

THE VISIBLE-VOLUME FUNCTION OF A SET OF CAMERAS IS CONTINUOUS, PIECEWISE RATIONAL, LOCALLY LIPSCHITZ, AND SEMI-ALGEBRAIC IN ALL DIMENSIONS

JÖRG RAMBAU

ABSTRACT. The visible-volume function assigns to a configuration of cameras and a flexible environment in a convex room the volume that can be supervised by the cameras. It is of interest to configure the cameras and the environment in such a way that the visible volume is maximized. Some methods of global optimization can take profit from desirable analytic properties of the visible-volume function. Earlier work has only considered this function in dimensions two and three or for static environments implicitly defined by level-set functions. In this paper it is shown, that the visible-volume function for a flexible environment modeled explicitly by a parametrized simplicial complex is continuous, piecewise rational, locally Lipschitz, and semi-algebraic in all dimensions. visible volume and discrete geometry and semi-algebraic sets and hyperplane arrangements and point configurations and oriented matroids, triangulations and piecewise rational and locally Lipschitz

1. INTRODUCTION

This paper reveals a number of analytic properties of the *visible-volume function* (formally defined below). For an adjustable camera system in a room with movable obstacles, this function assigns to each parameter combination the volume that can be supervised by the cameras. A natural goal is to maximize the visible-volume function over all adjustable parameters. This paper restricts attention to the analytic properties of the visible-volume function in an environment modeled by a simplicial complex with an adjustable embedding in some bounded convex room. The analytic properties are of interest, since they are important for the convergence guarantees for various numerical global optimization approaches.

The famous *art-gallery problem (AGP)* is closely related to the maximization of the visible volume. In the AGP one seeks for the minimal number of cameras with 360° view and their positions so that they together see the complete volume inside a simple two-dimensional polygon with n vertices. The famous simple proof in [10] of an earlier result in [6] shows – via a 3-coloring of a triangulation – that $\lfloor \frac{n}{3} \rfloor$ cameras suffice. However, the recent result in [1] shows that the AGP is as difficult as to decide whether an outer-description of a semi-algebraic set contains a real point. This implies that the analysis in this paper can probably not be simplified.

Optimizing coverage induced by camera placement has been investigated using various modeling techniques for a very long time, see [14] for a survey and [13] for a more recent thesis. The geometric computation of *shadows* on a surface (in the sense of day-to-day language) produced in three dimensions by a light at a certain position and some occluding 3D-objects dates back to at least 1977 [7]. That paper was motivated by the correct graphical

Date: October 10, 2022.

Preprint, accepted for publication in Discrete & Computational Geometry.

display of the illumination of architectural designs and similar issues in computer graphics. I mention this since the geometric description of a shadow of a simplex is a core component of this paper. In most applications today, the actual function to be optimized is usually called the *relevance function* and need not coincide with the visible volume. For example, points can have a weight according to their importance or distance to the camera. Still, common to most approaches is that the dependence of the visible volume on the camera system parameters must be captured in one way or another. This paper focuses on the visible-volume function and suggests a framework for its analysis in arbitrary dimensions.

In many practical instances like the one in [13], the visible-volume function must be considered as a black-box function, i.e., a closed-form formula is not available, the computation of a single function value is costly, and (sub)gradients – if they exist – cannot simply be calculated from formulas. For the global optimization of black-box functions, various methods have been proposed, see [4] for a more recent one. Although it is possible to apply such methods and hope for good results even if a convergence proof is not available, I find it mathematically desirable to know *a-priori* analytic properties of the function to be optimized that guarantee convergence. The criterion in [4] is, e.g., that the function is continuous, smooth on a dense subset of the domain, and locally Lipschitz or semi-algebraic.

In [12], the visible-volume function is analyzed for the case of an environment modeled by non-convex polygons in dimension two. In that paper, the independent variable of the visible-volume function is just the position of the camera, since it is assumed that it has a 360-degree view. The result is that the visible-volume function is continuous as well as smooth and locally Lipschitz in the camera position on a dense subset of the domain. Newer developments also consider environments in dimension three containing obstacles and more camera parameters like field of view and orientation of multiple cameras [13]. The visible-volume function is again continuous and smooth almost everywhere.

The only work I have found for general dimensions uses a different set-up based on the *level-set method* [15]: Inside a compact domain there are opaque solid obstacles implicitly given via a Lipschitz-continuous *level-set function*, a function that is negative in the interior and positive on the exterior of the obstacles [5]. A possible level-set function is the *signed-distance function* that assigns to each point the shortest distance to the boundary of some obstacle; it receives a negative sign if the point is inside the obstacle and a positive sign if it is outside. In that paper, the independent variables of the visible-volume function are again restricted to the positions of one or more cameras, i.e., the cameras possess omni-directional vision. Visibility of one point from another can be decided by checking whether the infimum of the level-set function on the connecting straight line (the *visibility function*) is positive. The visible-volume function then takes the form of an integral of the Heaviside function of the visibility function over the domain. It is shown that, under certain additional conditions on the visibility function (which the authors call “no fattening [of the visibility level sets]” [5, Page 5]), that the visible-volume function inherits Lipschitz-continuity from the level-set function. One may see a problem in the fact that the additional condition in that paper is a condition on the visibility function, which is not part of the input data.

Why bother at all about higher dimensions or a non-static environment? First, in dimension four the visible volume covered over time can be used to model video surveillance with time being the fourth dimension. Concerning the non-static environment the idea is that natural, usually adjustable camera parameters like field of view or direction can be modeled by obstacles inside the camera system (aperture blades). However, for this paper the true motivation was clearly curiosity: Is the dimension or a static environment relevant at all for some interesting analytic properties of the visible-volume function to hold, or is there a

general mathematical principle that guarantees an interesting result in all dimensions by a unified argument? What is the most explicit representation of the visible-volume function? Does it belong to a special class of functions?

How about generalizing the techniques from [12, 13]? In [12] so-called *inflection segments* are the basic building blocks of the analysis. Inflection segments are used to partition the interior of the non-convex polygon into regions in which the visible-volume function depends smoothly on the camera position. Generalizing their approach to higher dimensions would require to consider inflection hypersurfaces generating possibly non-convex polytopes as regions. This has in principle been carried out in a different language in [13, Tables 2.3–2.5] for dimension three by an extensive classification of so-called *vertex-incidence surfaces*. Following those lines, the number of cases that would have to be considered in a similar table would be enormous even in dimension four. For general dimension this seems intractable. To replicate the full result in this paper with the level-set method in [5] one would have to construct an explicit parametrized level-set function for movable obstacles that are described explicitly via parametrized embeddings of simplicial complexes. At least for me it is not obvious how to proceed along those lines. Yet, it would be interesting to see the result of this purely analytical approach.

This paper focuses on finite, discrete, explicit descriptions of the environment. It is shown by a unified method in a fairly general setting with obstacles, explicitly defined as parametrized pure simplicial complexes of codimension-one, that the visible-volume function, depending on the positions of the cameras and the environment is indeed continuous, smooth almost everywhere, locally Lipschitz and semi-algebraic in all dimensions. The methodology belongs to the wide field of discrete geometry – a combination of the basic theory of oriented matroids [16, 3], the basic theory of triangulations of point configurations [9], and the basic theory of semi-algebraic sets [2]. The actual result is even more specific: The domain of the visible-volume function can be stratified into open cubical cells so that the visible-volume function is a smooth rational function on each cell with degree at most cubical in the dimension. All structural results follow from this explicit representation, which goes beyond the results in [5]. Even in dimension two and three this is, to the best of my knowledge, the most concrete description of the visible-volume function.

