Chapter 13

Duality, Inversion and Polarity

By Sariel Har-Peled, February 21, 2023[①]

"I think you're insane," he said.
"I am just outspoken. I simply say, 'A man is a sperm's way to producing another sperm.' That's merely practical."

A maze of death, Philip K. Dick

Duality is a transformation that maps lines and points into points and lines, respectively, while preserving some properties in the process. Despite its relative simplicity, it is a powerful tool that can dualize what seem like "hard" problems into easy dual problems. There are several alternative definitions of duality, but they are essentially similar, and we present one that works well for our purposes.

13.1. Duality of lines and points

Consider a line $\ell \equiv y = ax + b$ in two dimensions. It is parameterized by two constants a and b, which we can interpret, paired together, as a point in the parametric space of the lines. Naturally, this also gives us a way of interpreting a point as defining the coefficients of a line. Thus, conceptually, points are lines and lines are points.

Formally, the ***dual point*** to the line $\ell \equiv y = ax + b$ is the point $\ell^\star = (a, -b)$. Similarly, for a point $p = (c, d)$ its ***dual line*** is $p^\star \equiv y = cx - d$. Namely,

$$\begin{aligned} p = (a, b) \qquad &\implies \qquad p^\star \equiv y = ax - b,\\ \ell \equiv y = cx + d \qquad &\implies \qquad \ell^\star = (c, -d). \end{aligned}$$

We will consider a line $\ell \equiv y = cx + d$ to be a linear function in one dimension and let $\ell(x) = cx + d$.

A point p lies *above* a line ℓ if p lies vertically above ℓ. Formally, for a point $p = (a, b)$ and a line $\ell \equiv y = cx + d$ we have

$$\begin{array}{lccccc} & p \succ \ell & \equiv & p \textbf{\textit{ above }} \ell & \equiv & b > \ell(a) = ca + d,\\ \text{and} & p \prec \ell & \equiv & p \textbf{\textit{ below }} \ell & \equiv & b < \ell(a). \end{array}$$

A line ℓ ***supports*** a convex set $S \subseteq \mathbb{R}^2$ if it intersects S but the interior of S lies completely on one side of ℓ.

Basic properties. For a point $p = (a, b)$ and a line $\ell \equiv y = cx + d$, we have the following:

(P1) $p^{\star\star} = (p^\star)^\star = p$.

Proof: Indeed, $p^\star \equiv y = ax - b$ and $(p^\star)^\star = (a, -(-b)) = p$. ∎

(P2) The point p lies above (resp. below, on) the line ℓ if and only if the point $\ell^\star$ lies above (resp. below, on) the line $p^\star$. (Namely, a point and a line change their vertical ordering in the dual.)

Proof: Indeed, $p \succ \ell(a)$ if and only if $b > ca + d$. Similarly, $(c, -d) = \ell^\star \succ p^\star \equiv y = ax - b$ if and only if

$$-d > ac - b \iff b > ca + d,$$

and this is the above condition. ■

(P3) The vertical distance between p and ℓ is the same as that between $p^\star$ and $\ell^\star$.

Proof: Indeed, the vertical distance between p and ℓ is $|b - \ell(a)| = |b - (ca + d)|$. The vertical distance between $\ell^\star = (c, -d)$ and $p^\star \equiv y = ax - b$ is $|(-d) - p^\star(c)| = |-d - (ac - b)| = |b - (ca + d)|$. ■

(P4) The vertical distance $\delta(\ell, \hbar)$ between two parallel lines ℓ and $\hbar$ is the same as the length of the vertical segment $\ell^\star\hbar^\star$.

Proof: The vertical distance between $\ell \equiv y = ax + b$ and $\hbar \equiv y = ax + e$ is $|b - e|$. Similarly, since $\ell^\star = (a, -b)$ and $\hbar^\star = (a, -e)$, we have that the segment $\ell^\star\hbar^\star$ is indeed vertical and the vertical distance between its endpoints is $|(-b) - (-e)| = |b - e|$. ■

The missing lines. Consider the vertical line $\ell \equiv x = 0$. Clearly, ℓ does not have a dual point (specifically, its hypothetical dual point has an x-coordinate with infinite value). In particular, our duality cannot handle vertical lines. To visualize the problem, consider a sequence of non-vertical lines ℓ_i that converges to a vertical line ℓ. The sequence of dual points $\ell_i^\star$ is a sequence of points that diverges to infinity.

