

Citation for published version:

Hook, J 2019, 'Max-plus linear inverse problems: 2-norm regression and system identification of max-plus linear dynamical systems with Gaussian noise', *Linear Algebra and its Applications*, vol. 579, pp. 1-31.
<https://doi.org/10.1016/j.laa.2019.05.024>

DOI:

[10.1016/j.laa.2019.05.024](https://doi.org/10.1016/j.laa.2019.05.024)

Publication date:

2019

Document Version

Peer reviewed version

[Link to publication](#)

Publisher Rights

CC BY-NC-ND

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Max-plus Linear Inverse Problems: 2-norm regression and system identification of max-plus linear dynamical systems with Gaussian noise

James Hook

January 2, 2019

Abstract

In this paper we present new theory and algorithms for 2-norm regression over the max-plus semiring. As an application we also show how max-plus 2-norm regression can be used in system identification of max-plus linear dynamical systems with Gaussian noise. We also introduce and provide methods for solving a max-plus linear inverse problem with regularization, which can be used when the the original problem is not well posed.

1 Introduction

Max-plus algebra concerns the max-plus semiring $\mathbb{R}_{\max} = [\mathbb{R} \cup \{-\infty\}, \oplus, \otimes]$, with

$$a \oplus b = \max\{a, b\}, \quad a \otimes b = a + b, \quad \text{for all } a, b \in \mathbb{R}_{\max}. \quad (1)$$

A max-plus matrix is an array of elements from \mathbb{R}_{\max} and max-plus matrix multiplication is defined in analogy to the classical (i.e. not max-plus) case. For $A \in \mathbb{R}_{\max}^{n \times d}$ and $B \in \mathbb{R}_{\max}^{d \times m}$ we have $A \otimes B \in \mathbb{R}_{\max}^{n \times m}$ with

$$(A \otimes B)_{ij} = \max_{k=1}^d (a_{ik} + b_{kj}), \quad (2)$$

for $i = 1, \dots, n$, $j = 1, \dots, m$. Max-plus algebra has found a wide range of applications in operations research, dynamical systems and control, see for example [4, 8, 13] and the references therein. This paper concerns the following max-plus regression problem with $p = 2$. To the best of our knowledge we are the first to study this problem explicitly.

Problem 1.1. For $A \in \mathbb{R}_{\max}^{n \times d}$, $\mathbf{y} \in \mathbb{R}_{\max}^n$ and $p \geq 1$, we seek

$$\min_{\mathbf{x} \in \mathbb{R}_{\max}^d} \|A \otimes \mathbf{x} - \mathbf{y}\|_p. \quad (3)$$

In order for Problem 1.1 to be well defined we need to be able to measure the p -norm distance between arbitrary vectors in \mathbb{R}_{\max}^n , which we do as follows. For $\mathbf{x} \in \mathbb{R}_{\max}^n$ define $\text{support}(\mathbf{x}) \subseteq \{1, \dots, n\}$, by $i \in \text{support}(\mathbf{x}) \Leftrightarrow \mathbf{x}_i > -\infty$. Then for $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\max}^n$, we use the convention that

$$\|\mathbf{x} - \mathbf{y}\|_p = \begin{cases} \|\mathbf{x}_{\text{support}(\mathbf{x})} - \mathbf{y}_{\text{support}(\mathbf{y})}\|_p, & \text{if } \text{support}(\mathbf{x}) = \text{support}(\mathbf{y}), \\ \infty, & \text{otherwise,} \end{cases} \quad (4)$$

where $\mathbf{x}_{\text{support}(\mathbf{x})}$ is the sub-vector of \mathbf{x} formed from its finite entries.

1.1 System identification

Virtually every application of max-plus algebra in dynamical systems and control exploit its ability to model certain classically non-linear phenomena in a linear way, as illustrated in the following example.

Example 1.2. Consider a distributed computing system in which d processors update a computation in parallel. At each stage processor i must wait until it has received input from its neighboring processors before beginning its next local computation. Then after completing its local computation it must broadcast some output back to its neighboring processors. Define the vectors of update times $\mathbf{t}(0), \dots, \mathbf{t}(N) \in \mathbb{R}_{\max}^d$, by $\mathbf{t}(n)_i =$ the time at which processor i completes its n th local computation. These update times can be modeled by

$$\mathbf{t}(n+1) = M \otimes \mathbf{t}(n), \quad \text{for } n = 0, \dots, N-1, \quad (5)$$

where $M \in \mathbb{R}_{\max}^{d \times d}$ is the max-plus matrix given by

$$m_{ij} = \begin{cases} a_i + c_{ij}, & \text{if } j \in J_i, \\ -\infty, & \text{otherwise,} \end{cases} \quad (6)$$

where a_i is the time taken for processor i 's local computation, c_{ij} is the time taken for communication from processor j and processor i receives input from the processors $J_i \subset \{1, \dots, d\}$, for $i, j = 1, \dots, d$. The update rule (5) constitutes a max-plus linear dynamical system. By studying the max-plus algebraic properties of the matrix M we can now predict the behavior of the system, for example computing its leading eigenvalue to determine the average update rate of the computations iteration.

Using the petri-net formalism, such max-plus linear models can be derived for more complicated systems of interacting timed events [8, Chapter 7]. These linear models can be extended by introducing stochasticity, which in the above example could model random variability in the time taken for messages to pass through the computer network [9, 8, Chapter 11], by allowing the system to switch between one of several governing max-plus linear equations [16], or by including a controller input [13]. This approach has been used to model a wide variety of processes including the Dutch railway system [8, Chapter 8], mRNA translation [3] and the Transmission Control Protocol (TCP) [1].

In this context *forwards problems* arise by presupposing a dynamical systems model then asking questions about how its orbits must behave. Conversely an *inverse problem* is to infer a dynamical systems model from an empirical time series recording. In the control theory literature this inverse problem is referred to as *system identification*. For example in [6, 14, 17, 7] the authors present methods for system identification of stochastic max-plus linear control systems. These methods, which can be applied to a very wide class of system, with non-Gaussian noise processes, work by formulating a non-linear programming problem for the unknown system parameters, which is then solved using one of several possible standard gradient based algorithm. However, the resulting problems are necessarily non-smooth and non-convex, which makes the optimization difficult.

1.2 This paper's contribution

In this paper we introduce the max-plus 2-norm regression problem and show how system identification of max-plus linear dynamical systems with Gaussian noise can be expressed in this way.

Doing so enables us to apply existing theory from the max-plus linear algebra literature, along with some new theory developed in this paper, to better understand the geometric and combinatorial aspects of the system identification problem. In particular we are able to demonstrate the geometric cause of the non-convexity and prove that in general determining whether or not a point is a local minimum of the residual is an NP-hard problem.

We provide two algorithms. One is an exponential cost method that is guaranteed to return the global minimum. The other is based on Newton’s method and has cost $\Theta(nd)$ per iteration but cannot be guaranteed to converge even to a local minimum, although as demonstrated in the example problem we find that it tends to work well in practice.

We present a regularized version of the max-plus 2-norm regression problem, which may be useful in data analysis applications. We also present a method to solve this regularized problem which we name Iteratively Reshifted Least Squares since it bears a striking similarity to the well known Iteratively Reweighted Least Squares method, which is used to solve classical 1-norm regularized 2-norm regression problems.

