



Light in Power: A General and Parameter-free Algorithm for Caustic Design

Jocelyn Meyron, Quentin Mériqot, Boris Thibert

► To cite this version:

Jocelyn Meyron, Quentin Mériqot, Boris Thibert. Light in Power: A General and Parameter-free Algorithm for Caustic Design. ACM Transactions on Graphics, 2018, 37 (6), pp.1-13. <10.1145/3272127.3275056>. <hal-01570739v2>

HAL Id: hal-01570739

<https://hal.science/hal-01570739v2>

Submitted on 7 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Light in Power: A General and Parameter-free Algorithm for Caustic Design

JOCELYN MEYRON

Univ. Grenoble Alpes, CNRS, LJK

QUENTIN MÉRIGOT

Univ. Paris-Sud, CNRS, LMO

BORIS THIBERT

Univ. Grenoble Alpes, CNRS, LJK

February 7, 2019

Abstract

We present in this paper a generic and parameter-free algorithm to efficiently build a wide variety of optical components, such as mirrors or lenses, that satisfy some light energy constraints. In all of our problems, one is given a collimated or point light source and a desired illumination after reflection or refraction and the goal is to design the geometry of a mirror or lens which transports exactly the light emitted by the source onto the target. We first propose a general framework and show that eight different optical component design problems amount to solving a light energy conservation equation that involves the computation of visibility diagrams. We then show that these diagrams all have the same structure and can be obtained by intersecting a 3D Power diagram with a planar or spherical domain. This allows us to propose an efficient and fully generic algorithm capable to solve these eight optical component design problems. The support of the prescribed target illumination can be a set of directions or a set of points located at a finite distance. Our solutions satisfy design constraints such as convexity or concavity. We show the effectiveness of our algorithm on simulated and fabricated examples.




Figure 1: Our algorithm can be used to design mirrors and lenses that reflect or refract collimated or point light sources onto a prescribed distribution of light. *From left to right:* Three lenses that refract the three channels of a color image; Mirror that reflects a point light source (located inside the mirror); Fabricated mirror that reflects a collimated light source; Fabricated lens that refracts a collimated light source.

1. INTRODUCTION

The field of non-imaging optics deals with the design of optical components whose goal is to transfer the radiation emitted by a light source onto a prescribed target. This question is at the heart of many applications where one wants to optimize the use of light energy by decreasing light loss or light pollution. Such problems appear in the design of car beams, public lighting, solar ovens and hydroponic agriculture. This problem has also been considered under the name of *caustic design*, with applications in architecture and interior decoration [Finckh et al., 2010].

In this paper, we consider the problem of designing a wide variety of mirrors and lenses that satisfy different kinds of light energy constraints. To be a little bit more specific, in each problem that we consider, one is given a light source and a desired illumination after reflection or refraction which is called the target. The goal is to design the geometry of a mirror or lens which transports exactly the light emitted by the source onto the target. The design of such optical components can be thought of as an *inverse problem*, where the *forward problem* would be the simulation of the target illumination from the description of the light source and the geometry of

the mirror or lens.

In practice, the mirror or lens needs to satisfy aesthetic and pragmatic design constraints. In many situations, such as for the construction of car lights, physical molds are built by milling and the mirror or lens is built on this mold. Sometimes the optical component itself is directly milled. This imposes some constraints that can be achieved by imposing convexity or smoothness conditions. The convexity constraint is classical since it allows in particular to mill the component with a tool of arbitrary large radius. Conversely, concavity allows to mill its mold. Also, convex mirrors are easier to chrome-plate, because convex surfaces have no bumps in which the chrome would spuriously concentrate [Cork et al., 1977].

In this paper, we propose a generic algorithm capable of solving eight different caustic design problems, see Figures 1 and 2. Our approach relies on the relation between these problems and optimal transport. The algorithm is fully generic in the sense that it can deal with any of the eight caustic-design problems just by changing a formula, and can handle virtually any ideal light source and target. Our contributions are the following:

- We propose a general framework for eight different optical component design problems (i.e. four non-imaging problems, for which we can produce either concave or convex solutions). These problems amount to solving the same *light energy conservation* equation (see Sec. 3), which involves prescribing the amount of light reflected or refracted in a finite number of directions.
- We propose a single algorithm with no parameter capable to solve this equation for the eight different problems. We will see that, in order to solve this equation, we need to compute integrals over *visibility cells*, which can be obtained in all cases by intersecting a 3D Power diagram with a planar or spherical domain. The equation is then solved using a damped Newton algorithm.
- In all of the four non-imaging problems, we can construct either a concave or convex optical component, easing their fabrication. Several components can then be combined to produce a single caustic, providing resilience to small obstacles and providing degrees of freedom to control the shape of the optical system.
- We show that we can solve near-field problems (when the target is at a finite distance) by iteratively solving far-field problems (when the target is a set of directions) for the eight optical component design problems.

2. RELATED WORK

The field of non-imaging optics has been extensively studied in the last thirty years. We give below an overview of the main




Figure 2: **Four inverse problems arising in non-imaging optics.** In each case, the goal is to build the surface \mathcal{R} of a mirror or a lens. For each problem, we provide two solutions (for instance, convex or concave surfaces when the light source is collimated). *Top/Bottom*: Collimated light sources/Point light sources. *Left/Right*: Mirror/Lens design.

approaches to tackle several of the problems of this field. A survey on inverse surface design from light transport behaviour is provided by Patow and Pueyo, 2005 .

Energy minimization methods. Many different methods to solve inverse problems arising in non-imaging optics rely on variational approaches. When the energies to be minimized are not convex, they can be handled by different kind of iterative methods. One class of methods deals with stochastic optimization. Finckh et al., 2010 propose to represent the optical component (mirror or lens) as a \mathcal{C}^2 B-spline triangle mesh and to use stochastic optimization to adjust the heights of the vertices so as to minimize a light energy constraint. Note that this approach is very costly, since a forward simulation needs to be done at every step and the number of steps is very high in practice. Furthermore, using this method, lots of artifacts in the final caustic images are present. Stochastic optimization has also been used by Papas et al., 2011 to design reflective or refractive caustics for collimated light sources. At the center of the method is the Expectation Minimization algorithm

initialized with a Capacity Constrained Voronoi Tessellation (CCVT) using a variant of Lloyd’s algorithm [Lloyd, 1982]. The source is a uniform directional light and is modeled using an array of curved microfacets. The target is represented by a mixture of Gaussian kernel functions. This method cannot accurately handle low intensity regions and artifacts due to the discretization are present. Microfacets were also used by Weyrich et al., 2009 to represent the mirror. Due to the sampling procedure, this method cannot correctly handle smooth regions and does not scale well with the size of the target. More recently, Piovarči et al., 2017 used microgeometry to design directional screens which provide increased gain and brightness. In their approach, the screen is decomposed into many small patches, each patch reflecting a set of rays toward a prescribed cone of directions. Their problem for each patch is similar to the one we solve for the special case of a directional source and a target at infinity and corresponds to one pillow, see Section 7. Similarly to us, their approach is based on convex optimization and produce convex patches. They prescribe the areas of the facets, whereas we prescribe their measures. Their numerical approach relies on gradient descent rather than on Newton’s method, and only deals with collimated source and far-field target.

The approaches proposed by Kiser et al., 2013, Yue et al., 2014 and Schwartzburg et al., 2014 have in common that they first compute some kind of relationship between the incident rays and their position on the target screen and then use an iterative method to compute the shape of the refractive surface. Yue et al., 2014 use a continuous parametrization and thus cannot correctly handle totally black and high-contrast regions (boundaries between very dark and very bright areas). Yue et al., 2012 proposed to use sticks to represent the refractive surface. This allows to reduce production cost, to be more entertaining for the user since a single set of sticks can produce different caustic patterns. The main problem with this approach is the computational complexity since they need to solve a NP-hard assignment problem. The problem of designing lenses for collimated light sources has also been considered by Schwartzburg et al., 2014. They propose a method to build lenses that can refract complicated and highly contrasted targets. They first use optimal transport on the target space to compute a mapping between the refracted rays of an initial lens and the desired normals, then perform a post-processing step to build a surface whose normals are close to the desired ones.

