

Tropical linear regression and low-rank approximation — a first step in tropical data analysis

Yang Qi

INRIA and CMAP École Polytechnique

joint with Marianne Akian, Stéphane Gaubert, and Omar Saadi

AGATES

OCTOBER 6, 2022



- 1 Introduction
- 2 Tropical linear regression
- 3 Tropical low-rank approximation
- 4 Generalization

- 1 Introduction
- 2 Tropical linear regression
- 3 Tropical low-rank approximation
- 4 Generalization

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y);$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y;$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y; \quad 3 \odot 4 = 7.$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y; \quad 3 \odot 4 = 7.$$

- The neutral element for \oplus is “0” = $-\infty$, and for \odot is “1” = 0

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y; \quad 3 \odot 4 = 7.$$

- The neutral element for \oplus is “0” = $-\infty$, and for \odot is “1” = 0
- The operations are defined for matrices:

$$(A \oplus B)_{ij} = \max(A_{ij}, B_{ij})$$

$$(A \odot B)_{ij} = \max_k (A_{ik} + B_{kj})$$

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y; \quad 3 \odot 4 = 7.$$

- The neutral element for \oplus is “0” = $-\infty$, and for \odot is “1” = 0
- The operations are defined for matrices:

$$(A \oplus B)_{ij} = \max(A_{ij}, B_{ij})$$

$$(A \odot B)_{ij} = \max_k (A_{ik} + B_{kj})$$

- A scalar acts on vector by “ λx ” = $\lambda \odot x = \lambda + x$ entrywise.

The Tropical Semiring

- The tropical semiring is $(\mathbb{R}_{\max}, \oplus, \odot)$, with :

$$\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$$

$$x \oplus y = \max(x, y); \quad 3 \oplus 4 = 4,$$

$$x \odot y = x + y; \quad 3 \odot 4 = 7.$$

- The neutral element for \oplus is “0” = $-\infty$, and for \odot is “1” = 0
- The operations are defined for matrices:

$$(A \oplus B)_{ij} = \max(A_{ij}, B_{ij})$$

$$(A \odot B)_{ij} = \max_k (A_{ik} + B_{kj})$$

- A scalar acts on vector by “ λx ” = $\lambda \odot x = \lambda + x$ entrywise.
- Eigenvalues and eigenvectors of a matrix A :

$$A \odot v = \lambda + v$$

Applications

Tropical data arise naturally:

- Control theory Game theory Economics ...

¹L. Zhang, G. Naitzat, L.-H. Lim, *Tropical geometry of deep neural networks*, ICML, 2018

Applications

Tropical data arise naturally:

- Control theory Game theory Economics ...

A baby example:

- Given a direct-distances matrix $A = (a_{ij})$, find the longest distances.
- Consider $A^2 = A \odot A$, i.e.,

$$(A^2)_{ij} = \max_k (a_{ik} + a_{kj}),$$

which gives the longest path between i and j of length 2.

- $A + A^2 + \dots$ gives longest paths of all lengths.

¹L. Zhang, G. Naitzat, L.-H. Lim, *Tropical geometry of deep neural networks*, ICML, 2018

Applications

Tropical data arise naturally:

- Control theory Game theory Economics ...

A baby example:

- Given a direct-distances matrix $A = (a_{ij})$, find the longest distances.
- Consider $A^2 = A \odot A$, i.e.,

$$(A^2)_{ij} = \max_k (a_{ik} + a_{kj}),$$

which gives the longest path between i and j of length 2.

- $A + A^2 + \dots$ gives longest paths of all lengths.

Another example ¹:

$$\left\{ \begin{array}{l} \text{feedforward ReLU neural networks} \\ \text{with integer weights} \end{array} \right\} \simeq \{\text{tropical rational maps}\}$$

¹L. Zhang, G. Naitzat, L.-H. Lim, *Tropical geometry of deep neural networks*, ICML, 2018

Ground states of spin glasses ²

Finding the ground state of the Ising spin glass with the energy function

$$E(\{\sigma\}) = - \sum_{i < j} J_{ij} \sigma_i \sigma_j - \sum_i h_i \sigma_i,$$

where $\{\sigma\} \in \{\pm 1\}^N$ a configuration of N Ising spins.

For the partition function $Z = \sum_{\sigma} e^{-\beta E}$, the ground state energy

$$E^* = - \lim_{\beta \rightarrow \infty} (1/\beta) \ln Z = - \lim_{\beta \rightarrow \infty} (1/\beta) \ln \sum_{\sigma} \prod_{i < j} e^{\beta J_{ij} \sigma_i \sigma_j} \prod_i e^{\beta h_i \sigma_i}$$

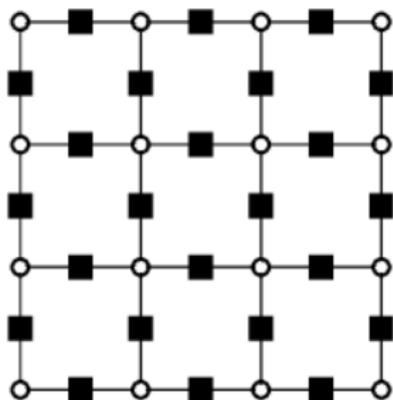
Taking the zero-temperature limit,

$$\lim_{\beta \rightarrow \infty} (1/\beta) \ln(e^{\beta x} + e^{\beta y}) = x \oplus y, \quad (1/\beta) \ln(e^{\beta x} \times e^{\beta y}) = x \odot y.$$

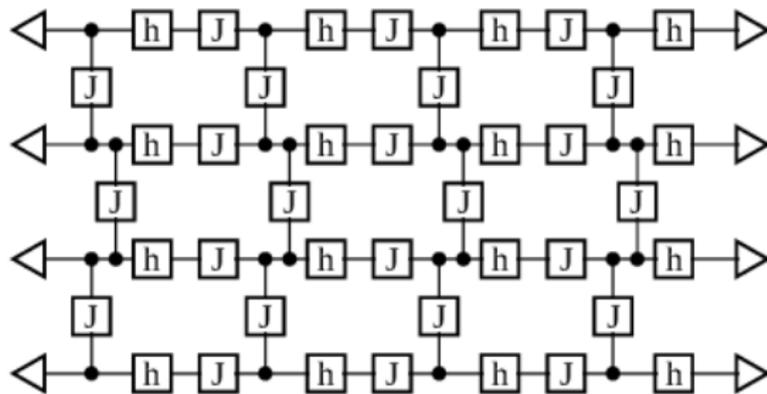
²J. Liu, L. Wang, P. Zhang, *Tropical tensor network for ground states of spin glasses*, Physical Review Letters, 2021

Tropical tensor networks

(a)



(b)



Example 1: function approximation

- X : compact metric
- Y : nonempty
- $\mathcal{C}(X)$: the space of continuous functions on X
- $\mathcal{B}(Y)$: the space of bounded functions on Y .
- Given $V : X \times Y \rightarrow \mathbb{R}$ bounded, find f and g such that

$$\inf_{f \in \mathcal{C}(X), g \in \mathcal{B}(Y)} \sup_{x \in X, y \in Y} |V(x, y) - f(x) - g(y)|$$

