

Energy-Bounded Caging: Formal Definition and 2D Energy Lower Bound Algorithm Based on Weighted Alpha Shapes

Jeffrey Mahler¹, Florian T. Pokorny¹, Zoe McCarthy¹, A. Frank van der Stappen³, Ken Goldberg²

Abstract—Caging grasps are valuable as they can be robust to bounded variations in object shape and pose and do not depend on friction. Full caging is useful but may not be necessary in cases where forces such as gravity are present (consider a stone in a cupped hand). This paper extends caging theory by defining energy-bounded cages under a constant potential energy field (such as gravity) based on the minimum energy required to escape. This paper also introduces Energy-Bounded-Cage-Analysis-2D (EBCA-2D), a sampling-based algorithm for planar analysis that takes as input a constant energy field specified as a function over poses, a polygonal object, and a configuration of rigid polygonal obstacles, and returns a lower bound on the minimum escape energy, which can be infinite when the object is fully caged. Building on recent results in collision detection and the computational geometric theory of weighted α -shapes, EBCA-2D is provably-correct and runs in time $O(N^2 + N \log(1/\Delta) + NV^3)$ time where N is the number of samples, Δ is an energy resolution used for binary search, and V is the total number of object and obstacle vertices. We implemented EBCA-2D and evaluated it with nine parallel-jaw gripper and four nonconvex obstacle configurations across six nonconvex polygonal objects. We found that the lower bounds returned by EBCA-2D are consistent with intuition and with an RRT* optimal motion planning algorithm that was unable to find escape paths with lower energy. EBCA-2D required an average of 3 minutes per problem on a single-core processor but has potential to be parallelized in a Cloud-based implementation. Additional proofs, data, and code are available at: <http://berkeleyautomation.github.io/caging/>.

I. INTRODUCTION

Consider a single movable object and a configuration of fixed obstacles. The object is caged if it cannot escape [22], [33]. When the obstacles describe the links of a robot gripper, they can restrain and, by moving, transport the caged object. Caging grasps are valuable as they can be robust to bounded variations in object shape and pose and do not depend on friction [1], [37], [47].

Full caging is useful but may not be necessary in cases where forces such as gravity are present. This paper extends caging theory by defining energy-bounded cages under a constant potential energy field (such as gravity) based on the minimum energy required to escape. For example, the objects in Fig. 1 would have to overcome gravitational forces to escape by lifting above the obstacles. However, computing the minimum escape energy may be challenging for arbitrary nonconvex objects and obstacles because there are typically




Fig. 1: Two energy-bounded cages of industrial parts (blue) by robotic grippers (black) under a gravitational field indicated by the center arrow. Neither object is caged by the classical definition, but both configurations are energy-bounded cages because the object must overcome the forces exerted by gravity to escape.

an uncountable number of paths that an object can take to escape a cage. This is also a problem in classical caging, as currently cages in 2D can only be verified under assumptions on the number of gripper fingers [1], [31], [42], the geometry of obstacles [37], or the geometry of the object [14], [43].

We present Energy-Bounded-Cage-Analysis-2D (EBCA-2D), the first sampling-based algorithm that can verify cages and energy-bounded cages for 2D nonconvex polygonal objects and an arbitrary number of nonconvex polygonal obstacles under a constant potential energy field. Our algorithm computes a lower-bound on the minimum escape energy using weighted α -shapes [11], [12], a discrete representation of the configuration space between the object and obstacles that has been used for proving path non-existence in motion planning [2], [30], [50]. We use weighted α -shapes to decompose the object configuration space into cells from a set of sampled object poses and a conservative estimate of their penetration depth. We then mark forbidden cells that lie strictly within the collision space or above an energy threshold and examine the connectivity of the free cells to prove the non-existence of object escape paths [30]. Finally, we lower bound the minimum escape energy by performing a binary search over energy levels, querying the connectivity of the free cells for each threshold. We evaluate our algorithm on a set of nine parallel-jaw gripper configurations and four configurations of nonconvex obstacles across six polygonal objects under gravity, and find that in each case an RRT* optimal path planner is not able to plan an escape path with lower energy than our estimated lower bound within 120 seconds of planning.

II. RELATED WORK

For surveys of the substantial literature on grasping, see Bicchi and Kumar [3] or Prattichizzo and Trinkle [32]. Many metrics for grasp quality follow one of two directions: wrench-space metrics or caging metrics. Grasp wrench space metrics measure the ability of a grasp to resist external forces and torques applied to a grasped object [15]. Wrench space

¹Department of Electrical Engineering and Computer Sciences; {jmahler, ftpokorny, zmccarthy}@berkeley.edu

²Department of Industrial Engineering and Operations Research; goldberg@berkeley.edu

¹⁻² University of California, Berkeley, USA

³Department of Information and Computing Sciences, Utrecht University, The Netherlands; A.F.vanderStappen@uu.nl

metrics have also been developed to measure the ability to resist task-specific wrenches on an object from point contacts [21], [20], [25], the set of perturbation wrenches that can be applied to an object [27], or robustness to uncertainty [17], [23], [28], [48].

While wrench space metrics depend on local properties of an object, caging metrics depend on the global geometry of an object and gripper. Early works defined a cage as a configuration of a “hand” (n points in a plane) such that a planar object could not be moved arbitrarily far away from the hand [22], [35]. Rimon and Blake [33] later characterized the space of caging hand configurations for a 1-parameter two-fingered gripping system with convex fingers. Rimon and Blake [34] also developed an algorithm to determine the maximal set of two-finger gripper configurations that cage an object, which was later extended to three fingers by Davidson and Blake [9]. Other research has presented algorithms for computing the set of caging configurations for grasps with two or three disc fingers on convex polygons [14], non-convex polygons [31], [42] and grasps on 3D polyhedra with fingers that can be decomposed into points [1].

Several works have studied the relation of caging configurations to uncertainty and to form closure grasps. Vahedi and van der Stappen [42] developed the concepts of squeezing and stretching cages for two-finger grippers, and showed that two-finger cages in the plane can always lead to a form closure grasp by either opening or closing the fingers. Rodriguez and Mason [36] extended this property to two finger cages of compact and contractible objects in arbitrary dimensions, and later generalized the link between caging and grasping to more than two fingers, showing that cages can be a useful waypoint to a form closure grasp of a polygonal object when the gripper stays in a sub- or super-level set of a gripper shape function [37]. Cages have also been shown experimentally to offer robustness to shape and pose uncertainty. Diankov et al. [10] found that caging grasps were empirically more successful than those ranked by local force closure metrics when manipulating articulated objects with handles. Other work has studied the robustness of caging grasps to object pose uncertainty [47] or uncertainty in object shape due to vision [39].