Monge-Ampère equations When the source and target lights are modeled by continuous functions, the problem amounts to solving a generalized Monge-Ampère equation, either in the plane for collimated light sources, or on the sphere

for point light sources. Let us explain this link more precisely for a collimated light source, assuming that the source rays are collinear to the constant vector $e_z = (0, 0, 1)$ and emitted from a horizontal domain $\Omega \subset \mathbb{R}^2 \times \{0\}$. We assume that the optical component is smooth and parameterized by a height-field function $\varphi : \Omega \rightarrow \mathbb{R}$ and denote by μ and ν the source and target measures. At every point $\varphi(x)$ of the optical component, the gradient $\nabla\varphi(x)$ encodes the normal, and we denote by $F(\nabla\varphi(x)) \in \mathbb{S}^2$ the direction of the ray that is reflected at $(x, \varphi(x))$ using Snell’s law. The conservation of light energy thus reads $\nu(A) = \mu((F \circ \nabla\varphi)^{-1}(A))$ for every set A . This is equivalent to having $\tilde{\nu}(B) = \mu((\nabla\varphi)^{-1}(B))$ for every set B , where $\tilde{\nu}(B) = \nu(F(B))$. When F and $\nabla\varphi$ are one-to-one (which is the case if the optical component is convex or concave) and μ and $\tilde{\nu}$ are modeled by continuous functions f and g , with the change of variable formula, the light energy conservation becomes equivalent to the following generalized Monge-Ampère equation

$$g(\nabla\varphi(x)) \det(D^2\varphi(x)) = f(x). \quad (1)$$

Similar equations are obtained for point light sources. The existence and regularity of their solutions, namely of the mirror or lens surfaces, have been extensively studied. When the light source is a point, this problem has been studied for mirrors [Caffarelli and Oliker, 2008; Caffarelli et al., 2008] and lenses [Gutiérrez and Huang, 2009] and when the light source is collimated one recovers (1) [Gutiérrez and Tournier, 2013]. We refer to the book of Gutiérrez, 2016 for an introduction to Monge-Ampère equations.

Optimal transport based methods in non-imaging optics In fact, the Monge-Ampère equations corresponding to the non-imaging problems considered in this paper can be recast as optimal transport problems. This was first observed by Wang, 2004 and Glimm and Oliker, 2003 for the mirror problem with a point light source. Many algorithms related to optimal transport have been developed to address non-imaging problems. For collimated sources, one could rely on wide-stencils finite difference schemes [Prins et al., 2013], or on numerical solvers for quadratic semi-discrete optimal transport [Mérigot, 2011; de Goes et al., 2012]. For point sources, there exist variants of the Oliker-Prussner algorithm for the mirror problem [Caffarelli et al., 1999] or the lens problem [Gutiérrez and Huang, 2009]. Both algorithms have a $O(N^4)$ complexity, restricting their use to small discretizations. A quasi-Newton method is proposed by de Castro et al., 2016 for point-source reflector design, handling up to 10^4 Dirac masses.

Finally we note that the approach of Schwartzburg et al., 2014 to build lenses also relies on optimal transport. However, the optimal transport step is used as a heuristic to estimate the

normals of the surface, and not to directly construct a solution to the non-imaging problem. A post-processing step is then performed by minimizing a non-convex energy composed of five weighted terms. In contrast, all the results presented in this article use no post-processing.

3. LIGHT ENERGY CONSERVATION

We present in this section several mirror and lens design problems arising in non-imaging optics. In all the problems, one is given a light source (emitted by either a plane or a point) and a desired illumination “at infinity” after reflection or refraction, which is called the target, and the goal is to design the geometry of a mirror or lens which transports the energy emitted by the source onto the target. We do not take into account multiple reflections or refractions. We show that even though the problems we consider are quite different from one another, they share a common structure that corresponds to a so-called generalized Monge-Ampère equation, whose discrete version is given by Equation (2). This section gathers and reformulates in a unified setting results about mirror and lens design for collimated and point light sources. We refer to the work of Caffarelli and Oliker, 2008, Prins et al., 2013, de Castro et al., 2016 and references therein.

3.1. Collimated light source

3.1.1 Convex mirror design

In this first problem, the light source is collimated, meaning that it emits parallel rays. The source is encoded by a light intensity function $\rho : \Omega \rightarrow \mathbb{R}$ over a 2D domain Ω contained in the (xy) plane $\mathbb{R}^2 \times \{0\} \subset \mathbb{R}^3$, and that all the rays are parallel and directed towards $e_z = (0, 0, 1)$. For simplicity, we will conflate \mathbb{R}^2 and $\mathbb{R}^2 \times \{0\}$. The desired target illumination is “at infinity”, and is described by a set of intensity values $\sigma = (\sigma_i)_{1 \leq i \leq N}$ supported on a finite set of directions $Y = \{y_1, \dots, y_N\}$ included in the unit sphere \mathbb{S}^2 . The problem is to find the surface \mathcal{R} of a mirror that reflects the source intensity ρ to the target intensity σ , see Figures 2 (top left) and 3.

Since the number of reflected directions Y is finite, the mirror surface \mathcal{R} is composed of a finite number of planar facets as illustrated in Figure 3. We will construct the mirror \mathcal{R} as the graph of a piecewise-linear convex function of the form $x \in \mathbb{R}^2 \mapsto \max_i \langle x | p_i \rangle - \psi_i$, where $\langle \cdot | \cdot \rangle$ denotes the dot product. The vectors $p_1, \dots, p_N \in \mathbb{R}^2$ and the elevations $\psi = (\psi_1, \dots, \psi_N) \in \mathbb{R}^N$ have to be determined. To choose p_i , we require that the plane $P_i = \{(x, \langle x | p_i \rangle) \mid x \in \mathbb{R}^2\}$ reflects rays with direction e_z towards the direction $y_i \in \mathbb{S}^2$. The downward pointing unit normal to P_i is given by the formula




Figure 3: **Convex Mirror for a collimated light source** (when $N = 16$). The mirror surface \mathcal{R} is the graph of a convex piecewise affine function. The support X_ρ of ρ is decomposed into visibility cells $(V_i(\psi))_{1 \leq i \leq N}$. Every vertical ray above a point $x \in X_\rho$ belongs to a cell $V_i(\psi)$, touches a plane with slope p_i and is reflected to the direction y_i .

$n_i = (p_i, -1)/(\|p_i\|^2 + 1) \in \mathbb{R}^3$, and Snell’s law of reflection gives us $y_i = e_z - 2\langle n_i | e_z \rangle n_i$. Solving for p_i and denoting $\mathbf{p}_{\mathbb{R}^2}$ the orthogonal projection onto $\mathbb{R}^2 \times \{0\}$, we get

$$p_i = -\mathbf{p}_{\mathbb{R}^2}(y_i - e_z)/\langle y_i - e_z | e_z \rangle.$$

Given a vector of elevations $\psi := (\psi_i)_{1 \leq i \leq N}$, we define the visibility cell $V_i(\psi)$

$$V_i(\psi) = \{x \in \mathbb{R}^2 \times \{0\} \mid \forall j, -\langle x | p_i \rangle + \psi_i \leq -\langle x | p_j \rangle + \psi_j\}.$$

By construction, for every $i \in \{1, \dots, N\}$, any vertical ray emanating from a point $x \in V_i(\psi)$ touches the mirror surface \mathcal{R} at an altitude $\langle x | p_i \rangle - \psi_i$, and is thus reflected towards direction y_i . Consequently, the amount of light reflected towards direction y_i equals the integral of ρ over $V_i(\psi)$. The *Collimated Source Mirror problem* (CS/Mirror) then amounts to finding elevations $\psi \in \mathbb{R}^N$ such that

$$\forall i \in \{1, \dots, n\} \quad \int_{V_i(\psi)} \rho(x) dx = \sigma_i. \quad (2)$$

By construction, a solution to Equation (2) provides a parameterization \mathcal{R}_ψ of a convex mirror that reflects the collimated light source ρ to the discrete target σ :

$$\mathcal{R}_\psi : x \in \mathbb{R}^2 \mapsto (x, \max_i \langle x | p_i \rangle - \psi_i) \in \mathbb{R}^3,$$

where $\mathbb{R}^2 \times \{0\}$ and \mathbb{R}^2 are identified. Notice that since the mirror is a graph over $\mathbb{R}^2 \times \{0\}$, the vectors y_i cannot be upward vertical. In practice we assume that every direction y_i belongs to the hemisphere $\mathbb{S}_-^2 := \{x \in \mathbb{S}^2, \langle x | e_z \rangle \leq 0\}$. Furthermore, we localize the position of the mirror by considering it only above the domain $X_\rho := \{x \in \mathbb{R}^2 \times \{0\}, \rho(x) \neq 0\}$.