Example 1: function approximation

- X : compact metric
- Y : nonempty
- $\mathcal{C}(X)$: the space of continuous functions on X
- $\mathcal{B}(Y)$: the space of bounded functions on Y .
- Given $V : X \times Y \rightarrow \mathbb{R}$ bounded, find f and g such that

$$\inf_{f \in \mathcal{C}(X), g \in \mathcal{B}(Y)} \sup_{x \in X, y \in Y} |V(x, y) - f(x) - g(y)|$$

will see later : $f + g$ is a best rank-one approximation of V

Example 2: invitation to tender markets

- A decision maker and n local companies

Example 2: invitation to tender markets

- A decision maker and n local companies
- For invitation $j \in [q]$, company $i \in [n]$ asks for the price p_{ij} (public)

Example 2: invitation to tender markets

- A decision maker and n local companies
- For invitation $j \in [q]$, company $i \in [n]$ asks for the price p_{ij} (public)
- Secret evaluation $0 < f_i \leq 1$ (**technical quality**)

Example 2: invitation to tender markets

- A decision maker and n local companies
- For invitation $j \in [q]$, company $i \in [n]$ asks for the price p_{ij} (public)
- Secret evaluation $0 < f_i \leq 1$ (**technical quality**)
- The decision maker minimizes his expected cost:

$$\min_{i \in [n]} p_{ij} f_i^{-1}$$

Equilibrium in invitation to tender markets

The prices constitute an equilibrium:

$\min_{i \in [q]} p_{ij} f_i^{-1}$ is achieved twice at least

Equilibrium in invitation to tender markets

The prices constitute an equilibrium:

$$\min_{i \in [q]} p_{ij} f_i^{-1} \text{ is achieved twice at least}$$

Indeed, if $p_{ij} f_i^{-1} < p_{kj} f_k^{-1}, \forall k \neq i$, then company i may raise its price and still wins the offer.

Equilibrium in invitation to tender markets

The prices constitute an equilibrium:

$$\min_{i \in [q]} p_{ij} f_i^{-1} \text{ is achieved twice at least}$$

Indeed, if $p_{ij} f_i^{-1} < p_{kj} f_k^{-1}, \forall k \neq i$, then company i may raise its price and still wins the offer.

Let $V_{ij} = -\log(p_{ij})$ and $a_i = \log(f_i)$, the equilibrium is:

$$\max_{i \in [n]} (V_{ij} + a_i) \text{ is achieved twice at least}$$

Equilibrium in invitation to tender markets

The prices constitute an equilibrium:

$$\min_{i \in [q]} p_{ij} f_i^{-1} \text{ is achieved twice at least}$$

Indeed, if $p_{ij} f_i^{-1} < p_{kj} f_k^{-1}, \forall k \neq i$, then company i may raise its price and still wins the offer.

Let $V_{ij} = -\log(p_{ij})$ and $a_i = \log(f_i)$, the equilibrium is:

$$\max_{i \in [n]} (V_{ij} + a_i) \text{ is achieved twice at least}$$

will see later: $\forall j \in [q], V_{\cdot j}$ lies in some hyperplane H_a

Given tropical data, we would like to do

- Tropical linear regression:
Given m points, find a hyperplane which best fits these data.
- Tropical low-rank approximation:
Given a tropical matrix, find a best low-rank approximation.

Overview

- 1 Introduction
- 2 Tropical linear regression**
- 3 Tropical low-rank approximation
- 4 Generalization

Tropical cones

- $C \subseteq (\mathbb{R}_{\max})^n$ is a **tropical (convex) cone** if $\forall x, y \in C, \forall \lambda \in \mathbb{R}_{\max}$:
 $\lambda + x \in C$ and $\max(x, y) \in C$.

Tropical cones

- $C \subseteq (\mathbb{R}_{\max})^n$ is a **tropical (convex) cone** if $\forall x, y \in C, \forall \lambda \in \mathbb{R}_{\max}$:
$$\lambda + x \in C \text{ and } \max(x, y) \in C.$$
- For a set $\mathcal{V} \subseteq (\mathbb{R}_{\max})^n$, the tropical cone generated by \mathcal{V} :
$$\text{Span}(\mathcal{V}) := \left\{ \sup_{v \in \mathcal{V}} (\lambda_v + v) \mid \lambda_v \in \mathbb{R}_{\max} \right\}.$$

Tropical cones

- $C \subseteq (\mathbb{R}_{\max})^n$ is a **tropical (convex) cone** if $\forall x, y \in C, \forall \lambda \in \mathbb{R}_{\max}$:

$$\lambda + x \in C \text{ and } \max(x, y) \in C.$$

- For a set $\mathcal{V} \subseteq (\mathbb{R}_{\max})^n$, the tropical cone generated by \mathcal{V} :

$$\text{Span}(\mathcal{V}) := \left\{ \sup_{v \in \mathcal{V}} (\lambda_v + v) \mid \lambda_v \in \mathbb{R}_{\max} \right\}.$$

- The **tropical projective space** $\mathbb{P}(\mathbb{R}_{\max})^n$ is the quotient

$$(\mathbb{R}_{\max}^n \setminus \{-\infty\}) / \sim,$$

where $x \sim y$ iff $x - y$ is a constant vector.

Tropical cones

- $C \subseteq (\mathbb{R}_{\max})^n$ is a **tropical (convex) cone** if $\forall x, y \in C, \forall \lambda \in \mathbb{R}_{\max}$:

$$\lambda + x \in C \text{ and } \max(x, y) \in C.$$

- For a set $\mathcal{V} \subseteq (\mathbb{R}_{\max})^n$, the tropical cone generated by \mathcal{V} :

$$\text{Span}(\mathcal{V}) := \left\{ \sup_{v \in \mathcal{V}} (\lambda_v + v) \mid \lambda_v \in \mathbb{R}_{\max} \right\}.$$

- The **tropical projective space** $\mathbb{P}(\mathbb{R}_{\max})^n$ is the quotient

$$(\mathbb{R}_{\max}^n \setminus \{-\infty\}) / \sim,$$

where $x \sim y$ iff $x - y$ is a constant vector.

- A natural metric on $\mathbb{P}(\mathbb{R}^n)$ is induced by **Hilbert's seminorm**:

$$\|z\|_H := \left(\max_{i \in [n]} z_i \right) - \left(\min_{i \in [n]} z_i \right).$$

Tropical cones

- $C \subseteq (\mathbb{R}_{\max})^n$ is a **tropical (convex) cone** if $\forall x, y \in C, \forall \lambda \in \mathbb{R}_{\max}$:

$$\lambda + x \in C \text{ and } \max(x, y) \in C.$$

- For a set $\mathcal{V} \subseteq (\mathbb{R}_{\max})^n$, the tropical cone generated by \mathcal{V} :

$$\text{Span}(\mathcal{V}) := \left\{ \sup_{v \in \mathcal{V}} (\lambda_v + v) \mid \lambda_v \in \mathbb{R}_{\max} \right\}.$$

- The **tropical projective space** $\mathbb{P}(\mathbb{R}_{\max})^n$ is the quotient

$$(\mathbb{R}_{\max}^n \setminus \{-\infty\}) / \sim,$$

where $x \sim y$ iff $x - y$ is a constant vector.