Due to the difficulty of computing the entire space of caging configurations for complex hand and object geometries, several recent works have studied heuristics for determining whether or not a single hand configuration cages an object such as leveraging holes in the object [38]. Makponyo et al. [29] introduced the concept of partial cage quality for a hand configuration, arguing that configurations that allow only rare escape motions may be successful in practice. The authors proposed a heuristic metric based on the length and curvature of escape paths generated by a motion planner. Wan et al. [46] determined cages for 2D polygons by mapping out the configurations in collision in a voxelized representation of the configuration space and checking connectivity. In comparison, we present a formal definition and metric of energy-bounded cages and formally prove that a cell decomposition of the 3D configuration space




Fig. 2: A total cage and an energy-bounded cage illustrated in the workspace (top) and translational configuration space (bottom) for a gravitational potential field. (Left) An object (blue) is caged by obstacles (black) if it cannot reach poses arbitrarily far away from its initial pose \mathbf{q}_0 . In configuration space this corresponds to a disconnection between the object initial pose \mathbf{q}_0 and the free space \mathcal{F} outside of the obstacles. (b) A gripper forms a u -energy-bounded cage of an object if all escape paths cross the u -superlevel set of a potential energy function $\mathcal{E} : SE(2) \rightarrow \mathbb{R}$ (yellow).

can be used to verify cages and energy-bounded cages.

Our work is also related to the problem of proving path existence and non-existence in the field of motion planning. When the free configuration space can be described by semi-algebraic functions, the free space can be analytically decomposed into cells to answer path existence queries [24]. However, such functions may not exist or computing such a decomposition may be prohibitively expensive, motivating alternative methods. Basch et al. [2] provided a quadratic-time algorithm to prove path non-existence of a polygon through a polygonal hole in an infinite wall. Zhang et al. [50] developed a method for approximately decomposing the free space and obstacle space for a robot into rectangular cells, labelling cells as being in collision using penetration depth computation, and searching for paths through cells in free space. McCarthy et al. [30], used configuration samples to approximate the collision space using α -shapes and present an algorithm that can verify path non-existence between two configurations. See [11] for a complete treatment of weighted α -shapes.

III. DEFINITIONS

A. Definitions

We consider the problem of caging a compact 2D polygonal object $\mathcal{O} \subset \mathbb{R}^2$ with V_o vertices by a fixed configuration of compact polygonal obstacles $\mathcal{G} \subset \mathbb{R}^2$ with V_g total vertices. We consider \mathcal{G} to be fixed in the environment and denote the object polygon in pose $\mathbf{q} \in SE(2)$ relative to its initial configuration \mathbf{q}_0 as $\mathcal{O}(\mathbf{q})$. Example obstacles \mathcal{G} include the end-effectors of a robotic gripper or parts of the environment such as walls or support surfaces. Note that we are interested in verifying cages for a fixed gripper configuration and currently do not jointly consider all possible configurations of the gripper.




Fig. 3: Illustration of the subset property of cages. (a) An object in pose \mathbf{q}_0 is caged because it cannot reach poses \mathbf{q}_f arbitrarily far away without crossing a forbidden region \mathcal{Z} . (b) If no path from \mathbf{q}_0 to \mathbf{q}_f exists without passing through a subset $\mathcal{Y} \subseteq \mathcal{Z}$, then the object must be caged. (c) However, we are not guaranteed to verify a cage using any subset $\mathcal{Y}' \subseteq \mathcal{Z}$ because \mathcal{Y}' may not block all paths from \mathbf{q}_0 to \mathbf{q}_f .

B. Caging

Consider a forbidden subset \mathcal{Z} of a topological space describing the configuration space \mathcal{C} of an object:

Definition 3.1: Let \mathcal{C} be a path-connected non-compact topological space and $\mathcal{Z} \subset \mathcal{C}$. We call a point $\mathbf{x} \in \mathcal{C} \setminus \mathcal{Z}$ *caged by \mathcal{Z} in \mathcal{C}* if \mathbf{x} lies in a compact path connected component of $\mathcal{C} \setminus \mathcal{Z}$. This definition is illustrated in the left panel of Fig. 2.

We can verify cages using a sufficient condition for caging:

Lemma 3.1: Let $\mathcal{Y} \subseteq \mathcal{Z} \subset \mathcal{C}$. If $\mathbf{x} \in \mathcal{C} \setminus \mathcal{Z}$ is caged by \mathcal{Y} , then \mathbf{x} is caged by \mathcal{Z} .

Proof: $\mathcal{C} \setminus \mathcal{Z} \subseteq \mathcal{C} \setminus \mathcal{Y}$, which implies that any path in $\mathcal{C} \setminus \mathcal{Z}$ can be restricted to $\mathcal{C} \setminus \mathcal{Y}$. ■

This property is illustrated in Fig. 3. Thus if we can prove the caging condition for a subset \mathcal{Y} of the true set of interest \mathcal{Z} , then the result will hold for \mathcal{Z} .

In this work, we are interested in the case where $\mathcal{C} \subset SE(2)$ is a subset of rigid transformations of a rigid object \mathcal{O} in the plane and $\mathcal{Z} \subseteq SE(2)$ is the *collision space* of \mathcal{O} relative to \mathcal{G} [24]:

$$\mathcal{Z} = \{\mathbf{q} \in SE(2) \mid \text{int}(\mathcal{O}(\mathbf{q})) \cap \mathcal{G} \neq \emptyset\}.$$

Note that \mathcal{Z} is compact based on our assumptions. We denote by $\mathcal{F} = SE(2) \setminus \mathcal{Z}$ the *free configuration space*.

C. Energy-Bounded Caging

When the object can escape, we seek to quantify the amount of potential energy required for the object to escape. Let $\mathcal{E} : SE(2) \rightarrow \mathbb{R}$ be a potential energy function on the space of poses that is convex when restricted to \mathbb{R}^2 , such as gravity. Also define $\mathcal{E}^{-1}(X) = \{\mathbf{q} \in SE(2) \mid \mathcal{E}(\mathbf{q}) \in X\}$ for any subset $X \subseteq \mathbb{R}$. Given an energy threshold $u \in \mathbb{R}$, we denote by $\mathcal{Z}_u = \mathcal{Z} \cup \mathcal{E}^{-1}([u, \infty))$ the *u -energy forbidden space* and by $\mathcal{F}_u = SE(2) \setminus \mathcal{Z}_u$ the *u -energy admissible space*. Using the previous definitions, we formally introduce a the notion of an energy-bounded cage:

Definition 3.2: Let $\mathcal{E} : SE(2) \rightarrow \mathbb{R}$ be a function on the poses of \mathcal{O} relative to an initial pose $\mathbf{q}_0 \in SE(2)$. We call \mathcal{G} a *u -energy-bounded cage* of \mathcal{O} with respect to \mathcal{E} if the initial configuration $\mathbf{q}_0 \in SE(2)$ of \mathcal{O} lies in a compact path-connected component of \mathcal{F}_u .