Concave mirror. The same approach also allows the construction of concave mirrors, using a concave function of the form $x \mapsto \min_i \langle x | p_i \rangle + \psi_i$. This amounts to replacing the visibility cells by

$$V_i(\psi) = \{x \in \mathbb{R}^2 \times \{0\} \mid \langle x | p_i \rangle + \psi_i \leq \langle x | p_j \rangle + \psi_j \forall j\}.$$

In that case, a solution to Equation (2) provides a parametrization of a concave mirror $\mathcal{R}_\psi(x) = (x, \min_i \langle x | p_i \rangle + \psi_i)$ that sends the collimated light source ρ to the discrete target σ .

3.1.2 Convex lens design

In this second design problem, we are interested in designing lenses that refract a given collimated light source intensity ρ to a target intensity σ , see the top right diagram in Figure 2. We denote by n_1 the refractive index of the lens, by n_2 the ambient space refractive index and by $\kappa = n_1/n_2$ the ratio of the two indices. We assume that the rays emitted by the source are vertical and that the bottom of the lens is flat and orthogonal to the vertical axis. There is no refraction angle when the rays enter the lens, and we only need to build the top part of the lens.

By a simple change of variable, we show that this problem is equivalent to (CS/Mirror). More precisely, for every $y_i \in Y$, we now define p_i to be the slope of a plane that refracts the vertical ray e_z to the direction y_i . We define \mathcal{R} as the graph of a convex function of the form $x \mapsto \max_i \langle x | p_i \rangle - \psi_i$, where $\psi = (\psi_i)_{1 \leq i \leq N}$ is the set of elevations. A calculation similar to the (CS/Mirror) case gives the following expression:

$$p_i = -\text{p}_{\mathbb{R}^2}(y_i - \kappa e_z) / \langle y_i - \kappa e_z | e_z \rangle.$$

In that case, we define the visibility cell $V_i(\psi)$ to be the set of points $x \in \mathbb{R}^2 \times \{0\}$ that are refracted to the direction y_i :

$$V_i(\psi) = \{x \in \mathbb{R}^2 \times \{0\} \mid \forall j, -\langle x | p_i \rangle + \psi_i \leq -\langle x | p_j \rangle + \psi_j\}.$$

The *Collimated Source Lens problem* (CS/Lens) then amounts to finding weights $(\psi_i)_{1 \leq i \leq N}$ that satisfy (2). In that case the lens surface is parameterized by

$$\mathcal{R}_\psi : x \in \mathbb{R}^2 \mapsto (x, \max_i \langle x | p_i \rangle - \psi_i).$$

In practice, we choose the directions y_i in \mathbb{S}_+^2 and the mirror to be parameterized over the support X_ρ of ρ .

Concave lens. Note that we can also build concave lenses by considering parameterizations with concave functions of the form $x \mapsto \min_i \langle x | p_i \rangle + \psi_i$. Figure 4 illustrates a concave and a convex solution to the same non-imaging optics problem.




Figure 4: **Concave and convex lenses** for a uniform collimated light source and the same target.

3.2. Point light source

3.2.1 Concave mirror design.

In this second mirror design problem, all the rays are now emitted from a single point in space, located at the origin, and the light source is described by an intensity function ρ on the unit sphere \mathbb{S}^2 . As in the previous cases, the target is “at infinity” and is described by a set of intensity values $\sigma = (\sigma_i)_{1 \leq i \leq N}$ supported on the finite set of directions $Y = \{y_1, \dots, y_N\} \subset \mathbb{S}^2$. The problem we consider is to find the surface \mathcal{R} of a mirror that sends the light intensity ρ to the light intensity σ (Fig. 2, bottom left).

Following Caffarelli and Oliker, 2008, we build a concave surface \mathcal{R} that is composed of pieces of confocal paraboloids. More precisely, we denote by $P(y_i, \psi_i)$ the solid (i.e filled) paraboloid whose focal point is at the origin with focal distance ψ_i and with direction y_i . We define the surface \mathcal{R}_ψ as the boundary of the intersection of the solid paraboloids, namely $\mathcal{R}_\psi = \partial(\cap_i P(y_i, \psi_i))$. The visibility cell $V_i(\psi)$ is the set of ray directions $x \in \mathbb{S}^2$ emanating from the light source that are reflected in the direction y_i . Since each paraboloid $\partial P(y_i, \psi_i)$ is parameterized over the sphere by $x \mapsto \psi_i x / (1 - \langle x | y_i \rangle)$, one has [Caffarelli and Oliker, 2008]

$$V_i(\psi) = \left\{ x \in \mathbb{S}^2 \mid \forall j, \frac{\psi_i}{1 - \langle x | y_i \rangle} \leq \frac{\psi_j}{1 - \langle x | y_j \rangle} \right\}.$$

The *Point Source Mirror problem* (PS/Mirror) then amounts to finding (ψ_i) that satisfy the light energy conservation equation (2). The mirror surface is then parameterized by

$$\mathcal{R}_\psi : x \in \mathbb{S}^2 \mapsto \min_i \frac{\psi_i}{1 - \langle x | y_i \rangle} x.$$

In practice, we assume that the target Y is included in \mathbb{S}_+^2 , that the support X_ρ of ρ is included $\mathbb{S}_+^2 := \{x \in \mathbb{S}^2, \langle x | e_z \rangle \geq 0\}$, and that the mirror is parameterized over X_ρ .

One can also define the mirror surface as the boundary of the union (instead of the intersection) of a family of solid

paraboloids. Then, the visibility cells become

$$V_i(\psi) = \left\{ x \in \mathbb{S}^2 \mid \forall j, \frac{\psi_i}{1 - \langle x | y_i \rangle} \geq \frac{\psi_j}{1 - \langle x | y_j \rangle} \right\}$$

and a solution to Equation (2) provides a parameterization $\mathcal{R}_\psi(x) = x \max_i \psi_i / (1 - \langle x | y_i \rangle)$ of the mirror surface. Let us note that in this case the mirror is neither convex nor concave.

3.2.2 Convex lens design.

We now consider the lens design problem for a point light source. As in the collimated setting, we fix the bottom part of the lens. We choose a piece of sphere centered at the source, so that the rays are not deviated. Following Gutiérrez and Huang, 2009, the lens is composed of pieces of ellipsoids of constant eccentricities $\kappa > 1$, where κ is the ratio of the indices of refraction. Each ellipsoid $\partial E(y_i, \psi_i)$ can be parameterized over the sphere by $x \mapsto \psi_i x / (1 - \kappa \langle x | y_i \rangle)$. The visibility cell of y_i is then

$$V_i(\psi) = \left\{ x \in \mathbb{S}^2 \mid \forall j, \frac{\psi_i}{1 - \kappa \langle x | y_i \rangle} \leq \frac{\psi_j}{1 - \kappa \langle x | y_j \rangle} \right\}.$$

The *Point Source Lens problem* (PS/Lens) then amounts to finding weights $(\psi_i)_{1 \leq i \leq N}$ that satisfy (2). Note that the top surface of the lens is then parameterized by

$$\mathcal{R}_\psi : x \in \mathbb{S}^2 \mapsto \min_i \frac{\psi_i}{1 - \kappa \langle x | y_i \rangle} x.$$

In practice, we choose the set of directions y_i to belong to \mathbb{S}_+^2 and the lens to be parameterized over the support $X_\rho \subset \mathbb{S}_+^2$ of ρ .