- A natural metric on $\mathbb{P}(\mathbb{R}^n)$ is induced by **Hilbert's seminorm**:

$$\|z\|_H := (\max_{i \in [n]} z_i) - (\min_{i \in [n]} z_i).$$

- Given $A, B \subseteq \mathbb{P}(\mathbb{R}_{\max})^n$, the **one-sided Hausdorff distance**:

$$d_H(A, B) := \sup_{a \in A} \inf_{b \in B} \|a - b\|_H.$$

Tropical hyperplanes

Given $a \in \mathbb{P}(\mathbb{R}_{\max})^n$, we define the **tropical hyperplane**:

$$H_a := \{x \in (\mathbb{R}_{\max})^n \mid \max_{i \in [n]}(a_i + x_i) \text{ achieved at least twice}\} .$$

Tropical hyperplanes

Given $a \in \mathbb{P}(\mathbb{R}_{\max})^n$, we define the **tropical hyperplane**:

$$H_a := \{x \in (\mathbb{R}_{\max})^n \mid \max_{i \in [n]}(a_i + x_i) \text{ achieved at least twice}\} .$$

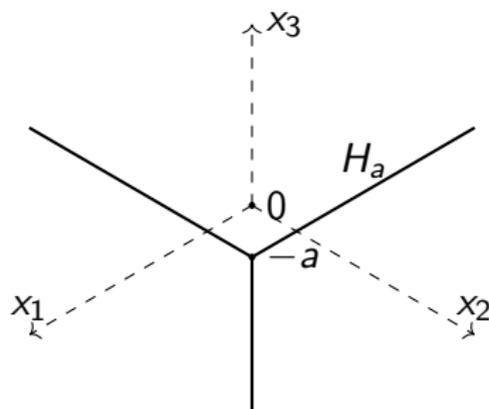


Figure: The hyperplane H_a with $a = (0, 0, 1)^\top$ and in $\mathbb{P}(\mathbb{R}_{\max})^3$.

Tropical hyperplanes

Given $a \in \mathbb{P}(\mathbb{R}_{\max})^n$, we define the **tropical hyperplane**:

$$H_a := \{x \in (\mathbb{R}_{\max})^n \mid \max_{i \in [n]}(a_i + x_i) \text{ achieved at least twice}\} .$$

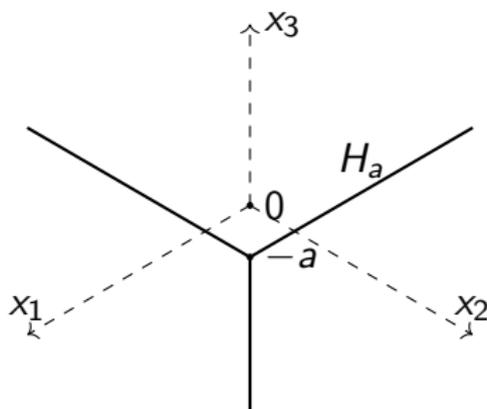


Figure: The hyperplane H_a with $a = (0, 0, 1)^\top$ and in $\mathbb{P}(\mathbb{R}_{\max})^3$.

Given a set $\mathcal{V} \subseteq \mathbb{R}_{\max}^n$, we solve **the tropical linear regression problem**:

$$\inf_{a \in \mathbb{P}(\mathbb{R}_{\max})^n} d_H(\mathcal{V}, H_a) .$$

For a subset \mathcal{V} of $(\mathbb{R}_{\max})^n$, the *inner radius* of \mathcal{V} is:

$$\text{in-rad}(\mathcal{V}) := \sup\{r \geq 0 \mid \exists b \in \mathbb{R}^n, B(b, r) \subseteq \text{Span}(\mathcal{V})\}.$$

For a subset \mathcal{V} of $(\mathbb{R}_{\max})^n$, the *inner radius* of \mathcal{V} is:

$$\text{in-rad}(\mathcal{V}) := \sup\{r \geq 0 \mid \exists b \in \mathbb{R}^n, B(b, r) \subseteq \text{Span}(\mathcal{V})\}.$$

Theorem (Akian–Gaubert–Q.–Saadi)

$$\inf_{a \in \mathbb{P}(\mathbb{R}_{\max})^n} d_H(\mathcal{V}, H_a) = \text{in-rad}(\mathcal{V}).$$

Question: how to compute $\text{in-rad}(\mathcal{V})$?

Mean Payoff Games

For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

Mean Payoff Games

For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

- Two players Min and Max

Mean Payoff Games

For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

- Two players Min and Max
- Starting from a state i , Min chooses $k \in [p]$ s.t. $V_{ik} \neq -\infty$

Mean Payoff Games

For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

- Two players Min and Max
- Starting from a state i , Min chooses $k \in [p]$ s.t. $V_{ik} \neq -\infty$
- Then Max chooses the next state $j \neq i$

Mean Payoff Games

For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

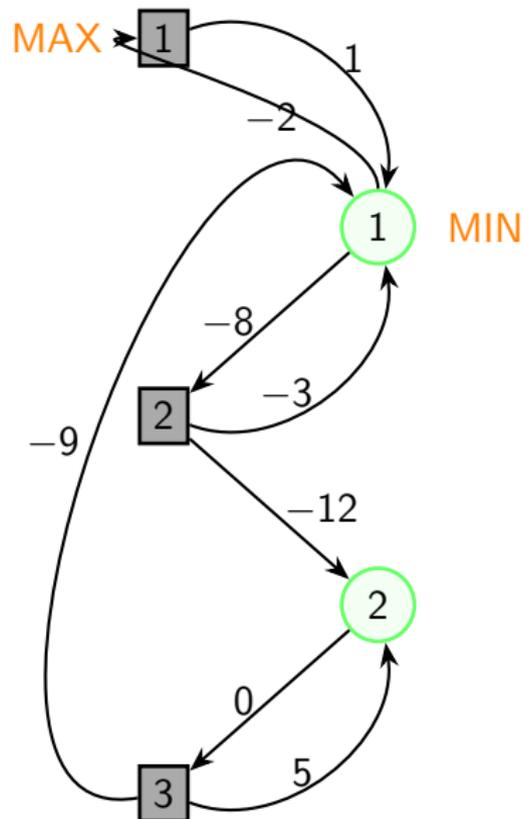
- Two players Min and Max
- Starting from a state i , Min chooses $k \in [p]$ s.t. $V_{ik} \neq -\infty$
- Then Max chooses the next state $j \neq i$
- $-V_{ik} + V_{jk}$ is the instantaneous payment made by Min to Max

Mean Payoff Games

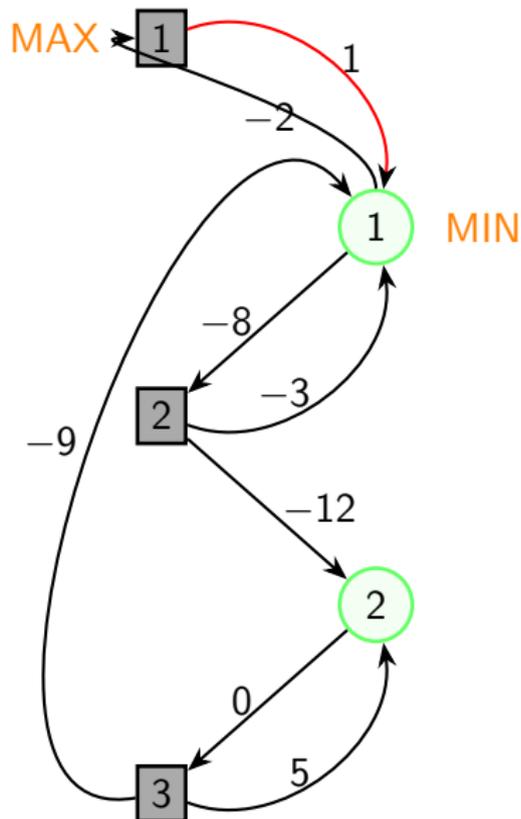
For $V \in (\mathbb{R}_{\max})^{n \times p}$, the *zero-sum two-player deterministic game*:

- Two players Min and Max
- Starting from a state i , Min chooses $k \in [p]$ s.t. $V_{ik} \neq -\infty$
- Then Max chooses the next state $j \neq i$
- $-V_{ik} + V_{jk}$ is the instantaneous payment made by Min to Max
- *Asymetry: Min can play tit for tat but Max cannot!*

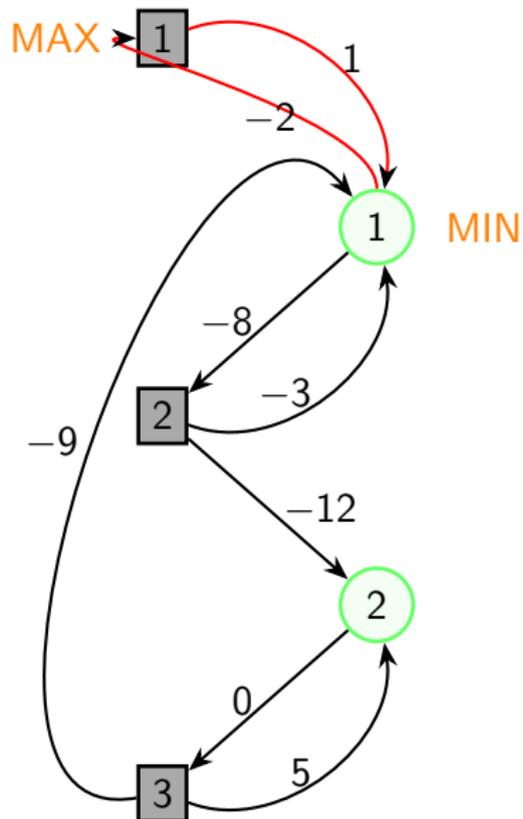
Baby example



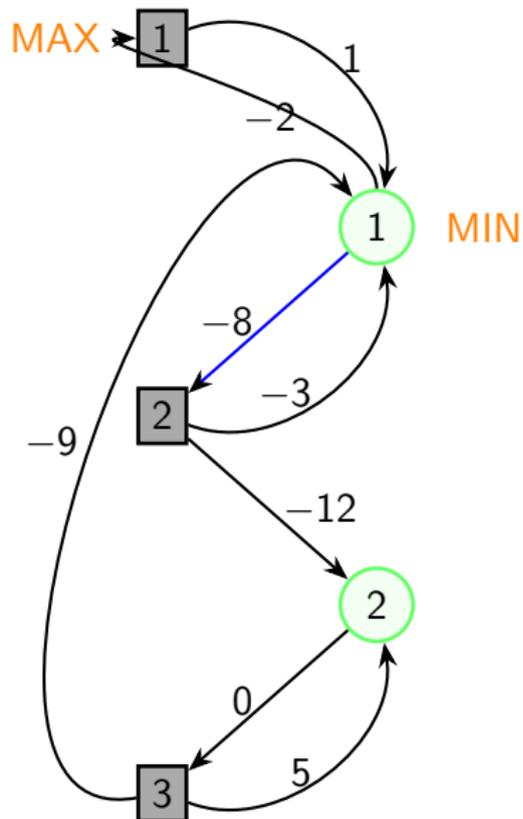
Baby example



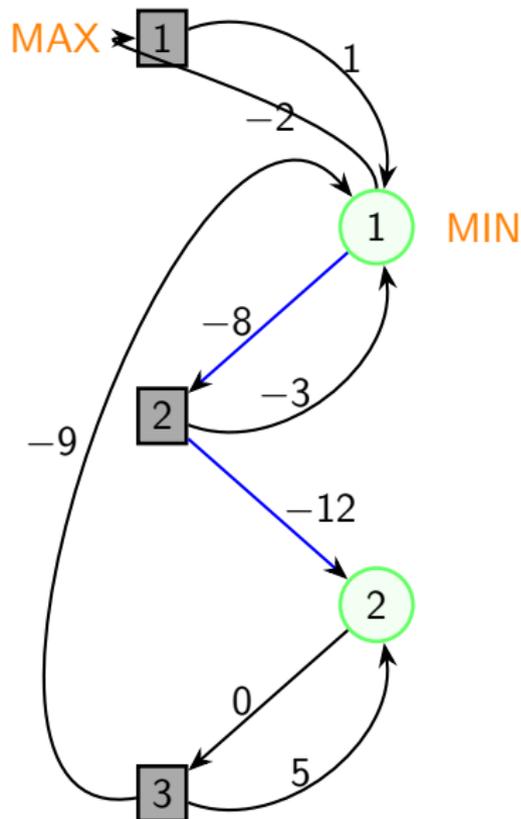
Baby example



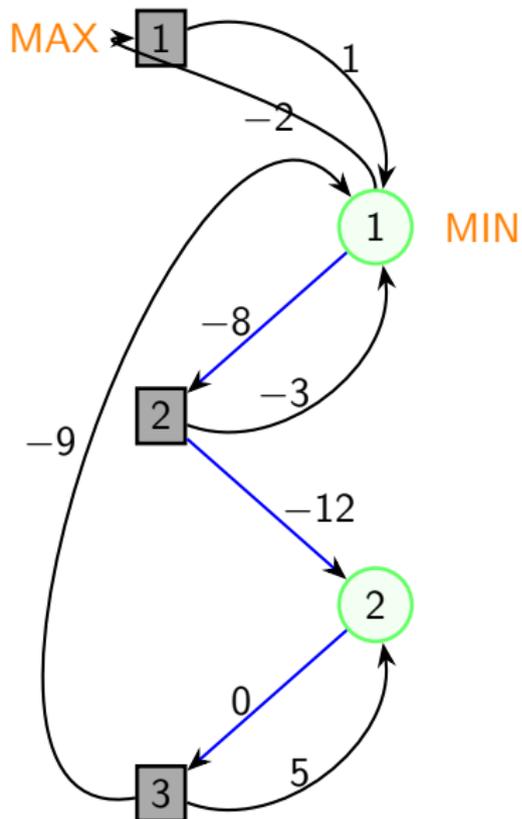
Baby example



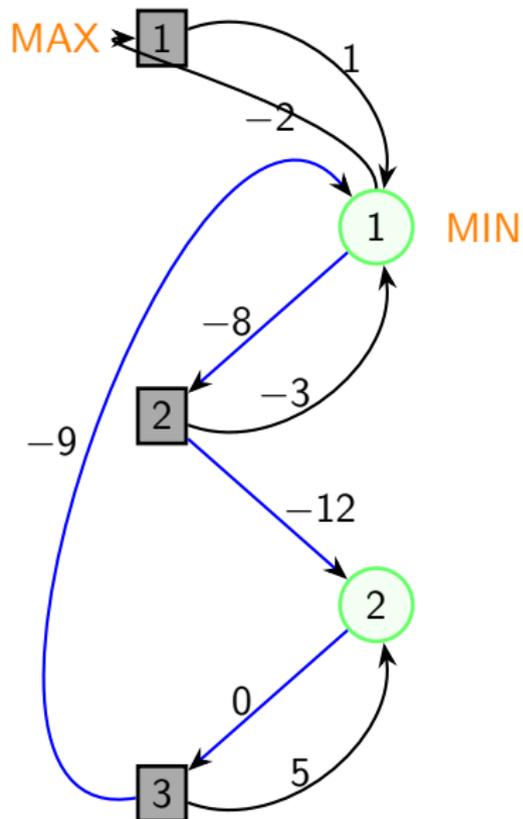
Baby example



Baby example



Baby example



After ℓ turns, under a strategy σ (resp. τ) of Min (resp. Max), if the sequence of actions is $i, k_1, i_1, \dots, k_\ell, i_\ell$, the total payment:

$$R_i^\ell(\sigma, \tau) = -V_{ik_1} + V_{i_1k_1} - V_{i_1k_2} - \dots + V_{i_\ell k_\ell}$$

The *value* v_i^ℓ of the game at horizon ℓ starting from i :

$$v_i^\ell := \min_{\sigma} \max_{\tau} R_i^\ell(\sigma, \tau) .$$

Shapley operator

Shapley operator $T : (\mathbb{R}_{\max})^n \rightarrow (\mathbb{R}_{\max})^n$,

$$T_i(x) = \min_{k \in [p], V_{ik} \neq -\infty} \left[-V_{ik} + \max_{j \in [n], j \neq i} (V_{jk} + x_j) \right], \quad i \in [n],$$

Theorem (Shapley)

$$v^\ell = T(v^{\ell-1}) = \dots = T^\ell(0).$$

Shapley operator

Shapley operator $T : (\mathbb{R}_{\max})^n \rightarrow (\mathbb{R}_{\max})^n$,

$$T_i(x) = \min_{k \in [p], V_{ik} \neq -\infty} \left[-V_{ik} + \max_{j \in [n], j \neq i} (V_{jk} + x_j) \right], \quad i \in [n],$$

Theorem (Shapley)

$$v^\ell = T(v^{\ell-1}) = \dots = T^\ell(0).$$

In the infinite horizon case, the *mean payoff vector*:

$$\chi(T) := \lim_{\ell \rightarrow \infty} T^\ell(0)/\ell.$$

Shapley operator

Shapley operator $T : (\mathbb{R}_{\max})^n \rightarrow (\mathbb{R}_{\max})^n$,

$$T_i(x) = \min_{k \in [p], V_{ik} \neq -\infty} \left[-V_{ik} + \max_{j \in [n], j \neq i} (V_{jk} + x_j) \right], \quad i \in [n],$$

Theorem (Shapley)

$$v^\ell = T(v^{\ell-1}) = \dots = T^\ell(0).$$

In the infinite horizon case, the *mean payoff vector*:

$$\chi(T) := \lim_{\ell \rightarrow \infty} T^\ell(0)/\ell.$$

- Complexity: $\text{NP} \cap \text{coNP}$.

Shapley operator

