, effectively getting "lost" in the wedge, i.e. preventing the appropriate map from taking place. This problem can be solved, however, by constructing the wedge as a collection of consecutive wedges of smaller angle, a sort of *fan* of interfaces as shown in Fig. S1(b). Increasing the number of such wedges then allows us to make the proportion of the "lost" rays arbitrary small.

Another solution involves symmetrically placing planar mirrors at the edge of the wedge, as shown in Fig. S1(c). As long as the angle μ between each mirror and the face of the wedge is chosen to lie in the range

$$\frac{\pi}{2} \leq \mu < \pi - \frac{\varepsilon}{2}, \tag{S1}$$

which is possible only if $\varepsilon/2 < \pi/2$, the mirrors ensure that *all* rays incident on the wedge emerge as desired. If $\varepsilon/2 \geq \pi/2$, a fan of wedges and mirrors can be used. Raytracing simulations confirm that this method works in 2D (see, for example, Fig. 2 in the main document) and in 3D (not shown).

The same ideas can be approximated ray-optically by replacing each interface between $n = +1$ and $n = -1$ with microstructured sheets that change the direction of transmitted light rays in the same way [1, 2].

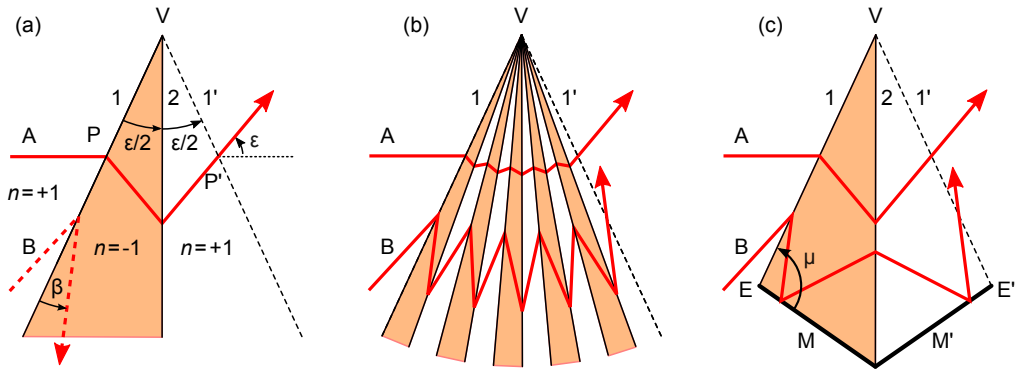

Fig. S1. Space-cancelling wedge based on negative refraction. (a) Half-plane 1, which is lying at the interface between the $n = +1$ and $n = -1$ media, gets mapped to half-plane 1′, the mirror image with respect to half-plane 2 of half-plane 1. Half-plane 1′ can alternatively be obtained by rotating half-plane 1 around the wedge edge V by an angle ε. Ray A (solid red line) is refracted as desired, but ray B (dashed red line) gets "lost" in the wedge. (b) A fan of N (here $N = 5$) wedges of $n = -1$ medium, each of angle $\varepsilon/(2N)$, refracts not only ray A as desired, but also ray B. Nevertheless, rays can still get lost in the wedges. (c) Mirrors M and M′, symmetrically positioned at the end of the $n = -1$ wedge and a corresponding $n = +1$ wedge, ensure that *all* light rays are redirected as desired. This requires the angle μ to be chosen such that $\pi/2 < \mu < \pi - \varepsilon/2$. Configuration (c) is preferred throughout this paper.

2. CANCELLING A WEDGE OF SPACE USING ABSOLUTE INSTRUMENTS

The second method for optically eliminating a wedge of space is based on a combination of absolute optical instruments and transformation optics. Let us explain the idea in detail. We want to map optically two faces ρ and σ of a 3D net of a 4D polyhedron, see Fig. S2(a), such as e.g. faces 1–1 in main text Fig. 3(a). To do this, we will use transformation optics within the wedge between planes ρ and σ outside the net. We denote by a the line where the faces ρ and σ meet, and by ε the angle between these faces outside of the net.

