Exercices: zonotopes and constrained zonotopes

General properties and examples

We recall the definition of zonotopes below:

Definition 1 (Zonotope) *An n-dimensional zonotope $\mathcal{Z}$ with center $c \in R^n$ and a vector $G = \left[g_1 \dots g_p\right] \in \mathbb{R}^{n,p}$ of p generators $g_j = (g_{ij})_{i=1,\dots,n} \in \mathbb{R}^n$ for $j = 1, \dots, p$ is defined as $\mathcal{Z} = \langle c, G\rangle = \{c + G\varepsilon \mid \|\varepsilon\|_\infty \leq 1\}$.*

In other words, for every dimension $1 \leq i \leq n$ we have the ith coordinate z_i of points $z \in \mathcal{Z}$ that belongs to the set:

$$z_i = \{c_i + \sum_{j=1}^{p} \boldsymbol{g}_{ij}\varepsilon_j \mid \varepsilon \in [-1, 1]^p\}$$

We now introduce constrained zonotopes, as zonotopes with linear constraints on the noise symbols ε_j:

Definition 2 (Constrained Zonotope) *An n-dimensional constrained zonotope $C\mathcal{Z}$ with center $c \in R^n$, a vector $G = \left[g_1 \dots g_p\right] \in \mathbb{R}^{n,p}$ of p generators $g_j \in \mathbb{R}^n$ for $j = 1, \dots, p$ and q constraints given by $H \in \mathbb{R}^{q,p}$ and $d \in \mathbb{R}^q$ is defined as $C\mathcal{Z} = \langle c, G, H, d\rangle = \{c + G\varepsilon \mid \|\varepsilon\|_\infty \leq 1,\ H\varepsilon \leq d\}$.*

In other words, a constrained zonotope is a zonotope with q constraints on the p noise symbols. These constraints can be used to refine the precision of the abstraction. Given that $H = (h_{ij})_{i=1,\dots,q;j=1,\dots,p}$, these constraints can be written as, for all $1 \leq k \leq q$:

$$\sum_{j=1}^{p} h_{kj}\varepsilon_j \leq d_k$$

Question 1 *Represent geometrically the zonotope*

$$\mathcal{Z} = \langle c, G\rangle = \left\langle \begin{pmatrix} 0 \\ 1/4 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1/2 & 1/4 \end{pmatrix} \right\rangle$$

in the (z_1, z_2) plane.

Answer 1 *The zonotope is shown at Figure 1*

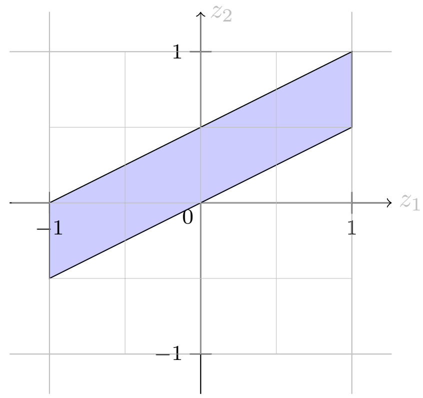

Figure 1: Answer Question 1

Question 2 *Represent geometrically, also in the (z_1, z_2) plane, the constrained zonotope*

$$C\mathcal{Z} = \langle c, G, H, d \rangle = \left\langle \begin{pmatrix} 0 \\ 1/4 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1/2 & 1/4 \end{pmatrix}, \begin{pmatrix} 1/2 & -1/4 \\ -1/2 & -1/4 \end{pmatrix}, \begin{pmatrix} 1/4 \\ 1/4 \end{pmatrix} \right\rangle$$

Hint: you should translate the constraints $H\varepsilon \leq d$ on the noise symbols ϵ_1 and ϵ_2 into constraints on the coordinates (z_1, z_2) of points $z \in C\mathcal{Z}$.

