The Tropical Geometry of Shortest Paths

Part 2: Tropical Polynomials and Hypersurfaces

Michael Joswig

TU Berlin

Braunschweig, 28 August 2025

Outline

1. Shortest Paths
2. Tropical Polynomials and Hypersurfaces
   - tropical polynomials
   - tropicalization
   - tropical determinant = linear assignment
   - tropical geometry vs optimization, more generally
3. Parametric Shortest Paths

Tropical Polynomials

We can consider (multivariate) tropical polynomials like

$$\begin{aligned} &4 \oplus 3X \oplus 4X^2 \oplus 2XY \oplus 6Y^2 \oplus \tfrac{9}{2}Y \\ &= \min(4, 3+X, 4+2X, 2+X+Y, 6+2Y, \tfrac{9}{2}+Y) \ . \end{aligned}$$

They can be added and multiplied tropically, to obtain another semiring: $\mathbb{T}[X, Y]$.

- via tropical evaluation a k-variate tropical polynomial $F \in \mathbb{T}[X_1, \ldots, X_k]$ defines a (continuous) piecewise linear map from $\mathbb{R}^k$ to $\mathbb{R}$
- the dome $\mathcal{D}(F) = \{(x, s) \in \mathbb{R}^{k+1} \mid s \leq F(x)\}$ is an unbounded polyhedron
- F tropically vanishes at $x \in \mathbb{R}^k$ $:\Longleftrightarrow$ minimum in evaluation $F(x)$ attained at least twice

Example: A Univariate Tropical Polynomial

$k = 1$

$$F(X) = (3 \odot X^3) \oplus (1 \odot X^2) \oplus (2 \odot X) \oplus 4$$

$$F(1) = \min(3+3, \boxed{1+2}, \boxed{2+1}, 4) = 3$$

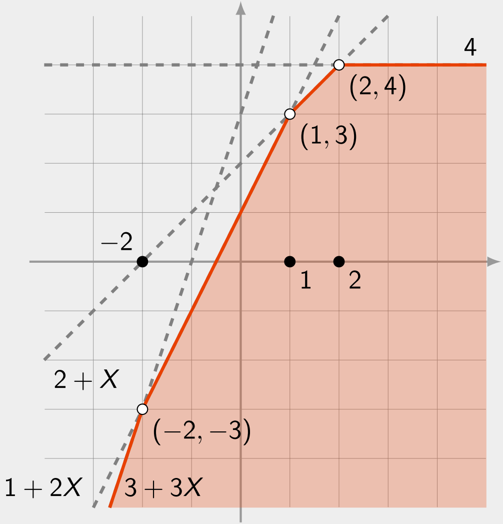

Regions of Linearity of (Bivariate) Tropical Polynomials

$k = 2$

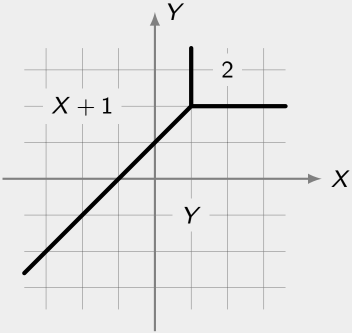

$\min(2, 1 + X, Y)$

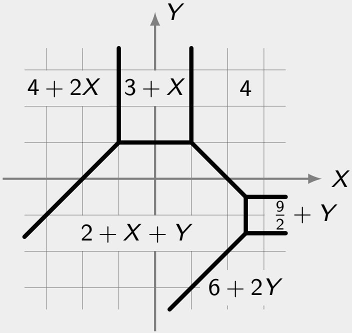

$\min(4, 3 + X, 4 + 2X, 2 + X + Y, 6 + 2Y, \frac{9}{2} + Y)$

Definition

The tropical hypersurface $\mathcal{T}(F)$ of a k-variate tropical polynomial F is the set of points $x \in \mathbb{R}^k$ where the minimum in the evaluation $F(x)$ is attained at least twice. F tropically vanishes.

Tropicalization

The ordinary polynomial

$$f \;=\; t^4 + t^3x + t^4x^2 + t^2xy + (t^6 + 23t^7)y^2 + t^{9/2}y$$

in, say $\mathbb{C}(t)[x,y]$, can be tropicalized to

$$\begin{aligned}\operatorname{trop}(f) \;&=\; 4 \oplus (3 \odot X) \oplus (4 \odot X^{\odot 2}) \oplus (2 \odot X \odot Y) \oplus (6 \odot Y^{\odot 2}) \oplus (\tfrac{9}{2} \odot Y)\\ &=\; \min(4, 3+X, 4+2X, 2+X+Y, 6+2Y, \tfrac{9}{2}+Y)\ ,\end{aligned}$$

where each coefficient $c \in \mathbb{K} = \mathbb{C}(t)$ is mapped to $\operatorname{val}(c)$ = lowest degree of t.