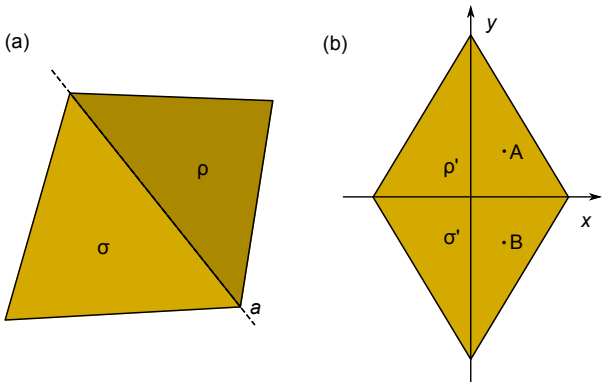

Fig. S2. Space-cancelling wedge based on absolute instruments. (a) The faces ρ and σ in physical space that are to be mapped on each other optically, and (b) the corresponding faces (polygons) in virtual space that lie in the plane XY. The point A is mapped optically to point B and vice versa by the refractive index, N.

We start with virtual space where we use Cartesian coordinates X, Y, Z; we define it as the half-space $Z > 0$ bounded by the plane XY. Within the half-plane $Z = 0; Y > 0$ we define a polygon ρ' of the same shape and size as the face ρ of the net, and similarly we define a polygon σ' within the half-plane $Z = 0; Y < 0$, see Fig. S2(b); the polygons ρ' and σ' can be transformed to each other by a rotation around the X-axis. This prepares virtual space to be transformed into the wedge between the net faces ρ and σ in physical space, whereby ρ' will be transformed into ρ and σ' will be transformed into σ.

Virtual space is filled with a suitable absolute-instrument refractive index $N(\vec{R})$ that maps optically the polygons ρ' and σ' to each other, see Fig. S2(b). Such an optical mapping can be achieved using the Lissajous lens [3], a device analogous to a 3D anisotropic harmonic oscillator in mechanics. If the frequency of such an oscillator in the X direction is chosen to be twice the frequency in Y and Z directions, there will be just a half an oscillation both in Y and Z per one full oscillation (period) in X. Therefore a ray entering virtual space at the point $\mathrm{A} = (X, Y, 0)$ with the wavevector (K_X, K_Y, K_Z) will leave it at the point $\mathrm{B} = (X, -Y, 0)$ with the wavevector $(K_X, -K_Y, -K_Z)$. The corresponding refractive index in virtual space $N(\vec{R})$ then follows from the general formula for Lissajous lens [3] and gets the form

$$N = \alpha \sqrt{1 - \frac{4X^2 + Y^2 + Z^2}{b^2}}\,, \tag{S2}$$

where α and b are constants that can be chosen freely according to our needs; b should be large enough to avoid negative index on the polygons ρ' and σ'.

Having the index in virtual space, we proceed to transforming virtual space to physical space. This is done by uniformly expanding the wedge of angle π between the half-planes of faces ρ' and σ' into the wedge of angle ε between the half-planes of faces ρ and σ; the X-axis is transformed into the axis a in physical space. Such a transformation is easiest described in cylindrical coordinates where only the azimuthal coordinate changes, being multiplied by the factor $\gamma \equiv \varepsilon/\pi$. This modifies the virtual-space medium into physical-space medium with permittivity and permeability tensors

$$\hat{\varepsilon}(\vec{r}) = N[\vec{R}(\vec{r})]^2 \mathrm{diag}\,(1/\gamma, \gamma, 1/\gamma)\,, \quad \hat{\mu} = \mathrm{diag}\,(1/\gamma, \gamma, 1/\gamma)\,, \tag{S3}$$

where we are assuming that the refractive index $N(\vec{R})$ in virtual space was realized purely dielectrically. We see that the mapping

from virtual to physical space introduces just a slight anisotropy to the medium.