Answer 2 *The constraints $H\varepsilon \leq d$ translate here into $z_2 \geq z_1$ and $\boldsymbol{z}_2 \geq 0$. $C\mathcal{Z}$ is obtained by intersecting the zonotope of Question 1 with these constraints, and is shown at Figure 2.*

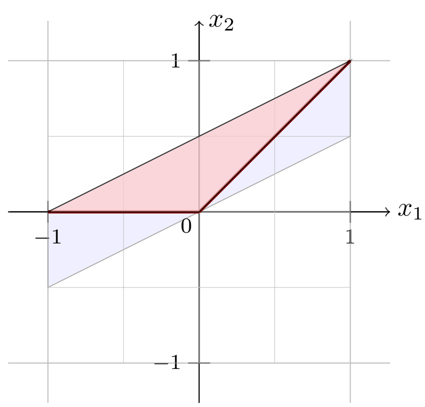

Figure 2: Answer Question 2

Question 3 *Is a constrained zonotope a polyhedra ? If yes, what is the constraint representation of the constrained zonotope $C\mathcal{Z}$ of Question 2? In that case also, what is the generator representation, as a polyhedron, of the constrained zonotope $C\mathcal{Z}$?*

Answer 3 *Constrained zonotopes are polyhedra since they are intersections of zonotopes (particular polyhedra) with polyhedric constraints, and polyhedra are stable under intersection.*

For $C\mathcal{Z}$ of Question 2:

- *two of the three faces are given by the constraints expressed in the (z_1, z_2) plane, we identified in Question 2, that is, $z_2 \geq z_1$ and $z_2 \geq 0$. The third face is given by the zonotopic generator $(1, 1/2)$ and is $z_1 \geq 2z_2 - 1$*
- *As a convex polyhedron, $C\mathcal{Z}$ is just a triangle, with three extreme points (i.e. generators) of coordinates $(-1, 0)$, $(0, 0)$ and $(1, 1)$.*

Question 4 *How can we compute, for a given constrained zonotope $C\mathcal{Z} = \langle c, G, H, d\rangle = \{c + G\varepsilon \mid \|\varepsilon\|_\infty \leq 1,\ H\varepsilon \leq d\}$, its projection onto coordinate z_i, $i = 1, \ldots, n$? This requires only a very brief answer.*

Answer 4 *By linear programming.*

Consider the concretization γ of constrained zonotopes $C\mathcal{Z} = \langle c, G, H, d\rangle$ to be

$$\gamma(C\mathcal{Z}) = \{c + G\varepsilon \mid \|\varepsilon\|_\infty \leq 1,\ H\varepsilon \leq d\}$$

Question 5 *Can two different constrained zonotopes have the same concretization? If you think so, please provide a small counter-example, otherwise, please write down a short argument.*

Given a set S in $\mathbb{R}^n$, is there always an abstraction of S as a constrained zonotope with minimal concretization? If you think so, please write down a short argument, otherwise please provide a small counter-example.

Answer 5 *Yes. For instance for $C\mathcal{Z}$ of Question 2, the constrained zonotope*

$$C\mathcal{Z}' = \langle c, G, H, d\rangle = \left\langle \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & -1 \\ 0 & -1 \\ -1 & 2 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\rangle$$

has the same concretization as $C\mathcal{Z}$.

For the second question, the answer is no. We run into the same problem as for general polyhedra. Consider e.g. the unit disc centered at 0 in the plane. We can always refine some outer-approximation of it by a constrained zonotope, by a smaller constrained zonotope that has one more face.

Affine transforms

We recall that zonotopes are closed under affine transformations: for $A \in \mathbb{R}^{m,n}$ and $b \in \mathbb{R}^m$ we can define $A\mathcal{Z} + b = \langle Ac + b, AG \rangle$ as the m-dimensional resulting zonotope.

