Uncertainty Envelopes *

Yaron Ostrovsky-Berman   Leo Joskowicz

Abstract

We introduce a new class of problems: computing the set of all the convex combinations of sites when their position is uncertain and depends linearly on shared parameters which vary according to a uniform distribution. The boundary of the set, called the uncertainty envelope, is useful in optimizing processes where there is uncertainty on the site positions and the objective functions. We provide upper bounds on the combinatorial complexity of the uncertainty envelope, and present the first efficient algorithm for its computation in the general case.

1   Introduction

Suppose an investor wishes to spend his budget on importing several products with fixed attributes (market value, import tax, maintenance, transportation costs, ...). Choosing a combination of products which maximizes his profit function is a standard optimization problem. However, when the product attributes are subject to change due to variation in market parameters (currency exchange rates, fuel and insurance costs, ...), there is uncertainty involved in any investment. The problem then becomes that of finding the best tradeoff between profit and risk.

This scenario is common to many fields involving uncertainty or change with the following model properties:

- *proportionality* - the attributes of the products are proportional to their portion of the budget.
- *divisibility* - the portion of the budget for each product is allowed to assume any fractional value.
- *linearity* - the attributes depend linearly on the market parameters.
- *independency* - the market parameters vary independently.

For optimization problems, the following additional properties hold:

- *uniformity* - the parameters vary according to a uniform distribution with fixed intervals.
- *multiple fuzzy objective functions* - there are multiple choices of the objective function form, where function coefficients have uncertainty that varies in a range.

*School of Engineering and Computer Science, The Hebrew University of Jerusalem, {yaronber,josko}@cs.huji.ac.il

We formalize the problem as follows. Let $S = \{s_1, s_2, \ldots, s_n\}$ be a set of sites (products, materials, components), and let $p = (p_1, p_2, \ldots, p_m)^T$ be the market parameters uncertainty vector, which is limited to the domain $\Delta = [-\delta_1, \delta_1] \times [-\delta_2, \delta_2] \times \ldots \times [-\delta_m, \delta_m]$. Each site s_i is associated with a nominal (current) value $b_i \in \mathbb{R}^d$ and an $m \times d$ sensitivity matrix A_i describing the attribute sensitivity to variations in p. The position of site s_i is given by $v_i(p) = A_i p + b_i$. We define the *uncertainty envelope* as the boundary of the union of the convex hulls of the sites when the uncertainty vector spans its domain:

$$\mathcal{E} = \partial\{\sum_{i=1}^{n} \lambda_i v_i(p) | \lambda_i \geq 0, \sum_{i=1}^{n} \lambda_i = 1, p \in \Delta\}$$

The uncertainty envelope provides a global view of the problem domain which is most useful in analyzing worst case variation of unknown or uncertain combination of the sites. Since there are possibly several objective functions for optimization, and each function may have uncertainty, the envelope allows sensitivity analysis on each of the solutions.

The class of problems we study depends on two key parameters: the number of sites n and the dimension d. In the following special cases, the problem has known solutions. When there is no uncertainty ($\Delta = \{0\}$), the envelope bounds the convex hull of $\{b_i\}_{i=1}^{n}$, the domain of the standard optimization problem. In one dimension ($d = 1$, any n), the envelope is a segment whose endpoints are determined by the maximum and minimum of $\{b_i \pm a_j^i \delta_j\}$, where a_j^i is the j^{th} entry of A_i. When there is only one site ($n = 1$, any d), the envelope bounds the d-zonotope defined by the sensitivity matrix A and the uncertainty domain Δ. Zonotopes play a crucial role in analyzing and computing uncertainty envelopes. When the sensitivity matrices of the sites are independent, that is the indices of their non-zero column vectors are mutually exclusive, the uncertainty envelope is the boundary of the convex hull of the uncertainty zonotopes of the sites. In the general case ($n, d \geq 2$), the envelope is the boundary of the volume swept by the zonotope as it spans the convex combinations of all the sites. For a solution to the special case $d = 2$, $n = 2$ see [4].