This construction can then be made for each pair of faces of the net that are to be identified. However, physical spaces for different such pairs overlap in general, which might lead to problems because the refractive index would have to be multi-valued. To avoid such problems, one has to choose the constant α large enough so that light rays undergo refraction when leaving the net (where refractive index equals unity) and entering physical space. A detailed analysis shows that it is always possible to choose α such that the rays explore only a portion of the "multivalued region"; this region can then be divided into disjoint parts, each being "used" for one pair of faces to be identified. This way the permittivity and permeability become single-valued, which solves the issue.

3. CANCELLING A WEDGE OF SPACE USING IDEAL THIN LENSES

The third approach for eliminating a wedge of space employs a combination of skew ideal thin lenses. Ideal thin lenses cannot be realised physically, but they can be useful in the initial design of instruments comprising physical lenses. Combinations of skew lenses are rarely used, but theoretical studies of structures comprising skew ideal lenses have uncovered interesting possibilities; for example, such structures can form transformation-optics devices [5], including omnidirectional cloaks [6].

We recently found that, in combinations of three skew ideal lenses, the relationship between object and image space can be a simple rotation by an arbitrary angle ε around the line where the planes of all three lenses intersect [7]. This implies that that such an ideal-lens combination is equivalent to a space-cancelling wedge of wedge angle ε.

Raytracing simulations, such as the one shown in Fig. S3, show that this is indeed the case, provided that the relevant light rays pass through all three lenses. But this is not the case for all light rays, resulting in a field-of-view limitation. This field-of-view limitation can be ameliorated with improved optical design, but it is unlikely that it can be removed entirely.

4. CANCELLING A WEDGE OF SPACE BY TRANSFERRING THE LIGHT FIELD

The SC wedge design that is perhaps easiest to realise in practice is inspired by light-field (or plenoptic) imaging based on lenslet (or microlens) arrays, in which information of both the position where light rays intersect a surface and the direction with which they do so is captured [8]. In our case, this information is captured in one face of the SC wedge and optically transferred, using optical fibres, to the other, optically identified, face of the SC wedge.

The scheme is presented in Fig. S4. Consider a lenslet L_1 located in one of the two optically identified faces of the SC wedge, and a parallel bundle of light rays incident on L_1. As the light rays are parallel, L_1 focusses them into the same point in its back focal plane, where one end of a thin optical fiber is located. Light rays enter this fiber only if they have passed through a position on the clear aperture of L_1 with a particular direction; the fibre therefore represents a particular combination of position and direction with which light was incident on the face of the SC wedge.

An identical lenslet, L_2, is placed into the corresponding position in the second of the optically identified faces of the SC wedge, and the other end of the optical fibre is placed in L_2's front focal plane such that L_2 collimates light exiting the fibre again. Altering the position of the end of the fibre in L_2's front focal plane alters the direction of the light; specifically, the fibre end can be placed such that the direction of the light emerging from L_2 is that with which light was incident on L_1, rotated by the deficit angle ε.

If the two lenslets L_1 and L_2 are parts of lenslet arrays covering the optically identified faces of the SC wedge, and if the "inner" focal planes of these two lenslet arrays are covered with the ends of an array of optical fibres, then the entire *light field* incident on plane of the first lenslet array is transferred to the plane of the second lenslet array.[1] Note that the setup is sym-

