Chapter 31

Duality

By Sariel Har-Peled, September 20, 2022[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.

31.1. 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,

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

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 $p = (a, b)$ lies ***above*** a line $\ell \equiv y = cx + d$ if p lies vertically above ℓ. Formally, we have that $b > \ell(a) = ca + d$. We will denote this fact by $p \succ \ell$. Similarly, the point p lies ***below*** ℓ if $b < \ell(a) = ca + d$, denoted by $p \prec \ell$.

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

(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

$$-d > ac - b \quad \Longleftrightarrow \quad b > ca + d,$$

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.

31.1.1. Examples

31.1.1.1. 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.

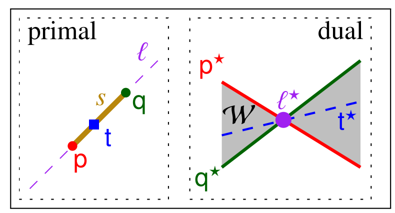

Consider now the point t as it moves along s. When it is equal to p (resp. q), then its dual line $t^\star$ is the line $p^\star$ (resp. $q^\star$). As t moves along s from p to q, its x-coordinate changes continuously, and hence the slope of its dual changes continuously from that of $p^\star$ to that 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.)

31.1.1.2. 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***

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

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, $\mathcal{E}_{\mathsf{L}}(x) = \mathcal{U}(x) - \mathcal{L}(x)$.

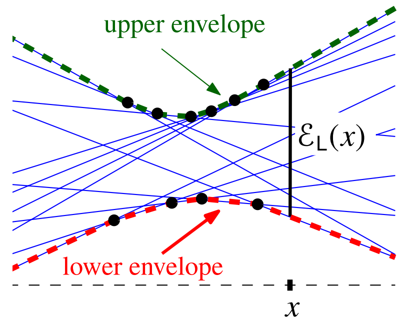

Computing the lower and/or upper envelopes can be useful. A line might represent a linear constraint, where the feasible solution must lie above this line. Thus, the feasible region is the region of points that lie above all the given lines. Namely, the region of the feasible solution is defined by the 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 31.1.1. *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 property (P4). ∎

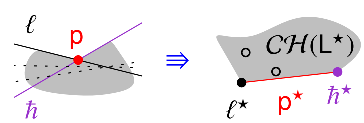

Thus, consider a vertex p of the upper envelope of the set of lines L. The point p is the intersection point of two lines ℓ and $\hbar$ of L (for the sake of simplicity of exposition, assume no other line of L passes through p). Consider the dual set of points $\mathsf{L}^\star$ and the dual line $p^\star$. 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 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 $\mathcal{P}$ 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 $\mathsf{L}^\star$.

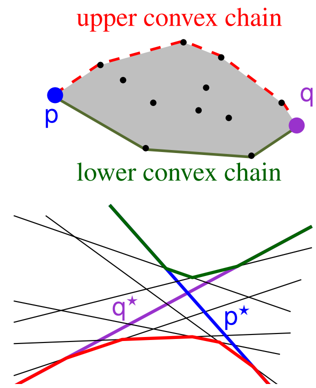

The lower chain of $\mathcal{CH}(\mathsf{L}^\star)$ corresponds to the upper envelope of L, and the upper chain corresponds to the lower envelope of L. Of special interest are the two x extreme points p and q of the convex hull. They are the dual of the two lines with the smallest/largest slopes in L. 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 31.1.2. *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. ■

31.2. 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:

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

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 1.1.1 and is left as an exercise.

Lemma 31.2.1. *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* $\mathcal{H}^\star$.

31.3. Bibliographical notes

The duality discussed here should not be confused with linear programming duality [Van97]. Although the two topics seem to be connected somehow, the author is unaware of a natural and easy connection.

A natural question is whether one can find a duality that preserves the orthogonal distances between lines and points. The surprising answer is no, as Exercise 1.2 testifies. It is not too hard to show using topological arguments that any duality must distort such distances arbitrarily badly.

Open Problem 31.3.1. Given a set P of n points in the plane and a set L of n lines in the plane, consider the minimum possible distortion of a duality (i.e., the one that minimizes the distortion of orthogonal distances) for P and L. What is the minimum distortion (of the duality) possible, as a function of n?

Formally, we define the distortion of the duality as

$$\max_{p \in P, \ell \in \mathsf{L}} \left(\frac{\mathrm{d}(p, \ell)}{\mathrm{d}(p^\star, \ell^\star)}, \frac{\mathrm{d}\big(p^\star, \ell^\star\big)}{\mathrm{d}(p, \ell)} \right).$$

A striking (negative) example of the power of duality is the work of Overmars and van Leeuwen [OL81] on the dynamic maintenance of convex hulls in the plane and the maintenance of the lower/upper envelopes of lines in the plane. Clearly, by duality, the two problems are identical. However, the authors (smart people indeed) did not observe it, and the paper is twice as long as it should be since the two problems are solved separately. (In defense of the authors this is an early work in computational geometry.)

Duality is heavily used throughout computational geometry, and it is hard to imagine managing without it. Results and techniques that use duality include bounds on k-sets/k-levels [Dey98], partition trees [Mat92], and coresets for extent measure [AHV04] (this is a random short list of relevant results and it is by no means exhaustive).