In this extended abstract, we provide upper bounds on the combinatorial complexity of uncertainty envelopes and present the first algorithm for their computation in the general case. The results are sum-

$dim \backslash \#sites$		1	2	3	$n > 3$
1	comb	$O(1)$			$O(1)$
	time	$O(m)$			$O(nm)$
2	comb	$\Theta(m)$	$O(\lambda_6(m^2))$		$O(n^2\lambda_6(n^2m^2))$
	time	$O(m\log m)$	$O(\lambda_6(m^2)\log(m))$		$O(n^2\lambda_6(n^2m^2))$
3	comb	$\Theta(m^2)$	$O(m^{6+\epsilon})$	$O(m^{8+\epsilon})$	$O((n^3m^4)^{2+\epsilon})$
	time	$\Theta(m^2)$	$O(m^{6+\epsilon})$	$O(m^{8+\epsilon})$	$O((n^3m^4)^{2+\epsilon})$
$d > 3$	comb	$\Theta(m^{d-1})$	$n \le d :\ O(d^d m^{d+n-2})^{d-1+\epsilon}$		$n > d :\ O(\binom{n}{d}d^d m^{2d-2})^{d-1+\epsilon}$
	time	$\Theta(m^{d-1})$	$n \le d :\ O(d^d m^{d+n-2})^{2d-4+\epsilon}$		$n > d :\ O(\binom{n}{d}d^d m^{2d-2})^{2d-4+\epsilon}$

Table 1: Combinatorial complexity of uncertainty envelopes and run time complexity of our algorithm. Except for $d > 3$, the algorithm's storage requirement is identical to the combinatorial complexity. The function $\lambda_s(m)$ is the upper bound on the complexity of a Davenport-Schnitzel sequence of order s on m symbols (nearly linear).

marized in Table 1. The algorithm identifies all the topologies assumed by the zonotope as it sweeps along the convex combinations of the sites, and then it sweeps individual facets that participate in topological changes. This is significantly better than the straightforward approach of sweeping the convex hull of the sites along trajectories determined by the faces of the uncertainty domain hyperbox Δ, which has exponential complexity in the number of parameters m.

This abstract is organized as follows. Section 2 reviews zonotope topology and computation. In Section 3, we describe an arrangement of surfaces which encodes the topology of the zonotope over all convex combinations of the sites. In Section 4 we use the arrangement to compute a small subset of the faces of hyperbox Δ which contribute to the uncertainty envelope. The general algorithm is described in Section 5. Section 6 described future work and open problems. Throughout this abstract we assume that the input is in general position, that is the the sensitivity matrix columns of the sites are pairwise linearly independent.

2 Zonotope topology

We now describe uncertainty envelopes for one site ($n = 1$). As the j^{th} parameter spans the interval $[-\delta_j, \delta_j]$, the site moves along a translate of the line segment $conv\{-a_j\delta_j, a_j\delta_j\}$, where a_j is the j^{th} column of the sensitivity matrix A. Since the parameters affect the site independently, the envelope is the boundary of the Minkowsky sum of the m line segments, translated by the nominal value b. This is the definition of a *zonotope*, a convex centrally symmetric polytope [1]. We review relevant zonotope properties below.

There is a geometric correspondence between zonotopes in $\mathbb{R}^d$ and hyperplane arrangements in $\mathbb{R}^{d-1}$. The correspondence is realized by the following transform: let $H_j = \{x | \langle x, a_j \rangle = 0\}$, the hyperplane normal to a_j and passing through the origin, and let $H_0 = \{x | x_d = 1\}$. Transform each column vector $a_j = (a_{j1}, a_{j2}, \ldots, a_{jd})^\perp$ to the $(d-2)$-dimensional hyperplane $h_j = H_j \cap H_0$ whose equation is:

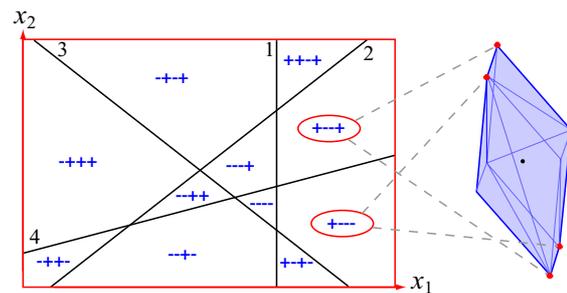$$a_{j1}x_1 + a_{j2}x_2 + \ldots + a_{j(d-1)}x_{d-1} + a_{jd} = 0 \quad (1)$$

Figure 1: The 3D zonotope (right) and its topology arrangement (left) corresponding to the sensitivity matrix $A = \begin{pmatrix} 1 & -2 & 1 & -1 \\ 0 & 2 & 1 & 3 \\ -1 & -1 & 2 & 5 \end{pmatrix}$. The j^{th} line equation is determined by substituting a_j into Eqn. 1. The marked neighboring 2-cells correspond to the neighboring zonotope vertices.