13.1.1. Examples

13.1.1.1. Segments and wedges

Consider a segment $s = pq$ that lies on a line ℓ. Observe, that the dual of a point $t \in \ell$ is a line $t^\star$ that passes through the point $\ell^\star$ (by (P2) above). Specifically, the two lines $p^\star$ and $q^\star$ define two double wedges. Let $\mathcal{W}$ be the double wedge that does not contain the vertical line that passes through $\ell^\star$; see the figure on the right.

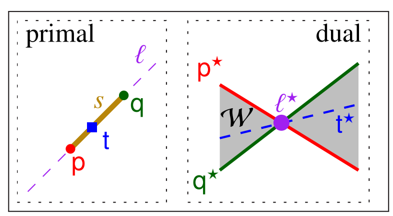

Consider now the point t as it moves along s. When it is equal to p (resp. q), then its dual line $t^\star$ is the line $p^\star$ (resp. $q^\star$). As t moves along s from p to q, its x-coordinate changes continuously, and hence the slope of its dual changes continuously from that of $p^\star$ to that of $q^\star$. Furthermore, all these dual lines must all pass through the point $\ell^\star$. As such, as t moves from p to q, the dual line $t^\star$ sweeps over the double wedge $\mathcal{W}$. Note that the x-coordinate of t during this process is in the interval $[\min(x(p), x(q)), \max(x(p), x(q))]$; namely, it is bounded. As such, the double wedge being swept over is the one that does not include the vertical line through $\ell^\star$.

What about the other double wedge? It represents the two rays forming $\ell \setminus s$. The vertical line through $\ell^\star$ represents the singularity point at infinity where the two rays are "connected" together. Thus, as t travels along one of the rays of $\ell \setminus s$ (say starting at q), the dual line $t^\star$ becomes steeper and steeper, till it becomes vertical. Now, the point t "jumps" from the "infinite endpoint" of this ray to the "infinite endpoint" of the other ray. Now, as t travels down the other ray, the dual line $t^\star$ continues to rotate from its current vertical position, sweeping over the rest of the double wedge, till t reaches p.[2] (The reader who feels uncomfortable with notions like "infinite endpoint" can rest assured that the author feels the same way. As such, this should be taken as an intuitive description of what's going on and not as a formally correct one. This argument can be formalized by using the projective plane.)

13.1.1.2. Convex hull and upper/lower envelopes

Consider a set L of lines in the plane. The minimization diagram of L, known as the ***lower envelope*** of L, is the function $\mathcal{L}_{\mathsf{L}} : \mathbb{R} \to \mathbb{R}$, where we have $\mathcal{L}(x) = \min_{\ell \in \mathsf{L}} \ell(x)$, for $x \in \mathbb{R}$. Similarly, the ***upper envelope*** of L is the function $\mathcal{U}(x) = \max_{\ell \in \mathsf{L}} \ell(x)$, for $x \in \mathbb{R}$. The ***extent*** of L at $x \in \mathbb{R}$ is the vertical distance between the upper and lower envelopes at x; namely, $\mathcal{E}_{\mathsf{L}}(x) = \mathcal{U}(x) - \mathcal{L}(x)$.

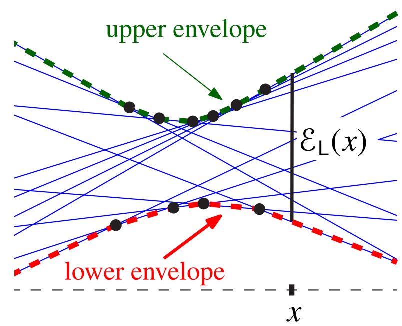

Computing the lower and/or upper envelopes can be useful. A line might represent a linear constraint, where the feasible solution must lie above this line. Thus, the feasible region is the region of points that lie above all the given lines. Namely, the region of the feasible solution is defined by the upper envelope of the lines.