An important remark is in place here: the analytic properties of the visible-volume function in the main result cannot be derived via splitting it into the composition of the “visible-subset function” (with the Hausdorff metric on the set of subsets) and the “volume-of-subsets function”. This is simply because the “visible-subset function” is not even continuous. See Figure 1 for visual evidence in dimension two. The same figure shows that if one considers an arbitrary positive amount of diffraction, then the visible-volume function itself becomes discontinuous: Indeed, while moving the camera position, in the first moment one can look through the obstacles at all, a cone will be visible with interior angle of at least the diffraction angle. Hence, in order to prove the main result straight-line geometric optics has to be assumed and exploited. Moreover, this shows that the results in this paper are in sharp contrast to the discontinuous dependence of image features on the view point, as is investigated in Computer Vision (see, e.g., [8]). Thus, the combination of methods in this paper is tailor-made specifically to the straight-line-optics visible-volume function.

2. PRELIMINARIES

This paper uses elementary notions and facts from the theories of semi-algebraic sets, oriented matroids, and triangulations. This section gives a brief overview over these tools in




FIGURE 1. The “visible-subset function” is not even continuous: an arbitrarily small left-movement of \mathbf{p} (which is located on the line from 7 to 9) to \mathbf{p}' makes points visible that are “far away” from the points visible from \mathbf{p} ; the increase of the visible-volume, however, will be arbitrarily small for a sufficiently small movement; the continuity of the visible volume is lost as soon as there is a constant-degree diffraction at corners

a phrasing adapted to the purposes of this paper. In order to distinguish the known from the novel, the terms “notion” and “fact” are used instead of “definition” and “theorem”.

Semi-algebraic sets (over the reals) arise as solution sets of polynomial equations and inequalities and enjoy a number of convenient closure and decomposition properties.

Notion 1 (Semi-algebraic sets and functions). A subset $S \subseteq \mathbb{R}^d$ is a *basic semi-algebraic set (over the reals)*, if there are finitely many polynomials

$$(1) \quad (p_i)_{i \in I}, (q_j)_{j \in J}: \mathbb{R}^d \rightarrow \mathbb{R}$$

with coefficients in \mathbb{R} so that

$$(2) \quad S = \{x \in \mathbb{R}^d : p_i(x) = 0 \forall i \in I, q_j(x) > 0 \forall j \in J\}.$$

A *semi-algebraic set (over \mathbb{R})* is a finite union of basic semi-algebraic sets [2, Section 2.3].

For semi-algebraic sets $S \subseteq \mathbb{R}^d$ and $T \subseteq \mathbb{R}^k$ a function $f: S \rightarrow T$ is a *semi-algebraic function* if its graph

$$(3) \quad \text{Graph}(f) := \{(x, y) \in S \times T : f(x) = y\}$$

is a semi-algebraic set [2, Section 2.5.2]. The class of semi-algebraic functions that are k -times continuously differentiable is denoted by S^k for $k = 0, 1, \dots, \infty$.

Fact 1 (Basic properties of semi-algebraic sets). Semi-algebraic sets are closed under finite union, finite intersection, and set complement [2, Section 2.3]. Moreover, projections of semi-algebraic are semi-algebraic [2, Section 2.4] as well as the closure and the interior [2, Proposition 3.1].

The composition of semi-algebraic functions is semi-algebraic. Moreover, the pointwise addition and multiplication are semi-algebraic [2, Propositions 2.102 and 2.103]. In particular, polynomials and rational functions defined on semi-algebraic sets are semi-algebraic.

The following, specialized version of the stratification theorem for semi-algebraic sets provides the basis for the piecewise description of the visible-volume function.

Fact 2 (Stratification theorem for semi-algebraic sets and functions). Every semi-algebraic set can be stratified into finitely many strata such that each stratum is S^∞ -homeomorphic

to an open cube, and the boundary of each stratum is a finite union of strata of lower dimensions [2, Theorem 5.38 (rephrased)].

Oriented matroids encode the combinatorics of finite geometric structures like finite point or vector configurations or finite arrangements of affine or linear hyperplanes. An oriented matroid may be defined in various ways all of which can be transformed into one another. In this paper, only *realizable oriented matroids* are used so that it is enough to define them via the geometric properties of point configurations or arrangements of hyperplanes (see [3] for a more general treatment). In the following, the *rank* is defined as $r := d + 1$.

Notion 2 (Oriented matroid of a hyperplane arrangement). Let $\mathcal{H} = (\mathbf{H}_i)_{i \in I}$ be a finite linear hyperplane arrangement in \mathbb{R}^r with at least one *basis*, i.e., a r -subset of hyperplanes with linearly independent normals labeled by a label set I (multiple hyperplanes are allowed, i.e., distinct labels label identical hyperplanes). The arrangement is *oriented* if for each hyperplane \mathbf{H} we have fixed the open positive halfspace \mathbf{H}^+ and the open negative halfspace \mathbf{H}^- . We denote by the zero-set \mathbf{H}^0 the hyperplane itself. A *signature* $\sigma : I \rightarrow \{-, 0, +\}$ is a function that assigns to each label a sign from $\{-, 0, +\}$, where “+” and “-” are interpreted as “+1” and “-1”. Each such signature induces a *region* \mathbf{R}_σ of the arrangement via

$$(4) \quad \mathbf{R}_\sigma := \bigcap_{i \in I} \mathbf{H}_i^{\sigma(i)},$$

which may be empty. A *covector* of \mathcal{H} is a signature σ whose region \mathbf{R}_σ is non-empty. A covector with an inclusion maximal number of zero-components is called a *cocircuit* of \mathcal{H} .

Two hyperplane arrangements with label set I are *combinatorially equivalent* if they have identical sets of covectors (or equivalently, identical sets of cocircuits). The combinatorial equivalence class of \mathcal{H} is called *the oriented matroid of \mathcal{H}* . Any choice of a *hyperplane at infinity* yields an *affine oriented matroid* where only the covectors are considered that are on the positive side of the hyperplane at infinity. The intersection of all hyperplanes not at infinity with the positive hemisphere defined by the hyperplane at infinity followed by a stereographic projection yields an affine hyperplane arrangement in dimension $d = r - 1$. For a matrix $W = (\mathbf{a}_i^\top | -b_i)_{i \in I} \in \mathbb{R}^{I \times r}$ the oriented affine hyperplane arrangement induced by $H_i^\sigma = \{\mathbf{x} \in \mathbb{R}^d : \text{sign}(\mathbf{a}_i^\top \mathbf{x} - b_i) = \sigma\}$ for all $\sigma \in \{-, 0, +\}$ and $i \in I$ is denoted by $\mathcal{H}(W)$. The oriented matroid of it is denoted by $\mathbf{M}_{\mathcal{H}}(W)$. [16, Section 7.1.3 (rephrased)].

Notion 3 (Oriented matroid of a point configuration). Let $\mathbf{A} = (\mathbf{v}_j)_{j \in J}$ be a finite point configuration in \mathbb{R}^d with at least one *basis*, i.e., a r -subset of points forming a d -dimensional simplex, labeled by a label set J (multiple points are allowed, i.e., distinct labels can label identical points). Each signature $\sigma : J \rightarrow \{-, 0, +\}$ on J defines a set of hyperplanes \mathcal{H}_σ consisting of all affine hyperplanes \mathbf{H} in \mathbb{R}^d with $\mathbf{v}_j \in \mathbf{H}^{\sigma(j)}$ for all $j \in J$. A *covector* of \mathbf{A} is a signature σ for which \mathcal{H}_σ is non-empty. A covector with an inclusion-maximal number of zero components is called a *cocircuit* of \mathbf{A} .

Two point configurations with label set J are *combinatorially equivalent* if they have identical sets of covectors (or equivalently, identical sets of cocircuits). The combinatorial equivalence class of \mathbf{A} is called *the oriented matroid of \mathbf{A}* . For a matrix $V = \begin{pmatrix} \mathbf{v}_j \\ 1 \end{pmatrix}_{j \in J} \in \mathbb{R}^{r \times J}$ the point configuration with points $(\mathbf{v}_j)_{j \in J}$, is denoted by $\mathbf{A}(V)$. The oriented matroid of it is denoted by $\mathbf{M}_{\mathbf{A}}(V)$ [16, Section 7.1.2 (rephrased)].