When u can be arbitrarily large, we obtain the standard notion of caging of a polygonal object \mathcal{O} relative to a collection of fixed obstacle polygons \mathcal{G} [22]. Fig. 2 illustrates




Fig. 4: Illustration of EBCA-2D, our algorithm to lower bound the minimum escape energy u^* for an obstacle configuration \mathcal{G} (best viewed in color). 1) First, we randomly sample object poses in collision, embed the poses in \mathbb{R}^3 , and compute the penetration depth r_i for each pose, which corresponds to a ball within the collision space (red). 2) The union of balls centered at the samples conservatively approximates the collision space \mathcal{Z} . We then decompose the configuration space into cells by computing the weighted Delaunay triangulation D from the points and use the weighted α -shape \mathcal{A} at $\alpha = 0$ (white triangles) to approximate \mathcal{Z} . 3) Finally, we search for the smallest u such that the object can escape by finding a set of forbidden cells $\mathcal{V}_u \subseteq \mathcal{Z}_u$ and checking the connectivity of $D \setminus \mathcal{V}_u$. Blue and green indicate connected components, while yellow indicates poses such that $\mathcal{E}(\mathbf{q}) > u$.

both the classical notion of caging and a u -energy-bounded cage with respect to the gravitational potential energy \mathcal{E} . To rank energy-bounded cages, we introduce the notion of the minimum escape energy for an object \mathcal{O} and gripper configuration \mathcal{G} :

Definition 3.3: The *minimum escape energy* for an object \mathcal{O} and obstacle configuration \mathcal{G} , denoted u^* , is the smallest value of u such that \mathcal{G} is not a u -energy-bounded cage of \mathcal{O} . The rest of this work is dedicated to computing a lower bound on the minimum escape energy for a configuration \mathcal{G} and \mathcal{O} .

IV. METHODOLOGY

We now detail EBCA-2D, our algorithm for lower bounding the minimum escape energy for an object \mathcal{O} and obstacle configuration \mathcal{G} , which is illustrated in Fig. 4. We first generate N samples of object poses in collision $\mathcal{Q} = \{\mathbf{q}_1, \dots, \mathbf{q}_N\}$, embed the samples into \mathbb{R}^3 to form a set \mathcal{X} , and compute a conservative estimate of the penetration depth \mathcal{R} for each embedded pose. We then use weighted α -shapes to construct a cell decomposition of the convex hull of \mathcal{X} and mark forbidden cells that lie strictly within the collision space or above an energy threshold and examine the connectivity of the free cells to prove the non-existence of object escape paths [30]. We then use binary search to find the highest value of u for which for a no path exists in our cell decomposition, thus lower-bounding the minimum escape energy.

A. Verifying Cages in $SE(2)$

Given a set of N sampled poses in collision $\mathcal{Q} = \{\mathbf{q}_1, \dots, \mathbf{q}_N\}$ where each $\mathbf{q}_i \in SE(2)$, the first step of our algorithm is to embed the samples in \mathbb{R}^3 . Let \mathbf{z} be the center of rotation of \mathcal{O} and $\rho = \max_{\mathbf{v} \in \mathcal{O}} \|\mathbf{v} - \mathbf{z}\|_2$ be the

maximum moment arm of \mathcal{O} . Then let $\pi : \mathbb{R}^3 \rightarrow SE(2) = \mathbb{R}^2 \times \mathbb{S}^1$ be the covering map defined by $\pi(x, y, z) = (x, y, (z/\rho) \bmod 2\pi)$, for $(x, y, z) \in \mathbb{R}^3$. We map from poses to the covering space with an inverse map $\pi_n^{-1} : SE(2) \rightarrow \mathbb{R}^3$ defined by $\pi_n^{-1}(x, y, \theta) = (x, y, \rho\theta + 2\pi n)$ for $n \in \mathbb{Z}$ [7]. Given $R \in \mathbb{Z}$, a fixed number of rotations to embed, our lifted set of pose samples is $\mathcal{A} = \{\hat{\mathbf{q}}_{i,n} = \pi_n^{-1}(\mathbf{q}_i) \mid \mathbf{q}_i \in \mathcal{Q}, n \in \{-R, \dots, 0, \dots, R\}\}$.

We relate path existence in the covering space to cages in the configuration space by means of the following result [18]:

Theorem 4.1: Let $\mathcal{Y} \subset \mathbb{R}^3$ be a bounded subset. Let $\mathbf{q}_0 \in SE(2)$ such that $\mathbf{q}_0 \in \pi(\text{Conv}(\mathcal{Y})) \setminus \pi(\mathcal{Y})$ and let $\hat{\mathbf{q}}_0$ be any point such that $\pi(\hat{\mathbf{q}}_0) = \mathbf{q}_0$. If there exists no continuous path from $\hat{\mathbf{q}}_0 \in \mathbb{R}^3$ to $\partial\text{Conv}(\mathcal{Y}) \subset \mathbb{R}^3$, then $\mathbf{q}_0 \in SE(2)$ is caged by $\pi(\mathcal{Y})$ in $SE(2)$.

Proof: Suppose the contrary. Since \mathbf{q}_0 is not caged by $\pi(\mathcal{Y})$ in $SE(2)$ and \mathbb{S}^1 is compact, there exists a continuous escaping path $\gamma(t) : [0, 1] \rightarrow SE(2)$ such that $\gamma(0) = \mathbf{q}_0 = (x_0, y_0, \theta_0)$ and $\gamma(1) = (x_1, y_1, \theta_1)$ where $\|(x_0, y_0) - (x_1, y_1)\|_2 > \text{diam}(\text{Conv}(\pi(\mathcal{Y})))$ [7]. By the properties of the covering map π , there exists a lifting of γ to a covering path $\hat{\gamma} : [0, 1] \rightarrow \mathbb{R}^3$ with $\hat{\gamma}(0) = \hat{\mathbf{q}}_0$ and $\pi(\hat{\gamma}(t)) = \gamma(t)$ for all $t \in [0, 1]$, where $\|\hat{\gamma}(0) - \hat{\gamma}(1)\|_2 > \text{diam}(\text{Conv}(\mathcal{Y}))$. Hence by the continuity of $\hat{\gamma}(t)$ there exists a smallest $t_0 \in (0, 1)$ such that $\hat{\gamma}(t_0) \in \partial\text{Conv}(\mathcal{Y})$ and $\hat{\gamma}([0, t_0]) \subset \mathbb{R}^3 \setminus \mathcal{Y}$. This contradicts our supposition that no continuous path exists from $\hat{\mathbf{q}}_0$ to $\partial\text{Conv}(\mathcal{Y})$. ■

This result implies that a lifting of the u -energy forbidden space $\hat{\mathcal{Z}}_u \subset \mathbb{R}^3$ such that $\pi(\hat{\mathcal{Z}}_u) = \mathcal{Z}_u$ can be used to check the existence of energy-bounded cages.

B. Approximating the u -Energy Forbidden Space \mathcal{Z}_u

It remains to construct a conservative approximation of the lifted u -energy forbidden space $\mathcal{V}_u \subseteq \hat{\mathcal{Z}}_u$ and to computationally prove path non-existence in the lifted space, which would prove an energy-bounded cage by Lemma 3.1 and Theorem 4.1. We first approximate the lifted collision space $\hat{\mathcal{Z}}$ by a set \mathcal{B} using a conservative estimate of penetration depth, then decompose the convex hull of \mathcal{B} into cells using weighted α -shapes, and finally form \mathcal{V}_u from cells lying strictly within $\hat{\mathcal{Z}}_u$.