The upper (and lower) envelope is a polygonal chain made out of two infinite rays and a sequence of segments, where each segment/ray lies on one of the given lines. As such, the upper envelop can be described as the sequence of lines appearing on it and the vertices where they change.

Developing an efficient algorithm for computing the upper envelope of a set of lines is a tedious but doable task. However, it becomes trivial if one uses duality.

Lemma 13.1.1. *Let L be a set of lines in the plane. Let $\alpha \in \mathbb{R}$ be any number, and let $\beta^- = \mathcal{L}_{\mathsf{L}}(\alpha)$ and $\beta^+ = \mathcal{U}_{\mathsf{L}}(\alpha)$. Let $p = (\alpha, \beta^-)$ and $q = (\alpha, \beta^+)$. Then:*

(i) The dual lines $p^\star$ and $q^\star$ are parallel, and they are both perpendicular to the direction $(\alpha, -1)$.

(ii) The lines $p^\star$ and $q^\star$ support $\mathcal{CH}(\mathsf{L}^\star)$.

(iii) The extent $\mathcal{E}_{\mathsf{L}}(\alpha)$ is the vertical distance between the lines $p^\star$ and $q^\star$.

Proof: (i) We have $p^\star \equiv y = \alpha x - \beta^-$ and $q^\star \equiv y = \alpha x - \beta^+$. These two lines are parallel since they have the same slope. In particular, they are parallel to the direction $(1, \alpha)$. But this direction is perpendicular to the direction $(\alpha, -1)$.

(ii) By property (P2), we have that all the points of $\mathsf{L}^\star$ are below (or on) the line $p^\star$. Furthermore, since p is on the lower envelope of L, it follows that $p^\star$ must pass through one of the points $\mathsf{L}^\star$. Namely, $p^\star$ supports $\mathcal{CH}(\mathsf{L}^\star)$ and it lies above it. A similar argument applies to $q^\star$.

(iii) This is a restatement of (P4). ■

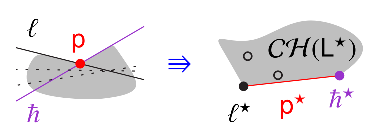

Thus, consider a vertex p of the upper envelope of the set of lines L. The point p is the intersection point of two lines ℓ and $\hbar$ of L (for the sake of simplicity of exposition, assume no other line of L passes through p). Consider the dual set of points $\mathsf{L}^\star$ and the dual line $p^\star$. Since p lies above (or on) all the lines of L, by property (P2), it must be that the line $p^\star$ lies below (or on) all the points of $\mathsf{L}^\star$. On the other hand (again by property (P2)), the line $p^\star$ passes through the two points $\ell^\star$ and $\hbar^\star$. Namely, $p^\star$ is a line that

[2]At this point t rests for awhile from this long trip of going to infinity and coming back.

supports the convex hull of $\mathsf{L}^\star$ and it passes through its vertices $\ell^\star$ and $\hbar^\star$. (The reader should verify that $\ell^\star$ and $\hbar^\star$ are indeed vertices of the convex hull.)

The convex hull of $\mathsf{L}^\star$ is a convex polygon Π which can be broken into two convex chains by breaking it at the two extreme points in the x direction (we are assuming here that L does not contain parallel lines, and as such the extreme points are unique). Note that such an endpoint is shared between the two chains and corresponds to a line that defines two asymptotes (one of the upper envelope and one, on the other side, for the lower envelope).

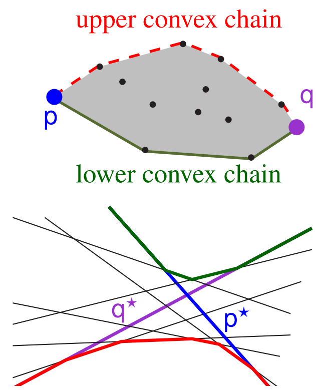

We will refer to this upper polygonal chain of the convex hull as the ***upper convex chain*** and to the lower one as the ***lower convex chain***. In particular, two consecutive segments of the upper envelope correspond to two consecutive vertices on the lower chain of the convex hull of $\mathsf{L}^\star$.

