TROPICAL GEOMETRY, LECTURE 11

JAN DRAISMA

Transversal intersections

- Definition: Two affine subspaces $w + L_1, w + L_2$ in $\mathbb{R}^n$ are said to *intersect em transversally* at w if $L_1 + L_2 = \mathbb{R}^n$.

  Two polyhedral complexes Σ_1, Σ_2 in $\mathbb{R}^n$ are said to intersect transversally at $w \in |\Sigma_1| \cap |\Sigma_2|$ if the affine spans of the unique cells $\sigma_1 \in \Sigma_1, \sigma_2 \in \Sigma_2$ having w in their relative interiors intersect transversally at w.

  Two tropical varieties $\mathrm{Trop}(X), \mathrm{Trop}(Y)$ are said to intersect transversally at $w \in \mathrm{Trop}(X) \cap \mathrm{Trop}(Y)$ if there exist polyhedral complexes Σ_1, Σ_2 with $|\Sigma_1| = \mathrm{Trop}(X)$ and $|\Sigma_2| = \mathrm{Trop}(Y)$ which intersect transversally at w.
- Theorem: let $X, Y \subseteq T^n$ closed subvarieties. If $\mathrm{trop}(X), \mathrm{trop}(Y)$ intersect transversally at $w \in \mathbb{R}^n$, then $w \in \mathrm{trop}(X \cap Y)$. In particular, if they intersect transversally everywhere, then $\mathrm{trop}(X \cap Y) = \mathrm{trop}(X) \cap \mathrm{trop}(Y)$.

  (Note that $\subseteq$ always holds!)
- Lemma: let $I, J \subseteq K[x_0, \ldots, x_n, y_0, \ldots, y_m]$ homogeneous and $w \in \mathbb{R}^{m+n+2}$. If $\mathrm{in}_w I$ is generated by $I' := (\mathrm{in}_w I) \cap k[x_0, \ldots, x_n]$ and $\mathrm{in}_w J$ is generated by $J' := (\mathrm{in}_w J) \cap k[y_0, \ldots, y_m]$, then $\mathrm{in}_w(I + J) = \mathrm{in}_w(I) + \mathrm{in}_w(J)$.

  [Note that $\supseteq$ always holds. Hence for any u we have

$$\mathrm{in}_u(\mathrm{in}_w(I + J)) \supseteq \mathrm{in}_u(\mathrm{in}_w(I) + \mathrm{in}_w(J)) \supseteq \mathrm{in}_u \mathrm{in}_w(I) + \mathrm{in}_u \mathrm{in}_w(J).$$

  If we can show that equality holds between the lhs and the rhs, then we are done since then all relevant ideals have the same Hilbert function.

  Note furthermore that $\mathrm{in}_u \mathrm{in}_w I$ equals the ideal I'' generated by $\mathrm{in}_{u'}(I')$, where u' is projection of u in R^{n+1}. Indeed, $\mathrm{in}_u \mathrm{in}_w I \supseteq I''$ and equality holds by comparing Hilbert functions: $\mathrm{in}_w I$ is the free $k[y_0, \ldots, y_m]$-module generated by I', and I'' is the free $k[y_0, \ldots, y_m]$-module generated by $\mathrm{in}_{u'} I'$.

  Pick u such that both the ideal on the lhs above and the two ideals on the rhs are monomial, and replace w by $w + \epsilon u$ for small $\epsilon > 0$. This reduces the lemma to the case where $\mathrm{in}_w(I + J), \mathrm{in}_w(I), \mathrm{in}_w(J)$ are monomial.

  Now suppose that $f \in I_d, g \in J_d$ are such that $\mathrm{in}_w(f + g)$ is a monomial $m = x^\alpha y^\beta$ not in $\mathrm{in}_w I$ and not in $\mathrm{in}_w J$. Hence neither $\mathrm{in}_w f$ nor $\mathrm{in}_w g$ contains the monomial m, and their sets of monomials is equal (or else we'd have $\mathrm{in}_w(f + g) = \mathrm{in}_w f + \mathrm{in}_w g$). Pick one of these monomials, say $m' = x^\gamma y^\delta$. Then there are $f_1 \in I$ and $g_1 \in J$ such that $\mathrm{in}_w f_1$ is a monomial $x^{\gamma'}$ dividing x^γ and $\mathrm{in}_w g_1$ is a monomial $y^{\delta'}$ dividing y^δ. Write

$$f = c x^{\gamma - \gamma'} y^\delta f_1 + f_2$$

  where c is the coefficient of $x^\gamma y^\delta$ in f, and similarly

$$g = d x^\gamma y^{\delta - \delta'} g_1 + g_2.$$

Then we have $v(c+d) > v(c) = v(d)$, or else $\mathrm{in}_w(f+g)$ would contain a nonzero constant times the monomial $x^\gamma y^\delta$. Thus $d = c(-1+a)$ with $v(a) > 0$.