The remainder of this paper is organized as follows. In Section 2 we introduce max-plus regression with a small example demonstrating the $p = 2$ and $p = \infty$ cases. In Section 2.1 we make a detailed study of the $p = 2$ case. In Subsection 2.2 we develop some necessary machinery for working with max-plus sets and functions. In Subsection 2.3 we show how any max-plus 2-norm regression problem can be reduced to a smaller problem with finite right hand side and at least one finite entry per row. In Subsections 2.3.1, 2.3.2, 2.3.3 and 2.3.4 we develop theory for the combinatorial and geometric aspects of the max-plus 2-norm regression problem before giving a small explicit example. In Section 2.2 we prove the NP-hardness result. In Sections 2.3 and 2.4 we introduce our two algorithms and in Section 3 we apply our new theory and algorithms to a systems identification example using both the original and regularized forms of the problem.

2 Max-plus regression

The column space of a max-plus matrix $A \in \mathbb{R}_{\max}^{n \times d}$ is simply the image of the matrix vector multiplication map

$$\text{col}(A) = \{A \otimes \mathbf{x} : \mathbf{x} \in \mathbb{R}_{\max}^d\}. \quad (7)$$

Just as in the classical case, the p -norm regression problem can be written as an optimization over the column space of the matrix

$$\min_{\mathbf{x} \in \mathbb{R}_{\max}^d} \|A \otimes \mathbf{x} - \mathbf{y}\|_p = \min_{\mathbf{z} \in \text{col}(A)} \|\mathbf{z} - \mathbf{y}\|_p. \quad (8)$$

Understanding the geometry of the column space is therefore key to understanding the regression problem. For more detail on max-plus linear spaces see [12] and the references therein.

Example 2.1. *Consider*

$$A = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

The column space of A is given by the union of two simplices

$$\text{col}(A) = \{[x_1, x_1 + 1, x_2 + 1]^\top : x_2 \leq x_1 \leq x_2 + 1\} \cup \{[x_2, x_1 + 1, x_2 + 1]^\top : x_1 \leq x_2 \leq x_1 + 1\}.$$

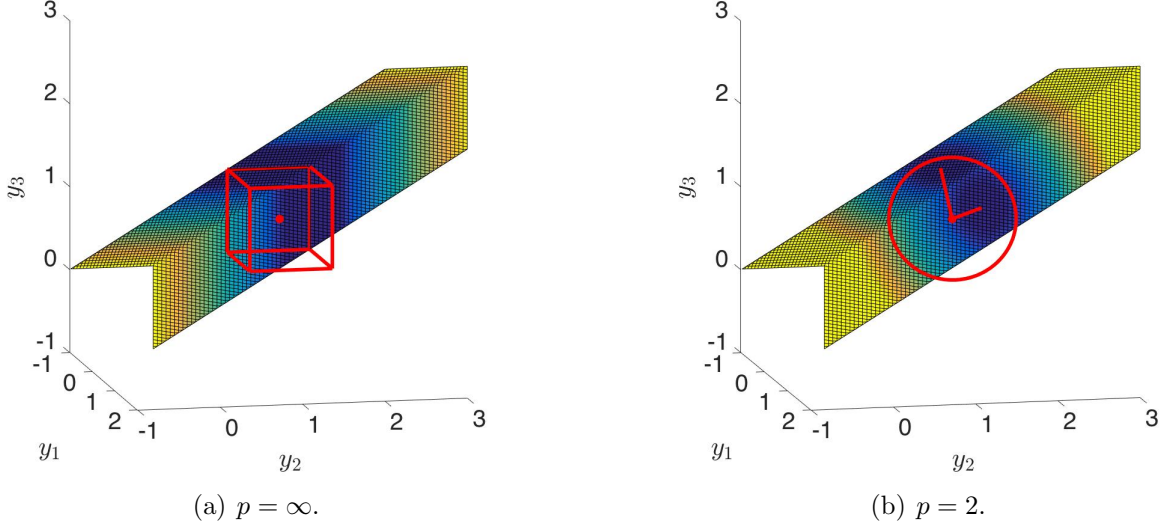


Figure 1: Column space view of max-plus regression.

Equivalently $\text{col}(A)$ is a prism with an L-shaped cross section

$$\text{col}(A) = \{\mathbf{l} + [\alpha, \alpha, \alpha]^\top : \mathbf{l} \in L, \alpha \in \mathbb{R}_{\max}\},$$

where

$$L = \{[0, t, 1]^\top : t \in [0, 1]\} \cup \{[0, 1, t]^\top : t \in [0, 1]\}.$$

Now consider Problem 1.1 with $p = \infty$. Figure 1 (a) displays the column space of A along with the target vector \mathbf{y} . We have also plotted the ball

$$B_\infty(\mathbf{y}, 1/2) = \{\mathbf{y}' \in \mathbb{R}_{\max}^3 : \|\mathbf{y}' - \mathbf{y}\|_\infty = 1/2\},$$

which is the smallest such ball that intersects $\text{col}(A)$. Therefore the minimum value of the residual is $1/2$ and the closest points in the column space are given by the L-shaped set

$$\arg \min_{z \in \text{col}(A)} \|z - \mathbf{y}\|_\infty = \{[1/2, 3/2, t]^\top : t \in [1/2, 3/2]\} \cup \{[1/2, t, 3/2]^\top : t \in [1/2, 3/2]\}.$$

Next consider Problem 1.1 with $p = 2$. Figure 1 (b) displays the column space of A along with the target vector \mathbf{y} . We have also plotted the ball

$$B_2(\mathbf{y}, 1/\sqrt{2}) = \{\mathbf{y}' \in \mathbb{R}_{\max}^3 : \|\mathbf{y}' - \mathbf{y}\|_2 = 1/\sqrt{2}\},$$

which is the smallest such ball that intersects $\text{col}(A)$. Therefore the minimum value of the residual is $1/\sqrt{2}$ and the closest points in the column space are given by

$$\arg \min_{z \in \text{col}(A)} \|z - \mathbf{y}\|_\infty = \{[1/2, 3/2, 1]^\top, [1/2, 1, 3/2]^\top\}.$$

We saw in Example 2.1 that the max-plus ∞ -norm regression problem could support multiple optimal solutions comprising a non-convex set. However it can be shown that the max-plus ∞ -norm regression problem is convex with respect to max-plus algebra [5] and that we can fairly

easily compute an optimal solution for it [4, Section 3.5]. We also saw that the max-plus 2-norm regression problem could support multiple isolated local minima and was therefore non-convex. However, because these local minima do not form a single, path connected set, this example also shows that the problem is max-plus non-convex.

2.1 2-norm regression

Since the max-plus 2-norm regression problem is non-convex in both classical and max-plus senses it is more difficult to solve than the ∞ -norm problem. However it may be more useful in application as 2-norm regression corresponds to computing a maximum likelihood estimate for a Bayesian inverse problem with Gaussian noise, as illustrated by the example in Section 3. In the remainder of this section we make a detailed study of the max-plus 2-norm regression problem.

2.2 Max-plus closure

Our approach to max-plus regression will be to develop classical mathematical description of all of the different max-plus objects that play an important role in the problem. In order to do this we need to be able to take closures of sets in \mathbb{R}_{\max}^n and extend or continue the definition of certain functions from \mathbb{R}^n to \mathbb{R}_{\max}^n .