The lower chain of $\mathcal{CH}(\mathsf{L}^\star)$ corresponds to the upper envelope of L, and the upper chain corresponds to the lower envelope of L. Of special interest are the two x extreme points p and q of the convex hull. They are the dual of the two lines with the smallest/largest slopes in L. These two lines appear on both the upper and lower envelopes of the lines and they contain the four infinite rays of these envelopes.

Lemma 13.1.2. *Given a set L of n lines in the plane, one can compute its lower and upper envelopes in $O(n \log n)$ time.*

Proof: One can compute the convex hull of n points in the plane in $O(n \log n)$ time. Thus, computing the convex hull of $\mathsf{L}^\star$ and dualizing the upper and lower chains of $\mathcal{CH}(\mathsf{L}^\star)$ results in the required envelopes. ■

13.2. Higher dimensions

The above discussion can be easily extended to higher dimensions. We provide the basic properties without further proof, since they are easy extensions of the two-dimensional case. A hyperplane $h \equiv x_\mathsf{d} = b_1 x_1 + \cdots + b_{\mathsf{d}-1} x_{\mathsf{d}-1} + b_\mathsf{d}$ in $\mathbb{R}^\mathsf{d}$ can be interpreted as a function from $\mathbb{R}^{\mathsf{d}-1}$ to $\mathbb{R}$. Given a point $p = (p_1, \ldots, p_\mathsf{d})$, let $h(p) = b_1 p_1 + \cdots + b_{\mathsf{d}-1} p_{\mathsf{d}-1} + b_\mathsf{d}$. In particular, a point p lies ***above*** the hyperplane h if $p_\mathsf{d} > h(p)$. Similarly, p lies ***below*** the hyperplane h if $p_\mathsf{d} < h(p)$. Finally, a point is on the hyperplane if $h(p) = p_\mathsf{d}$.

The ***dual*** of a point $p = (p_1, \ldots, p_\mathsf{d}) \in \mathbb{R}^\mathsf{d}$ is a hyperplane $p^\star \equiv x_\mathsf{d} = p_1 x_1 + \cdots p_{\mathsf{d}-1} x_{\mathsf{d}-1} - p_\mathsf{d}$, and the ***dual*** of a hyperplane $h \equiv x_\mathsf{d} = a_1 x_1 + a_2 x_2 + \cdots + a_{\mathsf{d}-1} x_{\mathsf{d}-1} + a_\mathsf{d}$ is the point $h^\star = (a_1, \ldots, a_{\mathsf{d}-1}, -a_\mathsf{d})$. Summarizing:

$$\begin{aligned} p = (p_1, \ldots, p_\mathsf{d}) \qquad &\Longrightarrow \qquad p^\star \equiv x_\mathsf{d} = p_1 x_1 + \cdots + p_{\mathsf{d}-1} x_{\mathsf{d}-1} - p_\mathsf{d} \\ h \equiv x_\mathsf{d} = a_1 x_1 + \cdots + a_{\mathsf{d}-1} x_{\mathsf{d}-1} + a_\mathsf{d} \qquad &\Longrightarrow \qquad h^\star = (a_1, \ldots, a_{\mathsf{d}-1}, -a_\mathsf{d}). \end{aligned}$$

In the following we will slightly abuse notation, and for a point $p \in \mathbb{R}^\mathsf{d}$ we will refer to $(p_1, \ldots, p_{\mathsf{d}-1}, \mathcal{L}_\mathcal{H}(p))$ as the point $\mathcal{L}_\mathcal{H}(p)$. Similarly, $\mathcal{U}_\mathcal{H}(p)$ would denote the corresponding point on the upper envelope of $\mathcal{H}$.

The proof of the following lemma is an easy extension of the proof of Lemma 13.1.1 and is left as an exercise.

Lemma 13.2.1. *For a point $p = (b_1, \ldots, b_\mathsf{d})$, we have the following:*

(A) $p^{\star\star} = p$.