Moreover, either $\mathrm{in}_w f, \mathrm{in}_w g$ have a single term and $\mathrm{trop}(f_2)(w) > \mathrm{trop}(f)(w)$ or else $\mathrm{in}_w f_2, \mathrm{in}_w g_2$ have the monomial m' fewer than $\mathrm{in}_w f$.

Now compute

$$\begin{aligned} f+g &= cx^{\gamma-\gamma'}y^{\delta-\delta'}(y^{\delta'}f_1 - x^{\gamma'}g_1 + ax^{\gamma'}g_1g_1) + f_2 + g_2 \\ &= cx^{\gamma-\gamma'}y^{\delta-\delta'}((y^{\delta'} - g_1)f_1 + (f_1 - x^{\gamma'})g_1 + ax^{\gamma'}g_1). \end{aligned}$$

Set

$$\begin{aligned} f' &:= cx^{\gamma-\gamma'}y^{\delta-\delta'}(y^{\delta'} - g_1)f_1 + f_2 \in I \\ g' &:= cx^{\gamma-\gamma'}y^{\delta-\delta'}((f_1 - x^{\gamma'})g_1 + ax^{\gamma'}g_1) + g_2 \in J. \end{aligned}$$

Then we have $f+g = f'+g'$ and either $\mathrm{trop}(f')(w) > \mathrm{trop}(f)(w)$ or else $\mathrm{in}_w f'$ has one monomial less than $\mathrm{in}_w f$.

By iterating this construction, we get a sequence of pairs $(f'', g'') \in I \times J$ that add up to $f+g$, and in each step either the value of $\mathrm{trop}(f'')(w)$ increases strictly or else it remains the same and the number of monomials in $\mathrm{in}_w f''$ decreases strictly. Clearly there are infinitely many steps of the first type. If the valuation happens to be discrete, then this contradicts the fact that $\mathrm{trop}(f'')(w) \leq \mathrm{trop}(f+g)(w)$. In the general case, there is a slightly more technical argument in the book.]

- Proof of the theorem: let Σ_1, Σ_2 be polyhedral complexes with support $\mathrm{trop}(X)$ and $\mathrm{trop}(Y)$, respectively, that intersect transversally at w, and let $\sigma_i \in \Sigma$ be the cells with w in their relative interiors. Their affine spans are $w + L_i$, where $L_i \subseteq \mathbb{R}^n$ is a subspace spanned by its integral points. Choose integral bases $a_1, \ldots, a_r \in \mathbb{Z}^n$ of $L_1 \cap L_2$ and extend to bases $a_1, \ldots, a_r, a_{r+1}, \ldots, a_s \in \mathbb{Z}^n$ of L_1 and $a_1, \ldots, a_r, a_{s+1}, \ldots, a_n \in \mathbb{Z}^n$ of L_2. Let $A : \mathbb{Z}^n \to \mathbb{Z}^n$ be the injective group homomorphism sending e_i to a_i; its image is a full-rank submodule of $\mathbb{Z}^n$, and the corresponding morphism $\phi : T^n \to T^n$ a finite morphism. Let $X' := \phi^{-1}(X)$ and $Y' := \phi^{-1}(Y)$ be the pre-images. By last week's work, $\mathrm{trop}(X) = \mathrm{trop}(\phi)\mathrm{trop}(X')$ and $\mathrm{trop}(Y) = \mathrm{trop}(\phi)\mathrm{trop}(Y')$, where $\mathrm{trop}(\phi)$ is the bijective linear map with matrix A. Moreover, we have $X' \cap Y' = \phi^{-1}(X \cap Y)$, hence $\mathrm{trop}(X \cap Y) = \mathrm{trop}(\phi)\mathrm{trop}(X' \cap Y')$. So it suffices to prove the theorem for X' and Y' instead of X and Y. Replace the latter by the former.