1) *Approximating the Collision Space Using Penetration Depth :* The 2D generalized penetration depth (GPD) $p : SE(2) \rightarrow \mathbb{R}$ between an object $\mathcal{O}(\mathbf{q}_i)$ in pose $\mathbf{q}_i = (x_i, y_i, \theta_i)$ and obstacle \mathcal{G} is defined as [51]:

$$p(\mathbf{q}_i) = \min_{\mathbf{q}_j \in SE(2)} \{d(\mathbf{q}_i, \mathbf{q}_j) \mid \text{int}(\mathcal{O}(\mathbf{q}_j)) \cap \mathcal{G} = \emptyset\}.$$

where $d : SE(2) \times SE(2) \rightarrow \mathbb{R}$ is a distance metric between poses. Following Zhang et. al [50], we use $d(\mathbf{q}_i, \mathbf{q}_j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} + \min_{m \in \mathbb{Z}} \rho|\theta_i - (\theta_j + 2\pi m)|$ which has the following important property:

Lemma 4.1: Let $r_i = r(\mathbf{q}_i) : SE(2) \rightarrow \mathbb{R}$ be an approximate solution to the above equation such that $r_i \leq p(\mathbf{q}_i)$ for all $\mathbf{q}_i \in \mathcal{C}$ and let $\mathbb{B}_r(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^3 : \|\mathbf{x} - \mathbf{y}\| \leq r\}$ be a standard Euclidean ball of radius r centered at $\mathbf{x} \in \mathbb{R}^3$. For




Fig. 5: Weighted α -shape construction and representation. (Top-left) The α -shape \mathcal{A} for $\alpha = 0$ is constructed from a set of Euclidean balls centered at points \mathcal{X} with radii $\mathcal{R} = \{r_1, \dots, r_N\}$ (Bottom-left) \mathcal{A} contains edges and triangles between the pairs and triplets of with a common intersection, respectively (bottom). (Top-right) As we increase the ball radius towards ∞ the set of balls becomes the power diagram of the point set, a generalization of the Voronoi diagram. (Bottom-right) The triangulation of the power diagram is the weighted Delaunay triangulation $D(\mathcal{X}, \mathcal{R})$ of the points, which contains the convex hull.

any embedded pose $\hat{\mathbf{q}}_i$, if $\pi(\hat{\mathbf{q}}_i) \in \mathcal{Z}$, then any embedded pose $\hat{\mathbf{q}}_j \in \mathbb{B}_{r_i}(\hat{\mathbf{q}}_i)$ also satisfies $\pi(\hat{\mathbf{q}}_j) \in \mathcal{Z}$.

A detailed version of the proof is given in the supplemental file at <http://berkeleyautomation.github.io/caging/>, and a similar proof is given by Zhang et al. [50]. For our set of pose samples $\mathcal{Q} \subset SE(2)$ with an associated lifting $\mathcal{X} \subset \mathbb{R}^3$ and associated GPD values $\mathcal{R} = \{r_{i,n} = r(\mathbf{q}_i) \mid \mathbf{q}_i \in \mathcal{Q}, n \in \{-R, \dots, 0, \dots, R\}\}$, define

$$\mathcal{B}(\mathcal{X}, \mathcal{R}) = \bigcup_{\mathbf{x} \in \mathcal{X}} \mathbb{B}_{r_{i,n}}(\hat{\mathbf{q}}_{i,n}).$$

It follows from Lemma 4.1 is that $\pi(\mathcal{B}(\mathcal{X}, \mathcal{R})) \subseteq \mathcal{Z}$.

In order to satisfy $r_i \leq p(\mathbf{q}_i)$, we use an algorithm by Zhang et. al [51] to lower bound the GPD between any two objects. The algorithm assumes a given convex decomposition of the two bodies [26], then computes the exact GPD between all possible pairs of convex pieces and takes the maximum over the GPD values between the pieces. The GPD between two convex bodies can be determined using the Gilbert-Johnson-Keerthi Expanding Polytope Algorithm (GJK-EPA) developed by van den Bergen [44] and implemented in libccd [16]. Each run of the GJK-EPA algorithm is $O(V)$ but can be modified to be near constant-time [6], where $V = V_o + V_g$ is the total number of vertices between \mathcal{O} and \mathcal{G} . There are up to $O(V^2)$ convex bodies to check, and thus the complexity of computing GPD is $O(V^3)$.

2) *Weighted α -Shapes:* Weighted α -shapes [11], [12], [13], illustrated in Fig. 5, represent a union of balls of varying radii by means of a simplicial complex whose vertices are given by the ball's centers. We use weighted α -shapes to decompose the configuration space into cells and to determine which cells belong strictly to the lifted collision space $\hat{\mathcal{Z}}$ and energy-bounded space $\hat{\mathcal{E}}^{-1}([u, \infty))$.

Weighted α -shapes are a type of simplicial complex [12], a key data-structure to represent a large collection of geometrically interesting spaces that generalize the notion of a graph

and a triangulation. Let $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_P\} \subset \mathbb{R}^3$ be a point set and $\mathcal{R} = \{r_1, \dots, r_P\}$ be positive scalars for each element of \mathcal{X} such that any subset of 4 points of \mathcal{X} are affinely independent. This is a weak condition since for uniformly sampled points this occurs with probability one [13]. The weighted Delaunay triangulation (WDT) of \mathcal{X} and \mathcal{R} is $D(\mathcal{X}, \mathcal{R}) = \{\sigma = \{\mathbf{x}_0, \dots, \mathbf{x}_k\} \mid \cap_{i=1}^k V_{\mathbf{x}_i}(\mathcal{X}, \mathcal{R}) \neq \emptyset \text{ and } 0 \leq k \leq 3\}$, where $V_{\mathbf{x}_i}(\mathcal{X}, \mathcal{R}) = \{\mathbf{y} \mid \|\mathbf{x}_i - \mathbf{y}\|^2 - r_i^2 \leq \|\mathbf{x}_j - \mathbf{y}\|^2 - r_j^2, \forall j \in \{1, \dots, P\}\}$ is the weighted Voronoi region for \mathbf{x}_i . The union of all simplices in $D(\mathcal{X}, \mathcal{R})$ is the convex hull $\text{Conv}(\mathcal{X})$ and when $r_i = 0$ for all i this reduces to the standard Delaunay triangulation of \mathcal{X} . The weighted α -shape $\mathcal{A} = \mathcal{A}(\mathcal{X}, \mathcal{R})$ at $\alpha = 0$ is a particular simplicial subcomplex of $D(\mathcal{X}, \mathcal{R})$ with several important properties:

Theorem 4.2 (Edelsbrunner et. al): Let $\mathcal{B}(\mathcal{X}, \mathcal{R}) = \bigcup_{i=1}^P \mathbb{B}_{r_i}(\mathbf{x}_i)$. Then $\mathcal{A}(\mathcal{X}, \mathcal{R})$ is homotopy equivalent to $\mathcal{B}(\mathcal{X}, \mathcal{R})$ and any k -simplex $\sigma = \{\mathbf{x}_{i,0}, \dots, \mathbf{x}_{i,k}\}$ in \mathcal{A} such that $0 \leq k \leq 3$ is completely contained in the union of balls $\bigcup_{j=0}^k \mathbb{B}_{r_{i,j}}(\mathbf{x}_{i,j})$.