(B) *The point p lies above (resp. below, on) the hyperplane h if and only if the point $h^\star$ lies above (resp. below, on) the hyperplane $p^\star$.*
(C) *The vertical distance between p and h is the same as that between $p^\star$ and $h^\star$.*
(D) *The vertical distance $\delta(h, g)$ between two parallel hyperplanes h and g is the same as the length of the vertical segment $h^\star g^\star$.*
(E) *Computing the lower and upper envelopes of $\mathcal{H}$ is equivalent to computing the convex hull of the dual set of points $\mathcal{H}^\star$.*

13.3. Application: Upper/lower envelopes in 3d of planes

Let H be a set of n plane in 3d (in general position, with no vertical planes, etc). We are interested in computing the upper or lower envelope of H. As a reminder, the lower envelope $\mathcal{L}_H$ and upper envelope $\mathcal{U}_H$ are the functions:

$$\forall p \in \mathbb{R}^2 \quad \mathcal{L}(p) = \min_{h \in H} h(p) \qquad \text{and} \qquad \mathcal{U}(p) = \max_{h \in H} h(p).$$

By duality $H^\star$ is a set of points, and arguing as above, the upper/lower envelope of H corresponds to the lower/upper parts of the convex-hull in 3d. Since the convex-hull in 3d can be computed in $O(n \log n)$ time, we readily get the following.

Theorem 13.3.1. *The upper/lower envelopes of a set of n planes in 3d can be computed in (expected) $O(n \log n)$ time.*

13.4. Inversion

Another useful transformation is ***inversion***. We quickly overview its properties (without proving anything). This is a useful to know transformation – not necessarily something we will be directly using.

For a point $p \neq \mathsf{o}$, its ***inversion***, through the unit circle, is the point $p^{-1} = p / \|p\|^2$, where o denotes the origin. Observe that p, p^{-1}, o are collinear, $\|p\| \, \|p^{-1}\| = 1$, and p and p^{-1} are on the same side of the origin on this line.

More generally, inversion can be defined in relation to any circle in the plane. If the circle C is centered at s and has radius r, then the inversion of a point p is the point p^{-1} lying on the line between p and s, such that $\|p^{-1}s\| \, \|ps\| = r^2$.

Let us go back to the simpler inversion through the unit-circle centered at the origin. The inversion maps a circle going through the origin to a line. It maps a circle that does not go through the origin to a circle. Note, that inversion maps the interior of the interior of the unit circle to be outside it, and vise versa.

The most striking property of inversions is that they are ***conformal***. Consider two curves σ, τ that intersection each other at a point p, and the angle between them at this point p is α. Then, the angle between the "inversion" curves $\sigma^{-1} = \{q^{-1} \mid q \in \sigma\}$ and τ^{-1} is also α at their intersection point p^{-1}. Proving this requires some work, so we omit it, as we are on principle against such things.

13.5. Polarity

An alternative to duality is polarity. It avoids the vertical line goes to infinity issue, and is especially useful when dealing with bounded sets.

13.5.1. Preliminaries

A direction in $\mathbb{R}^2$ can be represented as a unit vector in $\mathbb{R}^2$. The set of unit vectors (directions) in $\mathbb{R}^2$ is denoted by $\mathbb{S} = \left\{ p \in \mathbb{R}^2 \,\middle|\, \|p\| = 1 \right\}$.

Definition 13.5.1. For a line ℓ not passing through the origin, let $\hbar = \hbar(\ell)$ (resp. $\overline{\hbar} = \overline{\hbar}(\ell)$) be the (close) halfplane bounded by ℓ and containing (resp. not containing) the origin.

For a direction $v \in \mathbb{S}$ and a point $q \in \mathbb{R}^2$, let $\hbar_v(q)$ be the halfplane that is bounded by the line normal to direction v and passing through q, and that contains o.

Definition 13.5.2 (Extremal point, supporting line). For a set P of points in the plane, and a direction $v \in \mathbb{S}$, let p_v be the ***extremal point*** of P in the direction v. That is $p_v = \arg\max_{p \in \mathsf{P}} \langle v, p \rangle$. The point p_v is unique if v is not the outer normal of an edge of $\mathcal{CH}(\mathsf{P})$. Similarly, let ℓ_v be the *supporting line* of $\mathcal{CH}(\mathsf{P})$ normal to v and passing through p_v. Let $\hbar_v = \hbar(\ell_v)$ and $\overline{\hbar}_v = \overline{\hbar}(\ell_v)$. Observe that $\mathcal{CH}(\mathsf{P}) \subset \hbar_v$.

