International Journal of Computational Geometry & Applications

Far-field reflector problem under design constraints

JULIEN ANDRÉ
Laboratoire Jean Kuntzmann and Gipsa-Lab, Université Grenobles-Alpes, France

DOMINIQUE ATTALI
Gipsa-Lab and CNRS, Université Grenobles-Alpes, France

QUENTIN MÉRIGOT
Laboratoire Jean Kuntzmann and CNRS, Université Grenobles-Alpes, France

BORIS THIBERT
Laboratoire Jean Kuntzmann, Université Grenobles-Alpes, France

The far-field reflector problem is a well-known inverse problem arising in geometric optics. It consists in creating a mirror that reflects a given point light source to a prescribed target light at infinity. In this article, we study this problem under the common design constraint that the mirror is convex and is the graph of a polynomial over a given plane. We propose a method that iteratively improves the optical properties of the mirror surface while strictly fulfilling the design constraints. At each iteration, we first create an initial reflector by solving an optimal transport problem on the sphere. We then parameterize this reflector by the graph of a function over the plane. We test our algorithm with classical target lights at infinity and also show that our approach allows to create reflectors with more complex target lights.

Keywords: Reflector problem; optimal transport; nonimaging optic.

1. Introduction

The far-field reflector problem is a well-studied inverse problem in geometric optics. The inputs of this problem are the description of the light distribution emitted by a point light source located at the origin and a desired target distribution of light at infinity, that is, on the sphere of directions. The question is then to create a surface S which reflects the light emitted by the point light source into the target distribution of light.

For applications, the constructed reflector surface also needs to satisfy certain design constraints. In this article, we consider the following two constraints:

(C_g) The surface S should be the graph of a function over a fixed convex domain $\Omega \subseteq \mathbb{R}^2 \times \{0\}$. More precisely, for every u in Ω, the ray joining $(u, -\infty)$ to $(u, +\infty)$ should intersect the surface S exactly once.

(C_c) The surface S should be convex, in the sense that it is contained in the boundary of a convex domain of $\mathbb{R}^3$.

In addition to aesthetic reasons, these constraints are useful for building physical moulds for the reflectors, in particular for car lights. The convexity constraint (C_c) allows one to mill the mould exactly. Furthermore, it creates a reflector surface for which the chemical vapor deposition is easier. Indeed, once the reflector surface is built, some liquid aluminium is sprayed on the reflector surface for the reflection. Generally, the aluminised layer has the tendency to concentrate where the reflector contains small bumps or holes[4] and the (C_c) constraint is a way of avoiding these small bumps. The graph constraint (C_g) is a natural constraint that appears for the construction of car beams. To be more precise, a high beam is in general decomposed into several "pillows", each pillow being a reflector surface whose orthogonal projection onto a given plane is a simple two dimensional domain, such as a rectangle.[3] It is therefore natural to model the surface of each pillow by the graph of a function over a planar domain. The whole reflector, composed of several pillows, is then also the graph of a function. One may notice that this property allows one to be able to remove the mould after the fabrication of the reflector.

1.1. *Related work*

Caustics design. Many of the existing methods rely on a variational approach to solve the far-field reflector inverse problem. In these methods, one needs to choose a parameterization of the space of reflectors by a low-dimensional space $\mathbb{R}^k$, i.e. $\alpha \in \mathbb{R}^k \mapsto S_\alpha$. One also considers an objective function $L : \mathbb{R}^k \to \mathbb{R}$ such that $L(\alpha)$ measures the error between the desired light distribution at infinity and the light distribution that has been obtained after reflection on the reflector S_α. Starting with an initial guess α_0, these methods iteratively minimize the function L. In practice, the function L is highly non-convex, making the global minimum almost impossible to attain. One has to resort to stochastic optimization in order to reach a suitable local minimum. One may refer to Patow and Pueyo[15] for a survey on the reflector problem, including this type of techniques. Kinckh et al.[8] propose an implementation of these ideas capable of dealing with complex target light distributions, using a GPU-based raytracing technique and the stochastic algorithm SPSA. Note that these methods do not guarantee that the constructed reflector surface satisfies the convexity constraint (C_c). Partially related to our work, we should mention a method for designing prisms (and not reflectors) that produce prescribed caustic pattern in the particular case of a directional light source.[16] This method works in two-stage. It relies on optimal transport to construct an initial solution, which is then refined by local optimization. Note however, that this approach is restricted to directional light sources.

Monge-Ampère equations. When the two light distributions can be modeled by probability densities, the far-field reflector problem boils down to solving a Monge-

Ampère equation. This equation is a fully non-linear partial differential equation, whose existence, uniqueness and regularity has been well-studied.[2,18,10] However, because of this non-linearity, and because of the special geometric structure of this equation, there currently exists no numerical approach designed to handle the Monge-Ampère equation directly.

Optimal transport and supporting paraboloids methods. More recently, the reflector problem has been shown to be equivalent to an optimal transport problem on the sphere,[9,19] using a construction involving intersection of solid paraboloids. This optimal transport formulation allows to reformulate the reflector problem as a linear programming problem.[19] Furthermore, it can also be transformed into an unconstrained convex optimization problem, and solved efficiently for a discretization of the target distribution of light with up to $20k$ Dirac masses.[13,5] Additional details about this approach are given in Section 2. Moreover, the solution computed via the optimal transport formulation is a convex patch which satisfies the desired optical properties. Unfortunately, this surface, which is parameterized radially from the origin does not satisfy our design constraint: this surface is usually not the graph of a function over the convex domain $\Omega \subseteq \mathbb{R}^2$ prescribed in ($\mathrm{C_g}$).

1.2. *Contributions*

Our work stems from the following observation. Although supported paraboloid methods allow to find a solution to the reflector problem, the solution does not satisfy our main design constraint ($\mathrm{C_g}$). On the other hand, caustics design methods mentioned earlier can deal with ($\mathrm{C_g}$), but they are based on non-deterministic non-convex optimization methods, and do not guarantee the fact that the generated surface satisfies ($\mathrm{C_c}$). In this article, we propose a method that iteratively improves the optical properties of the surface patch while adhering strictly to the design constraints. This method relies on a fixed-point algorithm which alternates between the resolution of a discrete optimal transport that produces a radially parametrized surface, and the parameterization of this solution by the graph of a grid function.

We show that this method is able to handle complex distribution of light at infinity.

2. Far-field reflector problem and linear programming

In this section, we consider the far-field reflector problem without taking into account the design constraints. Since this is an inverse problem, we start by describing the direct problem more precisely.

2.1. *Direct problem: raytracing*

We consider a point light source, located at the origin of $\mathbb{R}^3$. The distribution of rays of this light source is described by a probability density ρ_0 on the sphere of directions, denoted by $\mathbb{S}_0$.

We are also given a smooth reflector surface $R \subseteq \mathbb{R}^3$. The direct problem consists in describing the distribution of rays reflected by R after a single reflection. For this, we assume that every ray $[0, x)$ starting from the origin and with direction $x \in \mathbb{S}_0$ intersects the surface R at most once, and we let n be the unit outward normal to the reflector at the intersection point $R \cap [0, x)$. We denote by $T_R(x)$ the direction of the reflected ray, given by Snell's law of reflection:

$$T_R(x) := x - 2\ (x \cdot n)\ n.$$

This direction is a unit vector on the sphere of directions $\mathbb{S}_\infty$, where the index ∞ indicates that this concerns the target illumination. The map T_R transports the probability density ρ_0 on $\mathbb{S}_0$ into a probability distribution ρ_∞ on $\mathbb{S}_\infty$.

2.2. *Reflector problem*

The *far-field reflector problem* is the corresponding inverse problem: given a target probability density ρ_∞ on $\mathbb{S}_\infty$, determine a surface R such that the probability distribution of ρ_0 after reflection coincides with ρ_∞. This problem can be formulated as an optimal transport problem[2,19] on the sphere between the probability densities ρ_0 and ρ_∞.