One can also choose to define the lens surface as the boundary of the union (instead of the intersection) of a family of solid ellipsoids. In that case, the visibility cells are given by

$$V_i(\psi) = \left\{ x \in \mathbb{S}^2 \mid \forall j, \frac{\psi_i}{1 - \kappa \langle x | y_i \rangle} \geq \frac{\psi_j}{1 - \kappa \langle x | y_j \rangle} \right\}$$

and a solution to Equation (2) provides a parameterization $\mathcal{R}_\psi(x) = x \max_i \psi_i / (1 - \kappa \langle x | y_i \rangle)$ of the lens surface. Let us note that in this case the lens is neither convex nor concave.

3.3. General formulation

Let X be a domain of either the plane $\mathbb{R}^2 \times \{0\}$ or the unit sphere \mathbb{S}^2 , $\rho : X \rightarrow \mathbb{R}$ a probability density and $Y = \{y_1, \dots, y_N\} \subset \mathbb{S}^2$ be a set of N points. We define the function $G : \mathbb{R}^N \rightarrow \mathbb{R}^N$ by

$$G_i(\psi) = \int_{V_i(\psi)} \rho(x) dx$$

where $G(\psi) = (G_i(\psi))_{1 \leq i \leq N}$ and $V_i(\psi) \subset X$ is the visibility cell of y_i , whose definition depends on the non-imaging problem. Using this notation, Equation (2) can be rephrased as finding weights $\psi = (\psi_i)_{1 \leq i \leq N}$ such that

$$\forall i \in \{1, \dots, N\}, \quad G_i(\psi) = \sigma_i. \quad (3)$$

Many other problems arising in non-imaging optics amount to solving equations of this form. For example, the design of a lens that refracts a point light source to a desired near-field target can also be modeled by a Monge-Ampère equation that has the same structure [Gutiérrez and Huang, 2009]. In this case, the visibility diagram correspond to the radial projection onto the sphere of pieces of confocal ellipsoids with non constant eccentricities and is not associated to an optimal transport problem.

4. VISIBILITY AND POWER CELLS

The main difficulty to evaluate the function G appearing in Equation (3) is to compute the visibility cells $V_i(\psi)$ associated to each optical modeling problem. We show in this section that the visibility cells have always the same structure, allowing us to build a generic algorithm in Section 5. We first need to introduce the notion of *Power diagram*.

Power diagrams. Let $P \subseteq \mathbb{R}^3 \times \mathbb{R}$ be a weighted point cloud, i.e. $P = \{(p_i, \omega_i)\}_{1 \leq i \leq N}$ with $p_i \in \mathbb{R}^2$ and $\omega_i \in \mathbb{R}$. The *Power cell* of the i th point p_i is given by

$$\text{Pow}_i(P) := \{x \in \mathbb{R}^3 \mid \forall j, \|x - p_i\|^2 + \omega_i \leq \|x - p_j\|^2 + \omega_j\}.$$

Power cells partition \mathbb{R}^3 into convex polyhedra up to a negligible set. Power diagrams are well-studied objects appearing in computational geometry [Aurenhammer, 1987], and can be computed efficiently in dimension 2 and 3. When all the weights are equal, one recover the usual *Voronoi diagram*.

Visibility diagram as a restricted Power diagram. We now show that in all the non-imaging problems of Section 3, the visibility cells are of the form

$$V_i(\psi) = \text{Pow}_i(P) \cap X. \quad (4)$$

For a collimated source, X denotes the plane $\mathbb{R}^2 \times \{0\}$ and for a point source, X is the unit sphere \mathbb{S}^2 . The expression of the weighted point cloud $P = \{(p_i, \omega_i)\}$ depends on the problem. We refer to Table 1 and the work of de Castro et al., 2016 for formulas in the (PS/Mirror) case, the other ones being obtained in a similar fashion. Let us show the derivation of

Table 1: Formulas for the weighted points used to define the Power cells in Equation (4) for the various problems. In the lens design problem, $\kappa > 0$ is the ratio of the indices of refraction, $\kappa > 1$ in the (PS/Lens) setting. Ccv means concave and Cvx convex. $\widetilde{\text{Ccv}}$ means that the optical component converges to a concave one when the discretization tends to infinity.

Type	Points	Weights
Cvx (CS/Mirror)	$p_i = -\frac{p_{\mathbb{R}^2}(y_i - e_z)}{\langle y_i - e_z e_z \rangle}$	$\omega_i = 2\psi_i - \ p_i\ ^2$
Ccv (CS/Mirror)	$p_i = \frac{p_{\mathbb{R}^2}(y_i - e_z)}{\langle y_i - e_z e_z \rangle}$	$\omega_i = 2\psi_i - \ p_i\ ^2$
Cvx (PS/Mirror)	$p_i = -\frac{y_i}{2 \ln(\psi_i)}$	$\omega_i = -\frac{1}{\ln(\psi_i)} - \frac{1}{4 \ln(\psi_i)^2}$
$\widetilde{\text{Ccv}}$ (PS/Mirror)	$p_i = y_i / (2 \ln(\psi_i))$	$\omega_i = \frac{1}{\ln(\psi_i)} - \frac{1}{4 \ln(\psi_i)^2}$
Cvx (CS/Lens)	$p_i = -\frac{p_{\mathbb{R}^2}(y_i - \kappa e_z)}{\langle y_i - \kappa e_z e_z \rangle}$	$\omega_i = 2\psi_i - \ p_i\ ^2$
Ccv (CS/Lens)	$p_i = \frac{p_{\mathbb{R}^2}(y_i - \kappa e_z)}{\langle y_i - \kappa e_z e_z \rangle}$	$\omega_i = 2\psi_i - \ p_i\ ^2$
Cvx (PS/Lens)	$p_i = -\kappa \frac{y_i}{2 \ln(\psi_i)}$	$\omega_i = -\frac{1}{\ln(\psi_i)} - \frac{\kappa^2}{4 \ln(\psi_i)^2}$
$\widetilde{\text{Ccv}}$ (PS/Lens)	$p_i = \kappa y_i / (2 \ln(\psi_i))$	$\omega_i = \frac{1}{\ln(\psi_i)} - \frac{\kappa^2}{4 \ln(\psi_i)^2}$

the formula in the (CS/Mirror) case, where the i th visibility cell is given by

$$\begin{aligned} V_i(\psi) &= \{x \in \mathbb{R}^2 \times \{0\} \mid \forall j, -\langle x | p_i \rangle + \psi_i \leq -\langle x | p_j \rangle + \psi_j\} \\ &= \{x \in \mathbb{R}^2 \times \{0\} \mid \forall j, \|x - p_i\|^2 + \omega_i \leq \|x - p_j\|^2 + \omega_j\}, \end{aligned}$$

where $\omega_i = 2\psi_i - \|p_i\|^2$. We conclude that the visibility cells for a convex mirror of the (CS/Mirror) problem are indeed given by (4), where the weighted point cloud is given by the first line of Table 1.

5. A GENERIC ALGORITHM

For each optical design problem, given a light source intensity function, a target light intensity function and a tolerance, Algorithm 1 outputs a triangulation of a mirror or a lens that satisfies the light energy conservation equation (2).

The main problem is to find weights ψ such that $G(\psi) = \sigma$ (see Equation (3)). This is done using a damped Newton algorithm similar to recent algorithms that have been shown to have a quadratic local convergence rate for optimal transport problems [Kitagawa et al., 2016] or for Monge-Ampère equations in the plane [Mirebeau, 2015]. A key point of this algorithm is to enforce the Jacobian matrix $DG(\psi)$ to always be of rank $N - 1$. To this purpose, we need to enforce all along the process that

$$\forall i \in \{1, \dots, N\}, G_i(\psi) > 0. \quad (5)$$

Indeed, first note that since G is invariant under the addition of a constant, the kernel of $DG(\psi)$ always contains the constant vector $(1, \dots, 1)$. Now note that if we have $G_i(\psi) = 0$, then

ALGORITHM 1: Mirror / lens construction

Input A light source intensity function ρ_{in} .
A target light intensity function σ_{in} .
A tolerance $\eta > 0$.

Output A triangulation \mathcal{R}_T of a mirror or lens.