Question 6 *Can affine transformations be also interpreted in an exact manner in contrained zonotopes? In that case, please define the affine transform of a constrained zonotope, otherwise give an short argument why this would not be the case.*

Answer 6 *For a constrained zonotope* $C\mathcal{Z} = \langle c, G, H, d \rangle = \{c + G\varepsilon \mid \|\varepsilon\|_\infty \leq 1,\ H\varepsilon \leq d\}$*, the linear transform defined by A and b as above is* $C\mathcal{Z}' = \langle Ac + b, AG, H, d \rangle = \{Ac + b + AG\varepsilon \mid \|\varepsilon\|_\infty \leq 1,\ H\varepsilon \leq d\}$*. This is an exact transformation.*

ReLU transforms

Different abstractions can be defined for the ReLU transform, among which the following one that we used in the course: let $[l_x, u_x]$ be the range reachable by component $\hat{x}$ of the input zonotope of the ReLU layer. When $l_x \leq 0$ and $u_x \geq 0$, we define the zonotope transformer for $\hat{y} = max(0, \hat{x})$ by

$$\hat{y} = \lambda \hat{x} - \frac{\lambda l_x}{2} - \frac{\lambda l_x}{2} \varepsilon_{new} \text{ with } \lambda = \frac{u_x}{u_x - l_x}. \tag{1}$$

Question 7 *Consider* $x_1 \in [-1, 1]$*, what is the zonotope abstraction of* (x_1, x_2) *for* $x_2 = ReLU(x_1)$ *using the abstraction of Equation (1) ?*

Answer 7 *This is exactly the one of Question 1.*

Question 8 *Consider again the contrained zonotope of Question 2. Is it a correct abstraction for* $x_2 = ReLU(x_1)$ *for* $x_1 \in [-1, 1]$*? Please give a short argument supporting your answer.*

Is it the best refinement, as a constrained zonotope, of the zonotope of Question 7? By refinement, we mean the following: $C\mathcal{Z}$ *is a refinement of* Z *if* $C\mathcal{Z}$ *has* Z *as underlying zonotope (hence just adding extra constraints).*

Answer 8 *Yes this is a correct abstraction, this is done by checking its concretization.*

Yes, this is the best refinement, as constrained zonotopes are particular polyhedra (in fact, we can represent all polytopes as constrained zonotopes), and as the best possible abstraction of $x_2 = ReLU(x_1)$ *as a polytope, which exists in that case, is the triangle we have pictured in Question 2.*

Question 9 *In view of the example of Question 8, define a ReLU transformer for constrained zonotopes refining the ReLU transformer for zonotopes by the addition of new constraints. Is it possible to make the transformer exact ?*

Answer 9 *Consider the constrained zonotope $\mathcal{Z}$ together with constraints, rewritten in the space of noise symbols ε, $\hat{y} \geq \hat{x}$ and $\hat{y} \geq 0$.*

It is not possible to make it exact. For instance, the graph in the (x_1, x_2) plane of the ReLU function, as in the example of Question 8 is not a convex polyhedra, and we have seen that constrained zonotopes are particular polyhedra.

Analyzing a small network

Consider the toy network of Figure 3, where for simplicity all biases are taken equal to zero, and the weights are represented on the edges:

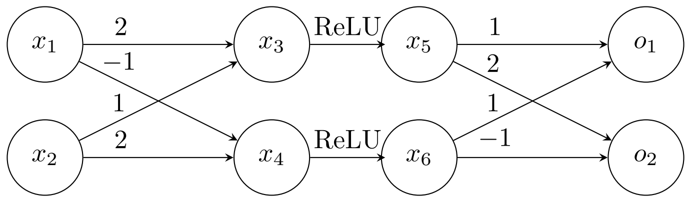

Figure 3: Toy network

Question 10 *We are interested in the local robustness of the network of Figure 3 around input $(1, 1)$.*

Using interval computations, is $[1-1/8, 1+1/8] \times [1-1/8, 1+1/8]$ a locally robust neighborhood of $(1, 1)$? For this neighborhood, we say that its radius (around $(1, 1)$, for the max distance) is 1/8.

What is the maximal robustness radius around $(1, 1)$ that can be proved for this neural net, using the interval abstraction ?