  Then we have achieved that $L_1 = \langle e_1, \ldots, e_s \rangle$ and $L_2 = \langle e_1, \ldots, e_r, e_{r+1}, \ldots, e_n \rangle$. Then, by an earlier argument, $\mathrm{in}_w I$ is homogeneous w.r.t. a $\mathbb{Z}^s$-grading, hence it is generated by polynomials in the variables $x_{s+1}, \ldots, x_n$. Similarly, $\mathrm{in}_w J$ is generated by polynomials in the variables $x_{r+1}, \ldots, x_s$. After homogenising $I \cap K[x_1, \ldots, x_n]$ using a variable x_{n+1} and letting I' be the ideal that this generates in $K[x_0, \ldots, x_{n+1}]$, and homogenising $J \cap K[x_1, \ldots, x_n]$ using a variable x_0, and letting J' be the ideal that this generates in $K[x_0, \ldots, x_{n+1}]$, we have that $\mathrm{in}_{(0,w,0)} I'$ is generated by polynomials in $x_{s+1}, \ldots, x_{n+1}$ and $\mathrm{in}_{(0,w,0)} J'$ is generated by polynomials in $x_0, x_{r+1}, \ldots, s_s$. Hence we find

$$\mathrm{in}_{(0,w,0)}(I' + J') = \mathrm{in}_{(0,w,0)} I' + \mathrm{in}_{(0,w,0)} J'$$

by the lemma. Setting $x_0 = x_{n+1} = 1$ we find

$$\mathrm{in}_w(I + J) = \mathrm{in}_w I + \mathrm{in}_w J.$$

The two ideals on the left are proper ideals generated by polynomials in disjoint sets of variables $x_{s+1}, \ldots, x_n$ and $x_{r+1}, \ldots, x_s$. Picking points $p \in (k^*)^{n-s}$ satisfying the first set of equations and $q \in (k^*)^{s-r}$ satisfying the second set of equations, we find that $(1, q, p)$, where there are r ones, satisfies both, so the sum is a proper ideal. Hence $w \in \mathrm{trop}(I + J)$.

- Example: $X = V(x + y + z + 1)$ and $Y = V(t^{-10}x + t^{-1}y + tz + 1)$ have transversely intersecting tropical varieties. It follows that these polynomials form a tropical basis for the ideal that they generate.