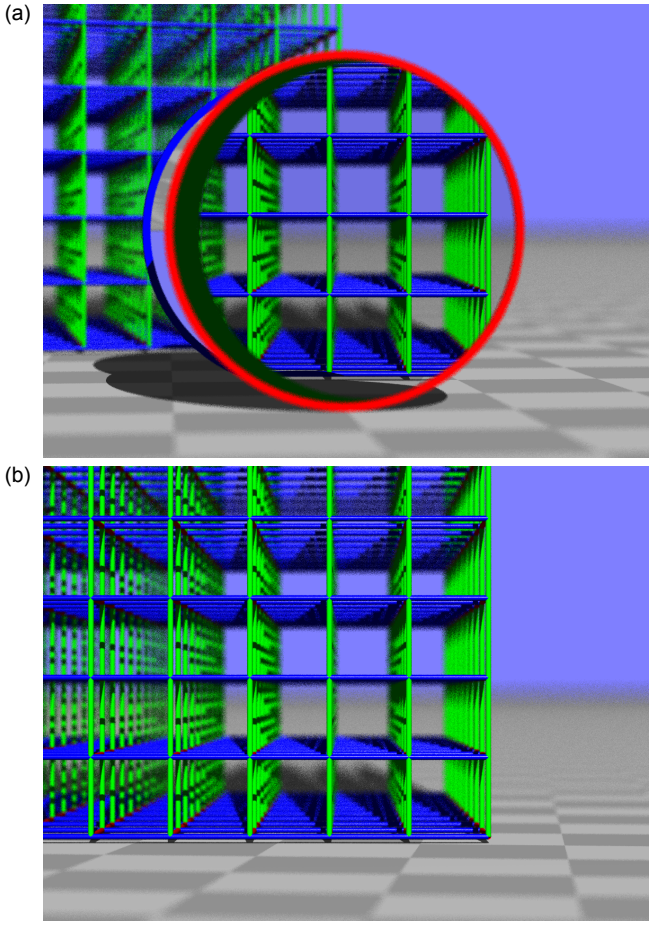

Fig. S3. Cancellation of a wedge of space with ideal thin lenses. (a) A lattice consisting of green and blue cylinders is placed behind three ideal thin lenses, of which only the closest lens, framed by a red ring, is seen directly. In the right part of the closest lens, the lattice is seen through all three lenses and appears rotated by 12° around a vertical axis. (b) For comparison, the simulation was repeated without the lenses but with the camera rotated by −12° around the same axis. The part of the image seen through all three lenses in (a) looks identical to the corresponding part in (b). Both images are ray-optics simulations, performed with an extended version of our scientific raytracer Dr TIM [4].

[1] It was recently shown that a coherent fibre bundle *on its own* can transmit light-field information [9]; however, this information cannot currently easily be converted back into a physical light field.

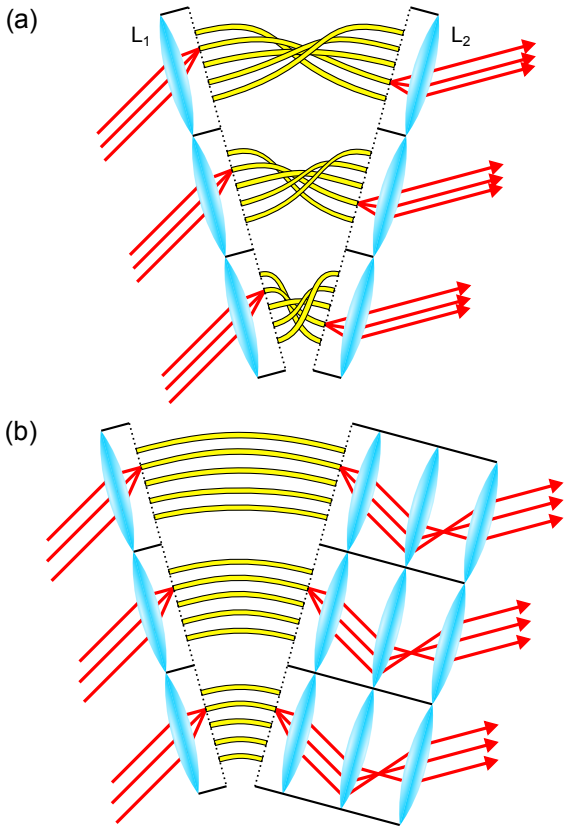

Fig. S4. Light-field transfer between optically identified faces of a SC wedge. (a) Two arrays of lenslets (cyan) are placed in the optically identified faces. The ends of an array of optical fibres (yellow) are placed in the "inner" focal planes (dotted lines) of the two lenslet arrays; baffles (solid black lines) ensure that light from only one lenslet in each array reaches a particular optical fibre. Each fibre therefore corresponds to one particular combination of a position (restricted to the aperture of the corresponding lenslet) and direction with which light rays (solid red lines) are incident on the faces. (b) It is possible to use a coherent fibre bundle, but to ensure that the light rays exit the setup with the correct direction, further optical components, here an afocal combination of two identical lenses, are required. The faces of the SC wedge are then formed by the outermost lens arrays.