Step 1 Initialization (Section 5.1)

$T, \rho \leftarrow \text{DISCRETIZATION_SOURCE}(\rho_{in})$
 $Y, \sigma \leftarrow \text{DISCRETIZATION_TARGET}(\sigma_{in})$
 $\psi^0 \leftarrow \text{INITIAL_WEIGHTS}(Y)$

Step 2 Solve Equation (3): $G(\psi) = \sigma$ (Section 5.2)

$\psi \leftarrow \text{DAMPED_NEWTON}(T, \rho, Y, \sigma, \psi^0, \eta)$

Step 3 Construct a triangulation \mathcal{R}_T of \mathcal{R} (Section 5.3)

$\mathcal{R}_T \leftarrow \text{SURFACE_CONSTRUCTION}(\psi, \mathcal{R}_\psi)$

the corresponding visibility cell $V_i(\psi)$ is empty, which implies that $\nabla G_i(\psi) = 0$ (the gradient being taken with respect to ψ). This is because the gradient of G_i involves integral on the boundary $\partial V_i(\psi)$, as shown for instance by Kitagawa et al., 2016 in Theorem 1.3. Hence, if $G_i(\psi) = 0$, then the rank of $DG(\psi)$ is at most $N - 2$ which prevents from using the Damped Newton method. Our method consists of three steps, described in Algorithm 1:

- **Initialization** (Sec. 5.1): We first discretize the source density into a piecewise affine density and the target one into a finitely supported measure. Then, we construct initial weights ψ^0 satisfying $\forall i, G_i(\psi^0) > 0$.
- **Damped Newton** (Sec 5.2): We construct a sequence ψ^k following Algorithm 2 until $\|G(\psi^k) - \sigma\|_\infty \leq \eta$. The main difficulty here is to evaluate $G(\psi^k)$ and $DG(\psi^k)$.
- **Surface construction** (Sec 5.3): Finally, we convert the solution $\psi^k \in \mathbb{R}^N$ into a triangulation. Depending on the non-imaging problem, this amounts to approximating an intersection (or union) of half-spaces (or solid paraboloids, or ellipsoids) by a triangulation.

5.1. Initialization

Discretization of light intensity functions Our framework allows to handle any kind of collimated or point light source or target light intensity functions. It can be for example any positive function on the plane or the sphere (depending on the problem) or a greyscale image, which we see as piecewise affine function. We first approach the support of the source density ρ by a triangulation T and assume that the density $\rho : T \rightarrow \mathbb{R}^+$ is affine on each triangle. We then normalize ρ by dividing it by the total integral $\int_T \rho(x) dx$.

ALGORITHM 2: Damped Newton method for $G(\psi) = \sigma$

Input The source $\rho : T \rightarrow \mathbb{R}^+$ and target $\sigma = \sum_i \sigma_i \delta_{y_i}$; an initial vector ψ^0 and a tolerance $\eta > 0$.

Step 1 Transformation to an Optimal Transport problem

If $X = \mathbb{R}^2 \times \{0\}$, then $\tilde{\psi}^0 = \psi^0$ (and $\tilde{G} = G$).

If $X = \mathbb{S}^2$, then $\tilde{\psi}^0 = (\ln(\psi_i^0))_{1 \leq i \leq N}$ (and $\tilde{G} = (G_i \circ \exp)_{1 \leq i \leq N}$).

Step 2 Solve the equation: $\tilde{G}(\tilde{\psi}) = \sigma$

Initialization: $\varepsilon_0 := \min [\min_i G_i(\psi^0), \min_i \sigma_i] > 0$,
 $k := 0$.

While $\|\tilde{G}(\tilde{\psi}^k) - \sigma\|_\infty \geq \eta$

- Compute $d_k = -DG(\tilde{\psi}^k)^+(\tilde{G}(\tilde{\psi}^k) - \sigma)$
- Find the smallest $\ell \in \mathbb{N}$ s.t. $\tilde{\psi}^{k,\ell} := \tilde{\psi}^k + 2^{-\ell} d_k$ satisfies

$$\begin{cases} \min_i \tilde{G}_i(\tilde{\psi}^{k,\ell}) \geq \varepsilon_0 \\ \|\tilde{G}(\tilde{\psi}^{k,\ell}) - \sigma\|_\infty \leq (1 - 2^{-(\ell+1)}) \|\tilde{G}(\tilde{\psi}^k) - \sigma\|_\infty \end{cases}$$

- Set $\tilde{\psi}^{k+1} = \tilde{\psi}^k + 2^{-\ell} d_k$ and $k \leftarrow k + 1$.

Return $\psi := (\tilde{\psi}_i^k)_{1 \leq i \leq N}$ if $X = \mathbb{R}^2 \times \{0\}$ or
 $\psi := (\exp(\tilde{\psi}_i^k))_{1 \leq i \leq N}$ if $X = \mathbb{S}^2$.

Similarly, the target light intensity function can also be any discrete probability measure. If the user provides an image, one can transform it into a discrete measure of the form $\sigma = \sum_i \sigma_i \delta_{y_i}$ using Lloyd's algorithm or more simply by taking one Dirac mass per pixel. We do the latter in all experiments. The target measure is also normalized by dividing with the discrete integral $\sum_i \sigma_i$. We need $\min_i \sigma_i > 0$ for the damped Newton algorithm, but this is not a restriction: if $\sigma_i = 0$, we simply remove the corresponding Dirac mass δ_{y_i} , thus ensuring that no light is sent towards y_i .

Choice of the initial family of weights ψ^0 . As mentioned at the beginning of this section, we need to ensure that at each iteration all the visibility cells have non-empty interiors. In particular, we need to choose a set of initial weights $\psi^0 = (\psi_i^0)_{1 \leq i \leq N}$ such that the initial visibility cells are not empty.

- For the collimated light sources cases (CS/Mirror) and (CS/Lens), we see that if we choose $\psi_i^0 = \|p_i\|^2/2$ then $\omega_i = 0$, where p_i is obtained using the formulas of the Section 4. Then, the visibility diagram becomes a Voronoi diagram, hence $p_i \in V_i(\psi^0)$.

- For the Point Source Mirror (PS/Mirror) case, an easy calculation shows that if we choose $\psi_i^0 = 1$, then $-y_i \in V_i(\psi^0)$.
- For the Point Source Lens (PS/Lens) case, we can show that if we also choose $\psi_i^0 = 1$, then $y_i \in V_i(\psi^0)$.

Note that the previous expressions for ψ^0 ensure that $G_i(\psi^0) = \rho(V_i(\psi^0)) > 0$ only when the support X_ρ of the light source is large enough. As an example in the (PS/Mirror) case, if $-y_i \notin X_\rho$, then we may have $G_i(\psi^0) = 0$. To handle this difficulty, we use a linear interpolation between ρ and a constant density supported on a set that contains the $(-y_i)$'s. This strategy also works for the (CS/Mirror), (PS/Lens) and (CS/Lens) cases.

5.2. Damped Newton algorithm

When the light source is collimated (*i.e.* $X = \mathbb{R}^2 \times \{0\}$), the problem is known to be an optimal transport problem in the plane for the quadratic cost, the function G is the gradient of a concave function, its Jacobian matrix DG is symmetric and $DG \leq 0$. Moreover, if $G_i(\psi) > 0$ for all i and if X_ρ is connected, then the kernel of DG is spanned by $\psi = \text{cst}$. This ensures the convergence of the damped Newton algorithm [Kitagawa et al., 2016] presented as Algorithm 2, where A^+ denotes the *pseudo-inverse* of the matrix A . Practically, taking the pseudo-inverse of $DG(\tilde{\psi}^k)$ guarantees that the mean of the $\tilde{\psi}^k$ remains constant. In practice, we remove a line and a column of the matrix to make it full rank.

When the light source is a point source, we make the change of variable $\tilde{\psi} = \ln(\psi)$ and $G = G \circ \exp$, so that $G(\psi) = \sigma$ if and only if $\tilde{G}(\tilde{\psi}) = \sigma$. This change of variable turns the optical component design problem into an optimal transport problem, ensuring that \tilde{G} is the gradient of a concave function and that $D\tilde{G}$ is symmetric negative [de Castro et al., 2016], thus easily invertible. In the (PS/Mirror) problem with convex mirrors, the damped Newton algorithm is also provably converging [Kitagawa et al., 2016].