Theorem (Einsiedler, Kapranov & Lind 2006)

Let $f \in \mathbb{K}[x_1,\dots,x_n]$ and $x \in \mathbb{K}^n$ with $f(x) = 0$. Then $\operatorname{trop}(f)$ tropically vanishes at $\operatorname{val}(x)$. Up to passing to the topological closure the converse holds, too.

The Linear Assignment Problem

$k = n^2$

Problem

Given n soccer players and n positions, what is the best formation?

$$A = \begin{pmatrix} 0 & -1 & 2 \\ 0 & -2 & -2 \\ 0 & 2 & 0 \end{pmatrix} \in \mathbb{T}^{3\times 3}$$

matrix of "errors" $A = (a_{ij}) \in \mathbb{T}^{n\times n}$

- assignment = choice of coefficients, one per column/row

$$\text{best} = \min_{\omega\in\mathrm{Sym}(n)} a_{1,\omega(1)} + a_{2,\omega(2)} + \cdots + a_{n,\omega(n)}$$
$$= \bigoplus_{\omega\in\mathrm{Sym}(n)} a_{1,\omega(1)} \odot a_{2,\omega(2)} \odot \cdots \odot a_{n,\omega(n)}$$

Definition (tropical determinant)

$\mathrm{tdet} = \mathrm{trop}(\det)$

Bipartite Perfect Matchings

For $A = (a_{ij}) \in \mathbb{T}^{n\times n}$ define undirected bipartite graph

$$B(A) = (V, E)$$

on $V = [n] \sqcup [n]$, where $\{i, j'\}$ is an edge if $a_{ij} \neq \infty$.

- matching = collection of edges such that each node is covered at most once
- matching is perfect $\iff$ each node covered exactly once
- linear assignment = minimum weight maximal bipartite matching

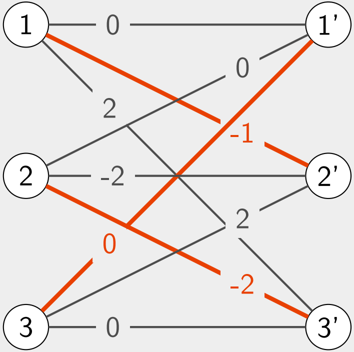

$$A = \begin{pmatrix} 0 & -1 & 2 \\ 0 & -2 & -2 \\ 0 & 2 & 0 \end{pmatrix}$$

Hungarian Method

Kuhn 1955

input: matrix $A \in \mathbb{T}^{n \times n}$
output: matching in $B(A)$ of minimum weight among all matchings of maximal size
$\mu \leftarrow \emptyset$
repeat
    $U_\mu \leftarrow$ nodes in $[n]$ not covered by μ
    $W_\mu \leftarrow$ nodes in $[n']$ not covered by μ
    $B_\mu \leftarrow$ directed graph with node set $[n] \sqcup [n']$,
    edges with weights induced by A, directed from $[n]$ to $[n']$,
    except for those in μ, which are reversed, with negated weights
    if there is a path from U_μ to W_μ in B_μ **then**
        $\pi \leftarrow$ edge set of shortest one among these
        $\mu \leftarrow \mu \triangle \pi$
until no path from U_μ to W_μ exists in B_μ
return μ

complexity:
$n \cdot$ cost for shortest path

Example

$\mu_0 = \emptyset$, and hence $U_0 = \{1, 2, 3\}$, $W_0 = \{1', 2', 3'\}$

1. pick $\mu_1 = \{(2, 2')\}$, one edge of lowest weight -2
   - $U_1 = \{1, 3\}$ and $W_1 = \{1', 3'\}$, and edge $(2, 2')$ is reversed with negated weight
   - directed path $(1, 2'), (2', 2), (2, 3')$ has (minimal) total weight $-1 + 2 - 2 = -1$
2. $\mu_2 = \mu_1 \triangle \{(1, 2'), (2, 2'), (2, 3')\} = \{(1, 2'), (2, 3')\}$
   - shortest path from $U_2 = \{3\}$ to $W_2 = \{1'\}$ has single edge $(3, 1')$
3. unique minimum weight perfect matching

$$\mu_3 \; = \; \mu_2 \triangle \{(3, 1')\} \; = \; \{(1, 2'), (2, 3'), (3, 1')\}$$

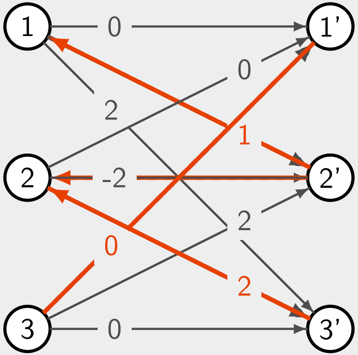

Intermediate Complexity Analysis

Complexity of Hungarian method: $n \cdot$ cost for shortest path; i.e., $O(n^4)$ with Floyd–Warshall.