Remark 1. In this paper, oriented matroids are identified with their *realization spaces* in the space of matrices, which is no problem for *realizable* oriented matroids.

The following is immediate from the definition:

Fact 3 (Finiteness property of oriented matroids). For finite label sets I and J there exists only a finite number of distinct oriented matroids of hyperplane arrangements and distinct oriented matroids of point configurations labeled by I and J , respectively.

In the theory of oriented matroids, it does not make a difference whether one considers the oriented matroid of a hyperplane arrangement or the oriented matroid of a point configuration. Moreover, in addition to covectors and cocircuits, there are other structures that specify an oriented matroid. For example, prescribing the set of covectors of a hyperplane arrangement or a point configuration is equivalent to prescribing all determinants of r -many rows of W or r -many columns of V (the *chirotope* of the oriented matroid) [16, Section 7.2.3]. This imposes polynomial equations and inequalities of degree r among the entries of the matrices. Therefore, each set of matrices inducing mutually combinatorially equivalent affine hyperplane arrangements or point configurations, respectively, is a semi-algebraic set [16, Section 7.3.2].

Fact 4 (Realization spaces of oriented matroids are semi-algebraic). Each combinatorial equivalence class in $\mathbb{R}^{I \times r}$ of all matrices W with mutually combinatorially equivalent hyperplane arrangements $\mathcal{H}(W)$ is a semi-algebraic set. Each combinatorial equivalence class in $\mathbb{R}^{r \times J}$ of all matrices V with mutually combinatorially equivalent point configurations $\mathbf{A}(V)$ is a semi-algebraic set.

Triangulations of finite point configurations are dissections of the convex hull of the point configuration into simplices all of whose vertices are contained in the point configuration so that any two simplices intersect properly. More formally:

Notion 4 (Triangulations of point configurations). Let \mathbf{A} be a finite point configuration in \mathbb{R}^d of affine dimension d labeled by a label set J (multiple points are allowed, i.e., distinct labels can label identical points). For an r -subset S of J the subconfiguration \mathbf{S} is the subconfiguration of points in \mathbf{A} with labels in S . If \mathbf{S} has affine dimension d , it is *spanning*. In this case, S is called an (abstract) d -simplex, and \mathbf{S} is called a (geometric) d -simplex. Two d -simplices (abstract or geometric) *intersect properly in \mathbf{A}* if the relative interiors of any two geometric faces of them is empty. A *triangulation of \mathbf{A}* is a family \mathcal{T} of properly intersecting abstract d -simplices in J so that the union of all corresponding geometric simplices equals $\text{conv } \mathbf{A}$ [9, Definition 2.3.1 (rephrased)].

The main connection of triangulations to oriented matroids is the following: whether or not a family of r -subsets of J forms a triangulation of \mathbf{A} only depends on the oriented matroid of \mathbf{A} , not on the exact coordinates.

Fact 5 (Triangulations depend only on oriented matroids). Combinatorially equivalent point configurations have identical sets of triangulations [9, Corollary 4.1.44 (specialized to triangulations)]. In particular, the set of all point configurations on the same label set having identical sets of triangulations is, as the union of (realization spaces of) oriented matroids of point configurations, a semi-algebraic set.

3. PROBLEM STATEMENT AND MAIN RESULT

In this section, the main objects of study are introduced, the notation is fixed, and the main problem is formalized.

Let d be the ambient dimension, and let $\mathbf{P} \subset \mathbb{R}^d$ be a d -polytope, which is called the *room* (see Figure 2 left).

Let J be a set, and let \mathcal{F} be an abstract simplicial complex on J of pure codimension one, the *abstract environment complex*. Moreover, let $\mathbf{e}: J \rightarrow \mathbf{P}$ be an embedding




FIGURE 2. The room and an intersection-free embedding of an environment in the room

of the vertices of \mathcal{F} into the room. Given such an embedding \mathbf{e} , we call the point configuration $\mathbf{A}(\mathbf{e}) := \mathbf{e}(J)$ the *vertices of the geometric environment complex*. For a $k-1$ -simplex $F = \{j_1, j_2, \dots, j_k\} \in \mathcal{F}$ in the abstract environment complex we obtain its *geometric embedding* $\mathbf{F}(\mathbf{e}) := \text{conv}\{\mathbf{e}(j_1), \mathbf{e}(j_2), \dots, \mathbf{e}(j_k)\}$. The *geometric environment complex* in the room is then given by $\mathcal{F}(\mathbf{e}) := \{\mathbf{F}(\mathbf{e}) : F \in \mathcal{F}\}$. Its geometric realization $|\mathcal{F}(\mathbf{e})| := \bigcup_{F \in \mathcal{F}} \mathbf{F}(\mathbf{e})$ is the union of all geometric embeddings of its simplices. We call an embedding \mathbf{e} *intersection-free* if any two geometric embeddings of open simplices of any dimension have empty intersection (see Figure 2 right for an example). In particular, for intersection-free embeddings all geometric simplices $\mathbf{F}(\mathbf{e})$ have dimension $\text{card}(F) - 1$, where $\text{card}(F)$ is the cardinality of the abstract simplex F .

Let $\mathbf{E} \subset (\mathbf{P})^J$ be any semi-algebraic subset of all intersection-free embeddings, equipped with the standard subspace topology on $(\mathbb{R}^d)^J$.

Remark 2. The subset of intersection-free embeddings is a semi-algebraic set: Each intersection point \mathbf{x} of two open geometric simplices $\mathbf{F}(\mathbf{e})$ and $\mathbf{F}'(\mathbf{e})$ has two representations in terms of the strictly positive barycentric coordinates of the simplices. More specifically, the geometric simplices $\mathbf{F}(\mathbf{e}) = \text{conv}\{\mathbf{e}(j_1), \mathbf{e}(j_2), \dots, \mathbf{e}(j_k)\}$ and $\mathbf{F}'(\mathbf{e}) = \text{conv}\{\mathbf{e}(j'_1), \mathbf{e}(j'_2), \dots, \mathbf{e}(j'_\ell)\}$ have an interior intersection point if and only if the following holds:

$$(5) \quad \exists \lambda_1, \lambda_2, \dots, \lambda_k, \mu_1, \mu_2, \dots, \mu_\ell \in \mathbb{R} :$$

$$(6) \quad \lambda_1, \lambda_2, \dots, \lambda_k, \mu_1, \mu_2, \dots, \mu_\ell > 0,$$

$$(7) \quad \lambda_1, \lambda_2, \dots, \lambda_k, \mu_1, \mu_2, \dots, \mu_\ell < 1,$$

$$(8) \quad \lambda_1 + \lambda_2 + \dots + \lambda_k = 1,$$

$$(9) \quad \mu_1 + \mu_2 + \dots + \mu_\ell = 1,$$

$$(10) \quad \lambda_1 \mathbf{e}(j_1) + \mathbf{e}(j_2) + \dots + \lambda_k \mathbf{e}(j_k) = \mu_1 \mathbf{e}(j'_1) + \mathbf{e}(j'_2) + \dots + \mu_\ell \mathbf{e}(j'_\ell).$$

The set of all \mathbf{e} for which this condition is true for a given pair of geometric simplices is consequently the projection of a semi-algebraic set, which is a semi-algebraic set. Moreover, the set of all not intersection-free \mathbf{e} is the union of the semi-algebraic sets over all pairs of simplices, thus a semi-algebraic set. The set of all intersection-free embeddings \mathbf{e} is, as the complement of this semi-algebraic set, again a semi-algebraic set. Note that all further considerations do not depend on the fact that we only consider intersection-free embeddings. \square




FIGURE 3. A camera configuration formed by position and aperture blades in the room

Remark 3. We do not model the usual camera parameters “field of view” or “orientation” separately because they can be simulated by movable “aperture blades” in the geometric environment complex.