Computation of G and DG By Section 4, the visibility cells $V_i(\psi)$ can be computed by intersecting a certain 3D Power diagram with a triangulation T of the support X_ρ of ρ . Such intersection can be computed using the algorithm developed by Lévy, 2015 . Then,

$$G_i(\psi) = \int_{V_i(\psi)} \rho(x) dx$$

can be computed using first order quadrature formulas. The computation of DG is done using forward-mode automatic differentiation, where we store the gradient of $G_i(\psi)$ as a sparse vector. Note that this works quite efficiently since all numbers that occur in the computation of $G_i(\psi)$ depend only on

ALGORITHM 3: Optical component design for a NF target.

Input The source $\rho : T \rightarrow \mathbb{R}^+$ and target $\sigma = \sum_{i=1}^N \sigma_i \delta_{z_i}$; an initial vector ψ^0 and two tolerances $\eta, \eta_{NF} > 0$.

Initialization $\forall i, c_i^0 = O$

While $\|c^{k+1} - c^k\|_1 / N > \eta_{NF}$

- Compute $v_i^k = \mathcal{R}_{\psi^k}(c_i^k)$
 - Set $y_i^k = (z_i - v_i^k) / \|z_i - v_i^k\|$
 - Solve $\psi^{k+1} \leftarrow \text{SOLVE_FF}(T, \rho, Y^k, \sigma, \eta)$
 - Update c_i^{k+1} to be the centroid of $V_i(\psi^{k+1})$
-

the values ψ_j where j is such that (i, j) are neighbors in the visibility diagram, i.e. $V_i(\psi) \cap V_j(\psi) \neq \emptyset$.

Linear system Since $D\tilde{G}$ is sparse and symmetric negative, we solve the linear systems using preconditioned conjugate gradient.

5.3. Surface construction

In the last step of Algorithm 1, we build a triangulation of the mirror or lens surface. The input is a family of weights ψ solving Equation (3) and the parameterization function \mathcal{R}_ψ whose formula is given in Section 3 and depends on the eight different cases. We triangulate each visibility cell by taking the convex hull of the vertices of its boundary. A vertex of the triangulation will belong to at least one visibility cell. For each vertex, we can compute exactly the normal to the (continuous surface) using Snell's law since we know the incident ray and the corresponding reflected/refracted direction y_i .

6. FINITE-DISTANCE CAUSTICS

In this section, we show that we can solve the eight optical component design problems when the target is at a finite distance. This setting is called *near-field* (NF) in contrast with the previous *far-field* (FF) case that deals with targets at infinity. This setting is interesting since in most applications, one wants the focused image to be at a finite distance and not at infinity.

To be more precise, given one of the eight optical component design problems mentioned above with a target illumination $\sigma = \sum_{i=1}^N \sigma_i \delta_{z_i}$ supported on a set of points $Z = \{z_1, \dots, z_N\} \subset \mathbb{R}^3$, we propose an algorithm that iteratively solves a FF problem, namely Equation (2), and converges to a solution of the NF problem.

6.1. Algorithm

The procedure consists in solving a sequence of FF problems that quickly converges to a solution of the NF one. Details can be found in Algorithm 3. In this algorithm, $\text{SOLVE_FF}(T, \rho, Y, \sigma, \eta)$ denotes an algorithm that solves the FF problem between a source $\rho : T \rightarrow \mathbb{R}^+$ and a target $\sigma = \sum_i \sigma_i \delta_{y_i}$ supported on $Y \subset \mathbb{S}^2$ for a numerical error η . It can for instance be Step 2 of Algorithm 1. This algorithm is used to produce all the lenses and mirrors of the article, except for the first image of Figure 15.

6.2. Convergence analysis

It is clear that when a fixed point is reached in Algorithm 3, the corresponding weight vector ψ is a good approximation of the NF problem. In practice, the algorithm converges very quickly: in all our examples, after 6 iterations, we get an error η_{NF} of 10^{-6} , see Figure 5.

We have no guarantee on the convergence of the discretization. However, note that the set of reflected (or refracted) rays emanating from a visibility cell has a diameter proportional to the diameter of the visibility cell. This generates a blur whose diameter is also proportional to the diameter of the cells. In practice, we observe that these cells have a small diameter (see Figures 7 and 10) and indeed the blur is unnoticeable. The convergence is illustrated in Figure 6 where we show LuxRender renderings at iteration 1 (which corresponds to the FF setting), at iteration 2 and after 6 steps (when the error $\|c^{k+1} - c^k\|_1 / N$ becomes smaller than $\eta_{NF} = 10^{-6}$). In all the images, we look at the projection of the refracted illumination onto the target screen. One can see that the first one is distorted since it corresponds to the projection of an image supported on the sphere onto a planar screen. Starting from the second iteration of the algorithm, the image is not distorted anymore.

The convergence of this algorithm is also illustrated in Figure 15. On the first image of Figure 15, the optical component is composed of three lenses that solve the FF problem. We observe that the LuxRender rendering creates three shifted copies of the same image. To be more precise, the translation between the 3 images is the same as the one between the lenses i.e. if each lens has a width of 1 then each image is also shifted by 1. In the second image, the nine lenses solve the NF problem for an image at a finite distance. The fact that the images do superimpose show that the NF problem is accurately solved. The quality of the first image of Figure 1 also assesses that the NF problem is solved accurately.

Table 2: Running time and number of Newton steps in Algorithm 2 (FF target) for the (CS/Mirror) and the TRAIN target in the far-field setting.

size	time	# Newton steps
128 ²	9s	11 iterations
256 ²	38s	13 iterations
512 ²	245s	15 iterations
1024 ²	1598s	18 iterations
2048 ²	7538s	24 iterations




Figure 5: Mean error ($\|c^{k+1} - c^k\|_1 / N$) in Algorithm 3 for a uniform collimated light source. Red: (CS/Lens) with the TRAIN target; Blue: (CS/Mirror) with the TRAIN target. The y-axis is in logarithmic scale.

6.3. Performance

We report performance of algorithms 1 and 3 to solve the (CS/Mirror) case for the TRAIN target on a laptop with an i7 CPU. Table 2 shows the running times in the FF setting. We underline that the number of iterations in the Newton step remains low: in all our examples it varies from 10 to 20. This means that the computational cost of the method is concentrated in the computation of the functional G , its Jacobian matrix DG and the resolution of the linear system. We believe that there is much room for improvement in the first two steps, by optimizing the computation of visibility cells, and by using an explicit computation of DG instead of relying on automatic differentiation.

In the NF setting, with the same configuration, the results are summarized in Table 3. The total running time is the one after $k = 6$ iterations. Note that the running time of the second iteration is greater than the first one. Indeed, in Algorithm 3, we saw that we use the weights found at one step to initialize the next one. Since the target directions change greatly between the first and second steps, these weights do not provide a good initialization. To fix that, we automatically apply a simple perturbation to obtain good initial weights. Starting from the third step, the running times decrease as the target directions do not move a lot at each iteration.

Table 3: Running time in Algorithm 3 (NF target) for the (CS/Mirror) and the TRAIN target in the near-field setting for different discretizations.

size	$k = 1$	$k = 2$	$k = 3$	$k = 4$	Total ($k = 6$)
128 ²	9s	9s	6s	2s	31s
256 ²	38s	61s	38s	31s	228s
512 ²	245s	294s	240s	194s	1303s
1024 ²	1598s	2095s	1586s	1489s	9077s



Figure 6: Forward simulation of different iterations of Algorithm 3. From left to right: target image, 1st iteration; 2nd iteration; after 6 iterations. First row: Concave (CS/Lens) TRAIN. Second row: Convex (CS/Mirror) TRAIN.

7. RESULTS AND DISCUSSION

In this section, we present several numerical examples for the different problems previously described as well as some other applications. In the experiments, we take $\kappa = 1.5$. Unless stated otherwise, the light source is chosen uniform and the discretization of the target (number of Diracs N) is equal to the size of the image. The stopping criterion to Newton's algorithm (Algorithm 2) is set to $\eta = 10^{-8}$. Since the convergence of Algorithm 3 is always fast, we do not use here the stopping criteria η_{NF} and stop the algorithm after $k = 6$ iterations in all the examples presented hereafter.

7.1. Evaluation strategy

The output of our algorithm is a triangulation equipped with a normal at each vertex. In all the simulations, we use the **LuxRender** rendering engine, with Bidirectional Path Tracing combined with a Sobol sampler and the Fresnel coefficient is not taken into account.