This implies that if we use the set of embedded poses \mathcal{X} with radii given by the penetration depths \mathcal{R} , then the weighted alpha shape $\mathcal{A}(\mathcal{X}, \mathcal{R}) \subset \hat{\mathcal{Z}}$ and can be used to verify cages by Lemma 3.1 and Theorem 4.1.

3) *Approximating the Potential Energy Superlevel Set:* It remains to find a subcomplex of $D(\mathcal{X}, \mathcal{R})$ such that $\pi(\mathbf{x}) \in \mathcal{E}^{-1}([u, \infty))$ for any \mathbf{x} in the subcomplex.

Lemma 4.2: For any k -simplex $\sigma \in D(\mathcal{X}, \mathcal{R})$, let $\mathcal{E}(\sigma) = \min_{\mathbf{x} \in \sigma} \mathcal{E}(\pi(\mathbf{x}))$. Furthermore, let $\mathcal{P}_u(\mathcal{X}, \mathcal{R}) = \{\sigma \in D(\mathcal{X}, \mathcal{R}) \mid \mathcal{E}(\sigma) > u\}$. Then $\mathcal{P}_u(\mathcal{X}, \mathcal{R})$ is a subcomplex of $D(\mathcal{X}, \mathcal{R})$ and $\mathcal{P}_u(\mathcal{X}, \mathcal{R}) \subseteq \mathcal{E}^{-1}([u, \infty))$.

Proof: Fix a simplex $\sigma_k \in \mathcal{P}_u(\mathcal{X}, \mathcal{R})$. Any face σ_j of σ_k is also a member of $\mathcal{P}_u(\mathcal{X}, \mathcal{R})$ by the minimum over the energy function. Furthermore, by the convexity of \mathcal{E} when restricted to \mathbb{R}^2 , any point \mathbf{y} in $\text{Conv}(\sigma_k)$ satisfies $\mathcal{E}(\pi(\mathbf{y})) \geq \mathcal{E}(\sigma_k) \geq u$, and therefore $\mathbf{y} \in \mathcal{E}^{-1}([u, \infty))$. ■

A result of this Lemma and Theorem 4.2 is that the *u-energy forbidden subcomplex* \mathcal{V}_u satisfies:

$$\mathcal{V}_u(\mathcal{X}, \mathcal{R}) = \mathcal{A}(\mathcal{X}, \mathcal{R}) \cup \mathcal{P}_u(\mathcal{X}, \mathcal{R}) \subset \hat{\mathcal{Z}}_u.$$

C. Verifying Path Non-Existence

We can now verify *u-energy-bounded cages* by showing that no path exists from the embedding of the object pose $\hat{\mathbf{q}}_0$ to $\partial D(\mathcal{X}, \mathcal{R})$ in $D(\mathcal{X}, \mathcal{R}) \setminus \mathcal{V}_u(\mathcal{X}, \mathcal{R})$ by Theorem 4.1. We use Algorithm 1, a modified version of the algorithm by McCarthy et al. [30], to verify that no escape paths exist. As shown by McCarthy et al. [30], the worst-case runtime to verify path non-existence is $O(N^2)$, where N is the number of sampled points, and is dominated by the construction of the weighted Delaunay triangulation $D(\mathcal{X}, \mathcal{R})$. Given $D(\mathcal{X}, \mathcal{R})$, Algorithm 1 takes $O(N)$ time in the worst case because each simplex in $D(\mathcal{X}, \mathcal{R})$ must be classified to construct a disjoint set structure.

Theorem 4.3: If \mathcal{V}_u is any subcomplex of $D(\mathcal{X}, \mathcal{R})$ in \mathbb{R}^3 such that $\hat{\mathbf{q}}_0 \in \text{Conv}(\mathcal{X}) \setminus \mathcal{V}_u$ and Algorithm 1 returns True, then there exists no continuous path from $\hat{\mathbf{q}}_0$ to $\partial \text{Conv}(\mathcal{X})$ in $D(\mathcal{X}, \mathcal{R}) \setminus \mathcal{V}_u$.

The proof is a slight modification of the main Theorem of [30] and is given in the supplemental file. Therefore, if Algorithm 1 returns true when run with \mathcal{V}_u as defined in Section IV-B.3, then we are guaranteed that \mathcal{V}_u forms a *u-energy bounded cage* of $\hat{\mathbf{q}}_0$.

```

1 Input: Lifted initial pose  $\hat{\mathbf{q}}_0$ , weighted Delaunay triangulation
    $D(\mathcal{X}, \mathcal{R})$ , u-Energy Forbidden Subcomplex  $\mathcal{V}_u$ 
Result: True if  $\mathcal{V}_u$  cages  $\mathcal{O}$  in pose  $\pi(\mathbf{x})$ , False otherwise
// Init free subcomplex and boundary
2  $\mathcal{U} = \{\sigma_i \mid \sigma_i \in D(\mathcal{X}, \mathcal{R}) \setminus \mathcal{V}_u, |\sigma_i| = 3\}$ ;
3  $\mathcal{W} = \{\sigma_j \mid \sigma_j \in \partial D(\mathcal{X}, \mathcal{R}) \setminus \mathcal{V}_u, |\sigma_j| = 2\}$ ;
// Compute connected components
4  $\mathcal{Q} = \text{DisjointSetStructure}(\mathcal{U} \cup \mathcal{W})$ ;
5 for  $\sigma_i \in \mathcal{U} \cup \mathcal{W}$  do
6   | for  $\sigma_j \in \text{Neighbors}(\sigma_i, D(\mathcal{X}, \mathcal{R}) \setminus \mathcal{V}_u)$  do
7   |   | if  $\sigma_i \cap \sigma_j \notin \mathcal{V}_u$  then
8   |   |   |  $\mathcal{Q}.\text{UnionSets}(\sigma_i, \sigma_j)$ ;
9   |   end
10 end
// Check connectivity
11  $\sigma_0 = \text{Locate}(\hat{\mathbf{q}}_0, D(\mathcal{X}, \mathcal{R}))$ ;
12 for  $\sigma_i \in \mathcal{W}$  do
13   | if  $\mathcal{Q}.\text{SameSet}(\sigma_0, \sigma_i)$  then
14   |   | return False;
15 end
16 return True;

```

Algorithm 1: Verifying *u*-Energy-Bounded Cages

D. Lower-Bounding the Minimum Escape Energy

We determine a lower bound to u^* by searching over values of u that form an energy-bounded cage. EBCA-2D, our full algorithm for computing the a lower bound, is given in Algorithm 2. EBCA-2D generates N samples of poses in collision \mathcal{Q} with penetration depths \mathcal{R} over the collision space using rejection sampling, embeds the poses in \mathbb{R}^3 using to form a set \mathcal{X} , constructs a weighted Delaunay triangulation $D(\mathcal{X}, \mathcal{R})$ and alpha shape $\mathcal{A}(\mathcal{X}, \mathcal{R})$ from the samples, and finds an approximation \hat{u} to u^* using binary search, where on each iteration we check for an energy-bounded cage using Algorithm 1. The complexity of Algorithm 2 is $O(N^2 + N \log(1/\Delta) + NV^3)$, where the $O(N^2)$ term is due to constructing $D(\mathcal{X}, \mathcal{R})$, the $O(N \log(1/\Delta))$ term is due to running Algorithm 1 for every iteration of binary search [30], and the $O(NV^3)$ term is due to the computation of the GPD for N pose samples.