For a real number ψ, let $\hbar_v \ominus \psi$ and $\overline{\hbar}_v \ominus \psi$ be the halfplanes formed by translating $\hbar_v$ and $\overline{\hbar}_v$, respectively, towards the origin by distance ψ.

13.5.2. Back to polarity

We use the *polarity transform*, which maps a point $p = (a, b) \neq \mathsf{o}$ to the line

$$p^{\odot} \equiv ax + by - 1 = 0 \equiv \langle p, (x, y) \rangle - 1 = 0 \equiv \left\langle p, (x, y) - \frac{p}{\|p\|^2} \right\rangle = 0.$$

Namely, the line $p^{\odot}$ is orthogonal to the vector $\mathsf{o}p$, and the closest point on $p^{\odot}$ to the origin is p^{-1}. Geometrically, a point p is being mapped to the line passing through the inverted point p^{-1} and orthogonal to the vector $\mathsf{o}p^{-1}$. Similarly, for a line ℓ, its *polar* point $\ell^{\odot}$ is q^{-1}, where q is the closest point to the origin on ℓ. Observe that $(\ell^{\odot})^{\odot} = \ell$ and $(p^{\odot})^{\odot} = p$ for any line ℓ and any point p.

If a point p lies on a line ℓ then $\ell^{\odot} \in p^{\odot}$. If p lies in the halfplane $\overline{\hbar}(\ell)$ (by Definition 13.5.1, we have $\mathsf{o} \notin \overline{\hbar}(\ell)$) if and only if $p^{\odot}$ intersects the segment $\mathsf{o}\ell^{\odot}$, see Figure 13.5.1 (left). Recall that $\rlap{/}{\delta}(\mathsf{L}) = \bigcap_{\ell \in \mathsf{L}} \hbar(\ell)$. Set $\mathsf{P}^{\odot} = \{ p^{\odot} \mid p \in \mathsf{P} \}$ and $\rlap{/}{\delta} = \rlap{/}{\delta}(\mathsf{P}^{\odot})$. Then the polygon $\rlap{/}{\delta}$ is the polar of $\mathcal{CH}(\mathsf{P})$, namely:

(I) If $p \in \mathsf{P}$ is a vertex of $\mathcal{CH}(\mathsf{P})$ then $p^{\odot}$ contains an edge of $\rlap{/}{\delta}$, see Figure 13.5.1 (right).

(II) The polar of line ℓ missing (resp. intersecting) $\mathcal{CH}(\mathsf{P})$ is a point lying in (resp. out) $\rlap{/}{\delta}$.

(III) For a point $p \in \mathcal{CH}(\mathsf{P})$, $\rlap{/}{\delta} \subset \hbar(p^{\odot})$.

Consider any direction $u \in \mathbb{S}$. Let p_u be the extremal point of P in direction u, and let ℓ_u be the corresponding supporting line, see Definition 13.5.2. The point $\ell_u^{\odot}$ lies on the edge of $\rlap{/}{\delta}$ supported by $p_u^{\odot}$, and $\ell_u^{\odot} / \|\ell_u^{\odot}\| = u$.