Eqn. 1 transforms the generators of the zonotope, $a_j \in \mathbb{R}^d$, to a hyperplane arrangement in $\mathbb{R}^{d-1}$ which has the same topology as the zonotope in the following sense. The sign vector $\sigma = (\sigma_1, \ldots, \sigma_m) \in \{-1, 0, 1\}^m$ of an arrangement cell containing a point x is determined according to the signs of $\langle x, a_j \rangle$ for $1 \le j \le m$. Each k-cell of the arrangement corresponds to two symmetric (antipodal) $(d - k - 1)$-cells of the zonotope. Specifically, the $(d - 1)$-cells of the arrangement with sign vector σ correspond to the vertices $v^+ = b + \sum_{j=1}^m a_j\delta_j\sigma_j$ and $v^- = b - \sum_{j=1}^m a_j\delta_j\sigma_j$, which achieve the maximum of $\langle x, v \rangle$ over x (equivalent to a direction) in the cell and v in the zonotope. The sign vectors of neighboring $(d - 1)$-cells differ in one entry only, and they correspond to neighboring vertices on the zonotope. Thus it is possible to compute the vertex representation of the zonotope in optimal $\Theta(m^{d-1})$ time. Figure 1 shows an example in 3D.

3 Swept zonotope topology

We now describe the topology of the zonotope as it is swept between the n sites. For clarity of explanation, we focus on the case $n = d$. In Section 5 we show how to solve for the general case.
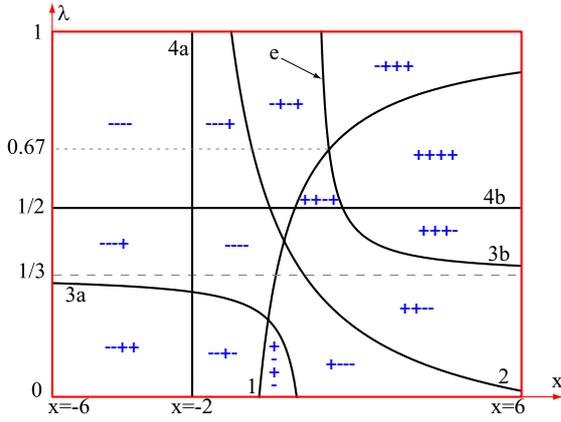
When $n = d$, the uncertainty envelope bounds the

Figure 2: Dynamic topology arrangement of $A_1 = \begin{pmatrix} 0.7 & 0.1 & -0.4 & -1 \\ 0.5 & -0.7 & 0.1 & -2 \end{pmatrix}$, $A_2 = \begin{pmatrix} 0 & 0.7 & 0.8 & 1 \\ -0.6 & 1 & -0.7 & 2 \end{pmatrix}$. Each of the curves is obtained by substituting the corresponding matrix columns into Eqn. 2. The curves corresponding to columns 3 and 4 demonstrate degeneracies which occur when a subset of the coefficients of x in Eqn. 2 are zero for some value of $\lambda \in \Lambda$.

union of $(d-1)$-simplices which are the convex hulls of sites whose positions vary within their respective zonotopes. Consider a convex combination of the sites defined by the coefficients $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$. The domain of the coefficients is $\Lambda = [0,1]^n \cap \{(\lambda \in \mathbb{R}^n | \sum_{i=1}^n \lambda_i = 1\}$. The position of a combination of sites is given by $w(p, \lambda) = \sum_{i=1}^n \lambda_i v_i(p) = \sum_{i=1}^n (\lambda_i A_i p + \lambda_i b_i)$. Let the combined sensitivity matrix be $\tilde{A}(\lambda) = \sum_{i=1}^n \lambda_i A_i$. Then the topology of the combined sites zonotope, defined by $\tilde{A}(\lambda)$ and denoted $Z(\lambda)$, changes as λ varies. Replacing a_{ji} in Eqn. 1 with the entries of $\tilde{A}(\lambda)$, we obtain:

$$\sum_{i=1}^n \lambda_i a_{j1}^i x_1 + \ldots \sum_{i=1}^n \lambda_i a_{j(d-1)}^i x_{d-1} + \sum_{i=1}^n \lambda_i a_{jd}^i = 0 \quad (2)$$

where a_{jk}^i denotes the k^{th} entry in the j^{th} column of matrix A_i. Since $\lambda_n = 1 - \sum_{i=1}^{n-1} \lambda_i$, this is an algebraic surface of degree two in the variables $(x_1, x_2, \ldots, x_{d-1}, \lambda_1, \lambda_2, \ldots, \lambda_{n-1})$. We denote the dimension of the embedding space by $D = 2(d-1)$.