We say that the sequence $(a_t \in \mathbb{R})_{t=1}^{\infty}$ converges to $\hat{a} \in \mathbb{R}_{\max}$, if either $\hat{a} \in \mathbb{R}$ and the sequence converges to \hat{a} in the usual sense, or $\hat{a} = -\infty$ and for all $m \in \mathbb{R}$, there exists t_0 such that $a_t \leq m$, for all $t \geq t_0$. This definition of convergence is extended componentwise to vectors in the obvious way. The max-plus closure of $X \subseteq \mathbb{R}^n$, denoted $\text{Cl}(X)$, consist of all such limit points of sequences in X . For example $\text{Cl}(\{0\}) = \{0\}$, $\text{Cl}(\{\mathbf{x} \in \mathbb{R}^2 : \mathbf{x}_1 > \mathbf{x}_2\}) = \{\mathbf{x} \in \mathbb{R}_{\max}^2 : \mathbf{x}_1 \geq \mathbf{x}_2\}$ and $\text{Cl}(\mathbb{R}^n) = \mathbb{R}_{\max}^n$.

A function $f : \mathbb{R}^n \mapsto \mathbb{R}$ is *monotonic* if $f(\mathbf{x}) \leq f(\mathbf{y})$, for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{x} \leq \mathbf{y}$, where \leq is the standard partial order on \mathbb{R}^n . A sequence $(\mathbf{a}_t \in \mathbb{R}^n)_{t=1}^{\infty}$ is *monotonically non-increasing* if $\mathbf{a}_{t+1} \leq \mathbf{a}_t$, for all $t = 1, 2, \dots$. Given a uniformly continuous monotonic function $f : \mathbb{R}^n \mapsto \mathbb{R}$, we *continue* its definition to $f : \mathbb{R}_{\max}^n \mapsto \mathbb{R}_{\max}$ using the rule that

$$f(\hat{\mathbf{a}}) = \lim_{t \rightarrow \infty} f(\mathbf{a}_t), \quad (9)$$

whenever the monotonically non-increasing sequence $(\mathbf{a}_t \in \mathbb{R}^n)_{t=1}^{\infty}$ converges to $\hat{\mathbf{a}} \in \mathbb{R}_{\max}^n$.

Theorem 2.2. *Let $f : \mathbb{R}^n \mapsto \mathbb{R}$ be a uniformly continuous monotonic function. Then the continuation $f : \mathbb{R}_{\max}^n \mapsto \mathbb{R}_{\max}$ is well defined.*

Proof. Suppose that $(\mathbf{a}_t \in \mathbb{R}^n)_{t=1}^{\infty}$ is a monotonically non-increasing sequences that converges to $\hat{\mathbf{a}} \in \mathbb{R}_{\max}^n$. Consider the sequences $(\mathbf{b}_t)_{t=1}^{\infty}$ and $(\mathbf{c}_t)_{t=1}^{\infty}$ defined by

$$(\mathbf{b}_t)_i = \begin{cases} (\hat{\mathbf{a}})_i & \text{if } i \in \text{support}(\hat{\mathbf{a}}), \\ -t & \text{otherwise.} \end{cases} \quad (\mathbf{c}_t)_i = \begin{cases} (\hat{\mathbf{a}})_i & \text{if } i \in \text{support}(\hat{\mathbf{a}}), \\ (\mathbf{a}_t)_i & \text{otherwise.} \end{cases}$$

Clearly $(\mathbf{b}_t)_{t=1}^{\infty}$ also converges to $\hat{\mathbf{a}}$ and $\lim_{t \rightarrow \infty} f(\mathbf{b}_t)$ exists in the sense of convergence introduced above because it is the limit of a non-increasing sequence. We now show that $(f(\mathbf{a}_t))_{t=1}^{\infty}$ has the same limit.

By uniform continuity, for all $\epsilon > 0$ there exists $\delta > 0$, such that $|f(\mathbf{x}) - f(\mathbf{y})| \leq \epsilon$, whenever $\|\mathbf{x} - \mathbf{y}\|_2 \leq \delta$. By the definition of convergence given above, for all such $\epsilon, \delta > 0$ and $m \in \mathbb{N}$ there

exists $T(\epsilon, m)$ such that

$$(\mathbf{a}_t)_i \leq \begin{cases} (\hat{\mathbf{a}})_i + \frac{\delta}{\sqrt{n}} & \text{if } i \in \text{support}(\hat{\mathbf{a}}), \\ -m & \text{otherwise,} \end{cases}$$

for all $i = 1, \dots, n$ and all $t \geq T(\epsilon, m)$. Note that $\|\mathbf{c}_t - \mathbf{a}_t\|_2 \leq \delta$, so that $|f(\mathbf{c}_t) - f(\mathbf{a}_t)| \leq \epsilon$, and $\mathbf{c}_t \leq \mathbf{b}_m$, for all $t \geq T(\epsilon, m)$. Therefore

$$f(\mathbf{a}_t) \leq f(\mathbf{b}_m) + \epsilon$$

and taking the limit $\epsilon \rightarrow 0$ and the limit $m \rightarrow \infty$ we obtain

$$\lim_{t \rightarrow \infty} f(\mathbf{a}_t) \leq \lim_{t \rightarrow \infty} f(\mathbf{b}_t). \quad (10)$$

Next consider the sequence $(B(t) \in \mathbb{N})_{t=1}^\infty$, defined by

$$B(t) = \max \left\{ 0, -\min_{i=1}^n \lfloor (\mathbf{a}_t)_i \rfloor \right\},$$

where $\lfloor \cdot \rfloor$ denotes the integer floor. Then the subsequence $(\mathbf{b}_{B(t)})_{t=1}^\infty$ satisfies $\mathbf{b}_{B(t)} \leq \mathbf{a}_t$, for all $t = 1, 2, \dots$. Therefore

$$\lim_{t \rightarrow \infty} f(\mathbf{a}_t) \geq \lim_{t \rightarrow \infty} f(\mathbf{b}_{B(t)}) = \lim_{t \rightarrow \infty} f(\mathbf{b}_t),$$

which together with (10) completes the proof. \square

2.3 Reduction to finite form

In this section we show that any instance of Problem 1.1 can either, be reduced to a problem with finite right hand side and at least one finite entry per row, or does not admit any solution with finite residual.

We say that $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}_{\max}^n$ are a *finite form*, if $\text{support}(\mathbf{y}) = \{1, 2, \dots, n\}$ and $\max_{j=1}^d a_{ij} > -\infty$, for all $i = 1, \dots, n$. Also define the residual $R: \mathbb{R}_{\max}^d \mapsto \mathbb{R} \cup \{+\infty\}$, by

$$R(\mathbf{x}) = \|A \otimes \mathbf{x} - \mathbf{y}\|_2^2 / 2.$$

Theorem 2.3. *If $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}_{\max}^n$ are a finite form, then $R(\mathbf{x}) < \infty$, for all $\mathbf{x} \in \mathbb{R}^d$.*

Proof. From definition (4) and the fact that $\text{support}(\mathbf{y}) = \{1, 2, \dots, n\}$, we have that $R(\mathbf{x}) < \infty$ whenever

$$\text{support}(A \otimes \mathbf{x}) = \{1, 2, \dots, n\}.$$

Since each row of A contains at least one finite entry, there exists $c: \{1, \dots, n\} \mapsto \{1, \dots, d\}$, such that $a_{ic(i)} > -\infty$, for all $i = 1, \dots, n$. Then for any $\mathbf{x} \in \mathbb{R}^d$, we have

$$(A \otimes \mathbf{x})_i = \max_{j=1}^d a_{ij} + \mathbf{x}_j \geq a_{ic(i)} + \mathbf{x}_{c(i)} > -\infty,$$

for all $i = 1, \dots, n$. So that $R(\mathbf{x}) < \infty$. \square

We say that $j \in \{1, \dots, d\}$ is *col-admissible*, if

$$i \notin \text{support}(\mathbf{y}) \Rightarrow a_{ij} = -\infty.$$

We say that $i \in \text{support}(\mathbf{y})$ is *row-admissible*, if

$$\exists \text{ col-admissible } j : a_{ij} > -\infty.$$

Let \mathcal{C} denote the set of all col-admissible indices and \mathcal{R} denote the set of all row-admissible indices.