2.3. *Reflector supported by a finite set of paraboloids*

Here, we describe the construction of a reflector that reflects light at the origin only in a finite set of directions Y and define its reflection map. Following the approach of Caffarelli and Oliker,[2] such a reflector surface is made up of pieces of paraboloids focused at the origin with axis directions in Y. Remark that it is natural to use such paraboloids since, by Snell's law of reflection, each paraboloid reflects light originating from its focal point back into its axis direction. To construct our reflector, we thus use one paraboloid for each direction in Y.

Consider the paraboloid focused at the origin, with focal distance $\lambda/2 \geq 0$ and axis direction $y \in Y$. We denote by $P_y(\lambda)$ the convex body whose boundary is this paraboloid. For a direction x in the unit sphere $\mathbb{S}_0$, the intersection between the ray $[0, x)$ and the paraboloid is the point $\lambda x/(1 - x \cdot y)$.

Given a family of focal distances $\boldsymbol{\lambda} = (\lambda_y)_{y \in Y}$, we define the reflector $R_{\boldsymbol{\lambda}}$ as the surface that bounds the intersection of convex bodies $P_y(\lambda_y)$, *i.e.*

$$R_{\boldsymbol{\lambda}} = \partial \left(\bigcap_{y \in Y} P_y(\lambda_y) \right).$$

By construction as an intersection, the reflector $R_{\boldsymbol{\lambda}}$ reflects an incident ray $[0, x)$ into the direction y in Y if and only if the ray $[0, x)$ hits the paraboloid with axis in direction y first; see Figure 1. We consider that if the ray $[0, x)$ hits the reflector at the intersection of two or more paraboloids as illustrated in Figure 1 on the right, then the reflector sends back a ray in the direction of every paraboloid hit by $[0, x)$.

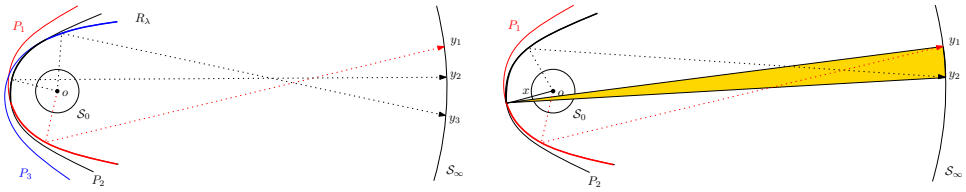

Fig. 1: (Left) A reflector $R_{\boldsymbol{\lambda}}$ composed of three paraboloids, illuminating the directions y_1, y_2 and y_3. (Right) In our model, a ray $[0x)$ that intersects the reflector surface exactly at the intersection of two paraboloids P_1 and P_2 may be reflected to the two directions y_1 and y_2. (In fact, in a more refined model, one could consider that there exists a one-parameter family of possible directions of reflection, which interpolate between y_1 and y_2. Each of these directions corresponds to a vector in the normal cone to the surface at the intersection point.)

It follows that the reflection map is best described using a multivalued function, which associates to a direction $x \in X \subset \mathbb{S}_0$ the set of directions:

$$T_{\boldsymbol{\lambda}}(x) = \arg\min_{y \in Y} \frac{\lambda_y}{1 - x \cdot y}. \tag{1}$$

In the above equation, we thus adopt the convention that $\arg\min_{y\in Y} f(y)$ is the set of values of Y at which $f(y)$ attains its smallest value. This equation can be seen as a generalization of Snell's law of reflection to a restricted class of non-smooth reflectors (see Figure 1 for an alternative generalization).

2.4. *Discretization of the reflector problem*

Here, we describe a possible discretization of the reflector problem, which consists in discretizing the two density probabilities ρ_0 and ρ_∞ as follows. We approximate them respectively by the two discrete probability measures μ and ν given by:

$$\mu = \sum_{x \in X} \mu_x \ \delta_x \quad \text{and} \quad \nu = \sum_{y \in Y} \nu_y \ \delta_y,$$

where δ_u is the dirac mass at u. The approximation of a density probability by a discrete probability measure refers to the quantization problem.

To state the reflector problem, we first need to introduce the notion of transport plan, which is a central notion in the optimal transport theory and was introduced by Kantorovich.[17] A *transport plan* between two measures μ and ν supported on finite sets X and Y is a matrix $(\pi_{x,y})_{x\in X, y\in Y}$ with non-negative entries i.e.

$$\forall x \in X, \ \forall y \in Y \quad \pi_{x,y} \geq 0. \tag{2}$$

The quantity $\pi_{x,y}$ measures the quantity of mass sent from x to y by the transport

plan. The law of conservation of mass then imposes the following linear constraints:

$$\begin{cases} \forall x \in X & \displaystyle\sum_{y \in Y} \pi_{x,y} = \mu_x \\ \forall y \in Y & \displaystyle\sum_{x \in X} \pi_{x,y} = \nu_y. \end{cases} \tag{3}$$

In particular, one has $\sum_x \mu_x = \sum_y \nu_y$. The set of transport plans between the two measures μ and ν is denoted $\Gamma(\mu, \nu)$. It is now possible to state the reflector problem. The intuition behind this formulation is given after the definition.

Definition 1 (Discrete reflector problem). A reflector $R_{\boldsymbol{\lambda}}$ is a solution to the discrete far field reflector problem between two measures μ and ν if there exists a transport plan π^* between μ and ν such that

$$\pi^*_{x,y} \neq 0 \Longrightarrow y \in T_{\boldsymbol{\lambda}}(x). \tag{FF}$$

In other words, π^* is supported on the graph of the multivalued reflection map $T_{\boldsymbol{\lambda}}$.

In this definition, the condition that π^* is a transport plan between μ and ν physically translates to the conservation of energy. The condition expressed in Equation (FF) means that this transport plan is geometrically feasible: if a part of the light emitted by x is sent to y (that is, $\pi^*_{x,y} \neq 0$), then the ray $[0, x)$ has to be reflected in the direction y by the reflector $R_{\boldsymbol{\lambda}}$ (that is, $y \in T_{\boldsymbol{\lambda}}(x)$). We recall that $T_{\boldsymbol{\lambda}}(x)$ is multivalued, meaning that the light emitted from x can be slit into several directions (as illustrated on the right of Figure 1). We will see at the end of Section 2.4 that this definition is easier to understand when the transport plan is a map, that is when for every point x there exists a single point y in Y with $\pi_{x,y} > 0$.

Furthermore, as shown in the lemma below, Equation (FF) is related to the notion of optimal transport: it implies that the transport plan π^* is optimal for the cost $c(x, y) = -\ln(1 - x \cdot y)$.

Lemma 1. *Let $R_{\boldsymbol{\lambda}}$ be a solution to the reflector problem and $\pi^* \in \Gamma(\mu, \nu)$ be a transport plan satisfying Equation (FF). Then*

$$\pi^* \in \arg\min \left\{ \sum_{x \in X} \sum_{y \in Y} c(x, y) \pi_{x,y}, \ \pi \in \Gamma(\mu, \nu) \right\},$$

where c is the cost function given by $c(x, y) = -\ln(1 - x \cdot y)$.

The proof is included in the appendix. Remark that this lemma implies that an optimal transport plan π^* minimizes a linear function under the linear constraints (2)–(3). Hence it is a solution of a linear programming problem.

Relation with the assignment problem. In the particular case where the sizes of the two point sets X and Y have the same cardinality M and the probability

measures μ and ν are uniform (i.e. $\mu_x = \nu_y = 1/M$), the transport plan π^* can be chosen so that the map $T_{\boldsymbol{\lambda}}$ is one-to-one.[17] The formulation of the discrete reflector problem is therefore much easier to state: a reflector $R_{\boldsymbol{\lambda}}$ is a solution to the reflector problem **(FF)** if and only if the map $T_{\boldsymbol{\lambda}}$ induces a bijection from X to Y.