The arrangement of the m surfaces is called the *dynamic topology arrangement* (DTA), because it encodes the topology of the zonotope of the combined site as its center moves along $conv\{b_i\}$. Every D-cell of the DTA represents values of x for which the sign vector is the same, defining two antipodal vertices of the zonotope $Z(\lambda)$ for λ within the cell. The intersection of two D-cells corresponds to two antipodal edges of the zonotope. With the general position assumption, an intersection of k D-cells which is also the intersection of exactly $k-1$ surfaces corresponds to two antipodal $(k-1)$-cells of the zonotope. The complexity of the entire DTA is $\Theta(m^D)$. Figure 2 shows a two dimensional DTA.
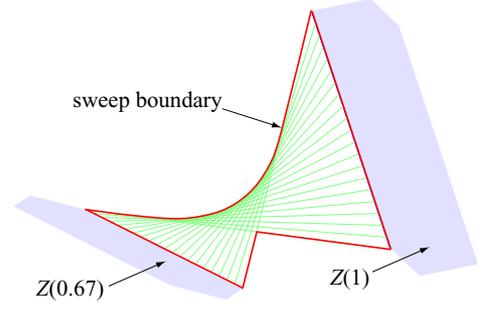

Figure 3: Sweeping the line segment corresponding to the edge e in Figure 2 through the subspace $P_C = (-\delta_1, \delta_2, 0, \delta_4) + conv\{-e_3, e_3\}$. Note that only a subset of Λ needs to be swept (from 0.67 to 1). The solid polygons are the zonotopes at $\lambda = 0.67$ and $\lambda = 1$. The segment is shown at 20 intermediate values of p_3. The closed curve bounds the swept area.

4 The sweep subspace

The uncertainty envelope is the boundary of the union of all the instances of the simplex determined by $p \in \Delta$. Although it is sufficient to include only instances with $p \in \partial \Delta$, the boundary has an exponential number of cells m.

We now show how the DTA reduces this number to a polynomial in m. When λ traces a path in its domain Λ, the zonotope undergoes a general sweep in Euclidean space. Weld and Leu [5] show that the volume swept by a compact d-manifold in $\mathbb{R}^d$ undergoing a general sweep is equal to the union of the volumes swept by its boundary and one location of the compact d-manifold in the sweep. The boundary facets of the zonotope generally correspond to the intersection of d D-cells of the DTA, i.e. $(d-1)$-cells of the arrangement. Points incident on the facet differ only in the value of $d-1$ out of m parameters in the uncertainty vector p. The indices of these parameters are exactly those with zero sign in the $(d-1)$-cell of the DTA corresponding to the facet.

Let F be a facet of the zonotope and let C be the DTA cell it corresponds to. Let $\sigma_C = (\sigma_1, \ldots, \sigma_m)$ be the sign vector of C, and let I_C be the set of indices with zero sign in σ_C. We define the sweep subspace of the cell as $P_C = (\sigma_1 \delta_1, \ldots \sigma_m \delta_m) + \sum_{i \in I_C} conv\{-e_i \delta_i, e_i \delta_i\}$, where $\{e_i\}$ are the standard basis vectors, and the plus and summation signs are Minkowski additions of sets. Since a point in F is attained by some value of $\lambda \in \Lambda$ and $p \in P_C$, the sweeping of the facet through all values of Λ is equivalent to sweeping the simplex through all the values of P_C. Therefore, to obtain the uncertainty envelope, it suffices to sweep the simplex through the values of P_C defined by all the $(d-1)$-cells of the DTA. For the computation and approximation of swept volumes see e.g. [3, 5]. Figure 3 shows an example in 2D.