Shapley operator $T : (\mathbb{R}_{\max})^n \rightarrow (\mathbb{R}_{\max})^n$,

$$T_i(x) = \min_{k \in [p], V_{ik} \neq -\infty} \left[-V_{ik} + \max_{j \in [n], j \neq i} (V_{jk} + x_j) \right], \quad i \in [n],$$

Theorem (Shapley)

$$v^\ell = T(v^{\ell-1}) = \dots = T^\ell(0).$$

In the infinite horizon case, the *mean payoff vector*:

$$\chi(T) := \lim_{\ell \rightarrow \infty} T^\ell(0)/\ell.$$

- Complexity: $\text{NP} \cap \text{coNP}$.
- **Zwick and Paterson [1996]**: Value iteration (pseudo-polynomial time).

Shapley operator

Shapley operator $T : (\mathbb{R}_{\max})^n \rightarrow (\mathbb{R}_{\max})^n$,

$$T_i(x) = \min_{k \in [p], V_{ik} \neq -\infty} \left[-V_{ik} + \max_{j \in [n], j \neq i} (V_{jk} + x_j) \right], \quad i \in [n],$$

Theorem (Shapley)

$$v^\ell = T(v^{\ell-1}) = \dots = T^\ell(0).$$

In the infinite horizon case, the *mean payoff vector*:

$$\chi(T) := \lim_{\ell \rightarrow \infty} T^\ell(0)/\ell.$$

- Complexity: $\text{NP} \cap \text{coNP}$.
- **Zwick and Paterson [1996]**: Value iteration (pseudo-polynomial time).
- Our case: $\chi(T) \leq 0$.

Equivalence between tropical linear regression and MPG

The *spectral radius* of T is defined as

$$\rho(T) = \sup\{\lambda \in \mathbb{R} \cup \{-\infty\} \mid \exists u \in (\mathbb{R}_{\max})^n, u \neq -\infty, T(u) = \lambda + u\} .$$

Theorem (Akian–Gaubert–Q.–Saadi)

$$\min_{a \in \mathbb{P}(\mathbb{R}_{\max})^n} d_H(V, H_a) = -\rho(T) = \text{in-rad}(V)$$

Moreover,

- if $T(a) \geq \rho(T) + a$, then H_a is optimal
- if $T(b) \leq \rho(T) + b$, then $B(-b, -\rho(T))$ is optimal

Equivalence between tropical linear regression and MPG

The *spectral radius* of T is defined as

$$\rho(T) = \sup\{\lambda \in \mathbb{R} \cup \{-\infty\} \mid \exists u \in (\mathbb{R}_{\max})^n, u \not\equiv -\infty, T(u) = \lambda + u\} .$$

Theorem (Akian–Gaubert–Q.–Saadi)

$$\min_{a \in \mathbb{P}(\mathbb{R}_{\max})^n} d_H(V, H_a) = -\rho(T) = \text{in-rad}(V)$$

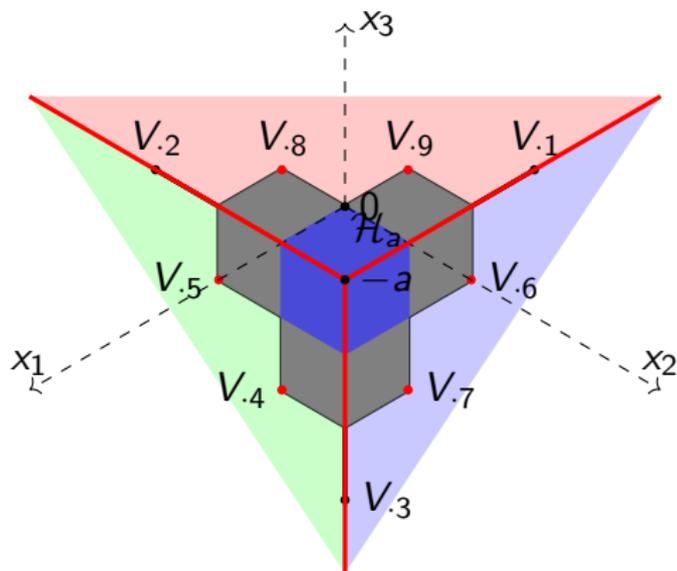
Moreover,

- if $T(a) \geq \rho(T) + a$, then H_a is optimal
- if $T(b) \leq \rho(T) + b$, then $B(-b, -\rho(T))$ is optimal

Corollary

The tropical linear regression problem is polynomial-time equivalent to the problem of solving a mean payoff game.