In that case, similarly as in Lemma 1, one can show that $R_{\boldsymbol{\lambda}}$ is a solution to the reflector problem if and only if the map $T_{\boldsymbol{\lambda}}$ minimizes the quantity

$$\sum_{x \in X} c(x, T(x)),$$

over all bijections $T : X \to Y$.[19] Therefore, the reflector problem is equivalent to a linear assignment problem, for which there exists many efficient algorithms.[1]

2.5. *Linear programming formulation*

We saw in Lemma 1 that the transport plan π^* associated to a solution $R_{\boldsymbol{\lambda}}$ to the reflector problem **(FF)** is a solution of a linear programming problem. In fact, we can show that the solution of the dual of this linear programming problem encodes a parametrization of a solution $R_{\boldsymbol{\lambda}}$ of **(FF)**. More precisely, one has:[19]

Theorem 1. *Let $\boldsymbol{a} = (a_x)_{x \in X}$ and $\boldsymbol{b} = (b_y)_{y \in Y}$ be a solution of the linear programming problem*

$$\max_{\boldsymbol{a}, \boldsymbol{b}} \left\{ \sum_{x \in X} a_x \ \mu_x + \sum_{y \in Y} b_y \ \nu_y, \ \forall x, y \ a_x + b_y \leq c(x, y) \right\}.$$

Then the surface $R_{\boldsymbol{\lambda}}$ solves reflector problem, with

$$\boldsymbol{\lambda} = (\exp(b_y))_{y \in Y}.$$

In particular, there always exists a solution $R_{\boldsymbol{\lambda}}$ to the discrete reflector problem. Moreover, for any positive number h, the reflector $R_{h\boldsymbol{\lambda}}$ is homothetic to $R_{\boldsymbol{\lambda}}$, and is also a solution to the reflector problem.

3. Far-field reflector problem for a patch

If the density probabilities ρ_0 and ρ_∞ are continuous, then up to homothety centered at the origin, there is a unique convex solution to the far field reflector problem.[2]

While mathematically pleasing, this uniqueness makes the global solutions too rigid for practical applications. In practice, the reflecting surface is decomposed into patches. In the remainder of the article, we consider only one of these patches. The purpose of our patch is to capture a certain fraction of the light emitted by the source ρ_0 and to reflect it into a prescribed target density ρ_∞. In typical applications, the density ρ_∞ is the indicator function of a domain of $\mathbb{S}_\infty$. Since each patch captures only a fraction of the light emitted by the source, we multiply the target density by a constant factor so as to satisfy conservation of total intensity (Eq. (4)).

Parameterized patch Now, we assume that the patch is parameterized by a rectangle Ω contained in the plane $H = \mathbb{R}^2 \times \{0\}$. Consider a finite-dimensional vector space $\mathcal{F}(\Omega)$ of functions on Ω, which are smooth or piecewise smooth. Given ϕ in $\mathcal{F}(\Omega)$, we denote S_ϕ the *graph* of ϕ, i.e.

$$S_\phi = \{(\boldsymbol{u}, \phi(\boldsymbol{u}));\ \boldsymbol{u} \in \Omega\} \subseteq \mathbb{R}^3.$$

Define $\pi_{\mathbb{S}}(x) = x/\left\|x\right\|$ for a point $x \in \mathbb{R}^3 \setminus \{0\}$. By construction, the surface patch S_ϕ captures all the light rays emitted by the source in a direction that lies in the set $\pi_{\mathbb{S}}(S_\phi)$ (see Figure 2). The total intensity reflected by S_ϕ is given by $\int_{\pi_{\mathbb{S}}(S_\phi)} \rho_0(x)\mathrm{d}x$.

Patch reflector problem Find $\phi \in \mathcal{F}(\Omega)$ such that intensity ρ_0 after reflection by the reflector S_ϕ coincides with $\kappa_\phi \cdot \rho_\infty$, where the constant κ_ϕ is determined by conservation of intensity:

$$\kappa_\phi = \left(\int_{\pi_{\mathbb{S}}(S_\phi)} \rho_0(x)\mathrm{d}x\right) / \left(\int_{\mathbb{S}_\infty} \rho_\infty(x)\mathrm{d}x\right). \tag{4}$$

Remark that the constant κ_ϕ depends on the surface S_ϕ. We have seen in Theorem 1 that the far-field reflector problem can be reformulated as a linear programming problem. It is not straightforward to adapt this formulation to the patch-reflector problem. In the next section, we therefore present a heuristic to handle this more difficult problem.

4. Solving the reflector problem with design constraints

In this section, we present a method for building a patch S_ϕ that reflects the light coming from the source with probability density ρ_0 towards the target with probability density $\kappa_\phi\rho_\infty$. We suppose that the target density ρ_∞ is approximated by a discrete probability measure ν defined on a finite point set Y with size M. The input data is thus a pair (ρ_0, ν) whose two entries are the source density and a discretization of the target density.

Iterative algorithm We approach ϕ iteratively, guided by the following idea. As we modify ϕ so that S_ϕ satisfies the desired optical properties, the part $\pi_{\mathbb{S}}(S_\phi)$ of the sphere $\mathbb{S}_0$ that illuminates S_ϕ evolves and the description of the source and target densities used to approximate S_ϕ must be updated accordingly.

Precisely, we build a sequence of functions ϕ_0, ϕ_1, …, ϕ_k and stop when we reach a fixed point, that is, when $\phi_k = \phi_{k-1}$. In practice, we end the process when ϕ_k and ϕ_{k-1} are sufficiently close. Experimentally, we observe that it suffices to take $k = 3$ or 4. During the course of the algorithm, we use the following discrete representation of the map $\phi_i : \Omega \to \mathbb{R}$. We sample the domain Ω uniformly by a point set U with size N and encode the function ϕ_i by recording for each point $\boldsymbol{u} \in U$ the value of the function $\phi_i(\boldsymbol{u})$ and the value of its gradient $\nabla\phi_i(\boldsymbol{u})$. Indeed, we shall see that this information is all we need to construct ϕ_{i+1} from ϕ_i. Since Ω is a rectangle, the point set U can easily be precomputed by taking for instance

the vertices of a discrete regular grid that approximates Ω and whose convex hull is equal to Ω.

Initially, we let S_{ϕ_0} be a paraboloid focused at the origin with focal distance λ_0 and whose axis is directed towards the center y_0 of the target. Precisely, we take $y_0 = \frac{1}{M}\sum_{y\in Y} y$ and adjust λ_0 so that the part of S_{ϕ_0} which projects orthogonally onto Ω lies entirely above Ω, *i.e.* so that $\phi_0(\boldsymbol{u}) \geq 0$ for all $u \in \Omega$. We then start modifying ϕ_0 iteratively. At step $i \in \{1, \ldots, k\}$ in the process, we construct $\phi = \phi_i$ from $\phi^- = \phi_{i-1}$ in four stages described respectively in Sections 4.1 to 4.4. The pseudo-code of our iterative algorithm is given below. Properties of the fixed point are described in Section 4.5. In Section 4.6 we will explain how to turn the output ϕ of this algorithm into a function ψ which belongs to a given finite-dimensional vector space of functions $\mathcal{F}(\Omega)$.

Algorithm 1 Compute ϕ with error ε.