Lemma 2.4. Let $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}_{\max}^n$. Then

$$\text{support}(\mathbf{x}) \subseteq \mathcal{C} \quad (11)$$

is a necessary condition for $R(\mathbf{x}) < \infty$.

Proof. Suppose there exists $j \in \{1, \dots, d\}$ such that $j \notin \mathcal{C}$ but $\mathbf{x}_j > -\infty$. By the definition of col-admissibility there must exist $i \in \{1, \dots, n\}$, such that $\mathbf{y}_i = -\infty$ and $a_{ij} > -\infty$. Therefore

$$(A \otimes \mathbf{x})_i = \max_{k=1}^d a_{ik} + \mathbf{x}_k \geq a_{ij} + \mathbf{x}_j > -\infty$$

and since $\mathbf{y}_i = -\infty$ we have $R(\mathbf{x}) = \infty$. □

Theorem 2.5. Let $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}_{\max}^n$.

1. If $\mathcal{R} \neq \text{support}(\mathbf{y})$, then $R(\mathbf{x}) = \infty$, for all $\mathbf{x} \in \mathbb{R}_{\max}^d$.
2. If $\mathcal{R} = \text{support}(\mathbf{y})$, then whenever $\mathbf{x} \in \mathbb{R}^d$ satisfies (11), we have

$$R(\mathbf{x}) = \|A_{\mathcal{R}\mathcal{C}} \otimes \mathbf{x}_{\mathcal{C}} - \mathbf{y}_{\mathcal{R}}\|_2.$$

where $A_{\mathcal{C},\mathcal{R}}$ denotes the sub-matrix of A formed from the rows in \mathcal{R} and columns in \mathcal{C} , likewise for the sub-vector $\mathbf{y}_{\mathcal{R}}$. Moreover $A_{\mathcal{C},\mathcal{R}}, \mathbf{y}_{\mathcal{R}}$ are a finite form.

Proof. For the first part first suppose that (11) is not satisfied then $R(\mathbf{x}) = \infty$. Now suppose that (11) is satisfied. From $\mathcal{R} \neq \text{support}(\mathbf{y})$, we have that there exists i with $\mathbf{y}_i > -\infty$ and $i \notin \mathcal{R}$. From the definition of row-admissibility, we have $a_{ij} = -\infty$, for all $j \in \mathcal{C}$. So that $(A \otimes \mathbf{x})_i = -\infty$, unless $\mathbf{x}_k > -\infty$, for some $k \notin \mathcal{C}$, and this would violate (11). Therefore $(A \otimes \mathbf{x})_i = -\infty$ and $\mathbf{y}_i > -\infty$, so $R(\mathbf{x}) = \infty$.

For the second part, from (11) and the definition of col-admissibility we have that

$$\text{support}(A \otimes \mathbf{x}) \subseteq \text{support}(\mathbf{y}).$$

So from the assumption $\mathcal{R} = \text{support}(\mathbf{y})$ and (4) we have

$$R(\mathbf{x}) = \|A \otimes \mathbf{x} - \mathbf{y}\|_2 = \|(A \otimes \mathbf{x})_{\mathcal{R}} - \mathbf{y}_{\mathcal{R}}\|_2,$$

and applying (11) a second time we have

$$R(\mathbf{x}) = \|A_{\mathcal{R}\mathcal{C}} \otimes \mathbf{x}_{\mathcal{C}} - \mathbf{y}_{\mathcal{R}}\|_2.$$

For the last part. From the assumption $\mathcal{R} = \text{support}(\mathbf{y})$, we have that every entry in $\mathbf{y}_{\mathcal{R}}$ is finite. Now consider an arbitrary row of $A_{\mathcal{R}\mathcal{C}}$. This will be the i th row of A , for some $i \in \mathcal{R}$. From the definition of row-admissibility we have that there exists $j \in \mathcal{C}$, such that $a_{ij} > -\infty$, and this will correspond to a finite entry in the chosen row of $A_{\mathcal{R}\mathcal{C}}$. Therefore $A_{\mathcal{C},\mathcal{R}}, \mathbf{y}_{\mathcal{R}}$ are a finite form. □

Theorem 2.5 tells us that given any max-plus 2-norm regression problem, either no solution with finite residual exists, or the problem can be reduced to a smaller problem in finite form, for which Theorem 2.3 guarantees that finite solutions with finite residuals exist.

2.3.1 Patterns of support and their domains

For $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{x} \in \mathbb{R}_{\max}^d$ define the *pattern of support* $\text{pattern}(\mathbf{x}) = (P_1, \dots, P_n) \in \mathcal{P}(\{1, \dots, d\})^n$, by

$$j \in P_i \Leftrightarrow a_{ij} + \mathbf{x}_j = (A \otimes \mathbf{x})_j, \quad (12)$$

for $i = 1, \dots, n$. Define the *domain* of a pattern $P \in \mathcal{P}(\{1, \dots, d\})^n$, by

$$X(P) = \{\mathbf{x} \in \mathbb{R}_{\max}^d : \text{pattern}(\mathbf{x}) = P\}. \quad (13)$$

We say that a pattern P is *feasible*, if $X(P) \cap \mathbb{R}^d \neq \emptyset$. For a pattern P define the binary relation $\tilde{\bowtie}_P$ on $\{1, \dots, d\}$, by $j \tilde{\bowtie}_P k$, if and only if $j, k \in P_i$, for some $i = 1, \dots, n$ and define \bowtie_P to be the transitive closure of $\tilde{\bowtie}_P$. Let $|\bowtie_P|$ be the number of equivalence classes of \bowtie_P . If P is a feasible pattern then $X(P)$ is a set of dimension $|\bowtie_P|$. From [5, Cor. 25] we have that

$$|\{\text{feasible } P : |\bowtie_P| = k\}| \leq \frac{(n + d - k - 1)!}{(n - k)! \cdot (d - k)! \cdot (k - 1)!}, \quad (14)$$

for $k = 1, \dots, d$, with equality for all k in the generic case of a matrix A with rows/cols in general position.

Define the ordering \preceq on $\mathcal{P}(\{1, \dots, d\})^n$, by $P \preceq P'$, if and only if $P_i \subseteq P'_i$, for all $i = 1, \dots, n$, with a strict inequality if at least one inclusion is a strict inclusion. Then the boundary of the domain $X(P)$ is given by $\cup_{\{P' : P \prec P'\}} X(P')$ and the closure by $\text{Cl}(X(P)) = \cup_{\{P' : P \preceq P'\}} X(P')$.