Definition 1 (Feasible Camera Configuration). The set

$$(11) \quad \mathbf{D} := \{(\mathbf{p}, \mathbf{e}) \in \mathbb{R}^d \times (\mathbb{R}^d)^J : \mathbf{e} \in \mathbf{E}, \mathbf{p} \in \mathbf{P} \setminus |\mathcal{F}(\mathbf{e})|\}$$

is the *set of feasible camera configurations*.

Remark 4. By Fact 1, \mathbf{D} is a semi-algebraic set.

The *affine hull* of a geometric simplex \mathbf{S} is the smallest affine subspace containing \mathbf{S} . For a geometric facet $\mathbf{F} \in \mathcal{F}(\mathbf{e})$, and a point $\mathbf{x} \in \mathbf{P} \setminus \text{aff } \mathbf{F}(\mathbf{e})$ the polyhedron $\text{cone}_{\mathbf{F}}(\mathbf{x})$ is defined as the *cone* and $\text{pyram}_{\mathbf{F}}(\mathbf{x})$ is defined as the *pyramid* over \mathbf{F} with apex \mathbf{x} . For a subset $\mathbf{X} \subset \mathbf{P}$, the notation $\overline{\mathbf{X}}$ is used for the topological closure and \mathbf{X}° for the topological interior.

Definition 2 (Shadow). For $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$ and $F \in \mathcal{F}$ so that let

$$(12) \quad \mathbf{S}_F(\mathbf{p}, \mathbf{e}) := \begin{cases} \overline{\text{cone}_{F(\mathbf{e})}(\mathbf{p}) \setminus \text{pyram}_{F(\mathbf{e})}(\mathbf{p})} & \text{if } \mathbf{p} \in \mathbf{P} \setminus \text{aff } F(\mathbf{e}), \\ \emptyset & \text{otherwise.} \end{cases}$$

be the *shadow of F with respect to the feasible camera configuration (\mathbf{p}, \mathbf{e})* .

The *total shadow of the environment with respect to the feasible camera configuration $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$* is defined as

$$(13) \quad \mathbf{S}(\mathbf{p}, \mathbf{e}) := \bigcup_{F \in \mathcal{F}} \mathbf{S}_F(\mathbf{p}, \mathbf{e}).$$

The intersection of the face or total shadow with the room is called the *cut-off (face or total) shadow*.

Figures 4 and 5 illustrate the cut-off shadows.

Parts of the room may be in the shadow of more than one face. The face shadow is a polyhedron, and its supporting hyperplanes are spanned by the face and the facets of the face together with the camera position. Note that if there is diffraction, then even in dimension two the invisible points are *not anymore* just the union of all face-shadows: adjacent segments in the environment would formally have a cone behind their intersection point that is not in the union of their shadows. However, that cone is clearly invisible in




FIGURE 4. Two cut-off face shadows: one for a face of the aperture blades and one for a face of the environment; some points in the room are in both shadows (darker)




FIGURE 5. The cut-off total shadows induced by two different camera configurations

reality. Thus, even for arbitrarily small positive diffraction (for which the main theorem is false) the construction of this paper fails at this point.

Definition 3 (Visible Volume). The *visible-volume function in room \mathbf{P} on the set of feasible camera configurations \mathbf{D}* is defined as the function

$$(14) \quad V(\cdot, \cdot): \begin{cases} \mathbf{D} & \rightarrow \mathbb{R}_{\geq 0} \\ (\mathbf{p}, \mathbf{e}) & \mapsto \text{vol}_d(\mathbf{P} \setminus \mathbf{S}(\mathbf{p}, \mathbf{e})) = \text{vol}_d(\mathbf{P}) - \text{vol}_d(\mathbf{S}(\mathbf{p}, \mathbf{e}) \cap \mathbf{P}) \end{cases}$$

The goal is to find analytic properties of the visible-volume function. In particular, for randomized local-search maximization methods it is of interest whether or not the visible-volume function is continuous and C^∞ almost everywhere.

Theorem 1 (Main Theorem). *The visible-volume function $V(\cdot, \cdot): \mathbf{D} \rightarrow \mathbb{R}_{\geq 0}$ is continuous. Moreover, there is a stratification of \mathbf{D} into strata S^∞ -diffeomorphic to open cubes with the following property: Restricted to the closure of each stratum the visible-volume function $V(\cdot, \cdot)$ is a rational function of degree at most d^3 .*

Since the graph of $V(\cdot, \cdot): \mathbf{D} \rightarrow \mathbb{R}_{\geq 0}$ is the union of the graphs of $V(\cdot, \cdot)$ restricted to the strata and $V(\cdot, \cdot)$ is rational on each stratum, $V(\cdot, \cdot)$ is semi-algebraic. Additionally, the only possibly non-smooth parts of the feasible domain are the boundaries between different strata. Since $V(\cdot, \cdot)$ is bounded and continuous, each rational function extends

to the boundary of each stratum, where it is equal to $V(\cdot, \cdot)$. Since rational functions are continuously differentiable on their complete domain, the one-sided directional derivatives of $V(\cdot, \cdot)$ restricted to the closure of each stratum are bounded for all directions into the stratum. Now, each stratum is in the boundary of only finitely many strata, so that the one-sided directional derivatives are bounded everywhere uniformly for each direction into the feasible domain. This implies that $V(\cdot, \cdot)$ is locally Lipschitz. Therefore, we have:

Corollary 1. *The visible-volume function $V(\cdot, \cdot)$ is piecewise smooth, in particular C^∞ almost everywhere. Moreover it is semi-algebraic and locally Lipschitz.*

In the next section, we will prove this theorem in eight steps by an explicit construction that generates a concrete formula for $V(\cdot, \cdot)$ via a (fairly fine) triangulation of the room. Note that the whole (convex!) room will be subdivided into convex regions and then triangulated. Thus, it will never be necessary to triangulate non-convex objects (which may fail starting in dimension three [9, Section 3.6]).

4. PROOF OF THE MAIN THEOREM

This section is devoted to the proof of the main theorem. Recall that $r := d + 1$ is the *rank* of the room \mathbf{P} . The proof idea is as follows: The set of feasible camera configurations can be partitioned into finitely many equivalence classes given by combinatorial equivalence of the induced hyperplane arrangement of all “visibility boundaries” and their intersection points (derived from oriented-matroid theory). Inside each of these finitely many equivalence classes the visible subset can be triangulated by the same (combinatorial) triangulation. Thus, in such an equivalence class the visible-volume function varies like the sum of the positive simplex-volumes over all visible simplices in the triangulation, which is essentially a determinant in the coordinates of the vertices of the simplices (derived from triangulation theory). Tracing back the data of these coordinates to the set of feasible camera configurations yields a rational function. Each equivalence class is a semi-algebraic set (derived from oriented-matroid theory and semi-algebraic-set theory) that can be further stratified into cubical cells (derived from semi-algebraic-set theory). Note that the final stratification is necessary, since a complete equivalence class can be complicated, e.g., not connected [16, Section 7.3.2]. The formal proof is structured into eight steps.

Step 1: The first part shows that the visible-volume function is continuous, which follows immediately from the definition via the face shadows.

Lemma 1. *The visible-volume function is continuous.*

Recall Figure 1 shows that the visible-subset function is not continuous, so that this result specifically holds for the volume, not for the visible subset of the room.