$\phi \leftarrow$ Parametrization$(\partial P_{y_0}(\lambda_0), U)$ {see Section 4.4}
repeat
  $\phi^- \leftarrow \phi$
  $(\mu, F) \leftarrow$ Discretization(ρ_0, U, ϕ^-) {see Section 4.1}
  $\boldsymbol{\lambda} \leftarrow$ FarFieldReflector$(\mu, \kappa_\mu \nu)$ {see Section 4.2}
  $h \leftarrow$ OptScaleFactor$(\boldsymbol{\lambda}, \mu, F, U)$ {see Section 4.3}
  $\phi \leftarrow$ Parametrization$(R_{h\boldsymbol{\lambda}}, U)$ {see Section 4.4}
until $\|\phi - \phi^-\| \leq \varepsilon$
return ϕ

Basic properties

Before describing the four stages of step i, we make a simple observation. Let $e = (0, 0, 1)$ be the unit vector pointing upward and $\boldsymbol{u} \in \Omega$. Consider the half-line starting at point $\boldsymbol{u}$ in direction e. Let $s_{\boldsymbol{u}}$ be the intersection point of the half-line with the surface S_ϕ and $n_{\boldsymbol{u}}$ the upward normal to S_ϕ at point $s_{\boldsymbol{u}}$. For the sake of simplicity and when clear from the context, we will use the same notation to denote $\nabla\phi(u) \in \mathbb{R}^2$ and $(\nabla\phi(u), 0) \in \mathbb{R}^3$. Given the pair $(\phi(\boldsymbol{u}), \nabla\phi(\boldsymbol{u}))$, we can deduce[7] the pair $(s_{\boldsymbol{u}}, n_{\boldsymbol{u}})$ using $s_{\boldsymbol{u}} = (\boldsymbol{u}, \phi(\boldsymbol{u}))$ and

$$n_{\boldsymbol{u}} = \frac{e - \nabla\phi(\boldsymbol{u})}{\sqrt{1 + \|\nabla\phi(\boldsymbol{u})\|^2}}. \tag{5}$$

Conversely, given the pair $(s_{\boldsymbol{u}}, n_{\boldsymbol{u}})$, we can recover the pair $(\phi(\boldsymbol{u}), \nabla\phi(\boldsymbol{u}))$ using $\phi(\boldsymbol{u}) = (s_{\boldsymbol{u}} - \boldsymbol{u}) \cdot e$ and

$$\nabla\phi(\boldsymbol{u}) = e - \frac{n_{\boldsymbol{u}}}{n_{\boldsymbol{u}} \cdot e}. \tag{6}$$

4.1. *Discretization of the source density*

Let $\phi^- = \phi_{i-1}$ be the function computed at step $i-1$. The goal of Stage 1 is to compute a discrete probability measure μ that approximates the restriction of the source density ρ_0 to $\pi_{\mathbb{S}}(S_{\phi^-})$. The support of μ is obtained by lifting points $\boldsymbol{u} \in U$ on the reflector S_{ϕ^-} and projecting the lifted points $(\boldsymbol{u}, \phi^-(\boldsymbol{u}))$ radially on the sphere $\mathbb{S}_0$. Introducting the map $F(\boldsymbol{u}) = \pi_{\mathbb{S}}(\boldsymbol{u}, \phi^-(\boldsymbol{u}))$, the support of μ is $F(U)$ and, by construction, a subset of $\pi_{\mathbb{S}}(S_{\phi^-})$. We then set

$$\mu \;=\; \sum_{x \in \mathrm{Supp}(\mu)} \mu_x \; \delta_x,$$

where the weight μ_x of point x is defined as follows:

$$\mu_x = \rho_0(F(\boldsymbol{u}))\; \frac{1}{\|s_u\|^3} \; \left| \frac{n_{\boldsymbol{u}} \cdot s_u}{n_{\boldsymbol{u}} \cdot e} \right|. \tag{7}$$

This choice for μ_x follows from a change of variable formula, and is explained in the Appendix. Finally, let us denote by DISCRETIZATION the function that takes as input (ρ_0, U, ϕ^-) and outputs the measure μ supported on $\mathrm{Supp}(\mu) := F(U)$ and the map $F : U \to \mathrm{Supp}(\mu)$. Note that the map F will only be used in Section 4.3.

4.2. *Parameterizing radially the reflector from the origin*

In Stage 2, we approximate the constant κ_ϕ defined in Equation (4) with

$$\kappa_\mu = \left(\sum_{x \in \mathrm{Supp}(\mu)} \mu_x \right) / \left(\sum_{y \in Y} \nu_y \right)$$

and solve the optimal transport problem between μ and $\kappa_\mu \nu$ as explained in Section 2. We get a 1-parameter family of reflector surfaces $\{R_{h\boldsymbol{\lambda}}\}_{h>0}$. Each surface $R_{h\boldsymbol{\lambda}}$ is made up of pieces of paraboloids with focal points at the origin and can be deduced from $R_{\boldsymbol{\lambda}}$ by applying a homothety with center the origin and ratio $h > 0$. We let FARFIELDREFLECTOR be the function that takes as input $(\mu, \kappa_\mu \nu)$ and outputs the set of focal distances $\boldsymbol{\lambda}$.

4.3. *Choosing the scale factor for focal distances*

Ideally, we would like to select a scale factor h such that points of the surface $R_{h\boldsymbol{\lambda}}$ hit by rays in directions $\mathrm{Supp}(\mu)$ lie above Ω. This is important to convert the radial solution $R_{h\boldsymbol{\lambda}}$ into the graph of some function $\phi : \Omega \to \mathbb{R}$. As we cannot be sure that such a scale factor exists, we use instead the following heuristic; see Figure 2. For $\boldsymbol{u} \in U$, let $r_{\boldsymbol{u}}$ be the point of $R_{\boldsymbol{\lambda}}$ hit by a light ray from the source in the direction $F(\boldsymbol{u})$ and let $\boldsymbol{u'}$ be the orthogonal projection of $r_{\boldsymbol{u}}$ onto the plane H. We select the parameter h which minimizes the quantity $\sum_{\boldsymbol{u} \in U} \|\boldsymbol{u} - h\boldsymbol{u'}\|^2$. This is a second-degree polynomial in h, and the minimal value is attained for

$$h = \frac{\sum_{\boldsymbol{u} \in U} \boldsymbol{u} \cdot \boldsymbol{u'}}{\sum_{\boldsymbol{u} \in U} \|\boldsymbol{u'}\|^2}.$$

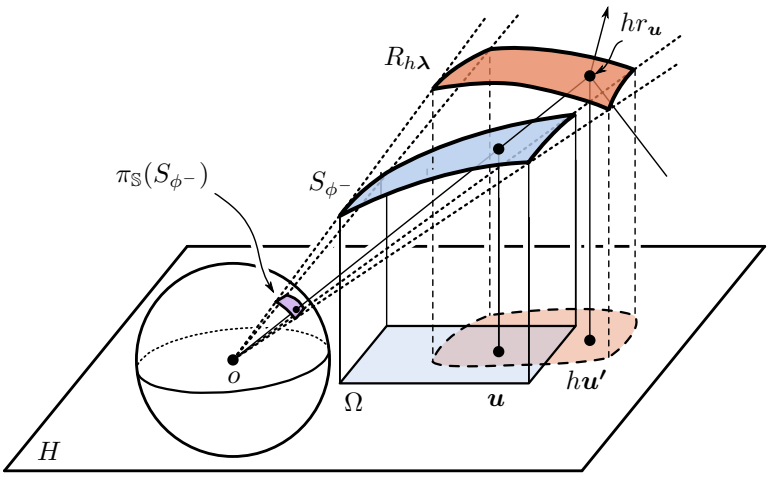

Fig. 2: When rays from the origin in directions $\pi_{\mathbb{S}}(S_{\phi^-})$ strike $R_{h\boldsymbol{\lambda}}$, they define a patch whose orthogonal projection onto the plane H does not necessarily coincides with Ω.

We let OptScaleFactor be the function that takes as input $(\boldsymbol{\lambda}, \mu, F, U)$ and returns the scale factor h.