Also define the *feasibility matrix* by $F_P \in \mathbb{R}_{\max}^{d \times d}$, by

$$f_{jk} = \begin{cases} 0, & \text{for } j = k, \\ \max\{-\infty, \max\{a_{ik} - a_{ij} : j \in P_i : i = 1, \dots, n\}\}, & \text{otherwise.} \end{cases} \quad (15)$$

We will need to quickly review some related results to support the following Theorem. For a max-plus matrix $B \in \mathbb{R}_{\max}^{d \times d}$, the *maximum cycle mean* of B is defined by

$$\lambda(B) = \max_{\zeta} \frac{W(\zeta)}{L(\zeta)}, \quad (16)$$

where the maximum is taken over cycles $\zeta = (\zeta(1) \mapsto \dots \mapsto \zeta(k) \mapsto \zeta(1)) \subset \{1, \dots, d\}$. The *weight* of a cycle is the sum of its edge weights $W(\zeta) = b_{\zeta(1)\zeta(2)} + \dots + b_{\zeta(k-1)\zeta(k)} + b_{\zeta(k)\zeta(1)}$ and the *length* of a cycle is its total number of edges $L(\zeta) = k + 1$. The *Klene star* of $B \in \mathbb{R}_{\max}^d$ is defined by

$$B^* = \lim_{t \rightarrow \infty} (I \oplus B \oplus B^{\otimes 2} \oplus \dots \oplus B^{\otimes t}), \quad (17)$$

where $I \in \mathbb{R}_{\max}^d$ is the max-plus identity matrix, with zeros on the diagonal and minus infinities off of the diagonal. From [4, Prop. 1.6.10 and Thm. 1.6.18] we have that if $\lambda(B) \leq 0$, then B^* exists and $\{x \in \mathbb{R}_{\max}^d : B \otimes x = x\} = \text{col}(B^*)$ and that if $\lambda(B) > 0$, then B^* does not exist and $\{x \in \mathbb{R}_{\max}^d : B \otimes x = x\} \cap \mathbb{R}^d = \emptyset$.

For $B \in \mathbb{R}_{\max}^{d \times d}$ define $\bar{B} \in \mathbb{R}_{\max}^d$ to be the arithmetic mean of the columns of B . It follows from [15, Thm. 3.3] that, provided $\lambda(B) \leq 0$, we have $\overline{(B^*)} \in \text{relint}(\text{col}(B^*))$.

Theorem 2.6. *For $A \in \mathbb{R}_{\max}^{n \times d}$ and $P \in \mathcal{P}(\{1, \dots, d\})^n$ we have*

$$\text{Cl}(X(P)) = \{\mathbf{x} \in \mathbb{R}_{\max}^d : F_P \otimes \mathbf{x} = \mathbf{x}\}.$$

Moreover, the pattern P is feasible, if and only if $\lambda(F_P) = 0$ and in the case where P is feasible, we have

$$X(P) = \text{relint}(\text{col}(F_P^*))$$

and $\overline{(F_P^*)} \in X(P)$.

Proof. First note that $\mathbf{x} \in \text{Cl}(X(P))$, if and only if $(A \otimes \mathbf{x})_i = a_{ij} + \mathbf{x}_j$, for all $j \in P_i$, for all $i = 1, \dots, n$, which is equivalent to $\max_{k=1}^d (a_{ik} + \mathbf{x}_k) \leq a_{ij} + \mathbf{x}_j$, for all $j \in P_i$, for all $i = 1, \dots, n$, which is equivalent to $F_P \otimes \mathbf{x} \leq \mathbf{x}$ and since F_P has zeros on its diagonal this is equivalent to $F_P \otimes \mathbf{x} = \mathbf{x}$. Next from [4, Prop. 1.6.10 and Thm. 1.6.18], we have that $\text{Cl}(X(P)) \cap \mathbb{R}^d$ is non-empty, if and only if $\lambda(F_P) \leq 0$ and since F_P has zeros on the diagonal this is equivalent to the condition $\lambda(F_P) = 0$. In the case that $\lambda(F_P) = 0$ we also have $\text{Cl}(X(P)) = \text{col}(F_P^*)$. Then note that

$$X(P) = \cup_{\{P' : P \prec P'\}} X(P') \setminus \cup_{\{P' : P \prec P'\}} X(P')$$

is equal to $\text{col}(F_P^*)$ minus its boundary, which is precisely $\text{relint}(\text{col}(F_P^*))$. The final result follows immediately from [15, Thm. 3.3]. \square

If there are m equivalence classes in \bowtie_P then label them arbitrarily with $\{1, \dots, m\}$ and define $c : \{1, \dots, d\} \mapsto \{1, \dots, m\}$, such that $c(j) = k$, if and only if j is in the k th equivalence class. Let $\#_k$ denote the number of elements in the k th equivalence class and define $C \in \mathbb{R}^{d \times m}$, by

$$c_{jc(j)} = \frac{1}{\sqrt{\#_{c(j)}}}, \quad (18)$$

for $j = 1, \dots, d$ and all other entries equal to zero. Then for any $\mathbf{x}_P \in \text{Cl}(X(P)) \cap \mathbb{R}^d$, we have that

$$\mathcal{A}(X(P) \cap \mathbb{R}^d) = \{C\mathbf{h} + \mathbf{x}_P : \mathbf{h} \in \mathbb{R}^m\} \quad (19)$$

is the smallest affine subspace of \mathbb{R}_{\max}^d containing $X(P) \cap \mathbb{R}^d$. We could choose $\mathbf{x}_P = \overline{(F_P^*)}$ but also need to consider the case where \mathbf{x}_P represents the current state of one of the algorithms that we detail later. We call $\text{Cl}(\mathcal{A}(X(P) \cap \mathbb{R}^d))$ the *extended domain* of P .

2.3.2 Local maps and images

Now define the *subpattern* p of P , by $p_i = \{\min(P_i)\}$, for $i = 1, \dots, n$. Since each element of a subpattern contains only a single element we can treat subpatterns like vectors and will write p_i to mean the unique element of the set p_i . Also define $L \in \{0, 1\}^{n \times d}$, by $l_{ip(i)} = 1$ and all other entries equal to zero. Then define the *local mapping* $A_P : \mathbb{R}_{\max}^d \mapsto \mathbb{R}_{\max}^n$, for real $\mathbf{x} \in \mathbb{R}^d$ by

$$A_P(\mathbf{x}) = L\mathbf{x} + \mathbf{y}_P, \quad (20)$$

where $(\mathbf{y}_P)_i = a_{ip(i)}$, for $i = 1, \dots, n$ and continue $A_P(\cdot)$ to \mathbb{R}_{\max}^n via (9). It follows that $\text{support}(\mathbf{y}_P) = \{1, \dots, n\}$, whenever A, \mathbf{y} are a finite form and P is a feasible pattern. Note that $A_P(\mathbf{x}) = A \otimes \mathbf{x}$, for all $\mathbf{x} \in \text{Cl}(X(P))$.

Define the *image* $Y(P) = A_P(X(P))$. Then we have

$$\text{col}(A) = \bigcup_P Y(P), \quad (21)$$

where the union is taken over all feasible patterns. We have that

$$\mathcal{A}(Y(P) \cap \mathbb{R}^n) = \{LCh + Lx_P + \mathbf{y}_P : \mathbf{h} \in \mathbb{R}^m\}, \quad (22)$$

is the smallest affine subspace of \mathbb{R}^n containing $Y(P) \cap \mathbb{R}^n$. We call $\text{Cl}\left(\mathcal{A}(Y(P) \cap \mathbb{R}^n)\right)$ the *extended image* of P .