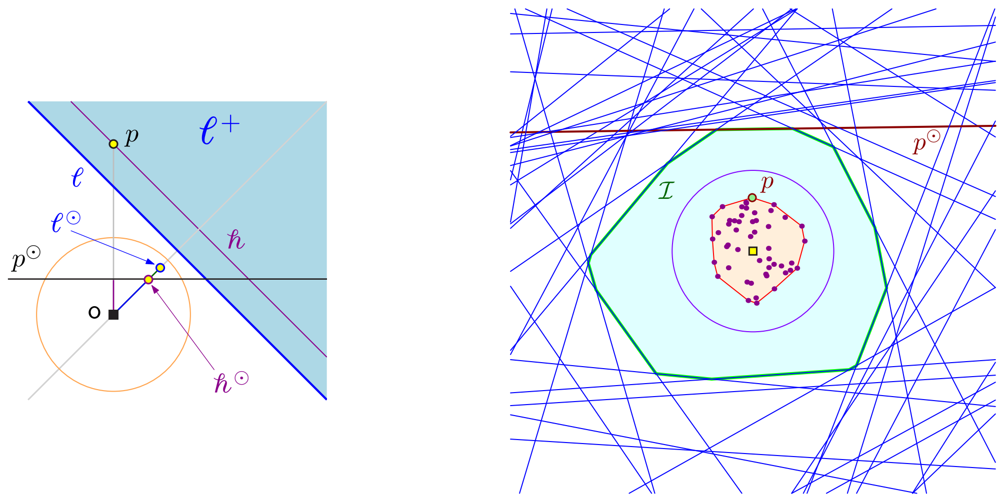

Figure 13.5.1: Left: A point p lies in the halfplane $\overline{\hbar}(\ell) \iff p^\odot$ intersects the segment $\mathsf{o}\ell^\odot$.
Right: A convex hull of a point set, and the corresponding "polar" polygon formed by the intersection of halfplanes.

13.6. Exercises

Exercise 13.6.1 (Duality of the extent). Prove Lemma 13.2.1

Exercise 13.6.2 (No duality preserves orthogonal distances). Show a counterexample proving that no duality can preserve (exactly) orthogonal distances between points and lines.

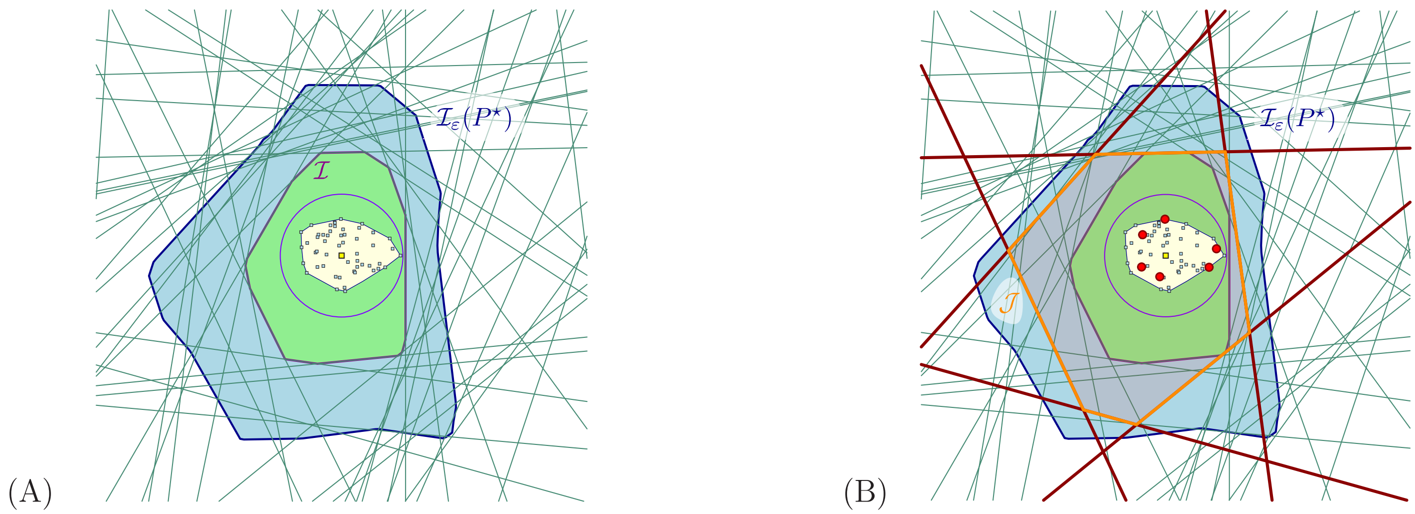

Figure 13.5.2: (A) $\mathcal{CH}(\mathsf{P})$, $\mathcal{I}(\mathsf{P}^\odot)$, $\mathcal{I}_\varepsilon(\mathsf{P}^\odot)$. (B) ε-kernel C and its polar $C^\odot$; $\mathcal{CH}(\mathsf{P}) \subseteq \mathcal{I}\mathsf{P}^\odot \subseteq \mathcal{I}_\varepsilon(C^\odot)$.