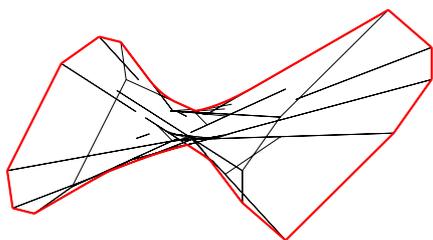|	
1.	Construct the dynamic topology arrangement $\mathcal{A}$.
2.	Compute sign vector of an arbitrary $(d-1)$-cell $C_{init} \in \mathcal{A}$.
3.	Traverse the $(d-1)$-cells of $\mathcal{A}$ in BFS order, starting from C_{init}:
	a. Update the sign vector of current cell C: $\sigma_C = (\sigma_1, \ldots, \sigma_m)$.
	b. Find indices I_C of zero signs in σ_C, defining sweep subspace P_C.
	c. Initialize swept volume $\mathcal{V}$ with simplex corresponding to arbitrary vertex $p \in P_C$.
	d. For each $i \in I_C$:
	- Sweep $\mathcal{V}$ from current $p = (p_1, \ldots, p_i, \ldots, p_m)$ to $p_{-i} = (p_1, \ldots, -p_i, \ldots, p_m)$.
	- Set $\mathcal{V}$ to new swept volume, set $p = p_{-i}$.
	e. Insert the boundary of $\mathcal{V}$ into arrangement of surfaces $\mathcal{H}$.
4.	Compute uncertainty envelope as inner boundary of the outer d-cell of $\mathcal{H}$.

Table 2: Algorithm for computing uncertainty envelopes

Figure 4: The arrangement $\mathcal{H}$ of swept boundaries of the input of Figure 2 with $b_1 = (0, 0)$, $b_2 = (1, 0)$, $\delta_1 = \delta_2 = 0.1$. The uncertainty envelope is the thick closed curve bounding $\mathcal{H}$. Notice that many of the curves do not contribute to the outer cell of the arrangement.

5 Algorithm

Table 2 presents the algorithm for computing the uncertainty envelope, based on the properties described in the previous sections. The algorithm starts by computing the DTA defined by Eqn. 2. It then traverses the DTA, iterating on the $(d-1)$-cells. For each cell C, it computes its sign vector, which defines the sweep subspace P_C. It then computes the volume swept by the simplex when p spans P_C. The collection of surfaces which bound the swept volumes of all the simplices form an arrangement $\mathcal{H}$, whose outer cell defines the uncertainty envelope.

The complexity of the algorithm, shown in Table 1, depends predominantly on the combinatorial and computational complexity of arrangements and their substructures. For $d > 3$, the current best algorithm for arrangement traversal and single cell computation is not optimal. For the most recent survey of arrangements, see [2].

Planar envelopes: The curves bounding the swept areas of line segments consist of line segments and sections of parabolas (Figure 3). Figure 4 shows an example.

3D envelopes: The swept volume of a line segment in space is a hyperbolic paraboloid. The envelope of a swept plane in space is a ruled and developable surface [5]. Thus, the sweep boundary of a triangle consists of both types of surfaces, which are swept again in the second iteration of step 3d to produce surfaces of $\mathcal{H}$.

Cases of $n \neq d$: When $n < d$ the DTA has dimension $d + n - 2$, and the algorithm sweeps $(n-1)$-simplices. For $n > d$, the uncertainty envelope is the boundary of the volume swept by the convex hull of the n sites when p spans Δ. Since the facets of the convex hull change during the sweep, we must consider all $\binom{n}{d}$ possible facets, and apply steps $1 - 3$ of the algorithm for each facet before applying step 4.

6 Conclusion

While the number of swept volumes depends on the complexity of the DTA, not all the N sweep boundaries inserted in step $3e$ contribute to the outer cell. Constructing an example in which $O(N)$ surfaces participate in the envelope, or proving a better bound, is an open problem. Furthermore, the properties of the outer cell of $\mathcal{H}$ suggest that the worst case bound on its complexity may be lower.

For efficiency of computation, the swept volumes in 3D can be approximated by polyhedrons [3]. Finally, the use of spatial data structures may reduce the number of facets considered when $n > d$.

References

[1] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[2] D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 24, pages 529–562. CRC Press LLC, Boca Raton, FL, 2004.

[3] Y. J. Kim, G. Varadhan, M. C. Lin, and D. Manocha. Fast swept volume approximation of complex polyhedral models. In *Symposium on Shape modeling and Applications*, 2003.

[4] Y. Ostrovsky-Berman and L. Joskowicz. Tolerance envelopes of planar mechanical parts. In *9th ACM Symposium on Solid Modeling and Applications*, pages 135–143, Genova, Italy, 2004.

[5] J. D. Weld and M. C. Leu. Geometric representation of swept volumes with application to polyhedral objects. *International Journal of Robotics Research*, 9(5):105–117, 1990.