2.3.3 Normal projections

For a feasible pattern P define the *normal projection* map $\Phi(P, \cdot) : \mathbb{R}_{\max}^n \mapsto \mathbb{R}_{\max}^n$, by

$$\Phi(P, \mathbf{y}) = \arg \min\{\|\mathbf{y} - \mathbf{y}'\|_2 : \mathbf{y}' \in \text{Cl}\left(\mathcal{A}(Y(P) \cap \mathbb{R}^n)\right)\}. \quad (23)$$

Then for real $\mathbf{y} \in \mathbb{R}^n$ we have

$$\Phi(P, \mathbf{y}) = LCh^* + Lx_P + \mathbf{y}_P, \quad (24)$$

where the normal equations [2] give

$$\mathbf{h}^* = ((LC)^\top LC)^\dagger (LC)^\top (\mathbf{y} - Lx_P - \mathbf{y}_P). \quad (25)$$

Note that $LC \in \mathbb{R}^{n \times m}$ with

$$(LC)_{i \ c(p(i))} = \frac{1}{\sqrt{\#_{c(p(i))}}} \quad (26)$$

and all other entries equal to zero. Hence we have that

$$\mathbf{h}_k^* = \sqrt{\#_k} \times \overline{\{(\mathbf{y} - Lx_P - \mathbf{y}_P)_i : c(p(i)) = k\}}, \quad (27)$$

for $k = 1, \dots, m$ such that $\{i : c(p(i)) = k\} \neq \emptyset$, and $\mathbf{h}_k^* = 0$ otherwise, and where the overline in (27) indicates taking the mean. Define the equivalence relation $\hat{\bowtie}_P$ on $\{1, \dots, n\}$, by $i \hat{\bowtie}_P j$, if and only if $p(i) \bowtie_P p(j)$. Then we have

$$\Phi(P, \mathbf{y})_i = \overline{\{(\mathbf{y} - Lx_P - \mathbf{y}_P)_j : i \hat{\bowtie}_P j\}} + (Lx_P + \mathbf{y}_P)_i, \quad (28)$$

for $i = 1, \dots, n$. We continue $\Phi(P, \cdot)$ to \mathbb{R}_{\max}^n via (9). We say that $\Phi(P, y)$ is *admissible*, if $\Phi(P, \mathbf{y}) \in \text{Cl}(Y(P))$.

Also define

$$A_P^{-1}(\Phi(P, \mathbf{y})) = \{\mathbf{x} \in \text{Cl}\left(\mathcal{A}(X(P) \cap \mathbb{R}^d)\right) : A_P(\mathbf{x}) = \Phi(P, \mathbf{y})\} \quad (29)$$

where for real $y \in \mathbb{R}^n$, we have

$$A_P^{-1}(\Phi(P, \mathbf{y})) = \text{Cl}(Ch^* + C \ker(LC) + x_P). \quad (30)$$

where

$$(Ch^*)_j = \overline{\{(\mathbf{y} - Lx_P - \mathbf{y}_P)_i : c(p(i)) = c(j)\}}, \quad (31)$$

for $j = 1, \dots, d$ and

$$C \ker(LC) = \text{span}\{\underline{e}_j \in \mathbb{R}_{\max}^d : j \in \{1, \dots, d\}/\text{support}(P)\}, \quad (32)$$

where $\text{support}(P) \subset \{1, \dots, d\}$ is the *support* of the pattern P , defined by $\text{support}(P) = \cup_{i=1}^n P_i$.

Define the *closest local minimum* map $\Psi(P, \mathbf{y}, \cdot) : \mathbb{R}_{\max}^d \mapsto \mathbb{R}_{\max}^d$, for real $\mathbf{y} \in \mathbb{R}^n$ by

$$\Psi(P, \mathbf{y}, \mathbf{x}) = \arg \min \{ \|\mathbf{x} - \mathbf{x}'\|_2 : \mathbf{x}' \in A_P^{-1}(\Phi(P, \mathbf{y})) \}. \quad (33)$$

Then for real $\mathbf{x} \in \mathbb{R}^d$, we have

$$\Psi(P, \mathbf{y}, \mathbf{x})_j = \begin{cases} (Ch^*)_j & \text{for } j \in \text{support}(P), \\ x_j & \text{otherwise,} \end{cases} \quad (34)$$

for $j = 1, \dots, d$. For all $\mathbf{y} \in \mathbb{R}^n$ we continue $\Psi(P, \mathbf{y}, \cdot)$ to \mathbb{R}_{\max}^d via (9), then for all $\mathbf{x} \in \mathbb{R}_{\max}^d$ we continue $\Psi(P, \cdot, \mathbf{x})$ to \mathbb{R}_{\max}^n via (9).

Theorem 2.7. *Let $A \in \mathbb{R}_{\max}^{n \times d}$, $\mathbf{y} \in \mathbb{R}_{\max}^n$ be a finite form and let $P \in \mathcal{P}(\{1, \dots, d\})^n$ be a feasible pattern. Then the normal projection $\Phi(P, \mathbf{y})$ is admissible, if and only if $\Psi(P, \mathbf{y}, \underline{-\infty}) \in \text{Cl}(X(P))$, where $\underline{-\infty} \in \mathbb{R}_{\max}^d$ is a vector with all entries equal to $-\infty$.*

Proof. First suppose that $\Psi(P, \mathbf{y}, \underline{-\infty}) \in \text{Cl}(X(P))$. By definition we have that $A \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) = \Phi(P, \mathbf{y})$ and from the supposition we have that $A \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) \in \text{Cl}(Y(P))$. Therefore $\Phi(P, \mathbf{y}) \in \text{Cl}(Y(P))$, or in words, $\Phi(P, \mathbf{y})$ is admissible.

Conversely suppose that $\Phi(P, \mathbf{y})$ is admissible. Note that from Theorem 2.6 we have that $\Psi(P, \mathbf{y}, \underline{-\infty}) \in X(P)$, if and only if

$$F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) = \Psi(P, \mathbf{y}, \underline{-\infty}). \quad (35)$$

From the supposition we have that there exists $\hat{\mathbf{x}} \in \text{Cl}(X(P))$, such that $A \otimes \hat{\mathbf{x}} = \Phi(P, \mathbf{y})$. Next note that since the feasibility matrix has zeros on its diagonal it suffices to show that

$$F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) \leq \Psi(P, \mathbf{y}, \underline{-\infty}).$$

Using the fact that $\Psi(P, \mathbf{y}, \underline{-\infty})$ is the infimum of $A_P^{-1}(\Phi(P, \mathbf{y}))$ with respect to the standard partial order, that $(F_P \otimes \mathbf{x})_k$ is a monotonically non-decreasing function of \mathbf{x} , for all $k \in \text{support}(P)$, and (34), we have

$$\left((F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty})) \right)_k \leq \left(F_P \otimes \hat{\mathbf{x}} \right)_k = \hat{\mathbf{x}}_k = (Ch^*)_j = \Psi(P, \mathbf{y}, \underline{-\infty})_k,$$

for $k \in \text{support}(P)$. Finally note from (34), that for $j \in \{1, \dots, d\} \setminus \text{support}(P)$, we have

$$\left((F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty})) \right)_j = \Psi(P, \mathbf{y}, \underline{-\infty})_j.$$

Therefore (35) holds and $\Psi(P, \mathbf{y}, \underline{-\infty}) \in \text{Cl}(X(P))$. □

2.3.4 Residual surface

For $A \in \mathbb{R}_{\max}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}_{\max}^n$ recall the *residual* $R : \mathbb{R}_{\max}^d \mapsto \mathbb{R}$, with $R(\mathbf{x}) = \|A \otimes \mathbf{x} - \mathbf{y}\|_2^2/2$. For a feasible pattern P define the *local residual* $R_P : \mathbb{R}_{\max}^d \mapsto \mathbb{R}$, by

$$R_P(\mathbf{x}) = \|A_P(\mathbf{x}) - \mathbf{y}\|_2^2/2. \quad (36)$$

For real $\mathbf{x} \in \mathbb{R}^d$ we have

$$R_P(\mathbf{x}) = \|L\mathbf{x} + \mathbf{y}_P - \mathbf{y}\|_2^2/2,$$

Note that for $\mathbf{x} \in \text{Cl}(X(P))$, we have $R(\mathbf{x}) = R_P(\mathbf{x})$. Hence R is piecewise quadratic.