metric, and so the transfer works both ways. Ambiguity about the lenslet through which a light ray entering a particular fibre has passed can be removed by inserting an array of baffles between the lenslet arrays and their relevant focal planes (see Fig. S4(a)), but this restricts the field of view of each lenslet. Furthermore, both the positions and directions are discretised in this light-field-transfer process, such that the overall number of fibres equals the space-bandwidth product of the system [10].

The array of fibres is simpler — and almost certainly easier to realise — if the relative position of the ends of each fibre is the same at both ends of the array. Such a fibre array is known as a *coherent fibre bundle*. However, if the coherent fibre bundle simply connects corresponding points in the inner focal planes of two lenslet arrays, light rays emerge with a direction that is rotated by 180° around the normal of the second lenslet-array plane relative to the desired direction. This can be fixed either by twisting the ends of each individual bundle of fibres that connects corresponding lenslets by 180° (as is shown in Fig. S4(a)), or by adding further optical components (additional lenses in the case shown in Fig. S4(b)) that rotate the light-ray direction by 180° to one end of the device.

Note that these ideas are very closely related not only to light-field capture [8] and display [11] devices, but also to Shack-Hartmann sensors [12], integral photography [13], integral imaging [14], generalised refraction using confocal lenslet arrays [15], moiré magnifiers [16], Gabor superlenses [17, 18], and digital integral cloaking [19].

5. RAY-TRACING SIMULATIONS

The ray-tracing simulations in this paper have been performed using an extended version of the open-source [20] raytracer Dr TIM [4]. (In the case of frames (g), (h), (j) and (k) of Fig. 2 in the main text, the images were manipulated in graphics software.) This software, packaged as executable Java Archive (JAR) files that enable interactive exploration of the nets of different manifolds, is available in Ref. [21], together with files detailing the parameters used in each simulation.

Dr TIM uses standard rendering raytracing, tracing backwards from the camera to the scene, and ultimately to a light source. This is efficient, as every ray that is being traced is one that ultimately contributes to the image, which is not generally the case when tracing rays that start at a light source.

When we created Dr TIM (then simply TIM [22]; the title "Dr" was conferred in 2014, when TIM gained significantly enhanced academic abilities [4]), we were guided by the following ideas:

1. Complex components can be simulated with different levels of idealisation.

2. Dr TIM is both a research tool and a dissemination tool.

The first idea has been applied to SC wedges, which Dr TIM can simulate in different ways. The most straightforward of these is to trace rays through the detailed structure of alternating wedges of positive and negative refractive index. This gives the results closest to any experimental realisation. In principle, it is possible to impedance-match these wedges, making transmission through the interfaces separating the wedges lossless. By constructing each null-space wedge from a large number of wedges of alternating negative and positive refractive index, almost all light rays are transmitted through the null-space wedge as intended. Because of the large number of wedges involved, simulations through such detailed structures are slow. Almost exactly the same result can be achieved by simulating these structures as ideal null-space wedges. Internally, this utilises Dr TIM's teleporting surface property [22]. This is what has been done in the simulations shown in this paper.

The second idea is the reason why we wrote Dr TIM in Java, which, once suitably compiled, can run on most computer systems without further adaptation. This then enables dissemination of the research software with publications to increase the transparency of the research, which is further enhanced by the ability to access Dr TIM's code [20].