Theorem 4.4: Let u^* denote the minimum escape energy for object \mathcal{O} and gripper configuration \mathcal{G} . Let \hat{u} be the result of running Algorithm 2 with \mathcal{O} and \mathcal{G} . Then $\hat{u} \leq u^*$.

Proof: By Lemma 3.1 and the subset properties of $\mathcal{A}(\mathcal{X}, \mathcal{R})$ from Lemma 4.1 and Theorem 4.2 we are guaranteed that if our algorithm terminates when checking $u = \infty$, then the object is truly caged. It remains to show that for all iterations of the binary search, the gripper configuration \mathcal{G} is a u_ℓ -energy-bounded cage. This is true for iteration 0, as the initial value satisfies $u_\ell \leq \mathcal{E}(\mathbf{q}_0)$. Furthermore, if the lower bound is updated to $u_\ell = u_m$ then \mathcal{G} is a u_m -energy-bounded cage of \mathcal{O} by Theorem 4.3, Lemma 4.1, and Lemma 4.2. ■

```

1 Input: Polygonal robot gripper  $\mathcal{G}$ , Polygonal object  $\mathcal{O}$ ,
   Number of pose samples  $N$ , Number of rotations  $R$  for
    $SE(2)$  lifting, Binary search resolution  $\Delta$ 
   Result:  $\hat{u}$ , a lower bound on the minimum escape energy  $u^*$ 
   // Sample poses in collision
2  $\mathcal{Q} = \emptyset, \mathcal{R} = \emptyset, \ell = \text{diam}(\mathcal{G}) + \text{diam}(\mathcal{O})$ ;
3  $\mathcal{W} = [-\ell, \ell] \times [-\ell, \ell] \times [0, 2\pi]$ ;
4 for  $i \in \{1, \dots, N_s\}$  do
5    $\mathbf{q}_i = \text{RejectionSample}(\mathcal{W})$ ;
6    $r_i = \text{LowerBoundPenDepth}(\mathbf{q}_i, \mathcal{O}, \mathcal{G})$ ;
7   if  $r_i > 0$  then
8      $\mathcal{Q} = \mathcal{Q} \cup \{\mathbf{q}_i\}, \mathcal{R} = \mathcal{R} \cup \{r_i\}$ ;
9 end
10  $\mathcal{X} = \{\pi_n^{-1}(\mathbf{q}_i) \mid \mathbf{q}_i \in \mathcal{Q}, n \in \{-N_r, \dots, N_r\}\}$ ;
   // Create alpha shape
11  $D(\mathcal{X}, \mathcal{R}) = \text{WeightedDelaunayTriangulation}(\mathcal{X}, \mathcal{R})$ ;
12  $\mathcal{A}(\mathcal{X}, \mathcal{R}) = \text{WeightedAlphaShape}(D(\mathcal{X}, \mathcal{R}), \alpha = 0)$ ;
   // Binary search for min escape energy
13 if  $\text{EnergyBoundedCage}(\mathbf{q}_0, D(\mathcal{X}, \mathcal{R}), \mathcal{A}(\mathcal{X}, \mathcal{R}))$  then
14   return  $\infty$ ;
15  $u_\ell = \min \mathcal{E}(\sigma_k)$  such that  $\sigma_k \in D(\mathcal{X}, \mathcal{R})$ ;
16  $u_u = \max \mathcal{E}(\sigma_k)$  such that  $\sigma_k \in D(\mathcal{X}, \mathcal{R})$ ;
17 while  $|u_u - u_\ell| > \Delta$  do
18    $u_m = 0.5 * (u_\ell + u_u)$ ;
19    $\mathcal{V}_{u_m} = \text{ForbiddenSubcomplex}(D(\mathcal{X}, \mathcal{R}), \mathcal{A}(\mathcal{X}, \mathcal{R}), u_m)$ ;
20   if  $\text{EnergyBoundedCage}(\mathbf{q}_0, D(\mathcal{X}, \mathcal{R}), \mathcal{V}_{u_m})$  then
21      $u_\ell = u_m$ ;
22   else
23      $u_u = u_m$ ;
24 end
25 return  $u_\ell$ ;

```

Algorithm 2: Energy-Bounded-Cage-Analysis-2D

V. EXPERIMENTS

To test our methods, we implemented EBCA-2D in C++ and evaluated the performance on a set of polygonal objects under a gravitational potential energy field. We used the CGAL library [8] to compute triangulations and α -shapes. For GPD computation we performed a convex decomposition of polygons using the algorithm of Lien et al. [26] and libccd [16] for the GJK-EPA algorithm. All experiments ran on a desktop with an Intel Core i7-4770K 350 GHz processor with 6 cores.

A. Energy-Bounded Cages Under Gravity

We ran our algorithm with $N = 200,000$ pose samples for varying obstacle configurations on a set of six polygonal parts. The parts were created by projecting 3D models from the YCB dataset [5] and 3DNet [49] onto a plane and triangulating the projection. We assumed a uniform mass density of $0.01\text{kg}/\text{cm}^2$ for each object, which we used to compute the mass M for each object. Each run of the algorithm took approximately 180 seconds to run, and more details on runtime can be found in Section V-B.

Fig. 6 shows the estimated normalized minimum escape energy $\hat{u}_n = \hat{u}/(Mg)$ for three parallel-jaw gripper configurations on each of three objects. To aid in visualization, we used RRT* implemented in OMPL [40] to plan an escape path to directly below the initial object pose, and we rendered the object in the pose along the solution path with maximum potential energy. The ranking of grasps by minimum escape




Fig. 6: Three example polygonal parts (blue) with three parallel-jaw configurations (black) for each object. Below each configuration is normalized minimum escape energy $\hat{u}_n = \hat{u}/Mg$ estimated by EBCA-2D with $N = 200,000$ pose samples under gravity, where M is the mass of the part. To visualize the output of EBCA-2D, we render the object translucently at the the highest point of an escape path found by an RRT* planner, with an arrow to indicate direction. We see that \hat{u}_n , which is the estimated minimum height that must be reached to escape, ranks the configurations for each object in the same order as the maximum height reached along the RRT* escape path.

energy matches our intuition, and appears to also match the ranking of grasps by the maximum energy reached along the RRT* visualization path. To evaluate the lower bound of Theorem 4.4, we also used RRT* to attempt to plan an object escape path over the set of collision-free poses with energy less than \hat{u} . In every case, the RRT* planner was not able to find an escape path with energy less than \hat{u} within 120 seconds.