4.4. *Parameterizing the reflector on Ω*

So far, the reflector surface $R = R_{h\boldsymbol{\lambda}}$ has been defined radially from the origin. In Stage 4, we want to express the radial surface R as the graph $\phi = \phi_i$ of a function on Ω, that is find $\phi : \Omega \to \mathbb{R}$ such that $S_\phi = R$. As explained before, it will be sufficient for our purposes to represent ϕ by an array that records for each point $\boldsymbol{u} \in U$ two pieces of information: the value of the function $\phi(\boldsymbol{u})$ and the value of its gradient $\nabla\phi(\boldsymbol{u})$. We use our observation to deduce this information from the pair $(s_{\boldsymbol{u}}, n_{\boldsymbol{u}})$, where $s_{\boldsymbol{u}}$ is the intersection of the half-line $[\boldsymbol{u}, e)$ with the reflector and $n_{\boldsymbol{u}}$ is the upward normal to the reflector at $s_{\boldsymbol{u}}$. We denote by Parametrization the function that takes as input the pair (R, U) and outputs the map ϕ. Note that this function is also used during initialization to convert the paraboloid $\partial P_{y_0}(\lambda_0)$ into the graph of a function ϕ_0.

4.5. *Reaching the fixed point*

When we reach a fixed point, that is, when $\phi = \phi^-$, the surface S_ϕ reflects the light rays emitted by the source with discrete probability measure μ towards infinity with discrete probability measure $\kappa_\mu\nu$. In other words, the surface S_ϕ is a solution for a discrete version of the reflector problem. The surface is encoded by the N set of pairs $\{(s_{\boldsymbol{u}}, n_{\boldsymbol{u}})\}_{\boldsymbol{u}\in U}$. By construction, the N points $s_{\boldsymbol{u}}$ are in convex position. More precisely, they all lie on the upper hull of $\{s_{\boldsymbol{u}}\}_{\boldsymbol{u}\in U}$, that is, on the part of the convex

hull of $\{s_{\boldsymbol{u}}\}_{\boldsymbol{u}\in U}$ that can be seen from $(0,0,+\infty)$. Notice that any surface passing through the points $s_{\boldsymbol{u}}$ and having outward normals $n_{\boldsymbol{u}}$ at $s_{\boldsymbol{u}}$ is also a solution for this discrete problem. Indeed, rays emitted by the source in directions $\mathrm{Supp}(\mu)$ hit the surface at points $s_{\boldsymbol{u}}$. For instance, the boundary of the intersection of half-spaces

$$H_{\boldsymbol{u}} = \{z \in \mathbb{R}^3 \mid (s_{\boldsymbol{u}} - z) \cdot n_{\boldsymbol{u}} \leq 0\}$$

is also a solution. Incidentally, we note that the algorithm terminates as soon as $\mu = \mu^-$, where μ^- denotes the discrete probability measure computed a step earlier.

4.6. *Adding more constraints on the shape of the reflector*

In this section, we describe briefly how to convert the solution $\phi = \phi_k$ into a function ψ that belongs to some finite-dimensional vector space $\mathcal{F}(\Omega)$. We suppose for that for every point $\boldsymbol{u} \in \Omega$, we are able to compute the two linear maps $\psi \in \mathcal{F}(\Omega) \mapsto \phi(\boldsymbol{u})$ and $\psi \mapsto \nabla\psi(\boldsymbol{u})$ in a certain basis. This is the case for tensor-product splines, and for other spaces of polynomial functions. Snell's law of reflection indicates that if we want the two surfaces S_ψ and S_ϕ to reflect lights in a similar way, then they must be close both in positions and normals. We thus introduce two energy terms. The first one is on positions:

$$E_{pos}(\psi) = \sum_{\boldsymbol{u}\in U} |\phi(\boldsymbol{u}) - \psi(\boldsymbol{u})|^2.$$

For the second energy term, we replace the constraint that normals are close by the constraint that gradients are close so as to get a convex quadratic energy:

$$E_{tan}(\psi) = \sum_{\boldsymbol{u}\in U} \|\nabla\phi(\boldsymbol{u}) - \nabla\psi(\boldsymbol{u})\|^2.$$

To balance the two energy terms, we introduce a weight $\omega \in [0,1]$ and choose the function ψ that minimizes the energy $E_\omega(\psi) = (1-\omega)E_{pos}(\psi) + \omega E_{tan}(\psi)$. The energy $E_\omega(\psi)$ is convex quadratic in ψ and its global minimum can therefore be found by solving a linear system.

5. Experiments

5.1. *Experimental set-up.*

Light densities. We begin by introducing the light densities we use on the two spheres, the one at the origin and the one at infinity. Recalling that $e = (0,0,1)$, the source density ρ_0 is Lambertian and defined by $\rho_0(x) = \max\{x \cdot e, 0\}$. At infinity, we consider different discrete probability measures ν defined as follows. Let $g : \mathbb{R}^2 \times \{-1\} \to \mathbb{R}$ be a probability density defined on the horizontal plane passing through the south pole $(0,0,-1)$. The measure with density g is approximated by the uniform measure over a finite set Y_g with a prescribed number of points using the variant of the k-means algorithm introduced in De Goes et al..[6] Projecting the

finite set Y_g onto the unit sphere $\mathbb{S}_\infty$, we thus get a discrete probability measure on $\mathbb{S}_\infty$ which we denote ν. The densities g we will use in our experiments are mainly indicator functions of various shapes drawn in the plane $\mathbb{R}^2 \times \{-1\}$. These shapes will include a disk, a lozenge, an annulus and the letters P and S drawn with a bold font. The corresponding discrete probability measure ν will be denoted respectively $\nu_\bigcirc$, $\nu_\diamond$, $\nu_\circledcirc$, $\nu_{\mathbf{P}}$ and $\nu_{\mathbf{S}}$.

Shape of reflectors

To build our reflectors, we call Algorithm 1 on the input (ρ_0, ν). This produces a function $\phi : U \to \mathbb{R}$ whose graph S_ϕ has the desired optical properties (see Section 4.5). We then approximate the resulting reflector S_ϕ using one of the two following forms. In both cases, the approximating patch will be denoted S_ψ.

PL patches. In this case, the patch S_ψ is constructed as the upper part of the convex hull of the N points $s_{\boldsymbol{u}} = (\boldsymbol{u}, \phi(\boldsymbol{u}))$. Since these points are in convex position, this amounts to interpolating linearly the values of $\phi : U \to \mathbb{R}$ on a triangulation of U adapted to this convex hull.

Bézier patches. In this case, we consider for the finite-dimensional vector space $\mathcal{F}(\Omega)$ the set of Bézier functions:

$$\psi_\alpha(\boldsymbol{u}) := \sum_{0 \le i \le d} \sum_{0 \le j \le d} \alpha_{i,j} B_i^d(u_1) B_j^d(u_2)$$

of degree $d \times d$ over Ω. The control points of the associated Bézier patch are the points of coordinates $(i/d, j/d, \alpha_{i,j})$. We then approximate the reflector surface with the Bézier patch S_ψ whose parametrization ψ minimizes the energy E_ω over $\mathcal{F}(\Omega)$ (see Section 4.6). In all our experiments we set $\omega = 1/2$, except for the third column of Figure 6. Remark that even though the points $\{s_{\boldsymbol{u}}\}_{\boldsymbol{u} \in U}$ are in convex position, there is no guarantee for the Bézier function ψ_α to be convex.[11] However, the simple approximation proposed in Section 4.6 seems to suffice in practice to get a convex solution.

5.2. *Results*

We ran our program on several examples. The resolution of the linear programming problem appearing in the optimal transport stage is the most time consuming step. Indeed, the complexity of the transportation problem is at least $\Omega(N^2)$, where N is the number of points. To give an idea, on a regular laptop, it takes about one minute for 4 iterations in the algorithm with $N = 2500$.

Visualization of the reflected intensity. We include two types of figures to visualize the optical properties of the resulting surface S_ψ. The most obvious is the **reflected intensity** that this patch reflects at infinity, and which is obtained by raytracing. The raytraced images are generated using the unbiased renderer LuxRender.