Geometric illustration



$$V = \begin{pmatrix} -3 & 0 & 0 & 1 & 1 & -1 & 0 & 0 & -1 \\ 0 & -3 & 0 & 0 & -1 & 1 & 1 & -1 & 0 \\ -1 & -1 & -4 & -2 & -1 & -1 & -2 & 0 & 0 \end{pmatrix}$$

Theorem (Akian–Gaubert–Q.–Saadi)

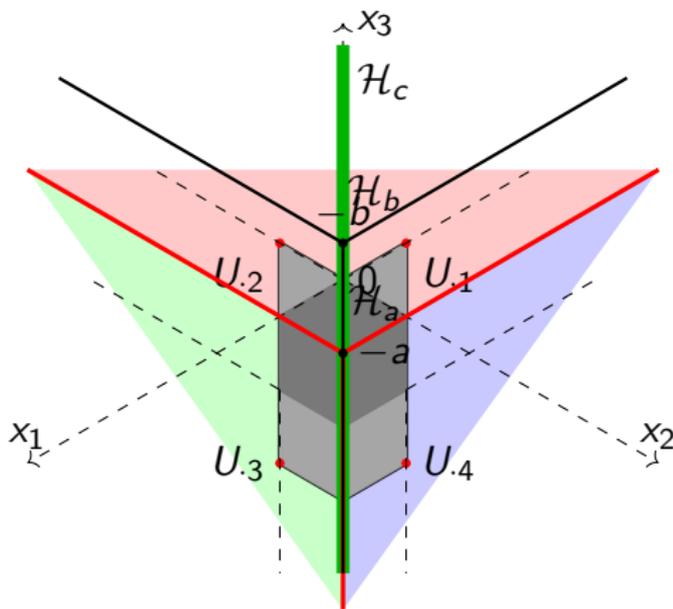
For $a \in \mathbb{R}^n$, the following assertions are equivalent:

- $T(a) = \rho(T) + a$;
- The hyperplane \mathcal{H}_a admits a “witness” point in each sector, i.e.,
 $\forall i \in [n], \exists k \in [p], v^{(k)} \in S_i(a)$ and $\text{dist}_H(v^{(k)}, \mathcal{H}_a) = \text{dist}_H(\mathcal{V}, \mathcal{H}_a)$.

Moreover, if these assertions hold, then

- $\rho(T) = -\text{dist}_H(\mathcal{V}, \mathcal{H}_a)$
- \mathcal{H}_a is an optimal solution of the tropical linear regression problem
- $B(-a, \text{dist}_H(\mathcal{V}, \mathcal{H}_a))$ is a Hilbert ball of maximal radius included in $\text{Sp}(\mathcal{V})$.

Nonunique optimal solutions



$\text{Col}(U)$ with multiple optimal hyperplanes and multiple optimal inner balls, but a unique optimal hyperplane with witness points in each sector.

Revisit to equilibrium in ITT

Equilibrium:

$$\min_{i \in [q]} p_{ij} f_i^{-1} \text{ is achieved twice at least}$$

By letting $V_{ij} = -\log(p_{ij})$ and $a_i = \log(f_i)$, the equilibrium is:

$$\max_{i \in [n]} (V_{ij} + a_i) \text{ is achieved at least twice,}$$

namely,

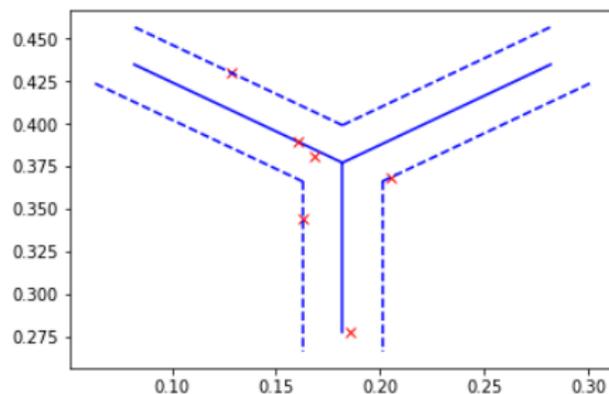
$$\forall j \in [q], \quad V_{\cdot j} \in H_a$$

In practice, we solve the **tropical linear regression problem**:

$$\min_{b \in \mathbb{P}(\mathbb{R}_{\max})^n} d_H(V, H_b).$$

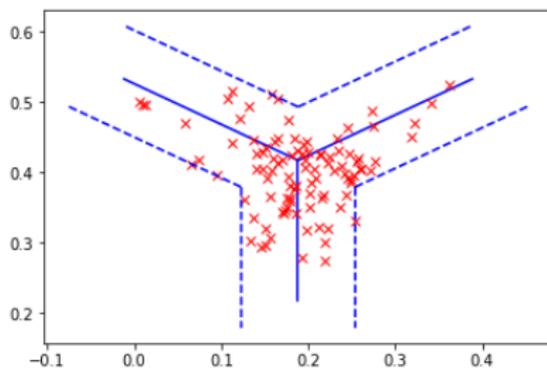
Example

	ind. houses	social housing	school	road	stadium	bridge	f	f^{reg}
Firm 1	1.02	3.21	<u>8.72</u>	26.2	69.8	<u>123</u>	1	1
Firm 2	0.81	2.65	7.49	20.3	<u>53.8</u>	106	0.8	0.81
Firm 3	<u>0.6</u>	<u>1.86</u>	5.5	<u>14.7</u>	41.8	76	0.6	0.605



Algorithm

- Goal: $T(v) = \rho(T) + v$
- Algorithm: projective Krasnoselkii-Mann value iteration algorithm
 - Given $\epsilon > 0$, start with $v^0 = (0, \dots, 0)^\top$.
 - If $\|T(v^k) - v^k\| \geq \epsilon$, let
$$\tilde{v}^{k+1} = T(v^k) - (\max_{i \in [n]} T(v^k)_i) e,$$
$$v^{k+1} = (1 - \beta)v^k + \beta\tilde{v}^{k+1},$$
where $e = (1, \dots, 1)^\top \in \mathbb{R}^n$ and $\beta \in (0, 1)$ fixed.



Theorem (Akian–Gaubert–Q.–Saadi)

Suppose that $V \in \mathbb{R}^{n \times p}$ is finite, and let

$$W := \max_{k \in [p]} \|V_{\cdot, k}\|_H .$$

Then, an ϵ -approximation of $\text{in-rad}(V)$, and vectors $v, z \in \mathbb{R}^n$ satisfying

$$B_H(v, \text{in-rad}(\mathcal{V}) - \epsilon) \subseteq \text{Span}(V)$$

and

$$d_H(\text{Span}(V), H_z) \leq \text{in-rad}(\mathcal{V}) + \epsilon$$

can be obtained in $O(npW/\epsilon)$ arithmetic operations.