We also ran our algorithm on a set of configurations with more than two nonconvex obstacles. Fig. 7 displays \hat{u}_n for four examples: capturing an object using a single rectangular jaw and ramp, bowl-shaped jaws pinning an object against a vertical wall, three rectangular jaws, and a robotic gripper on a doorknob inspired by [10]. Our algorithm is able to prove cages for configurations 3 and 4, and the ranking of configurations 1 and 2 by \hat{u}_n matches our intuition. Again, RRT* was not able to find an escape path within the set of collision-free poses with energy less than \hat{u} .

B. Sensitivity to Number of Pose Samples

We also studied the sensitivity of \hat{u} and the total runtime to the number of pose samples N used to approximate the collision space. The left panel of Fig. 8 shows the ratio of \hat{u} at $N = \{6.25, 12.5, 25, 50, 100, 200, 400\} \times 10^3$ to \hat{u} at $N = 400,000$ for configuration 1 for each of the objects in Fig. 6. Each ratio is averaged over 5 independent trials per value of N to smooth the effects of random initializations. We see that for less than about 25,000 samples the output tends to be $\hat{u} \approx 0$ because the collision space is not well-approximated, leading to “holes” in the algorithm’s representation of the collision space for lower y -coordinates. However, as N




Fig. 7: Four example configurations of polygonal parts (blue) and obstacle configurations (black) with varying shape and number of components. Under each configuration is the normalized minimum escape energy \hat{u}_n estimated by EBCA-2D with $N = 200,000$. We see that EBCA-2D verifies that configurations 3 and 4 are cages, both of which are challenging because of the nonconvexity of the parts and the nonconvex obstacle in configuration 4.




Fig. 8: (Left) The ratio of the minimum escape energy \hat{u} estimated by EBCA-2D for N pose samples to \hat{u} at $N = 400,000$ for $N = \{6.25, 12.5, 25, 50, 100, 200, 400\} \times 10^3$. The values plotted are for configuration 1 of each object in Fig. 6 and are averaged over 5 independent trials per value of N . We see that the object B, the “fattest” object, converges the fastest and object C, the object with the “thinnest” pieces, converges the slowest. (Right) The runtime of EBCA-2D in seconds broken down by component of the algorithm versus the number of pose samples $N = \{6.25, 12.5, 25, 50, 100, 200, 400\} \times 10^3$. Each datapoint is averaged over 5 independent trials per value of N and configuration 1 of the objects in Fig. 6. The scaling of the average runtime is approximately linear in N , and the runtime becomes dominated by the time to generate pose samples for large N .

becomes large, \hat{u} converges towards a nonzero value. Interestingly, object B, which is the “fattest” [45] converges the fastest, taking only about 50,000 samples to converge to within 90% of its value at $N = 400,000$. On the other hand, object C takes nearly 200,000 to converge to within 90% of its value at $N = 400,000$, as u^* depends on the very thin tips of the object in this configuration. This is possibly because u^* for object C depends on a very thin part of the collision space, which requires more samples to approximate.

The right panel of Fig. 8 shows the scaling of the runtime in seconds versus the number of pose samples N averaged over 5 independent trials for configuration 1 of the objects in Fig. 6. The runtime is broken down by component of the algorithm: pose sampling, approximating the configuration space using α -shapes, and the binary search over potential energies. We see that the total runtime for these shapes and obstacles is approximately linear in the number of pose samples N , with pose sampling taking the largest portion of the runtime. However, the amount of time to sample poses and the time to construct an approximation to the configuration space both appear to be slightly superlinear in N . These results suggest that runtime remains well below the worst case N^2 scaling with the number of samples in practice.

VI. DISCUSSION AND FUTURE WORK

We defined energy-bounded caging configurations and the minimum escape energy, or the minimum energy that external perturbations must exert on an object for it to escape a set of obstacles. We also developed Energy-Bounded-Cage-Analysis-2D (EBCA-2D), an algorithm to compute a lower bound on the minimum escape energy for 2D polygonal objects and obstacles using weighted α -shapes. Our experiments demonstrate that we are able to verify cages and suggest that our algorithm returns an intuitive lower bound for a set of nonconvex polygonal objects and gripper configurations.

Future work will investigate extensions of our algorithm to synthesize obstacle configurations that form energy-bounded cages and to analyze 3D objects and obstacles. One barrier to using our algorithm for synthesis is the runtime for analyzing a single configuration, which is largely dominated by pose sampling and GPD computation. To reduce runtime, future work will study adaptive sampling procedures to approximate the thin parts of the collision space with fewer samples, such as Gaussian sampling from motion planning [4], parallel implementations of sampling using Cloud-based implementations, and a dual approach to upper bound u^* using optimal motion planning or the connectivity of the free space. Furthermore, we will study alternative approaches to synthesis, such as searching over the full obstacle configuration space to generate all obstacle placements that require at least a given amount of energy for escape [37], [42].

While in principle the theory behind our approach can be generalized to 3D, a challenge for synthesizing and analyzing configurations in 3D is the increase in dimensionality of the configuration space from 3D to 6D. This increases the computational load to construct dense α -shapes [13] and may also increase the number of samples needed to approximate the configuration space. In future work, we will investigate alternative representations of the forbidden space such as Vietoris-Rips complexes [19], a sparser simplicial complex representation of point samples, or precomputed simplicial complexes that cover the configuration space [46]. Another difficulty is that scaling to 3D would require an embedding of $SE(3)$ into \mathbb{R}^6 , which is more challenging due to the topology of $SE(3)$ [7] and because no implementation of higher dimensional weighted α -shapes exists in common software such as CGAL [41]. Thus, we will also study Cloud-based construction of α -shapes.

VII. ACKNOWLEDGMENTS

This work is supported in part by the U.S. National Science Foundation under Award IIS-1227536, and by grants from Google and Cisco. We thank our colleagues who gave feedback and suggestions, in particular Animesh Garg, James Kuffner, Michael Laskey, Sanjay Krishnan, Stephen McKinley, Jia Pan, Sachin Patil, Nan Tian, and Jur van den Berg.