7.2. Results

Genericity. Our algorithm is able to solve eight different optical component design problems. We present for instance in Figure 7 four examples for which we display the visibility diagram of X_ρ as well as the optical component (lens or mirror) above it, a mesh of the optical component and a forward simulation using LuxRender.

Then, for the examples of Figures 8, 9, 11, and 12, we display the target distribution as an image; a mean curvature plot (blue represents low mean curvature and red high mean curvature) of the constructed mesh \mathcal{R}_T and a forward simulation using LuxRender.

High-contrast and complex target lights. We can handle any kind of target distribution. Figures 8 and 9 shows several examples of mirror design for respectively a collimated and a point light source. Note that we are able to construct mirrors for smooth images such as the TRAIN image as well as images with totally black areas (third and fourth rows). We are also able to handle target supported on *non-convex* sets such as the HIKARI and SIGGRAPH images. One can notice that since the area of the visibility cells are equal to the greyscale values of the image then the triangles have roughly the same size, implying that one can recognize the target image in the mesh of the surface, see Figure 7 for zooms on different meshes. The mean curvature plot shows the discontinuities in the surface which come from the black areas in the image. Figures 11 and 12 show the same kind of results for the lens design problems (CS/Lens) and (PS/Lens).

Non-uniform light sources. Our algorithms can be used with non-uniform light sources. Below, we compare the meshes that are generated in the (CS/Lens) case when the source is either uniform or a Gaussian (Fig. 10 left). Because of the higher concentration of light, the details of the triangulation are more concentrated in the middle in the Gaussian case (middle) than in the uniform case (right).

Convex / concave optical components. As shown in Section 3, for each problem, one can choose between two different parameterizations. For instance, for the (CS/Lens) problem, one can build a lens which is either concave or convex, see Figure 4 for an illustration of these differences. Note that in the (PS/Lens) setting (which corresponds to the last row of Figure 7 and Figure 12), the light source is not supported on the full hemisphere \mathbb{S}_+^2 but instead on a smaller part of it. Indeed, choosing a smaller support for μ enforces that \mathcal{R}_T is a graph above the plane instead of the hemisphere and thus avoids potential inter-refractions. Furthermore, since we




Figure 7: **Four non-imaging problems solved with Algorithm 1.** *From left to right:* visibility diagram on X_ρ (wireframe) with the optical component \mathcal{R} , Triangulation \mathcal{R}_T of \mathcal{R} ; forward simulation using LuxRender. *From top to bottom:* Convex Collimated Source Mirror; Concave Point Source Mirror; Concave Collimated Source Lens; Point Source Lens (union of ellipsoids).




Figure 8: **Convex Collimated Source Mirror** problem with a uniform light source for different target distributions. *From left to right:* target distribution, mean curvature of the mirror, forward simulation using LuxRender. Dimensions of images (top to bottom): 256^2 , 300^2 , 400^2 .




Figure 9: **Concave Point Source Mirror** problem for a uniform point light source with different target distributions. *From left to right:* target distribution, mean curvature of the mirror (top view), forward simulation using LuxRender. Dimensions of images (top to bottom): 256^2 , 300^2 , 400^2 .




Figure 10: **Triangulation \mathcal{R}_T for a non-uniform light source.** *From left to right:* non uniform collimated light source; mesh of the lens for this non-uniform light; mesh of the lens for a uniform light source.




Figure 11: **Concave Collimated Source Lens** with a uniform light source for different target distributions. *From left to right*: target distribution, mean curvature of the lens (top view), forward simulation using LuxRender. Dimensions of images (top to bottom): 256^2 , 300^2 , 400^2 .

parametrize the lens as a union of ellipsoids, it is neither convex nor concave. As for all figures, we have performed no post-processing on \mathcal{R}_T in order to emphasize the benefit of designing convex or concave optical components (convexity is a form of regularity). One also observes that when the lens is rotated with respect to the light source (Figure 14 and first row of Figure 19), or when the target screen is not at the right distance (Figure 18), the image is deformed in a monotonic and regular way. We believe this is due to the monotonicity properties of optimal transport and to the convexity/concavity properties of the optical components.

Comparison with previous work. Figure 13 compares the state of the art results obtained by Schwartzburg et al., 2014 (second column) and the LuxRender renderings obtained by our method (third column) on two target distributions for the (CS/Lens) case with a collimated uniform light source in the near-field setting. Although the results are comparable, one can notice, in the second column, the presence of small artifacts between the black and white regions, for instance around the rings (notably in the center). The contrast is more accurate with our convex lenses.




Figure 12: **Point Source Lens** with a uniform light source for different target distributions. The lens surface is the boundary of the union of filled ellipsoids, hence is not convex, nor concave. *From left to right*: target distribution, mean curvature of the lens (top view), forward simulation using LuxRender. Dimensions of images (top to bottom): 256^2 , 300^2 , 400^2 .




Figure 13: **Comparison with [Schwartzburg et al., 2014]** *From left to right:* target distribution; images obtained by [Schwartzburg et al., 2014] and taken from their article; our forward simulation using LuxRender. Last row: meshes of the two corresponding convex lenses: RINGS (left) and SIGGRAPH (right).




Figure 15: **Pillows and differences between FF and NF.** *Top:* The lens is composed of three pillows that solve the FF problem. Note in the last image the shift between the three projected images. *Middle:* A lens composed of nine pillows (each of them solving the NF problem) that refracts a uniform collimated light source; *Bottom:* The same lens with an obstacle in red.



Figure 14: **Stability under rotation of the lens.** LuxRender renderings in the (CS/Lens) setting for the TRAIN target while rotating the lens with respect to the direction of the collimated light source (0° / 5° / 15°).

Application to pillows This problem consists in decomposing the optical component (mirror or lens) into several smaller optical components that are called *pillows*, as illustrated in Figure 1, and are widely used in car headlight design. Each pillow independently satisfies a non-imaging problem with the same target light, but with a different source (since it receives only a portion of the light). Hence, the optical component made with all the pillows glued together is more reliable and allows for example to reduce the artifacts due to small occluders. Indeed if one object is in front of one or more lenses, the quality of the refracted image decreases but the image can still be recognized. Using pillows also gives some flexibility to the designer to improve the appearance and the volume occupied by the component. An example with 9 pillows can be found in Figure 15, and the effect of a small occluder. In practice, in order to avoid a shift between the nine simulated images, we solve the near-field problem (see Section 6).

Application to color images Using pillows, we can also target color images. Indeed, we can build one component for each of the Red, Green and Blue channels of an image. If we then place three lights (red, green and blue) in front of each component, using Algorithm 3 the 3 images will be perfectly align and thus produce the original color image, see the first image of Figure 1.

8. PHYSICAL PROTOTYPES

We built three lenses (see Figure 16) and two mirrors (see figures 17 and 1) corresponding to collimated light sources. The lenses are fabricated in PMMA (whose index of refraction is 1.49) and the mirrors in aluminum. All the lenses and mirrors have size $100mm \times 100mm$ and were milled in one pass on 3-axis CNC machines after milling the blank. For the Hikari lens, we choose to focalize the target image on a wall at 2 meters and the target is a square of size $600mm \times 600mm$. For the Train and Einstein signature lenses and the two mirrors, we choose to focalize the target image on a wall at 1 meter and the target is a square of size $300mm \times 300mm$. The five components were milled in one pass on the *DMG-DC100V* machine with a $10mm$ radius end-mill.

The milling process is very sensitive. For the lenses, the end-mill is a super finishing ball mill D10, 3 teeth and is following a concentric spiral trajectory. For the mirrors the end-mill is a PCD ball mill hooped D10, 2 teeth and is following a parallel scanning trajectory. We observe that the precision of the milling is not accurate enough: when the collimated light source is traversing a lens or reflected by a mirror with no sandpapering and polishing, the light is dispersed and we do not recognize the target (see Figure 19, second row). We had to sandpaper them by hand before polishing them with a polishing paste. This clearly damages the lens surface: there is a tradeoff between removing the artifacts due to the milling and smoothing too much the surface (see Figure 19 rows 3-5), thus damaging its refractive properties. Note that thanks to the convexity property (see Figure 19 row 4), the lens surface is quite regular and is more robust to sandpapering.