Figure 2: Domains of feasible patterns for the matrix A of Example 2.8.

Example 2.8. Consider

$$A = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} 0 \\ 0.5 \\ 0 \end{bmatrix}, \quad \mathbf{y}' = \begin{bmatrix} 0 \\ 1.5 \\ 2 \end{bmatrix}.$$

Since all entries in A , \mathbf{y} and \mathbf{y}' are finite both A, \mathbf{y} and A, \mathbf{y}' are finite forms. There are seven feasible patterns, their domains are displayed in Figure 2.

For $P = P(3) = (\{1\}, \{1\}, \{2\})$, the feasibility matrix is given by

$$F_P = \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix}.$$

From Theorem 2.6, since $\lambda(F_P) = 0$, we have that P is admissible and that its domain is given by $X(P) = \text{relint}(\text{col}(F_P^*))$. In this case $F_P^* = F_P$ and

$$\text{col}(F_P^*) = \{\mathbf{x} \in \mathbb{R}_{\max}^2 : x_2 + 1 \geq x_1 \geq x_2\}, \quad \text{relint}(\text{col}(F_P^*)) = \{\mathbf{x} \in \mathbb{R}_{\max}^2 : x_2 + 1 > x_1 > x_2\}.$$

The boundary of $X(P)$ is given by $X(P(2)) \cup X(P(4))$. These patterns are both feasible and

$$F_{P(2)} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \quad X(P(2)) = \text{relint}(\text{col}(F_{P(2)}^*)) = \text{col}(F_{P(2)}^*) = \{\mathbf{x} \in \mathbb{R}_{\max}^2 : x_1 = x_2 + 1\},$$

$$F_{P(4)} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad X(P(4)) = \text{relint}(\text{col}(F_{P(4)}^*)) = \text{col}(F_{P(4)}^*) = \{\mathbf{x} \in \mathbb{R}_{\max}^2 : x_1 = x_2\}.$$

For $P = (\{1, 2\}, \{2\}, \{1\})$, we have

$$F_P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Since $\lambda(F_P) = 1$, we have that P is not admissible.

Note that $\mathbf{z} \in \text{col}(A)$, if and only if $\mathbf{z} \otimes \alpha = [z_1 + \alpha, z_2 + \alpha, z_3 + \alpha]^\top \in \text{col}(A)$, for all $\alpha \in \mathbb{R}_{\max}$. Therefore $\text{col}(A)$ has translational symmetry in the $[1, 1, 1]^\top$ direction. Similarly for all of the pattern images and extended images. We can therefore study these objects by examining their image under the orthogonal projection $\Pi : \mathbb{R}_{\max}^3 \mapsto \{\mathbf{z} \in \mathbb{R}_{\max}^3 : [1, 1, 1]\mathbf{z} = 0\}$. This is the same idea as in the tropical projected space \mathbb{TP}^2 , which is usually taken to be a projection onto $\{\mathbf{y} \in \mathbb{R}_{\max}^3 : [1, 0, 0]\mathbf{y} = 0\}$. However the orthogonal projection that we use here is more convenient for analyzing the 2-norm regression problem. The projected pattern images, extended images and a sample of normal projections are displayed in Figure 3.

We return our attention to the pattern $P = P(3) = (\{1\}, \{1\}, \{2\})$ and set

$$\mathbf{x}_P = \overline{F_P^*} = \overline{F_P} = \begin{bmatrix} 0 \\ -0.5 \end{bmatrix}.$$

We have that \bowtie_P is the identity relation with equivalence classes $\{1\}$ and $\{2\}$, which we label as the 1st and 2nd classes respectively. So that $c(1) = 1$ and $c(2) = 2$ and C is the 2×2 identity matrix. The extended domain is therefore given by

$$\text{Cl}(\mathcal{A}(X(P) \cap \mathbb{R}^2)) = \text{Cl}(\{C\mathbf{h} + \mathbf{x}_P : \mathbf{h} \in \mathbb{R}^2\}) = \text{Cl}(\mathbb{R}^2) = \mathbb{R}_{\max}^2.$$

The subpattern of P is given by $p = P$, since P already has exactly one element in each subset. Written as a vector we have $p = [1, 1, 2]$. The matrix L and the vector \mathbf{y}_P are therefore given by

$$L = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{y}_P = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}.$$

The local map is given by $A_P(\mathbf{x}) = L\mathbf{x} + \mathbf{y}_P$. The extended image is given by

$$\begin{aligned} \text{Cl}(\mathcal{A}(Y(P) \cap \mathbb{R}^3)) &= \text{Cl}(\{LC\mathbf{h} + L\mathbf{x}_P + \mathbf{y}_P : \mathbf{h} \in \mathbb{R}^2\}) \\ &= \text{Cl}(\{\mathbf{z} \in \mathbb{R}^3 : z_2 = z_1 + 1\}) \\ &= \{\mathbf{z} \in \mathbb{R}_{\max}^3 : z_2 = z_1 + 1\}. \end{aligned}$$

and the image is given by

$$Y(P) = \{\mathbf{z} \in \mathbb{R}_{\max}^3 : z_2 = z_1 + 1, z_3 > z_1 > z_3 - 1\}.$$

The boundary of the image is given by $Y(P(2)) \cup Y(P(4))$, where

$$Y(P(2)) = \{\mathbf{z} \in \mathbb{R}_{\max}^3 : z_2 = z_1 + 1, z_3 = z_1\}, \quad Y(P(4)) = \{\mathbf{z} \in \mathbb{R}_{\max}^3 : z_2 = z_1 + 1, z_3 = z_1 + 1\}.$$

The normal projection is given by

$$\Phi(P, \mathbf{y}) = \begin{bmatrix} \overline{\{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_j : j \hat{\bowtie}_P 1\}} \\ \overline{\{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_j : j \hat{\bowtie}_P 2\}} \\ \overline{\{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_j : j \hat{\bowtie}_P 3\}} \end{bmatrix} + L\mathbf{x}_P + \mathbf{y}_P,$$

where

$$\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P = \begin{bmatrix} 0 \\ -0.5 \\ -0.5 \end{bmatrix}, \quad L\mathbf{x}_P + \mathbf{y}_P = \begin{bmatrix} 0 \\ 1 \\ 0.5 \end{bmatrix},$$

and since the equivalence relation $\hat{\simeq}_P$ has equivalence classes $\{1, 2\}$ and $\{3\}$, we have