REFERENCES

- [1] T. F. Allen, E. Rimon, and J. W. Burdick, "Two-finger caging of 3d polyhedra using contact space search," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2014, pp. 2005–2012.
- [2] J. Basch, L. J. Guibas, D. Hsu, and A. T. Nguyen, "Disconnection proofs for motion planning," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 2. IEEE, 2001, pp. 1765–1772.
- [3] A. Bicchieri and V. Kumar, "Robotic grasping and contact: A review," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. Citeseer, 2000, pp. 348–353.
- [4] V. Boor, M. H. Overmars, et al., "The gaussian sampling strategy for probabilistic roadmap planners," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 2. IEEE, 1999, pp. 1018–1023.
- [5] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols," *arXiv preprint arXiv:1502.03143*, 2015.
- [6] S. Cameron, "A comparison of two fast algorithms for computing the distance between convex polyhedra," 1996.
- [7] M. Caroli and M. Teillaud, "Computing 3d periodic triangulations," in *Algorithms-ESA 2009*. Springer, 2009, pp. 59–70.
- [8] T. K. F. Da, S. Lorient, and M. Yvinec, "3D alpha shapes," in *CGAL User and Reference Manual*, 4.6.2 ed. CGAL Editorial Board, 2015. [Online]. Available: <http://doc.cgal.org/4.6.2/Manual/packages.html#PkgAlphaShapes3Summary>
- [9] C. Davidson and A. Blake, "Caging planar objects with a three-finger one-parameter gripper," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 1998, pp. 2722–2727.
- [10] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner, "Manipulation planning with caging grasps," in *IEEE Conf. on Humanoid Robots*. IEEE, 2008, pp. 285–292.
- [11] H. Edelsbrunner, *Weighted alpha shapes*. University of Illinois at Urbana-Champaign, Department of Computer Science, 1992.
- [12] —, "The union of balls and its dual shape," in *Proceedings of the ninth annual symposium on Computational geometry*. ACM, 1993, pp. 218–231.
- [13] H. Edelsbrunner and E. P. Mücke, "Three-dimensional alpha shapes," *ACM Transactions on Graphics (TOG)*, vol. 13, no. 1, pp. 43–72, 1994.
- [14] J. Erickson, S. Thite, F. Rothganger, and J. Ponce, "Capturing a convex object with three discs," *IEEE Trans. Robotics*, vol. 23, no. 6, pp. 1133–1140, 2007.
- [15] C. Ferrari and J. Canny, "Planning optimal grasps," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 1992, pp. 2290–2295.
- [16] D. Fiser, "libccd - collision detection between convex shapes," <http://libccd.danfis.cz/>.
- [17] K. Y. Goldberg and M. T. Mason, "Bayesian grasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 1990, pp. 1264–1269.
- [18] A. Hatcher, *Algebraic topology*. Cambridge University Press, 2002.
- [19] J.-C. Hausmann et al., "On the vietoris-rips complexes and a cohomology theory for metric spaces," *Ann. Math. Studies*, vol. 138, pp. 175–188, 1995.
- [20] H. Kruger, E. Rimon, and A. F. van der Stappen, "Local force closure," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 4176–4182.
- [21] H. Kruger and A. F. van der Stappen, "Partial closure grasps: Metrics and computation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2011, pp. 5024–5030.
- [22] W. Kuperberg, "Problems on polytopes and convex sets," in *DIMACS Workshop on polytopes*, 1990, pp. 584–589.
- [23] M. Laskey, J. Mahler, Z. McCarthy, F. Pokorny, S. Patil, J. van den Berg, D. Kragic, P. Abbeel, and K. Goldberg, "Multi-arm bandit models for 2d sample based grasp planning with uncertainty," in *Proc. IEEE Conf. on Automation Science and Engineering (CASE)*. IEEE, 2015.
- [24] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [25] Z. Li and S. S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," *IEEE Journal of Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988.
- [26] J.-M. Lien and N. M. Amato, "Approximate convex decomposition of polygons," in *Proceedings of the twentieth annual symposium on Computational geometry*. ACM, 2004, pp. 17–26.
- [27] S. Liu and S. Carpin, "A fast algorithm for grasp quality evaluation using the object wrench space," 2015.
- [28] J. Mahler, S. Patil, B. Kehoe, J. van den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "Gp-gpis-opt: Grasp planning under shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2015.
- [29] T. Makapunyo, T. Phoka, P. Pipattanasomporn, N. Niparnan, and A. Sudsang, "Measurement framework of partial cage quality based on probabilistic motion planning," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2013, pp. 1574–1579.
- [30] Z. McCarthy, T. Bretl, and S. Hutchinson, "Proving path non-existence using sampling and alpha shapes," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 2563–2569.
- [31] P. Pipattanasomporn and A. Sudsang, "Two-finger caging of nonconvex polytopes," *IEEE Trans. Robotics*, vol. 27, no. 2, pp. 324–333, 2011.
- [32] D. Prattichizzo and J. C. Trinkle, "Grasping," in *Springer handbook of robotics*. Springer, 2008, pp. 671–700.
- [33] E. Rimon and A. Blake, "Caging 2d bodies by 1-parameter two-fingered gripping systems," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 1996, pp. 1458–1464.
- [34] —, "Caging planar bodies by one-parameter two-fingered gripping systems," *Int. J. Robotics Research (IJRR)*, vol. 18, no. 3, pp. 299–318, 1999.
- [35] E. Rimon and J. Burdick, "Mobility of bodies in contact. i. a new 2nd order mobility index for multiple-finger grasps," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 1994, pp. 2329–2335.
- [36] A. Rodriguez and M. T. Mason, "Two finger caging: squeezing and stretching," *Algorithmic Foundation of Robotics VIII*, vol. 57, 2010.
- [37] A. Rodriguez, M. T. Mason, and S. Ferry, "From caging to grasping," *Int. J. Robotics Research (IJRR)*, pp. 886–900, 2012.
- [38] J. A. Stork, F. T. Pokorny, and D. Kragic, "Integrated motion and clasp planning with virtual linking," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2013, pp. 3007–3014.
- [39] J. Su, H. Qiao, Z. Ou, and Z.-Y. Liu, "Vision-based caging grasps of polyhedron-like workpieces with a binary industrial gripper," 2015.
- [40] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, December 2012, <http://ompl.kavrakilab.org>.
- [41] The CGAL Project, *CGAL User and Reference Manual*, 4th ed. CGAL Editorial Board, 2015.
- [42] M. Vahedi and A. F. van der Stappen, "Caging polygons with two and three fingers," *Int. J. Robotics Research (IJRR)*, vol. 27, no. 11-12, pp. 1308–1324, 2008.
- [43] —, "On the complexity of the set of three-finger caging grasps of convex polygons," in *Proc. Robotics: Science and Systems (RSS)*. Citeseer, 2009.
- [44] G. Van Den Bergen, "Proximity queries and penetration depth computation on 3d game objects," in *Game developers conference*, vol. 170, 2001.
- [45] A. F. van der Stappen and M. H. Overmars, "Motion planning amidst fat obstacles," in *Proceedings of the tenth annual symposium on Computational geometry*. ACM, 1994, pp. 31–40.
- [46] W. Wan, R. Fukui, M. Shimosaka, T. Sato, and Y. Kuniyoshi, "A new grasping by caging solution by using eigen-shapes and space mapping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2013, pp. 1566–1573.
- [47] —, "Grasping by caging: A promising tool to deal with uncertainty," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 5142–5149.
- [48] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 557–562.
- [49] W. Wohlkinger, A. Aldoma, R. B. Rusu, and M. Vincze, "3dnet: Large-scale object class recognition from cad models," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 5384–5391.
- [50] L. Zhang, Y. J. Kim, and D. Manocha, "Efficient cell labelling and path non-existence computation using c-obstacle query," *Int. J. Robotics Research (IJRR)*, vol. 27, no. 11-12, pp. 1246–1257, 2008.
- [51] L. Zhang, Y. J. Kim, G. Varadhan, and D. Manocha, "Generalized penetration depth computation," *Computer-Aided Design*, vol. 39, no. 8, pp. 625–638, 2007.