31.3.1. Projective geometry and duality

Note that our duality did not work for vertical lines. This "missing lines phenomena" is inherent to all dualities, since the space of lines in the plane has the topology of an open Möbius strip which is not homeomorphic to the plane; that is, there is no continuous mapping between all lines in the plane and all points in the plane. Naturally, there are a lot of other possible dualities, and the one presented here is the one most useful for our purposes.

One way to overcome the limitations that not all lines or points are presented by the duality is to add an extra coordinate. Now a point is represented by a triplet (w, x, y), which represents the planar point $(x/w, y/w)$. Thus a point no longer has a unique representation, and for example the triplets $(1, 1, 1)$ and $(2, 2, 2)$ represent the same point $(1, 1)$. Specifically, the set of all points in three dimensions representing a point in the plane is a line in three dimensions passing through the origin.

Similarly, a line in the plane is now represented by a triplet $\langle A, B, C \rangle$ which corresponds to the (planar) line $A+Bx+Cy = 0$. Alternatively, consider this triplet as representing the plane $Aw+Bx+Cy = 0$ that passes through the origin.

Duality is now defined in a very natural way. Indeed, if the point (a, b, c) lies on the line (A, B, C), then we have that $Aa + Bb + Cc = 0$, which implies that the point represented by (A, B, C) lies on the line represented by (a, b, c).

We can now think about the plane as being the unit sphere in three dimensions. Geometrically, we can interpret a plane (i.e., two-dimensional line) as a great circle (i.e., the intersection of the respective

3D plane and the unit sphere) and a point by its representative point on the sphere. That is, a point (x, y) in the plane induces the line $\ell \equiv \left\{ c \cdot (1, x, y) \,\middle|\, c \in \mathbb{R} \right\}$, and its (say top) intersection point with the unit sphere represents the original point. Specifically, the two antipodal intersection points are considered to be the same. (That is, we "collapse" the unit sphere by treating antipodal points as the same point.) Now, we get a classical projective geometry, where any two distinct lines intersect in a single intersection point and any two distinct points define a single line.

This homogeneous representation has the beauty that all lines are now represented and duality is universal. Expressions for the intersection point of two lines no longer involve division, which makes life much easier if one wants to implement geometric algorithms using exact arithmetic. For example, this representation is currently used by the CGAL project [FGK+00] (a software library implementing basic geometric algorithms). Another advantage is that any theorem in the primal has an immediate dual theorem in the dual. This is mathematically elegant.

This duality is still not perfect, since now there is no natural definition for what a segment connecting two points is. Indeed, there are two portions of a great circle connecting a pair of points. There is an extension of this notion to add orientation. Thus $\langle 1, 1, 1 \rangle$ and $\langle -1, -1, -1 \rangle$ represent different lines. Intuitively, one of them represents one halfplane bounded by this line, and the other represents the other halfplane. Now, if one goes through the details carefully, everything falls into place and you can speak about segments (or precisely oriented segments), and so on.

This topic is presented quite nicely in the book by Stolfi [Sto91].

31.3.2. Duality, Voronoi diagrams, and Delaunay triangulations

Given a set P of points in $\mathbb{R}^{\mathsf{d}}$, its ***Voronoi diagram*** is a partition of space into cells, where each cell is the region closest to one of the points of P. The ***Delaunay triangulation*** of a point set is a graph which is planar (the planarity holds only for $\mathsf{d} = 2$, but a similar definition holds in any dimension) where two points are connected by a straight segment if there is a ball that touches both points and its interior is empty. It is easy to verify that these two structures are dual to each other in the sense of graph duality. Maybe more interestingly, this duality also has an easy geometric interpretation.

Indeed, given P, its Voronoi diagram boils down to the computation of the lower envelope of cones. Indeed, for a point $p \in P$ the function $f(x) = \|x - p\|$ has an image, which is a cone in three dimensions. Clearly, the minimization diagram of all these functions is the Voronoi diagram. This set of cones can be linearized and then the computation of the Voronoi diagram boils down to computing the lower envelope of hyperplanes, one hyperplane for each point of P. Similarly, the computation of the Delaunay triangulation of P can be reduced, after lifting the points to the hyperboloid, to the computation of the convex hull of the points. Specifically, the projection down of the lower part of the convex hull is the required triangulation. Thus, the two structures are dual to each other. The interested reader should check out [BCKO08].

31.4. Exercises

Exercise 31.1 (Duality of the extent). Prove Lemma 1.2.1

Exercise 31.2 (No duality preserves orthogonal distances). Show a counterexample proving that no duality can preserve (exactly) orthogonal distances between points and lines.

References

[AHV04] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *J. Assoc. Comput. Mach.*, 51(4): 606–635, 2004.

[BCKO08] M. de Berg, O. Cheong, M. J. van Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. 3rd. Santa Clara, CA, USA: Springer, 2008.

[Dey98] T. K. Dey. Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.*, 19(3): 373–382, 1998.

[FGK+00] A. Fabri, G.-J. Giezeman, L. Kettner, S. Schirra, and S. Schönherr. On the design of CGAL a computational geometry algorithms library. *Softw. – Pract. Exp.*, 30(11): 1167–1202, 2000.

[Mat92] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8: 315–334, 1992.

[OL81] M. H. Overmars and J. van Leeuwen. Maintenance of configurations in the plane. *J. Comput. Syst. Sci.*, 23: 166–204, 1981.

[Sto91] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, 1991.

[Van97] R. J. Vanderbei. *Linear programming: Foundations and extensions*. Kluwer, 1997.