Proof of Lemma 1. It is enough to show that the cut-off total shadow is continuous. Each face shadow intersected with the room, the *cut-off face-shadow*, is a polytope whose outer description is continuous in the camera configuration. Thus, its volume function is bounded and continuous. By the inclusion-exclusion formula for the volume measure, the volume of the total shadow equals the sum of all cut-off face-shadow volumes minus the sum of all volumes of intersections of two cut-off face shadows plus the sum of all volumes of intersections of three cut-off face shadows . . . plus/minus the volume of the intersection of all cut-off face shadows. All appearing intersections are intersections of polytopes. Thus, all volumes in the inclusion-exclusion formula are continuous. Thus, the visible-volume function as a finite sum of continuous functions is continuous. \square




FIGURE 6. The hyperplane arrangement of a camera configuration; the longitudinal visibility planes are those that meet in the camera position

Step 2: For any feasible camera configuration $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$ and a $d - 1$ -face F of the environment complex \mathcal{F} , consider the affine hyperplane $\mathbf{H}(\mathbf{F}(\mathbf{e}))$ spanned by the corresponding geometric face $\mathbf{F}(\mathbf{e})$. It is called the *transversal visibility plane of F with respect to (\mathbf{p}, \mathbf{e})* . Moreover, for any feasible camera configuration $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$ and a $d - 2$ -face F' of the environment complex \mathcal{F} with $\mathbf{p} \notin \text{aff} \mathbf{F}'(\mathbf{e})$ (only these bound a $d - 1$ -face producing a shadow), consider the affine hyperplane $\mathbf{H}(\mathbf{F}'(\mathbf{e}), \mathbf{p})$ spanned by the corresponding geometric face $\mathbf{F}'(\mathbf{e})$ and the camera position \mathbf{p} . It is called the *longitudinal visibility plane of F' with respect to (\mathbf{p}, \mathbf{e})* . Since the environment complex has pure dimension $d - 1$, all $d - 2$ -faces of the environment complex bound some $d - 1$ -face of the environment complex.

Now let $\mathcal{H}(\mathbf{p}, \mathbf{e})$ be the affine hyperplane arrangement of all transversal and longitudinal visibility planes together with all facet-supporting hyperplanes of the room. They form the “visibility boundaries” mentioned in the proof idea. We consider this arrangement labeled by a label set I . This way, hyperplanes spanned by distinct simplices are distinguished, even if they are identical as subsets of \mathbb{R}^d . The label set is independent of (\mathbf{p}, \mathbf{e}) , since the environment complex \mathcal{F} is independent of (\mathbf{p}, \mathbf{e}) by assumption. See Figure 6 for an illustration of the resulting hyperplane arrangement.

Since there are only finitely many $d - 1$ -faces and $d - 2$ -faces in \mathcal{F} , the arrangement $\mathcal{H}(\mathbf{p}, \mathbf{e})$ is finite. More specifically, the number of hyperplanes in $\mathcal{H}(\mathbf{p}, \mathbf{e})$ is bounded by r times the number of $d - 1$ -faces in \mathcal{F} in the extreme case where no two $d - 1$ -faces of \mathcal{F} are adjacent and where no geometric $d - 2$ -face of \mathcal{F} contains \mathbf{p} in its affine hull.

The hyperplane data is a polynomial function in the data of the camera configuration.

Lemma 2. A hyperplane \mathbf{H} in $\mathcal{H}(\mathbf{p}, \mathbf{e})$ can be represented by equations $\mathbf{a}_{\mathbf{H}}^{\top} \mathbf{x} - b_{\mathbf{H}} = 0$ where the coordinates of $\mathbf{a}_{\mathbf{H}}$ and $b_{\mathbf{H}}$ are multilinear polynomial functions of degree at most d in (\mathbf{p}, \mathbf{e}) .

Proof of Lemma 2. Recall that (\mathbf{p}, \mathbf{e}) represents the coordinates of the camera position and all vertices of the geometric environment complex. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_d \in \mathbf{A}(\mathbf{e}) \cup \{\mathbf{p}\}$ span a hyperplane in $\mathcal{H}(\mathbf{p}, \mathbf{e})$. Then, a point \mathbf{x} lies in the hyperplane if and only if it satisfies the




FIGURE 7. Each region of the hyperplane arrangement is either completely visible or completely invisible from the camera position

following equation:

$$(15) \quad \det \begin{pmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_d & \mathbf{x} \\ 1 & 1 & \cdots & 1 & 1 \end{pmatrix} = 0$$

The coefficients of the coordinates of \mathbf{x} as well as the absolute term are, using development of the determinant by the last column, determinants of $d \times d$ -matrices whose entries are coordinates of points in $\mathbf{A}(\mathbf{e}) \cup \{\mathbf{p}\}$. Thus, they are multilinear polynomials of degree at most d in the camera configuration. \square

For the following, fix such equations $\mathbf{a}_i^\top(\mathbf{p}, \mathbf{e})\mathbf{x} - b_i(\mathbf{p}, \mathbf{e}) = 0$ for all hyperplanes \mathbf{H}_i with $i \in I$ and denote the matrix with rows $(\mathbf{a}_i^\top(\mathbf{p}, \mathbf{e}) | -b_i(\mathbf{p}, \mathbf{e}))_{i \in I}$ by $W(\mathbf{p}, \mathbf{e})$.

Step 3: An *open region* $\mathbf{R}(\mathbf{p}, \mathbf{e})$ of $\mathcal{H}(\mathbf{p}, \mathbf{e})$ is a connected component of the complement $\mathbf{P} \setminus \bigcup_{i \in I} \mathbf{H}_i$ of the arrangement in the room. The closures $\overline{\mathbf{R}(\mathbf{p}, \mathbf{e})}$ of the open regions are intersections of closed halfspaces with the polytope \mathbf{P} and, therefore, convex polytopes. All open faces of these polytopes are also called regions. Denote by $\mathcal{R}(\mathbf{p}, \mathbf{e})$ the set of (open) regions of $\mathcal{H}(\mathbf{p}, \mathbf{e})$. See Figure 7 on the left for the regions in the running example.

Note that, by construction, each open region of the hyperplane arrangement is either completely contained in the total shadow (invisible from \mathbf{p}) or completely in the complement of the total shadow (visible from \mathbf{p}). This is because all supporting hyperplanes of facets of these shadows have been added to $\mathcal{H}(\mathbf{p}, \mathbf{e})$. Thus, the visible volume is the sum of volumes of all d -dimensional regions that are inside the room and outside the total shadow, which is the statement of the following lemma.

Lemma 3. *For a fixed feasible camera configuration $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$, the visible volume can be computed as follows:*

$$(16) \quad V(\mathbf{p}, \mathbf{e}) = \sum_{\substack{\mathbf{R}(\mathbf{p}, \mathbf{e}) \in \mathcal{R}(\mathbf{p}, \mathbf{e}) \\ \mathbf{R}(\mathbf{p}, \mathbf{e}) \subseteq \mathbf{P} \setminus \mathbf{S}(\mathbf{p}, \mathbf{e})}} \text{vol}_d(\mathbf{R}(\mathbf{p}, \mathbf{e})).$$

\square

Figure 7 on the right shows the visible volume, subdivided into visible regions.




FIGURE 8. Two regions and their covectors

Step 4: Associated with $\mathcal{H}(\mathbf{p}, \mathbf{e})$ is its oriented matroid (see Notion 2), determined by its finite set of signed covectors. Recall that this is a set of signatures $\sigma_{\mathbf{R}(\mathbf{p}, \mathbf{e})}: I \rightarrow \{-, 0, +\}$. Each signature in this set corresponds to an open region $\mathbf{R}(\mathbf{p}, \mathbf{e})$ in $\mathcal{R}(\mathbf{p}, \mathbf{e})$ of $\mathcal{H}(\mathbf{p}, \mathbf{e})$ so that

$$(17) \quad \mathbf{R}(\mathbf{p}, \mathbf{e}) := \{\mathbf{x} \in \mathbf{P} : \text{sign}(\mathbf{a}_i^\top(\mathbf{p}, \mathbf{e})\mathbf{x} - b_i(\mathbf{p}, \mathbf{e})) = \sigma_{\mathbf{R}(\mathbf{p}, \mathbf{e})}(i), i \in I\}.$$

Figure 8 illustrates the covectors of two full-dimensional regions. Each covector is indicated as a string of signs. Each position in the string corresponds to a hyperplane. The sign in that position is the signature value of that hyperplane.