Overview

- 1 Introduction
- 2 Tropical linear regression
- 3 Tropical low-rank approximation**
- 4 Generalization

Let $\phi : \mathbb{R}_{\geq 0}^p \rightarrow \mathbb{R}_{\geq 0}$ be a positive, continuous, order-preserving, positively homogeneous of degree 1 function. Given a norm or a semi-norm $\|\cdot\|$ on \mathbb{R}^n , define the map $\|\cdot\|_\phi$ on $\mathbb{R}^{n \times p}$ by

$$\forall A \in \mathbb{R}^{n \times p}, \quad \|A\|_\phi = \phi(\|A_{\cdot 1}\|, \dots, \|A_{\cdot p}\|) .$$

If ϕ is subadditive ($\phi(x + y) \leq \phi(x) + \phi(y)$), $\|A\|_\phi$ is a (semi-)norm.

Theorem

For $M \in (\mathbb{R}_{\max})^{n \times p}$ and $r \leq \min(n, p)$,

$$\begin{aligned} & \min_{A \in (\mathbb{R}_{\max})^{n \times r}, B \in (\mathbb{R}_{\max})^{r \times p}} d_{\phi, \|\cdot\|}(M, AB) \\ = & \min_{A \in (\mathbb{R}_{\max})^{n \times r}} \phi(d_{\|\cdot\|}(M_{\cdot 1}, \text{Col}(A)), \dots, d_{\|\cdot\|}(M_{\cdot p}, \text{Col}(A))) \end{aligned}$$

Outer radius

For a subset \mathcal{V} of $(\mathbb{R}_{\max})^n$, the *outer radius* of \mathcal{V} is:

$$\text{out-rad}(\mathcal{V}) := \inf\{r \geq 0 \mid \exists b \in \mathbb{R}^n, B(b, r) \supseteq \text{Span}(\mathcal{V})\}.$$

Theorem (Akian–Gaubert–Q.–Saadi)

For any matrix $V \in \mathbb{R}^{n \times p}$, we have

$$\min_{A \in \mathbb{R}^{n \times p}, \text{rank } A=1} \|V - A\|_{\infty} = \frac{1}{2} \text{out-rad}(\text{Span } V).$$

Question: how to compute $\text{out-rad}(\text{Span } V)$?

Reduction to eigenvalue problem

Given $A \in (\mathbb{R}_{\max})^{n \times n}$, define

$$A^* = I \oplus A \oplus A^2 \oplus \dots$$

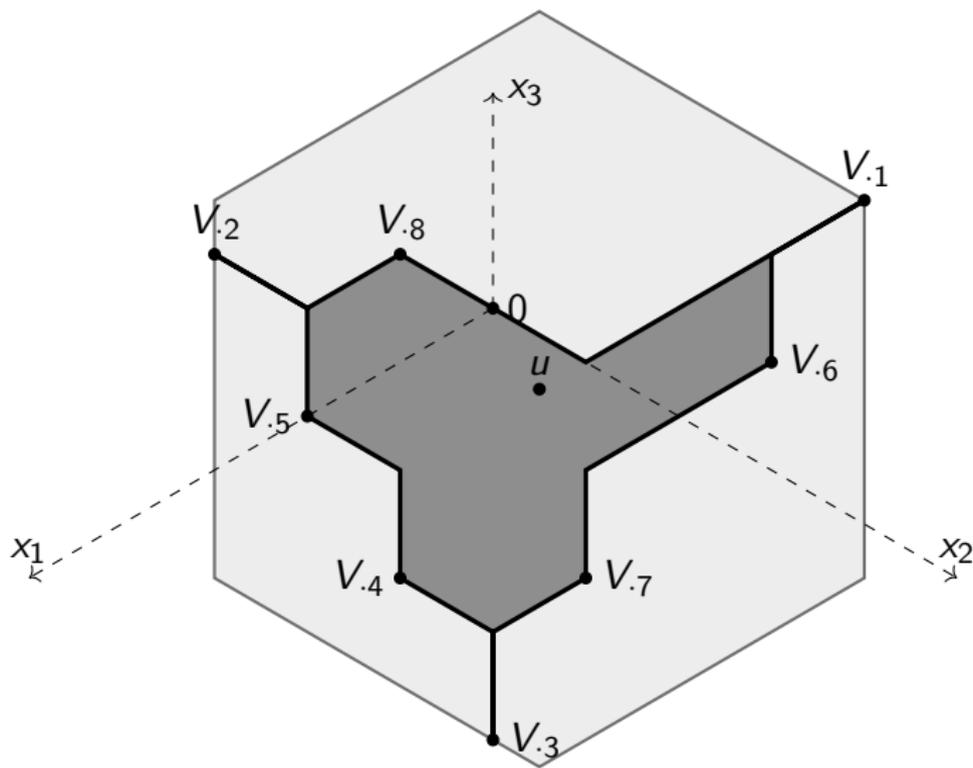
For $V \in \mathbb{R}^{n \times p}$, let $H = V \odot (-V^T) \in \mathbb{R}^{n \times n}$, where

$$H_{ik} = \max_{j \in [p]} (V_{ij} - V_{kj}), \quad i, k \in [n].$$

Theorem (Akian–Gaubert–Q.–Saadi)

H has a unique eigenvalue, which equals $\text{out-rad}(\text{Span } V)$. Moreover, the set of centers of all Hilbert outer balls of $\text{Span } V$ is the column space of $(-\lambda + H)^$.*

Example



Kernel approximations

X : compact metric, Y : nonempty, $\mathcal{C}(X)$: the space of continuous functions on X , $\mathcal{B}(Y)$: the space of bounded functions on Y .

Theorem (Akian–Gaubert–Q.–Saadi)

If $V : X \times Y \rightarrow \mathbb{R}$ bounded and $\{V(\cdot, y)\}_{y \in Y}$ equicontinuous, then

$$\inf_{f \in \mathcal{C}(X), g \in \mathcal{B}(Y)} \sup_{x \in X, y \in Y} |V(x, y) - f(x) - g(y)|$$

achieves an optimal solution, which is equal to one half of the tropical eigenvalue of H , where

$$H(x, z) = \sup_{y \in Y} (V(x, y) - V(z, y)) ,$$

and f is a tropical eigenvector of H .