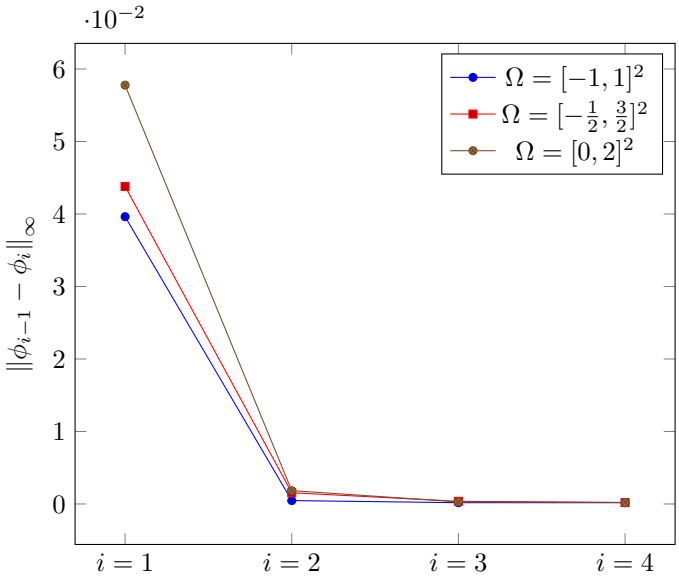

Fig. 3: Evaluation of the fixed point property. The target density is the lozenge, sampled with 400 points and the domains Ω are unit squares centered at three different locations. The maximum amount of change between two iterations reaches 10^{-3} in at most 4 iterations.

Our second type of visualization allows to better apprehend the deformations induced by the reflector surface. We start from a regular mesh of the rectangle domain Ω, obtained by decomposing the rectangle into squares and each square into two triangles. Each point of the mesh is first lifted to the surface S_ψ by ψ, then projected on $\mathbb{S}_0$ by $\pi_{\mathbb{S}}$, and finally reflected to $\mathbb{S}_\infty$ by T_{S_ψ}. Keeping the connectivity of the original mesh, this defines the **reflected mesh** on $\mathbb{S}_\infty$.

Fitting of normals When building a Bézier patch, the parameter ω introduced in the last step of the algorithm (Section 4.6) allows to interpolate between a fitting based only on the position of the reflector S_ϕ constructed by the optimal transport phase ($\omega = 0$) and a fitting based only on its normals ($\omega = 1$). The choice of this parameter has a very mild impact on the optical quality of the reconstructed surface, as long as it remains in the open interval $(0, 1)$. Setting $\omega = 0$, however, corresponds to completely disregarding the normals of the surface S_ϕ, and can degrade the optical properties of the reconstructed surface. The second and third columns of Figure 6 illustrate the difference between $\omega = 1/2$ (left) and $\omega = 0$ (right) when the target density is the disk.

Evaluation of the fixed point property. In this experiment, the target on $\mathbb{S}_\infty$ is a uniform sampling of a lozenge with 400 points. The domain Ω_i ($0 \leq i \leq 2$) is the square $[-1, 1]^2$ translated by the vector $(i/2, i/2)$. Figure 3 displays the maximum

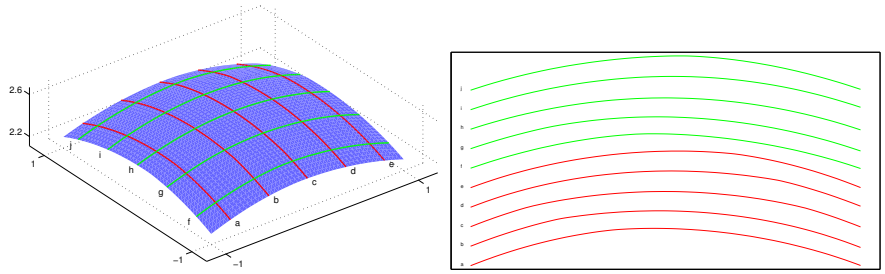

Fig. 4: Convexity of the constructed reflectors: (Top) PL patch obtained for the target density $\nu_{\mathbf{S}}$ supported by shape S, with 2500 points in the optimal transport stage. (Bottom) Convex curves extracted from the reflector surface.

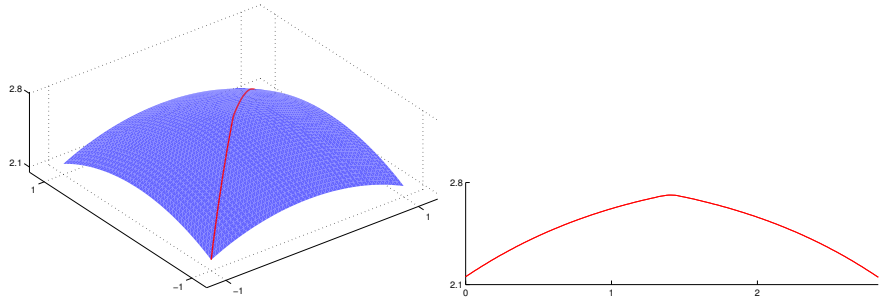

Fig. 5: Non-smooth reflector. The target probability measure $\nu_{\circledcirc}$ is supported by the ring. The number of points in the optimal transport stage is 2500.

change in ϕ between an iteration of the algorithm and the next one. One can observe that a fixed point is attained slightly faster for the centered domain Ω_0 than for the domains Ω_1 and Ω_2. In all three cases, the maximum change in ϕ reaches $\varepsilon := 10^{-3}$ in less than four iterations. This has to be compared to the value of ϕ. To give an idea, the order of magnitude of the initial function ϕ_0 is approximatively 2.5.

Smoothness of the solution. Depending on the target, the exact solution to the reflector problem may be non-smooth. As an indication for this, consider a cone with apex above the origin and directed towards the origin. The cone reflects the light from the origin in directions that lie within an annulus on $\mathbb{S}_\infty$. We thus expect the inverse problem to produce non-smooth solutions when targets are non-convex. And indeed, when the target measure is $\nu = \nu_{\circledcirc}$, our algorithm produces a non-smooth reflector which resembles a cone; see Figure 4. However, the approximation with a Bézier patch has the consequence of smoothing the result and degrading the optical properties. This is visible on Figure 7 where we can see that for small degrees, our reflector starts illuminating the hole inside the annulus. We can note that as the degree increases, the size of the triangles of the reflected mesh inside the

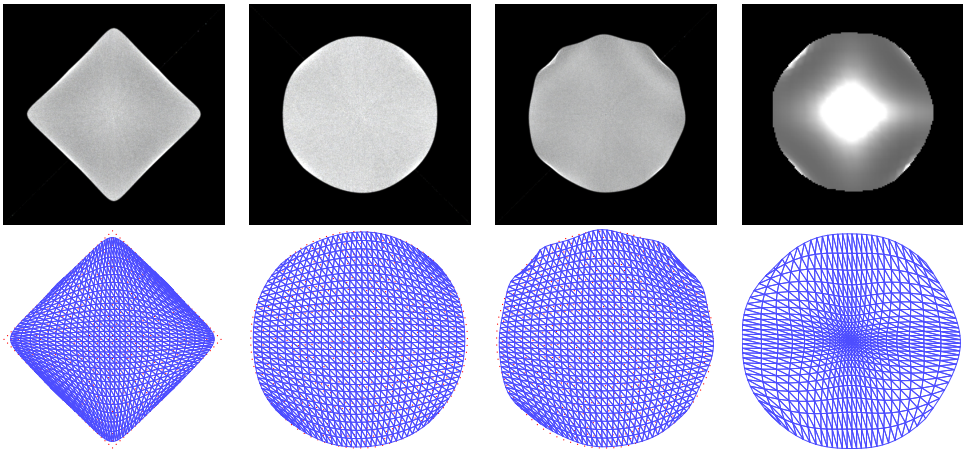

Fig. 6: Reflected intensities (Top) and meshes (Bottom) for three target densities: $\nu_\diamond$, $\nu_\circ$ and a non-uniform density on the disk. In these four cases, the reflector is a Bézier patch of degree 7×7, and the number of points in the optimal transport stage is 900 for the first three columns and 400 for the last one. The difference between columns two and three is the value of ω in the fitting step (§4.6): on the left, $\omega = 0.5$ and on the right $\omega = 0$.

annulus increases, which means that this part receives less and less intensity. This can be explained by the fact that the Bézier approximation is getting closer to the exact solution as the degree increases.