We define an equivalence relation “ $\sim_{\mathcal{H}}$ ” on \mathbf{D} where $(\mathbf{p}, \mathbf{e}) \sim_{\mathcal{H}} (\mathbf{p}', \mathbf{e}')$ if and only if $\mathcal{H}(\mathbf{p}, \mathbf{e})$ and $\mathcal{H}(\mathbf{p}', \mathbf{e}')$ are combinatorially equivalent (Notion 2). The equivalence class of (\mathbf{p}, \mathbf{e}) with respect to “ $\sim_{\mathcal{H}}$ ” is denoted by $M_{\mathcal{H}}(\mathbf{p}, \mathbf{e})$.

Let $\mathcal{M}_{\mathcal{H}}$ be the set (see Fact 3) of such equivalence classes in \mathbf{D} . This set is finite by Fact 3. Whether or not a region is totally visible or totally not visible only depends on $M_{\mathcal{H}}(\mathbf{p}, \mathbf{e})$, as the following crucial lemma establishes.

Lemma 4. *A region $\mathbf{R}(\mathbf{p}, \mathbf{e})$ is in the face shadow $\mathbf{S}_F(\mathbf{p}, \mathbf{e})$ of a $d - 1$ -simplex $F \in \mathcal{F}$ with respect to a feasible camera configuration (\mathbf{p}, \mathbf{e}) if and only if the following three conditions hold for $\sigma_{\mathbf{R}(\mathbf{p}, \mathbf{e})}$:*

- (i) *The vertices of $\mathbf{F}(\mathbf{e})$ and the camera position \mathbf{p} are cocircuits of $\mathcal{H}(\mathbf{p}, \mathbf{e})$.*
- (ii) *Let $i \in I$ be the label of the transversal visibility plane spanned by F . Then*

$$(18) \quad \sigma_{\mathbf{R}(\mathbf{p}, \mathbf{e})}(i) = -\sigma_{\mathbf{p}}(i)$$

- (iii) *Let $i \in I$ be the label of any longitudinal visibility plane spanned by \mathbf{p} and a $d - 2$ -face $F' \in \mathcal{F}$ of F and let $\mathbf{v}_{F \setminus F'}$ be the geometric vertex of $\mathbf{F}(\mathbf{p}, \mathbf{e})$ missing in $\mathbf{F}'(\mathbf{p}, \mathbf{e})$. Then*

$$(19) \quad \sigma_{\mathbf{R}(\mathbf{p}, \mathbf{e})}(i) = \sigma_{\mathbf{v}_{F \setminus F'}}(i)$$

Moreover, $\mathbf{R}(\mathbf{p}, \mathbf{e})$ is in the total shadow $\mathbf{S}(\mathbf{p}, \mathbf{e})$ if and only if the above conditions are satisfied for at least one $d - 1$ -face of \mathcal{F} .

Proof of Lemma 4. Assume first that $\mathbf{R}(\mathbf{p}, \mathbf{e})$ is contained in the face shadow of F . In particular, F has a non-empty face shadow. In particular, the camera position is affinely independent of the vertices of $\mathbf{F}(\mathbf{e})$ with which it spans a d -dimensional simplex. The

facet-supporting hyperplane of this simplex are, by construction, exactly the transversal and the longitudinal visibility hyperplanes associated to F , which are members of $\mathcal{H}(\mathbf{p}, \mathbf{e})$. This implies condition (i), and, hence, that conditions (ii) and (iii) are well-defined. In this case, since $\mathbf{R}(\mathbf{p}, \mathbf{e})$ is contained in the face shadow of F , it satisfies all facet defining inequalities of it. This implies conditions (ii) and (iii), which are a reformulation of the same fact in terms of cocircuit signatures.

Assume now that all three conditions are satisfied. In particular, (ii) and (iii) are well-defined. By those conditions, $\mathbf{R}(\mathbf{p}, \mathbf{e})$ satisfies all facet defining inequalities of the face shadow of F . Thus, it is contained in that face shadow.

The assertion about the total shadow follows from its definition as the finite union of face shadows. \square

Corollary 2. *Whether or not a region is in the total shadow does not change inside each $M_{\mathcal{H}}(\mathbf{p}, \mathbf{e})$.*

Regions inside the total shadow are called *invisible from \mathbf{p}* , not invisible regions are called *visible from \mathbf{p}* .

Step 5: Recall that the set $\mathcal{M}_{\mathcal{H}}$ is finite (Fact 3). Recall further that the set $\mathcal{M}_{\mathcal{H}}$ is the set of all equivalence classes of elements in \mathbf{D} , where the members of each equivalence class induce combinatorially equivalent hyperplane arrangements. The following lemma states that all equivalence classes are semi-algebraic sets.

Lemma 5. *For each $(\bar{\mathbf{p}}, \bar{\mathbf{e}}) \in \mathbf{D}$, the equivalence class $M_{\mathcal{H}}(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ is a semi-algebraic set. Moreover, there are only finitely many equivalence classes in \mathbf{D} .*

Proof of Lemma 5. The equivalence class $M_{\mathcal{H}}(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ of $(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ in \mathbf{D} can be written as follows:

$$(20) \quad M_{\mathcal{H}}(\bar{\mathbf{p}}, \bar{\mathbf{e}}) = \mathbf{D} \cap \text{proj}_{(\mathbf{p}, \mathbf{e})} \left\{ (\mathbf{p}, \mathbf{e}, W, \bar{W}) \in \mathbb{R}^d \times (\mathbb{R}^d)^J \times \mathbb{R}^{I \times r} \times \mathbb{R}^{I \times r} : \right. \\ \left. \begin{aligned} W &= W(\mathbf{p}, \mathbf{e}), \\ \bar{W} &= W(\bar{\mathbf{p}}, \bar{\mathbf{e}}), \\ M_{\mathcal{H}}(W) &= M_{\mathcal{H}}(\bar{W}) \end{aligned} \right\}.$$

Here \mathbf{D} is a semi-algebraic set by assumption. The projected set is semi-algebraic by Fact 4 and Lemma 2. The projection of it to the (\mathbf{p}, \mathbf{e}) -coordinates is a semi-algebraic set by Fact 1. The intersection with \mathbf{D} is semi-algebraic by Fact 1. Therefore, the equivalence class $M_{\mathcal{H}}(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ is a semi-algebraic set. Since there are only finitely many oriented matroids on a finite set of hyperplanes (Fact 3), there are only finitely many equivalence classes. \square

Step 6: The plan is now to compute the volumes of the regions not in the total shadow more explicitly. To this end, each region is triangulated. In order to keep control of the dependence on the camera configuration (\mathbf{p}, \mathbf{e}) , the same combinatorial type of triangulation should be kept for as many camera configurations as possible. A sufficient condition for a triangulation to stay feasible for the same point configuration with tweaked coordinates is that the oriented matroid of the point configuration does not change.

Note that the oriented matroid of a hyperplane arrangement does not uniquely determine the oriented matroid of its vertices as a point configuration, not even for each region separately.

To by-pass this problem, each equivalence class of the previous step is further subdivided. Denote by $\mathbf{C}(\mathbf{p}, \mathbf{e})$ the set of vertices of the polyhedral complex of regions of $\mathcal{H}(\mathbf{p}, \mathbf{e}) \cap \mathbf{P}$ and let $K(\mathbf{p}, \mathbf{e})$ be its label set. In order to guarantee a consistent labeling, we label each




FIGURE 9. The intersection points of the hyperplane arrangement are the vertices of the regions

point by the subset of labels of those hyperplanes it lies in, where intersection points outside the room are ignored. Note that intersection points inside and outside the room lie on opposite sides of at least one facet-supporting hyperplane of the room, leading to a different oriented matroid of the hyperplane arrangement. Thus, in each equivalence class $M_{\mathcal{H}}(\mathbf{p}, \mathbf{e})$ the subset of intersection points inside the room is constant. Figure 9 shows the point configuration of all intersection points inside the room.