Corollary (Akian–Gaubert–Q.–Saadi)

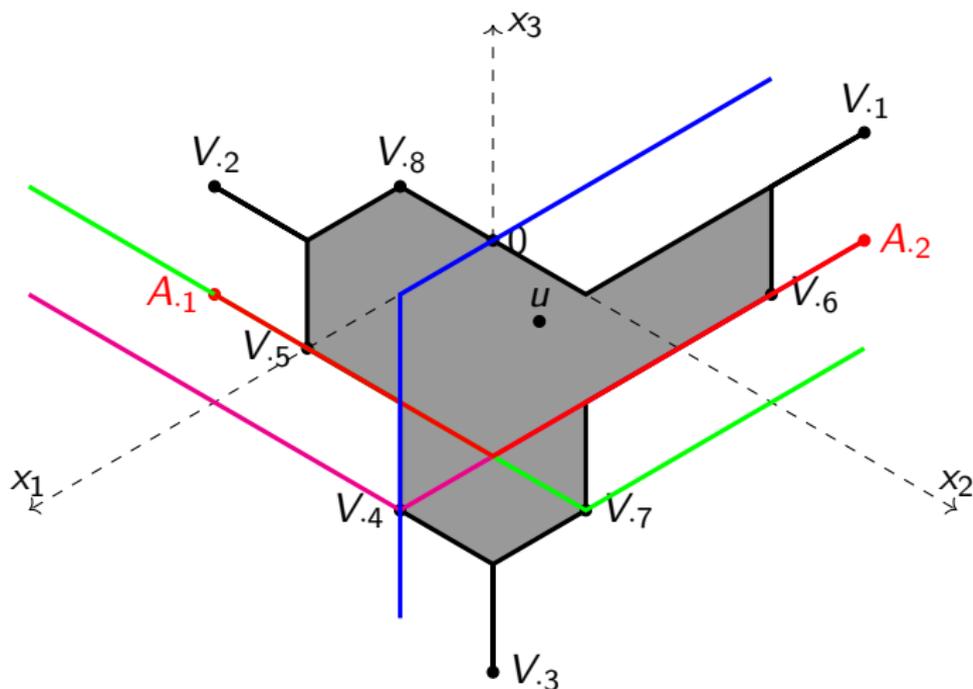
A best rank-two approximation of the matrix $V \in \mathbb{R}^{3 \times p}$ can be obtained by tropical linear regression:

$$\begin{aligned} \min_{A \in (\mathbb{R}_{\max})^{3 \times 2}, B \in (\mathbb{R}_{\max})^{2 \times p}} d(V, AB) &= \min_{A \in (\mathbb{R}_{\max})^{3 \times 2}} \max_{k \in [p]} d_H(V_{\cdot k}, \text{Col}(A)) \\ &= \max_{k \in [p]} d_H(V_{\cdot k}, \mathcal{H}_{a^*}) , \end{aligned}$$

and needs $O(p)$ arithmetic operations.

Provides us a way to compute a best rank-2 approximation for $V \in \mathbb{R}^{n \times p}$.

Example



Overview

- 1 Introduction
- 2 Tropical linear regression
- 3 Tropical low-rank approximation
- 4 Generalization**

General setting

- Let \mathbb{K} be the field of Puiseux series over \mathbb{C} , namely any $f \in \mathbb{K}$ looks like

$$f = \sum_{k=k_0}^{\infty} c_k T^{k/n},$$

where $c_k \in \mathbb{C}$ and $\text{val}(f) = k_0/n$ is the natural valuation.

- A projective variety $\mathcal{X} \subseteq \mathbb{K}\mathbb{P}^{N-1}$ is called *nondegenerate* if \mathcal{X} is not contained in any hyperplane.
- Let $X = \text{val}(\widehat{\mathcal{X}})$. Then the closure \overline{X} is a tropical variety in \mathbb{R}_{\min}^N .
- Denote by $\mathcal{B} = \{e_1, \dots, e_N\}$ the standard basis of \mathbb{R}_{\min}^N , where e_i is the vector whose i th entry is 0 and other entries are ∞ .
- Define $\text{Var}(\mathbb{K}\mathbb{P}^{N-1})$ to be

$$\{\mathcal{X} \subseteq \mathbb{K}\mathbb{P}^{N-1} \mid \mathcal{X} \text{ is (i) irreducible, (ii) nondegenerate, (iii) } \mathcal{B} \subseteq X\}.$$

Given $\mathcal{X} \in \text{Var}(\mathbb{KP}^{N-1})$,

- The minimum integer r such that $p = x_1 + \cdots + x_r$ is called the \mathcal{X} -rank of p and denoted by $\text{rank}_{\mathcal{X}}(p) = r$, where $x_1, \dots, x_r \in \widehat{\mathcal{X}}$.
- The minimum integer r such that $p \in \mathbb{R}_{\min}^N$ can be written as $p = \min\{x_1, \dots, x_r\}$ for some $x_1, \dots, x_r \in X$ is called the *factor X-rank*, or *Barvinok X-rank*, and written as $\text{Frank}_{\mathcal{X}}(p) = r$.
- The Kapranov X -rank of p is defined to be

$$\text{Krank}_{\mathcal{X}}(p) = \min\{r \mid \tilde{p} = x_1 + \cdots + x_r, \text{val}(\tilde{p}) = p, x_1, \dots, x_r \in \widehat{\mathcal{X}}\}.$$

Theorem

$$\text{Krank}_{\mathcal{X}}(p) \leq \text{Frank}_{\mathcal{X}}(p)$$

Matrix case

For $M \in \mathbb{R}_{\min}^{d \times n}$,

- Frank(M): smallest r for which

$$M = \min_{1 \leq i \leq r} \{u_i \odot v_i^\top\},$$

where $u_i \in \mathbb{R}_{\min}^d$, $v_i \in \mathbb{R}_{\min}^n$.

- Krank(M): smallest dimension of any tropical linear space containing the columns of M .
- Tropical rank Trank(M): largest r such that M has a nonsingular $r \times r$ minor, where N is *singular* if

$$\det(N) = \bigoplus_{\sigma \in \mathcal{S}_r} N_{1\sigma_1} \odot \cdots \odot N_{r\sigma_r} = \min\{N_{1\sigma_1} + \cdots + N_{r\sigma_r}\}$$

is attained at least twice.

Theorem (Develin–Santos–Sturmfels)

$$\text{Trank}(M) \leq \text{Krank}(M) \leq \text{Frank}(M)$$

Tropical typical ranks

- The integer r_g such that the set $\{p \in \mathbb{K}^N \mid \text{rank}_{\mathcal{X}}(p) = r_g\}$ contains a nonempty Zariski open subset of \mathbb{K}^N is called the *generic \mathcal{X} -rank*.
- If the set $\{p \in \mathbb{R}_{\min}^N \mid \text{Frank}_{\mathcal{X}}(p) = r\}$ contains a nonempty Euclidean open subset of \mathbb{R}_{\min}^N , then r is called a *tropical typical \mathcal{X} -rank*.

Theorem (Blekherman–Teitler)

The minimum real typical rank equals the generic rank.

Theorem (Akian–Gaubert–Q.)

- *If m and M are two tropical typical \mathcal{X} -ranks, then any integer between m and M is also a tropical typical \mathcal{X} -rank.*
- *The minimum tropical typical \mathcal{X} -rank is the generic \mathcal{X} -rank.*
- *The maximum tropical typical \mathcal{X} -rank is the maximum factor \mathcal{X} -rank.*

Consider $\mathbb{R}_{\min}^{n_1} \otimes \cdots \otimes \mathbb{R}_{\min}^{n_d}$, and assume $n_1 \leq \cdots \leq n_d$.

Example

The typical tropical tensor ranks are integers between the generic rank r_g and $n_1 \cdots n_{d-1}$.

- Tropical linear regression is equivalent to finding inner radius, which is polynomial-time equivalent to mean payoff game.
- Finding a best tropical rank-one approximation is equivalent to finding outer radius, which is equivalent to eigenvalue problem.
- A rank-2 approximation can be obtained by tropical linear regression.

Thank you for your attention!