REFERENCES

1. J. Courtial and J. Nelson, "Ray-optical negative refraction and pseudoscopic imaging with Dove-prism arrays," New J. Phys. **10**, 023028 (2008).

2. J. Courtial, "Ray-optical refraction with confocal lenslet arrays," New J. Phys. **10**, 083033 (2008).
3. A. J. Danner, H. L. Dao, and T. Tyc, "The Lissajous lens: a three-dimensional absolute optical instrument without spherical symmetry," Opt. Express **23**, 5716–5722 (2015).
4. S. Oxburgh, T. Tyc, and J. Courtial, "Dr TIM: Ray-tracer TIM, with additional specialist capabilities," Comp. Phys. Commun. **185**, 1027–1037 (2014).
5. J. Courtial, T. Tyc, J. Bělín, S. Oxburgh, G. Ferenczi, E. N. Cowie, and C. D. White, "Ray-optical transformation optics with ideal thin lenses makes omnidirectional lenses," Opt. Express **26**, 17872–17888 (2018).
6. J. Bělín, T. Tyc, M. Grunwald, S. Oxburgh, E. N. Cowie, C. D. White, and J. Courtial, "Ideal-lens cloaks and new cloaking strategies," submitted (2019).
7. J. Bělín, G. Ferenczi, and J. Courtial, "Ideal-lens image rotator," submitted for publication (2019).
8. T. Adelson and J. Y. A. Wang, "Single lens stereo with a plenoptic camera," IEEE Transactions on Pattern Analysis and Machine Intelligence **14**, 99–106 (1992).
9. A. Orth, M. Ploschner, E. R. Wilson, I. S. Maksymov, and B. C. Gibson, "Optical fiber bundles: Ultra-slim light field imaging probes," Science Advances **5**, eaav1555 (2019).
10. A. W. Lohmann, R. G. Dorsch, D. Mendlovic, Z. Zalevsky, and C. Ferreira, "Space–bandwidth product of optical signals and systems," J. Opt. Soc. Am. A **13**, 470–473 (1996).
11. J. Geng, "Three-dimensional display technologies," Adv. Opt. Photon. **5**, 456–535 (2013).
12. R. B. Shack and B. C. Platt, "Production and use of a lenticular hartmann screen," J. Opt. Soc. Am. **61**, 656 (1971).
13. G. Lippmann, "La photographie intégrale," C. R. Hebd. Acad. Sci. **146**, 446–451 (1908).
14. R. F. Stevens and T. G. Harvey, "Lens arrays for a three-dimensional imaging system," J. Opt. A: Pure Appl. Opt. **4**, S17–S21 (2002).
15. A. C. Hamilton and J. Courtial, "Generalized refraction using lenslet arrays," J. Opt. A: Pure Appl. Opt. **11**, 065502 (2009).
16. M. C. Hutley, R. Hunt, R. F. Stevens, and P. Savander, "The moiré magnifier," Pure Appl. Opt.: JEOS A **3**, 133–142 (1994).
17. C. Hembd-Sölner, R. F. Stevens, and M. C. Hutley, "Imaging properties of the Gabor superlens," J. Opt. A: Pure Appl. Opt. **1**, 94–102 (1999).
18. D. Gabor, "Optical system composed of lenticules," U.S. Patent 2,351,034 (June 13, 1944).
19. J. S. Choi and J. C. Howell, "Digital integral cloaking," Optica **3**, 536–540 (2016).
20. "Dr TIM, a highly scientific raytracer," https://github.com/jkcuk/Dr-TIM.
21. J. Courtial, D. G. Garcia, G. J. Chaplain, J. Bělín, T. Tyc, and C. Englert, "Simulation software and parameter files for optical simulation of curved spaces," figshare, https://doi.org/10.6084/m9.figshare.9863093 (2019).
22. D. Lambert, A. C. Hamilton, G. Constable, H. Snehanshu, S. Talati, and J. Courtial, "TIM, a ray-tracing program for METATOY research and its dissemination," Comp. Phys. Commun. **183**, 711–732 (2012).