We define an equivalence relation “ $\sim_{\mathbf{C}}$ ” on \mathbf{D} where $(\mathbf{p}, \mathbf{e}) \sim_{\mathbf{C}} (\mathbf{p}', \mathbf{e}')$ if and only if $\mathbf{C}(\mathbf{p}, \mathbf{e})$ and $\mathbf{C}(\mathbf{p}', \mathbf{e}')$ are combinatorially equivalent (Notion 3). The equivalence class of (\mathbf{p}, \mathbf{e}) with respect to “ $\sim_{\mathbf{C}}$ ” is denoted by $M_{\mathbf{C}}(\mathbf{p}, \mathbf{e})$.

Lemma 6. *The coordinates of the points in $\mathbf{C}(\mathbf{p}, \mathbf{e})$ are rational functions of degree d in the coordinates of the coefficient vectors and the right-hand sides of the hyperplanes in $\mathcal{H}(\mathbf{p}, \mathbf{e})$. Consequently, they are rational functions of degree at most d^2 in (\mathbf{p}, \mathbf{e}) .*

Proof of Lemma 6. Each point in $\mathbf{C}(\mathbf{p}, \mathbf{e})$ is the unique intersection of d hyperplanes in $\mathcal{H}(\mathbf{p}, \mathbf{e})$. Thus, each coordinate of such a point is, by Cramer’s rule, a quotient of $d \times d$ -determinants in the coefficients and right-hand sides of the hyperplanes. This is a rational function of degree at most d in the data of $\mathcal{H}(\mathbf{p}, \mathbf{e})$. Moreover, by composing this function with the polynomial function from Lemma 2, we obtain a rational function of degree at most d^2 , as claimed. \square

Define an equivalence relation on the set of all feasible camera configurations $(\mathbf{p}, \mathbf{e}) \in \mathbf{D}$ as follows: (\mathbf{p}, \mathbf{e}) and $(\mathbf{p}', \mathbf{e}') \in \mathbf{D}$ are *equivalent*, $(\mathbf{p}, \mathbf{e}) \sim (\mathbf{p}', \mathbf{e}')$, if and only if $(\mathbf{p}, \mathbf{e}) \sim_{\mathcal{H}} (\mathbf{p}', \mathbf{e}')$ and, additionally, $(\mathbf{p}, \mathbf{e}) \sim_{\mathbf{C}} (\mathbf{p}', \mathbf{e}')$. The set of all equivalence classes is denoted by \mathcal{M}_{\sim} . Thus, by Fact 5, equivalent camera configurations induce identical sets of (combinatorial types of) triangulations of $\mathbf{C}(\mathbf{p}, \mathbf{e})$ and identical sets of (combinatorial types of) regions. This, finally, is the partitioning of \mathbf{D} into equivalence classes mentioned in the proof idea.

Each equivalence class $M_{\sim}(\bar{\mathbf{p}}, \bar{\mathbf{e}}) \in \mathcal{M}_{\sim}$ of some $(\bar{\mathbf{p}}, \bar{\mathbf{e}}) \in \mathbf{D}$ is the intersection of \mathbf{D} with $M_{\mathcal{H}}(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ and $M_{\mathbf{C}}(\bar{\mathbf{p}}, \bar{\mathbf{e}})$. Analogously to the proof of Lemma 5 it can be seen that all three are semi-algebraic sets (by Fact 4 and Lemmas 2 and 6), and so is their intersection (Fact 1). Consequently, by Fact 2, each such equivalence class in \mathbf{D} and, hence, the semi-algebraic




FIGURE 10. The volumes of the visible regions can be computed by a triangulation of the intersection points refining the polyhedral subdivision by regions; the camera configuration on the right is better than the one on the left

set of feasible camera configurations \mathbf{D} itself can be stratified into strata S^∞ -diffeomorphic to open cubes of certain dimensions so that the boundary of each stratum is a finite union of smaller-dimensional strata and so that inside each stratum all camera configurations are equivalent. Moreover, the volumes of the finitely many max-dimensional strata add up to the volume of \mathbf{D} , and the not-max-dimensional strata form a set of measure zero.

Step 7: For each $M_\sim(\mathbf{p}, \mathbf{e}) \in \mathcal{M}_\sim$ refine the polyhedral complex of regions in $\mathcal{R}(\mathbf{p}, \mathbf{e})$ that are inside \mathbf{P} to a fixed triangulation of $\mathbf{C}(\mathbf{p}, \mathbf{e})$, e.g., by some generic regular refinement [9, Corollary 2.3.18]. Any such triangulation (seen as an abstract simplicial complex on $K(\mathbf{p}, \mathbf{e})$, see Notion 4) defines via the embedding \mathbf{e} and the position \mathbf{p} a valid triangulation for all feasible camera configurations equivalent to (\mathbf{p}, \mathbf{e}) , since the validity of a triangulation depends only on the oriented matroid of the point configuration to be triangulated (Fact 5). Call this triangulation $\mathcal{T}(M_\sim(\mathbf{p}, \mathbf{e}))$. Because $\mathcal{T}(M_\sim(\mathbf{p}, \mathbf{e}))$ refines $\mathcal{R}(\mathbf{p}, \mathbf{e})$ inside \mathbf{P} , the open geometric embedding of each simplex in $\mathcal{T}(M_\sim(\mathbf{p}, \mathbf{e}))$ is either completely visible or completely invisible from \mathbf{p} , and the visibility does not change inside an equivalence class.

Figure 10 shows a possible triangulation for the visible region and the resulting visible volume.

Step 8: Consider some feasible camera configuration $(\bar{\mathbf{p}}, \bar{\mathbf{e}}) \in \mathbf{D}$. Then the visible volume can be computed for all (\mathbf{p}, \mathbf{e}) in the equivalence class $M_\sim(\bar{\mathbf{p}}, \bar{\mathbf{e}})$ as follows :

$$\begin{aligned}
 V(\mathbf{p}, \mathbf{e}) &= \sum_{\substack{F \in \mathcal{T}(M_\sim(\bar{\mathbf{p}}, \bar{\mathbf{e}})): \\ \mathbf{F}^\circ(\mathbf{p}, \mathbf{e}) \subseteq \mathbf{P} \setminus \mathbf{S}(\mathbf{p}, \mathbf{e})}} \text{vol}_d(\mathbf{F}(\mathbf{p}, \mathbf{e})) \\
 (21) \quad &= \frac{1}{d!} \sum_{\substack{F = \{k_1, \dots, k_r\} \in \mathcal{T}(M_\sim(\bar{\mathbf{p}}, \bar{\mathbf{e}})): \\ \mathbf{F}^\circ(\mathbf{p}, \mathbf{e}) \subseteq \mathbf{P} \setminus \mathbf{S}(\mathbf{p}, \mathbf{e})}} \left| \det \begin{pmatrix} \mathbf{v}^{k_1} & \cdots & \mathbf{v}^{k_r} \\ 1 & \cdots & 1 \end{pmatrix} \right|.
 \end{aligned}$$

This is the sum of positive simplex volumes mentioned in the proof idea. Thus, the visible-volume function restricted to an equivalence class is a sum of determinants of $r \times r$ -matrices with a row of ones whose columns are given by the homogeneous coordinates of points in $\mathbf{C}(\mathbf{p}, \mathbf{e})$, thus a multilinear polynomial function of degree at most d in the coordinates

of the points in $\mathbf{C}(\mathbf{p}, \mathbf{e})$. By composition with earlier results, this is a rational function of degree at most d^3 in the camera configuration data. By continuity and boundedness of the visible-volume function everywhere in \mathbf{D} , the visible-volume function formula (21) extends to the boundary of each stratum. (Intuitively, some of the simplex volumes contained in the triangulation corresponding to the interior of the stratum degenerate to zero on some boundary stratum; these simplices are exactly those missing in the triangulation corresponding to the boundary stratum, which leads to the same volume.) Thus, the visible-volume function is a rational function of degree at most d^3 on the closure of each stratum. \square

Remark 5. Although all illustrations are two-dimensional, all constructions are correct in arbitrary spacial dimensions d . Note, though, that in dimension two the result can be achieved in an easier way, exploiting structures of dimension two that cannot be generalized to higher dimensions.