Convexity of the solution. By construction, Algorithm 1 produces reflectors S_ϕ which are convex (precisely, which are the graphs of convex functions ϕ). This property is preserved when approximating S_ϕ with PL patches S_ψ. This can be observed in Figure 5, where we plot the curves obtained by intersecting the reflector with vertical planes: all of intersection curves are convex.

Caustics. Some caustics patterns are observed, especially in Figure 8. This phenomenon seems to be related to the lack of regularity of the reflector surface. In Figure 8, the reflector surface is a triangulation and the reflected light seems to concentrate along some particular lines. This phenomenon does not occur when the reflector surface is a Bezier patch of low degree, and is therefore very smooth, see Figure 7.

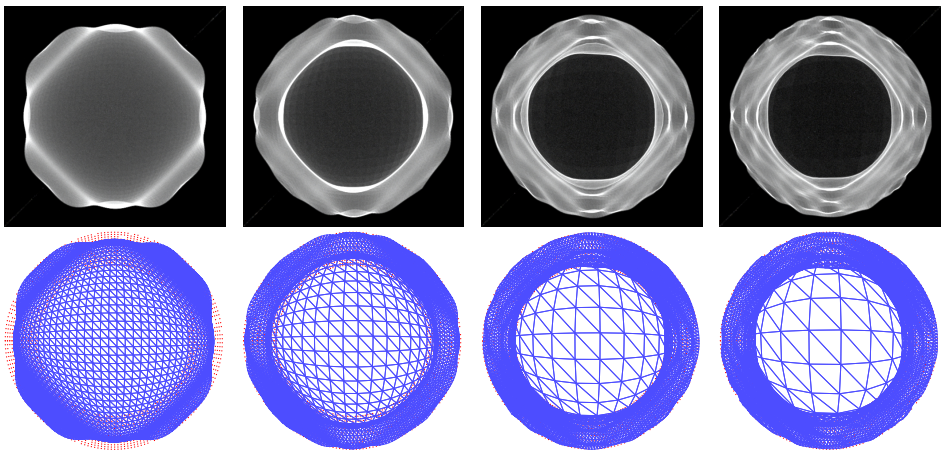

Fig. 7: Reflected intensities (Top) and meshes (Bottom) for a non-convex target, here the annulus. The exact solutions to the reflector problem are nonsmooth. From left to right, the degree of the Bézier surfaces used to define the reflector are 5×5, 9×9, 15×15 and 19×19. The red dots correspond to the support of $\nu_{\circledcirc}$. The number of points of the optimal transport stage is 2500. One can observe that when the degree is too low, the constructed surface tends to fill in the hole of the annulus.

6. Discussion

We close with a few remarks and open questions. We believe that the patch reflector problem would deserve a more detailed theoretical study, for probability measures with density. Under what conditions can we guarantee that a solution exists? We proposed a fixed point heuristic to find a solution and our heuristic seems to work well in practice: this fixed point approach could be a starting point for the theoretical analysis. If a solution exists, is it unique, and convex as in the far field reflector problem? Another issue concerns the choice of a fitting method for producing the final surface. We wonder whether more sophisticated methods would allow to avoid the formation of spurious caustic patterns. For instance, one could think of minimizing a higher order energy over a space of convexity-preserving interpolating spline surfaces. [14,12]

Appendix: Proofs

Proof of Lemma 1

Let x be a point of X. Taking the logarithm of a function does not change where the minimum is attained, and we can rewrite Equation (1) as

$$T_{\boldsymbol{\lambda}}(x) = \arg\min_{y \in Y} c(x, y) + \ln(\lambda_y).$$

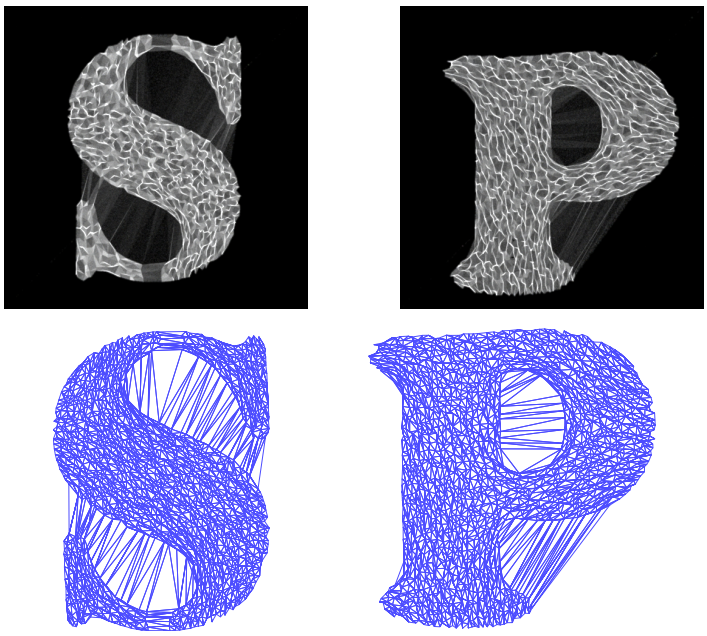

Fig. 8: Example of more difficult target densities supported on non-convex sets. In the two cases, the reflector is a triangulated surface and the number of points of the optimal transport stage is 2500.

Let $y_0 \in T_{\boldsymbol{\lambda}}(x)$. By Equation (FF), and since $y \mapsto c(x,y) + \ln(\lambda_y)$ is constant over $T_{\boldsymbol{\lambda}}(x)$, one has

$$\begin{aligned}
&\sum_{y\in Y}(c(x,y)+\ln(\lambda_y))\pi^*_{x,y}\\
&\quad=\sum_{y\in T_{\boldsymbol{\lambda}}(x)}(c(x,y)+\ln(\lambda_y))\pi^*_{x,y}\\
&\quad=(c(x,y_0)+\ln(\lambda_{y_0}))\sum_{y\in T_{\boldsymbol{\lambda}}(x)}\pi^*_{x,y}\\
&\quad=(c(x,y_0)+\ln(\lambda_{y_0}))\mu_x.
\end{aligned}$$

Let now $\pi \in \Gamma(\mu,\nu)$ be any transport plan. Then

$$\begin{aligned}
&(c(x,y_0)+\ln(\lambda_{y_0}))\mu_x\\
&=\sum_{y\in Y}(c(x,y_0)+\ln(\lambda_{y_0}))\pi_{x,y}\\
&\leq\sum_{y\in Y}(c(x,y)+\ln(\lambda_y))\pi_{x,y}
\end{aligned}$$

We finally deduce that

$$\sum_{y\in Y}(c(x,y)+\ln(\lambda_y))\pi^*_{x,y} \leq \sum_{y\in Y}(c(x,y)+\ln(\lambda_y))\pi_{x,y}.$$

Now, since the two transport plans π and π^* belong to $\Gamma(\mu,\nu)$, one has

$$\begin{aligned}\sum_{x\in X}\sum_{y\in Y}\ln(\lambda_y)\pi^*_{x,y} &= \sum_{y\in Y}\ln(\lambda_y)\nu_y \\ &= \sum_{x\in X}\sum_{y\in Y}\ln(\lambda_y)\pi_{x,y},\end{aligned}$$

which allows us to conclude.