- however, Floyd–Warshall solves all-pairs shortest paths, which is more than we need!

Strategy: Dijkstra's algorithm

- solves single source/target shortest path problem in $O(n^2)$ time
- however, requires nonnegative weights!
- address that issue later

Dijkstra's Algorithm (1959)

input: nonnegative matrix $A \in \mathbb{T}^{n \times n}$ and a node $t \in [n]$
output: vector of shortest path weights $\mathrm{wt}^*(\cdot, t)$ for $\Gamma(A)$

```
initialize vector p of length n with p(t) = 0 and p(v) = ∞ for v ≠ t
U ← [n]
while U contains node with finite p-value do
    v ← node in U whose p-value is minimal (and thus finite)
    U ← U − v
    foreach u ∈ δ⁻(v) do
        λ ← wt(u, v) + p(v)
        if λ < p(u) then  p(u) ← λ
return p
```

Example: Dijkstra's Shortest-Path Tree

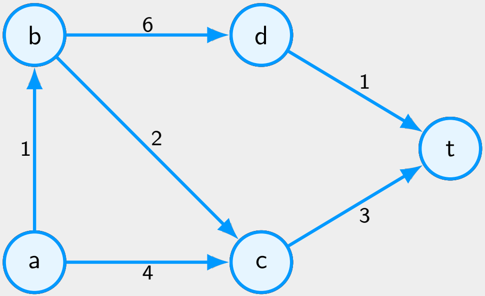

$p(a) = 6$

$p(b) = 5$

$p(c) = 3$

$p(d) = 1$

$p(t) = 0$

Final Complexity Analysis

Theorem

Suppose that $A \in \mathbb{T}^{n\times n}$ is nonnegative.
Then Dijkstra's algorithm computes one shortest path tree in the digraph $\Gamma(A)$ to obtain the (backward) potential $A^{*(t)} = \mathrm{wt}^*(\cdot, t)$ in $O(n^2)$ time.

- forward potentials from Dijkstra correspond to rows $\mathrm{wt}^*(u, \cdot)$
- $O(n^3)$ for the entire Kleene star A^*, just like Floyd–Warshall

Theorem

The complexity of computing $\mathrm{tdet}(A)$, via the Hungarian method and Dijkstra is $O(n^3)$.

What about negative coefficients?

Dijkstra With a Potential

Now let $A \in \mathbb{T}^{n\times n}$ be arbitrary, but no negative cycles in $\Gamma = \Gamma(A)$.
With $p \in Q(A)$ be a finite potential we adjust the weights to

$$\mathrm{wt}_p(u,v) \;:=\; \mathrm{wt}(u,v) - p(u) + p(v) \;. \tag{1}$$

- wt_p nonnegative weights on arcs of Γ

Observation

Any shortest path tree for the adjusted weight function wt_p is also a shortest path tree for the original weight function wt. The adjusted weight function gives the actual distance from any node, v, to the target, t, in the graph with the original weights by the formula

$$\mathrm{wt}^*(v,t) \;=\; \mathrm{wt}_p^*(v,t) + p(v) - p(t) \;. \tag{2}$$

Further Connections Between Tropical Geometry and Optimization

dramatically incomplete and biased

- Baldwin & Klemperer 2019; Tran & Yu 2019: product-mix auctions
  - Crowell & Tran 2016: mechanism design
  - J., Klimm & Spitz 2022: revenue maximization
- ordinary polyhedra and linear programs can be tropicalized
  - Akian, Gaubert & Gutermann 2012: mean-payoff games
  - Allamigeon, Benchimol, Gaubert & J. (2014, 2015, 2018, 2021); Allamigeon, Gaubert & Vandame (2022): complexity of interior point method
- Lin & Yoshida 2018: tropical Fermat–Weber problems
  - J. & Comăneci 2024: asymmetric tropical distance
- Gärtner & Jaggi 2006: tropical support vector machines
  - Tang, Wang & Yoshida 2020: application to phylogenetics
  - Zhang, Naitzat & Lim 2018; Montúfar, Ren & Zhang 2021: neural networks
- Murota 1996: M-convexity; ...; Brändén & Huh 2020: Lorentzian polynomials

Summary

- Ordinary polynomials can be tropicalized.
- Computing a tropical determinant is equivalent to solving a linear assignment problem.
- A k-variate tropical polynomials decomposes $\mathbb{R}^k$ into polyhedral regions, via evaluation.
- Dijkstra'a algorithm computes one shortest path tree toward a fixed node in $O(n^2)$, provided that the weights are nonnegative (or a potential is given).