Answer 10 *We compute the output of the neural net at $(1, 1)$. We get $x_3 = 3$, $x_4 = 1$, thus $x_5 = 3$, $x_6 = 1$ an $o_1 = 4$, $o_2 = 5$, hence class 2 is classified by this network ($o_2 \geq o_1$).*

We now carry on the same computation, by interval computations, for $x_1 = [7/8, 9/8]$ and $x_2 = [7/8, 9/8]$. We get $x_3 = [21/8, 27/8]$, $x_4 = [5/8, 11/8]$, thus $x_5 = [21/8, 27/8]$, $x_6 = [5/8, 11/8]$. Then $o_1 = [13/4, 19/4]$, $o_2 = [31/8, 49/8]$. We find that o_1 is not always greater than o_2, and similarly for o_2 wrt o_1. But in fact, an exact calculation would show that $o_2 \geq o_1$ (this initial box enjoys the inequalities $-x_1 + 2x_2 \geq 0$ and $3x_2 \leq 4x_1$ which is the connected domain containing $(1, 1)$ with $o_2 \geq o_1$).

We are carrying on the same computation starting with $x_1 = [1-r, 1+r]$ and $x_2 = [1-r, 1+r]$. We find $x_3 = [3-3r, 3+3r]$, $x_4 = [1-3r, 1+3r]$, and, supposing that $r \leq 1/3$, x_3 and x_4 are positive, hence $x_5 = x_3$ and $x_6 = x_4$. We can suppose this for r since that we have seen that already for $r = 1/7 \leq 1/3$, the interval computation will not permit to prove local robustness. Finally, $o_1 = [4-6r, 4+6r]$, $o_2 = [5-9r, 5+9r]$ and $o_2 - o1 = [1-15r, 1+15r]$. Finally, $o_2 - o_1 \geq 0$ is provable by interval arithmetic iff $r \leq 1/15$.

Question 11 *Compute the zonotope for each layer of the network of Figure 3 obtained using the zonotope abstraction with input domain* $(x_1, x_2) \in [2/3, 4/3] \times [2/3, 4/3] \wedge 3x_2 \leq 4x_1$*. As this input domain is not a zonotope, we are obliged to compute with, the input zonotope being given by the square* $[2/3, 4/3] \times [2/3, 4/3]$*.*

Can you use the zonotopic analysis to prove or disprove the property that for this input domain, we always have on the outputs $o_2 \geq o_1$ *?*

Answer 11 *Starting with* $[2/3, 4/3] \times [2/3, 4/3]$ *we cannot prove robustness with zonotopes, as we show now. We compute:* $x_1 = 1 + 1/3\varepsilon_1$, $x_2 = 1 + 1/3\varepsilon_2$, $x_3 = 3 + 2/3\varepsilon_1 + 1/3\varepsilon_2$, $x_4 = 1 - 1/3\varepsilon_1 + 2/3\varepsilon_2$*. The concretization of* x_3 *is positive so* $x_5 = x_3$*. Similarly for* x_6, $x_6 = x_4$ *since* x_4 *is always positive. Finally* $o_1 = 4 + 1/3\varepsilon_1 + \varepsilon_2$, $o_2 = 5 + 5/3\varepsilon_1$ *and* $o_2 - o_1 = 1 + 4/3\varepsilon_1 - \varepsilon_2 \in [-4/3, 10/3]$*: we cannot prove nor disprove* $o_2 \geq o_1$*.*

Question 12 *Now same question as Question 11, with constrained zonotopes instead of zonotopes.*

Answer 12 *This time we can prove robustness with constrained zonotopes: we get the same underlying zonotope but this time we have the constraint* $3x_2 \leq 4x_1$ *which translates into* $1+4/3\varepsilon_1 - \epsilon_2 \geq 0$ *which is exactly* $o_2 \geq o_1$ *as the calculation of* $o_2 - o_1$ *shows in the answer to Question 11.*