Formal justification of Equation (7)**.**

We assume that the map F is a smooth diffeomorphism from the two-dimensional domain Ω to a domain of the unit sphere $\mathbb{S}_0$. Recal that the jacobian $\mathrm{J}F(\boldsymbol{u})$ of F at a point $\boldsymbol{u}$ measures the distortion of area induced by the map F near $\boldsymbol{u}$, that is

$$\mathrm{J}F(\boldsymbol{u}) = \lim_{\varepsilon\to 0}\frac{\mathrm{area}(F(\mathrm{B}(\boldsymbol{u},\varepsilon)))}{\mathrm{area}(\mathrm{B}(\boldsymbol{u},\varepsilon))},$$

where $\mathrm{B}(\boldsymbol{u},\varepsilon)$ denotes the disk of radius ε centered at $\boldsymbol{u}$.

We also assume that the point set U is uniformly distributed on Ω in the following sense: there exists a partition of Ω into domains $(\mathcal{V}_u)_{u\in U}$ of equal area, such that $\mathcal{V}_u$ is contained in the disk $\mathrm{B}(\boldsymbol{u},\varepsilon)$. Then, the domains of the sphere $(F(\mathcal{V}_u))_{u\in U}$ also form a partition of $\pi_{\mathbb{S}}(S_{\phi^-})$. Therefore for $x=F(u)$, it is natural to take for μ_x the area of $F(\mathcal{V}_u)$, which equals, by the change of variable formula

$$\begin{aligned}\mathrm{area}(F(\mathcal{V}_u)) &= \int_{\mathcal{V}_u}\rho_0(F(\boldsymbol{w}))\ \mathrm{J}F(\boldsymbol{w})d\boldsymbol{w} \\ &= \mathrm{area}(\mathcal{V}_u)\rho_0(F(\boldsymbol{u}))\mathrm{J}F(\boldsymbol{u}) + o(\varepsilon). \\ &\simeq \frac{\mathrm{area}(\Omega)}{|U|}\rho_0(F(\boldsymbol{u}))\mathrm{J}F(\boldsymbol{u})\end{aligned}$$

Combining this equation with the Lemma below gives us Equation (7) up to a common multiplicative constant independent of $\boldsymbol{u}$.

Lemma 2. *The Jacobian of the map $F:\boldsymbol{u}\mapsto\pi_{\mathbb{S}}(\boldsymbol{u},\phi(\boldsymbol{u}))$ is given by the formula*

$$\mathrm{J}F(\boldsymbol{u}) = \frac{1}{\|s_{\boldsymbol{u}}\|^3}\left|\frac{n_{\boldsymbol{u}}\cdot s_{\boldsymbol{u}}}{n_{\boldsymbol{u}}\cdot e}\right|.$$

Proof. We put $f(\boldsymbol{u}) := s_{\boldsymbol{u}}$. By differentiating the map $F=\pi_{\mathbb{S}}\circ f$ and taking the jacobian, one has

$$\mathrm{J}F(\boldsymbol{u}) = |\det(\mathrm{D}\pi_{\mathbb{S}}(s_{\boldsymbol{u}})_{|n_{\boldsymbol{u}}^{\perp}})|\times \mathrm{J}f(\boldsymbol{u}),$$

where $\mathrm{D}\pi_{\mathbb{S}}(s_{\boldsymbol{u}})_{|n_{\boldsymbol{u}}^{\perp}}$ is the restriction of $\mathrm{D}\pi_{\mathbb{S}}(s_{\boldsymbol{u}})$ to the plane orthogonal to $n_{\boldsymbol{u}}$. Now, since the restriction of the projection $\pi_{\mathbb{S}}$ to the plane orthogonal to $n_{\boldsymbol{u}}$ is the

composition of an orthogonal projection onto the plane orthogonal to $s_{\boldsymbol{u}}$ and an homothety of scale $1/\|s_{\boldsymbol{u}}\|^2$, one has

$$\det(\mathrm{D}\pi_{\mathbb{S}}(s_{\boldsymbol{u}})_{|n_{\boldsymbol{u}}^{\perp}}) = \frac{1}{\|s_{\boldsymbol{u}}\|^2} \left(n_{\boldsymbol{u}} \cdot \frac{s_{\boldsymbol{u}}}{\|s_{\boldsymbol{u}}\|} \right).$$

Since $f(\boldsymbol{u}) = (\boldsymbol{u}, \phi(\boldsymbol{u}))$, one has $\mathrm{J}f(\boldsymbol{u}) = \sqrt{1 + \|\nabla\phi(\boldsymbol{u})\|^2}$. We deduce from Equation (6) that

$$\mathrm{J}f(\boldsymbol{u}) = \sqrt{1 + \left(e - \frac{n_{\boldsymbol{u}}}{n_{\boldsymbol{u}} \cdot e} \right)^2} = \frac{1}{|n_{\boldsymbol{u}} \cdot e|},$$

thus concluding the proof. □

References

1. Rainer E Burkard, Mauro Dell'Amico, and Silvano Martello. *Assignment problems*. Siam, 2009.
2. L. A. Caffarelli and V. Oliker. Weak solutions of one inverse problem in geometric optics. *Journal of Mathematical Sciences*, 154(1):39–49, 2008.
3. M.A. Campos, M. Cejnek, and M. Olivik. Automotive tail lamp with large rake angle, September 21 1999. US Patent 5,954,427.
4. F. Cork, D.F. Bettridge, and P.C. Clarke. Method of and mixture for aluminizing a metal surface, February 22 1977. US Patent 4,009,146.
5. Pedro Machado Manhães De Castro, Quentin Mérigot, and Boris Thibert. Intersection of paraboloids and application to Minkowski-type problems. *arXiv preprint arXiv:1403.0062*, 2014.
6. Fernando de Goes, Katherine Breeden, Victor Ostromoukhov, and Mathieu Desbrun. Blue noise through optimal transport. *ACM Transactions on Graphics*, 31(6):171, 2012.
7. Manfredo Perdigao Do Carmo and Manfredo Perdigao Do Carmo. *Differential geometry of curves and surfaces*, volume 2. Prentice-Hall Englewood Cliffs, 1976.
8. Manuel Finckh, Holger Dammertz, and Hendrik PA Lensch. Geometry construction from caustic images. In *Computer Vision–ECCV 2010*, pages 464–477. Springer, 2010.
9. T Glimm and V Oliker. Optical design of single reflector systems and the monge–kantorovich mass transfer problem. *Journal of Mathematical Sciences*, 117(3):4096–4108, 2003.
10. Pengfei Guan and Xu-Jia Wang. On a monge-ampere equation arising in geometric optics. *J. Differential Geom*, 48(2):205–223, 1998.
11. B. Jüttler. Surface fitting using convex tensor-product splines. *J. Comput. Appl. Math.*, 84(1):23–44, 1997.
12. Bert Jüttler. Surface fitting using convex tensor-product splines. *Journal of computational and applied mathematics*, 84(1):23–44, 1997.
13. Sergey A Kochengin and Vladimir I Oliker. Computational algorithms for constructing reflectors. *Computing and Visualization in Science*, 6(1):15–21, 2003.
14. Ming-Jun Lai. Convex preserving scattered data interpolation using bivariate¡ i¿ c¡/i¿¡ sup¿ 1¡/sup¿ cubic splines. *Journal of computational and applied mathematics*, 119(1):249–258, 2000.
15. Gustavo Patow and Xavier Pueyo. A survey of inverse surface design from light transport behavior specification. In *Computer Graphics Forum*, volume 24, pages 773–789. Wiley Online Library, 2005.

16. Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. High-contrast computational caustic design. *ACM Transactions on Graphics (TOG)*, 33(4):74, 2014.
17. C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2003.
18. Xu-Jia Wang. On the design of a reflector antenna. *Inverse problems*, 12(3):351, 1996.
19. Xu-Jia Wang. On the design of a reflector antenna II. *Calculus of Variations and Partial Differential Equations*, 20(3):329–341, 2004.