We can also observe some artifacts in the milling process. For instance, some corrugations are present in the lens (Figure 19 first row) and induce some artifacts in the projected image (Figure 16, first row). We observe that although the image are very contrasted, the projected image are very accurate. The boundary of the target is often slightly blurred and this is due to the boundary of the lens or mirror where the milling was less good. Our model do not take into account the different wavelengths of the white color and we observe on the boundary of the projected images a small chromatic aberration (the boundary is slightly blue).

9. CONCLUSION AND PERSPECTIVES

We presented a general framework for eight different optical component design problems satisfying light energy constraints. We proposed an efficient algorithm able to solve them whether the target is at infinity (far-field) or at a finite distance (near-field). The main limitation of the approach is the fact that we only deal with ideal light sources (a light bulb is for instance neither collimated nor punctual). Another limitation is that we do not account for self shadowing and internal reflections (although, this is not a problem in the situations we have encountered). In the future, we also want to try fabricating physical prototypes when the source is punctual. This problem becomes more difficult for the lenses since we also have to mill the two sides of the lens, and in particular the inner sphere. We also believe that the robustness and versatility of the proposed approach can make it a useful component for the design of heuristics able to deal with extended light sources and in computer graphics for caustic design.

Acknowledgements We would like to thank Alain Di Donato from the *GINOVA* platform, *S.MART DS* department for building the physical prototypes, as well as the GMP department of the Institute of Technology of Aix-en-Provence in France, and André Lieutier and Jimmy Cresson from Dassault-Systèmes for helping with the fabrication process, and Philippe Halot for taking the pictures. This work has been partially supported by the LabEx PERSYVAL-Lab (ANR-11-LABX-0025-01) and by ANR project MAGA (ANR-16-CE40-0014 MAGA).

REFERENCES

- [Aurenhammer, 1987] Aurenhammer, F. (1987). Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96.
- [Caffarelli et al., 1999] Caffarelli, L., Kochengin, S., and Oliker, V. (1999). On the numerical solution of the problem of reflector design with given far-field scattering data. *Contemporary Mathematics*, 226:13–32.
- [Caffarelli et al., 2008] Caffarelli, L. A., Gutiérrez, C. E., and Huang, Q. (2008). On the regularity of reflector antennas. *Annals of mathematics*, 167(1):299–323.
- [Caffarelli and Oliker, 2008] Caffarelli, L. A. and Oliker, V. (2008). Weak solutions of one inverse problem in geometric optics. *Journal of Mathematical Sciences*, 154(1):39–49.



Figure 16: **Fabricated lenses for a collimated light source.** From left to right: experimental setup, zoom on the target screen. From top to bottom: TRAIN, HIKARI, EINSTEIN'S SIGNATURE targets. Images are focused on a screen at 2 meters for the first two rows and 1 meter for the last one.



Figure 17: **Fabricated mirror for a collimated light source.** HIKARI target.



Figure 18: **Stability with respect to the depth of the focus plane.** The lens of HIKARI is designed to focus at a distance of 2 meters. The target screen is at different depths, top: 1.5 meters; bottom: 1m (left), 50cm (right).




Figure 19: **Fabrication process.** *First row:* lens for the TRAIN target with a zoom on milling errors (left). The lens is rotated (20 degrees) around the axis of the light source (right). *Second row:* Collimated light projected after reflection or refraction on the screen (rough mirror and rough EINSTEIN’S SIGNATURE) *Third row:* Rough mirror / sandpapered and polished mirror. *Fourth row:* Rough lens / sandpapered and polished lens. *Fifth row:* Sandpapering with water / Polishing by hand.

- [Cork et al., 1977] Cork, F., Bettridge, D., and Clarke, P. (1977). Method of and mixture for aluminizing a metal surface. US Patent 4,009,146.
- [de Castro et al., 2016] de Castro, P. M. M., Mérigot, Q., and Thibert, B. (2016). Far-field reflector problem and intersection of paraboloids. *Numerische Mathematik*, 134(2):389–411.
- [de Goes et al., 2012] de Goes, F., Breeden, K., Ostromoukhov, V., and Desbrun, M. (2012). Blue noise through optimal transport. *ACM Transactions on Graphics*, 31(6):171.
- [Finckh et al., 2010] Finckh, M., Dammertz, H., and Lensch, H. (2010). Geometry construction from caustic images. In *Proceedings of the 11th European conference on Computer vision: Part V*, pages 464–477. Springer-Verlag.
- [Glimm and Olikar, 2003] Glimm, T. and Olikar, V. (2003). Optical design of single reflector systems and the monge–kantorovich mass transfer problem. *Journal of Mathematical Sciences*, 117(3):4096–4108.
- [Gutiérrez, 2016] Gutiérrez, C. E. (2016). *The Monge-Ampère Equation*, volume 89. Birkhäuser.
- [Gutiérrez and Huang, 2009] Gutiérrez, C. E. and Huang, Q. (2009). The refractor problem in reshaping light beams. *Archive for rational mechanics and analysis*, 193(2):423–443.
- [Gutiérrez and Tournier, 2013] Gutiérrez, C. E. and Tournier, F. (2013). The parallel refractor. In *From Fourier Analysis and Number Theory to Radon Transforms and Geometry*, pages 325–334. Springer.
- [Kiser et al., 2013] Kiser, T., Eigensatz, M., Nguyen, M. M., Bompas, P., and Pauly, M. (2013). Architectural caustic—controlling light with geometry. In *Advances in Architectural Geometry 2012*, pages 91–106. Springer.

- [Kitagawa et al., 2016] Kitagawa, J., Mérigot, Q., and Thibert, B. (2016). A newton algorithm for semi-discrete optimal transport. *arXiv preprint arXiv:1603.05579*.
- [Lévy, 2015] Lévy, B. (2015). A numerical algorithm for 12 semi-discrete optimal transport in 3d. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1693–1715.
- [Lloyd, 1982] Lloyd, S. (1982). Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137.

-
- [Mérigot, 2011] Mérigot, Q. (2011). A multiscale approach to optimal transport. *Computer Graphics Forum*, 30(5):1583–1592.
- [Mirebeau, 2015] Mirebeau, J.-M. (2015). Discretization of the 3d monge-ampere operator, between wide stencils and power diagrams. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(5):1511–1523.
- [Papas et al., 2011] Papas, M., Jarosz, W., Jakob, W., Rusinkiewicz, S., Matusik, W., and Weyrich, T. (2011). Goal-based caustics. *Computer Graphics Forum*, 30(2):503–511.
- [Patow and Pueyo, 2005] Patow, G. and Pueyo, X. (2005). A survey of inverse surface design from light transport behavior specification. *Computer Graphics Forum*, 24(4):773–789.
- [Piovarči et al., 2017] Piovarči, M., Wessely, M., Jagielski, M., Alexa, M., Matusik, W., and Didyk, P. (2017). Directional screens. In *Proceedings of the 1st Annual ACM Symposium on Computational Fabrication*, page 1. ACM.
- [Prins et al., 2013] Prins, C., Boonkkamp, T., Roosmalen, v. J. J., IJzerman, W., and Tukker, T. (2013). A numerical method for the design of free-form reflectors for lighting applications. Technische Universiteit Eindhoven.
- [Schwartzburg et al., 2014] Schwartzburg, Y., Testuz, R., Tagliasacchi, A., and Pauly, M. (2014). High-contrast computational caustic design. *ACM Transactions on Graphics (TOG)*, 33(4):74.
- [Wang, 2004] Wang, X. (2004). On the design of a reflector antenna ii. *Calculus of Variations and Partial Differential Equations*, 20(3):329–341.
- [Weyrich et al., 2009] Weyrich, T., Peers, P., Matusik, W., and Rusinkiewicz, S. (2009). Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics*, 28(3):32.
- [Yue et al., 2012] Yue, Y., Iwasaki, K., Chen, B.-Y., Dobashi, Y., and Nishita, T. (2012). Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum*, 31(2pt3):575–582.
- [Yue et al., 2014] Yue, Y., Iwasaki, K., Chen, B.-Y., Dobashi, Y., and Nishita, T. (2014). Poisson-based continuous surface generation for goal-based caustics. *ACM Transactions on Graphics*, 33(3):31.