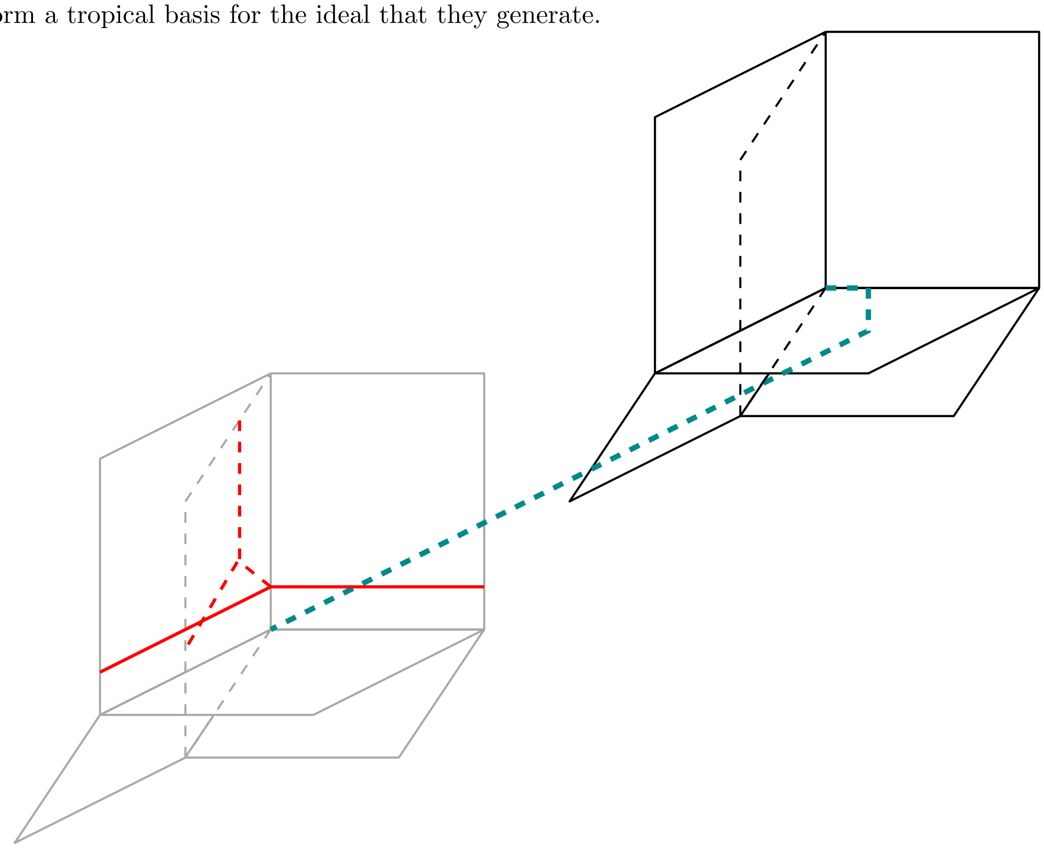

The four-point condition

- Definition: A finite metric is an $n \times n$-matrix (d_{ij}) satisfying $d_{ij} = d_{ji} \geq 0$ and $d_{ii} = 0$ and $d_{ij} + d_{jk} \geq d_{ik}$.
- Example: Glue a finite number of positive-length closed intervals together, and consider the result as an (infinite) metric space (X, δ) with shortest-path metric. Given a labelling $\phi : \{1, \ldots, n\} \to X$ we can define $d_{ij} := \delta(\phi(i), \phi(j))$. We say that (X, δ, ϕ) *realises* d.
- If X is a *tree*, then the d_{ij} satisfy the *four-point condition*: for any four distinct i, j, k, l the maximum of $\{d_{ij} + d_{kl}, d_{ik} + d_{jl}, d_{il} + d_{jk}\}$ is attained at least twice.

- Theorem: conversely, suppose that a finite metric satisfies the four-point condition, then there is a unique metric tree realising it in the above sense with the additional property that all leaves are labelled.

  [The proof idea is this: imagine the tree we're trying to find, and shorten edges leading to leaves until one becomes a point. Then remove that label, and proceed by induction.

  The basis of the induction, $n = 2$, is immediate.

  Suppose the statement is true for $n-1$, set $q := \min\{d_{kl} + d_{lm} - d_{km} \mid k, l, m \text{ distinct}\}$, and assume the minimum is attained in k, l, m.

  Define $e_{ij} := d_{ij} - q$. This is still a finite metric and satisfies the triangle inequality for k, l, m with equality. (Thanks to Arthur: to see that $e_{ij} \geq 0$, assume that $d_{ij} < q$ and choose an arbitrary $j' \neq i, j$. Then $q \leq d_{ij} + d_{jj'} - d_{ij'} < q + d_{jj'} - d_{ij'}$ so $d_{ij'} > d_{jj'}$. But the converse also holds by symmetry of the argument in i, j.) Moreover, e satisfies the four-point condition.

  By induction there is a metric tree T' realising the finite metric e on the labels $1, \ldots, \hat{l}, \ldots, n$. On the path from k to m in T, label the point at distance e_{kl} from k with l. This automatically has the right distance e_{lm} to m in T', as well—and (for uniqueness) it is the only place in T' where you can put l to match these distances.

  We now show that e_{il} is the distance from l to i in T' for all $i \neq k, m$, as well. Consider the subtree of T' spanned by k, i, l, m. After removing l, this tree splits into two or three connected components. Case 1: two components, without loss of generality with k, i in the same component. We know that the maximum of $\{e_{ik} + e_{lm}, e_{im} + e_{kl}, e_{il} + e_{km}\}$ is attained at least twice. The first two numbers are honest distances in T', and the second is larger than the first. Hence the last must equal the second, so that we find

  $$e_{im} + e_{kl} = e_{il} + e_{km} = e_{il} + e_{kl} + e_{lm}$$

  and hence $e_{il} = e_{im} - e_{lm}$, which is also the distance between i and l in T'.

  Case 2: three components. Suppose the maximum is attained by the first two. Then in particular $e_{im} + e_{kl} \geq e_{il} + e_{km} = e_{il} + e_{kl} + e_{lm}$, so $e_{im} \geq e_{il} + e_{lm}$, and by the triangle inequality equality must hold, and we're back in the previous case.

  Thus we have realised e by a tree T'. Now realise d by growing a new leaf edge leading to l of length $q/2$, and increasing the lengths of all leaf edges by this same number.]
- A version of this proof is called the *neighbour joining algorithm* in computational biology.