The construction allows for the following generalization: If there is more than one camera to be configured, we can construct the open regions visible and invisible from each camera separately. Each open region in the combined hyperplane arrangement is then a subset of a unique open region in each separate arrangement and, hence, for each camera either completely visible or invisible. We can, therefore, decide the exact subset of cameras from which it is visible completely. Therefore, the result extends to the volume covered by a given collection of subsets of cameras. In particular, this applies to the volume covered by at least/exactly/at most k cameras. Call the function that assigns to each feasible configuration of multiple cameras the volume visible by at least/exactly/at most k cameras the *at-least/exactly/at-most- k -fold-visible-volume function*, respectively. With this definition, the following corollary holds true.

Corollary 3. *For all $k \in \mathbb{N}$ the at-least/exactly/at-most- k -fold-visible-volume functions are continuous, piecewise rational, locally Lipschitz, and semi-algebraic in all dimensions.* \square

Remark 6. A short note on complexity estimates: If K is the number of facets of the environment complex and the room, then there are K transversal and at least $K(r-1)$ longitudinal visibility hyperplanes. With this, in general, one must consider roughly $\sum_{i=1}^r \binom{K}{i} \binom{K(r-1)}{r-i}$ regions contributing to potentially different summands over all strata. The number of strata is even more discouraging: even if each equivalence class constitutes a single stratum without further decomposition, then a-priori there can be roughly as many strata as there are different oriented matroids on arrangements of Kr hyperplanes, where $K(r-1)$ of them all intersect in a point. For example, in physical space, i.e., in rank 4, there are already 10,559,305 realizable simple acyclic oriented matroids on 8 elements (= hyperplanes) [11]. Even if all but $r=4$ many longitudinal hyperplanes are ignored because of their forced common intersection, this number of strata would be reached already by the at least $K=4$ facets of the room without any obstacles.

5. CONCLUSIONS

In this paper, the visible-volume function was shown to be continuous, piecewise rational, locally Lipschitz, and semi-algebraic in all dimensions. The set-up was chosen more general than in earlier work: the visible-volume function may depend on the positions of a set of cameras and the embedding of an environment in a convex room. It returns the volume that the cameras can supervise given the cameras' positions and the embedding of the environment. The environment can be an arbitrary abstract simplicial complex of pure

codimension one. The embedding and the camera positions can be restricted to an arbitrary semi-algebraic set. The actual result provides a piecewise smooth formula with a rational function on each piece. Given this knowledge, we can be sure that certain global blackbox optimization methods applied to the visible-volume function will converge (at least if the restrictions of the embedding and the camera positions are chosen to be tractable). Moreover, from the explicit piecewise formula, gradient information of the visible-volume formula can easily be obtained at each feasible camera configuration.

One contribution of this paper is to formalize the problem and construct suitable objects of discrete geometry in such a way that a novel combination of the basic theories of oriented matroids, triangulations, and semi-algebraic sets implies the result. This can be seen as a show-case how those theories play together, which may be interesting in its own right.

At first glance, it seems that the exploitation of the explicit piecewise formula in this paper for optimization is out of reach for sufficiently realistic environments. However, it is conceivable that large-scale methods from mixed integer linear programming, e.g., with dynamically generated selection variables for oriented matroids, might constitute a viable approach to find an exact global extremum of the visible volume function. Whether such methods can be implemented to show a satisfactory performance for practical problems or at least for medium-sized academic benchmark examples is an open question.

ACKNOWLEDGEMENT

The author thanks Maria Hänel and Christian Knauer for inspiring discussions.

REFERENCES

1. Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow, *The art gallery problem is $\exists \mathbb{R}$ -complete*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (New York, NY, USA), STOC 2018, Association for Computing Machinery, 2018, p. 65–73.
2. Saugata Basu, Richard Pollack, and Marie-Françoise Roy, *Algorithms in real algebraic geometry (algorithms and computation in mathematics)*, Springer-Verlag, Berlin, Heidelberg, 2006.
3. Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Gunter M. Ziegler, *Oriented matroids*, 2 ed., Encyclopedia of Mathematics and its Applications, Cambridge University Press, 1999.
4. James V. Burke, Frank E. Curtis, Adrian S. Lewis, Michael L. Overton, and Lucas E. A. Simões, *Gradient sampling methods for nonsmooth optimization*, Numerical Nonsmooth Optimization: State of the Art Algorithms (Cham) (Adil M. Bagirov, Manlio Gaudioso, Napsu Karmitsa, Marko M. Mäkelä, and Sona Taheri, eds.), Springer International Publishing, 2020, pp. 201–225.
5. Li-Tien Cheng and Yen-Hsi Tsai, *Visibility optimization using variational approaches*, Commun. Math. Sci. **3** (2005), no. 3, 425–451.
6. Vasek Chvátal, *A combinatorial theorem in plane geometry*, Journal of Combinatorial Theory, Series B **18** (1975), 39–41.
7. Franklin C. Crow, *Shadow algorithms for computer graphics*, Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques (New York, NY, USA), SIGGRAPH '77, Association for Computing Machinery, 1977, p. 242–248.
8. James Damon, Peter Giblin, and Gareth Haslinger, *Local features in natural images via singularity theory*, 1st ed., Springer Publishing Company, Incorporated, 2016.
9. Jesús De Loera, Jörg Rambau, and Francisco Santos, *Triangulations – Structures for Applications and Algorithms*, Algorithms and Computation in Mathematics, vol. 25, Springer, 2010.
10. Steve Fisk, *A short proof of Chvátal's watchman theorem*, Journal of Combinatorial Theory, Series B **24** (1978), no. 3, 374.
11. Komei Fukuda, Hiroyuki Miyata, and Sonoko Moriyama, *Complete enumeration of small realizable oriented matroids*, Discrete and Computational Geometry **49** (2013), no. 2, 359–381.
12. Anurag Ganguli, Jorge Cortés, and Francesco Bullo, *Maximizing visibility in nonconvex polygons: Nonsmooth analysis and gradient algorithm design*, SIAM J. Control and Optimization **45** (2005), 792 – 797 vol. 2.
13. Maria Hänel, *A matter of perspective - three-dimensional placement of multiple cameras to maximize their coverage*, Ph.D. thesis, University of Bayreuth, Bayreuth, Oktober 2015.

14. Aaron Mavrinac and Xiang Chen, *Modeling coverage in camera networks: A survey*, International Journal of Computer Vision **101** (2013), no. 1, 205–226.
15. Stanley Osher, *Level set methods*, Geometric Level Set Methods in Imaging, Vision, and Graphics (Stanley Osher and Nikos Paragios, eds.), Springer New York, New York, NY, 2003, pp. 3–20.
16. Jürgen Richter-Gebert and Günter M. Ziegler, *Oriented matroids*, Handbook of Discrete and Computational Geometry (Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, eds.), CRC Press, Inc., USA, 1997, p. 111–132.

JÖRG RAMBAU, UNIVERSITY OF BAYREUTH, TEL.: +49-921-55-7350
Email address: joerg.rambau@uni-bayreuth.de