$$\Phi(P, \mathbf{y}) = [-0.25, 0.75, 0]^\top.$$

We have $\text{support}(P) = \{1, 2\}$, so that $\ker(LC) = \{0\}$ and therefore

$$A_P^{-1}(\Phi(P, \mathbf{y})) = \text{Cl}(C\mathbf{h}^* + \mathbf{x}_P),$$

where

$$\mathbf{h}^* = \begin{bmatrix} \sqrt{1} \times \frac{\{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_i : c(p(i)) = 1\}}{\sqrt{1} \times \{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_i : c(p(i)) = 2\}} \\ \sqrt{1} \times \frac{\{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_i : c(p(i)) = 1\}}{\sqrt{1} \times \{(\mathbf{y} - L\mathbf{x}_P - \mathbf{y}_P)_i : c(p(i)) = 2\}} \end{bmatrix} = \begin{bmatrix} -0.25 \\ -0.5 \end{bmatrix},$$

so that $A_P^{-1}(\Phi(P, \mathbf{y})) = \{[-0.25, -1]^\top\}$. Since $A_P^{-1}(\Phi(P, \mathbf{y}))$ contains a single vector we have

$$\Psi(P, \mathbf{y}, \mathbf{x}) = \arg \min\{\|\mathbf{x} - \mathbf{x}'\|_2 : \mathbf{x}' \in A_P^{-1}(\Phi(P, \mathbf{y}))\} = [-0.25, -1]^\top,$$

for all $\mathbf{x} \in \mathbb{R}_{\max}^2$. Therefore $\Psi(P, \mathbf{y}, \underline{-\infty}) = [-0.25, -1]^\top$. Checking

$$F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) = \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix} \otimes \begin{bmatrix} -0.25 \\ -1 \end{bmatrix} = \begin{bmatrix} -0.25 \\ -1 \end{bmatrix} = \Psi(P, \mathbf{y}, \underline{-\infty}),$$

we verify by Theorem 2.7 that $\Phi(P, \mathbf{y})$ is admissible.

For the alternative target vector \mathbf{y}' we have $\Phi(P, \mathbf{y}') = [0.25, 1.25, 2]^\top$,

$$A_P^{-1}(\Phi(P, \mathbf{y}')) = C\mathbf{h}^* = [0.25, 1]^\top$$

and

$$F_P \otimes \Psi(P, \mathbf{y}', \underline{-\infty}) = \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix} \otimes \begin{bmatrix} 0.25 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} > \Psi(P, \mathbf{y}', \underline{-\infty}),$$

which verifies by Theorem 2.7 that $\Phi(P, \mathbf{y}')$ is not admissible. See Figure 3. Note how the normal projection of \mathbf{y} lies inside the column space of A but that for the alternative target \mathbf{y}' it does not.

2.4 NP-hardness of finding descent directions

For $A \in \mathbb{R}_{\max}^{n \times d}$, $\mathbf{y} \in \mathbb{R}_{\max}^n$ and $\mathbf{x} \in \mathbb{R}_{\max}^d$ we say that $\mathbf{z} \in \mathbb{R}^d$ is an *descent direction* for \mathbf{x} if there exists $\epsilon > 0$ such that

$$R(\mathbf{x} + \mu\mathbf{z}) < R(\mathbf{x}), \tag{37}$$

for all $0 < \mu \leq \epsilon$.

The following problem is equivalent to determining whether $\mathbf{x} \in \mathbb{R}_{\max}^d$ is a local minimum of R . In particular \mathbf{x} is a local minimum if and only if it does not have any descent directions.

Problem 2.9. Let $A \in \mathbb{R}_{\max}^{n \times d}$, $\mathbf{y} \in \mathbb{R}_{\max}^n$ and $\mathbf{x} \in \mathbb{R}_{\max}^d$. Does there exist a descent direction for \mathbf{x} ?

The following problem is known as the set covering problem.




Figure 3: For the problem of example Example 2.8. Projected pattern images and extended images for the matrix A . $Y_{1/2} = Y(P(1)) = \mathcal{A}(Y(P(1))) = Y(P(2)) = \mathcal{A}(Y(P(2)))$, $Y_3 = Y(P(3))$, $Y_4 = Y(P(4)) = \mathcal{A}(Y(P(4)))$, $Y_5 = Y(P(5))$ and $Y_{6/7} = Y(P(6)) = \mathcal{A}(Y(P(6))) = Y(P(7)) = \mathcal{A}(Y(P(7)))$. Target vectors \mathbf{y} and \mathbf{y}' with normal projections onto $\mathcal{A}(Y_3)$.

Problem 2.10. Let $F = \{F_i \subset \{1, \dots, n\} : i = 1, \dots, m\}$ be a family of subsets with $\cup_{i=1}^m F_i = \{1, \dots, n\}$ and let $1 < k < m$. Does there exist a subset $J \subset \{1, \dots, m\}$, such that $|J| \leq k$ and $\cup_{j \in J} F_j = \{1, \dots, n\}$?

We will now prove that any instance of Problem 2.10 can be reduced to an instance of Problem 2.9. Because the set covering problem is known to be NP-hard this suffices to show that Problem 2.9 is also NP-hard [10].

Lemma 2.11. Let $A \in \{0, -\infty\}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}^n$ be a finite form. Then $\mathbf{z} \in \mathbb{R}^d$ is a descent direction for the vector of zeros $\underline{0} \in \mathbb{R}_{\max}^d$, if and only if

$$\langle A \otimes \mathbf{z}, \mathbf{y} \rangle > 0. \quad (38)$$

Moreover, in the case that $\sum_{i=1}^n y_i = 0$, if $\underline{0} \in \mathbb{R}_{\max}^d$ has a descent direction $\mathbf{z} \in \mathbb{R}^d$ then it has a descent direction $\mathbf{z}' \in \{0, 1\}^d$.

Proof. Note that since $A \in \{0, -\infty\}^{n \times d}$, we have $A \otimes (\underline{0} + \mu \mathbf{z}) = \mu(A \otimes \mathbf{z})$, for all $\mu \geq 0$. From (37) we have that $\mathbf{z} \in \mathbb{R}^d$ is a descent direction for $\underline{0}$ if and only if there exists $\epsilon > 0$, such that

$$R(\underline{0} + \mu \mathbf{z}) = \|\mu(A \otimes \mathbf{z}) - \mathbf{y}\|_2^2 / 2 < R(\underline{0}) = \|\mathbf{y}\|_2^2 / 2$$

for all $0 < \mu \leq \epsilon$. Which is equivalent to

$$\lim_{\mu \rightarrow 0^+} \frac{\delta(\|\mu(A \otimes \mathbf{z}) - \mathbf{y}\|_2^2 / 2)}{\delta \mu} = -\langle A \otimes \mathbf{z}, \mathbf{y} \rangle < 0.$$

For the second part suppose that $\sum_{i=1}^n y_i = 0$ and that $\mathbf{z} \in \mathbb{R}^d$ is a descent direction for $\underline{0}$. Note that for any $a, b \in \mathbb{R}$ with $a > 0$, we have that $a\mathbf{z} + b\underline{1}$ is also a descent direction, where $\underline{1} \in \mathbb{R}^d$

is a vector of ones. Furthermore since $\sum_{i=1}^n y_i = 0$, we must have that $\min_{i=1}^d z_i < \max_{i=1}^d z_i$. So we can assume without loss of generality that $\min_{i=1}^d z_i = 0$ and $\max_{i=1}^d z_i = 1$. Now let σ be a permutation of $\{1, \dots, d\}$ such that

$$z_{\sigma(1)} \leq z_{\sigma(2)} \leq \dots \leq z_{\sigma(d)}.$$

And define the sequence $b_1 < b_2 < \dots < b_m$ such that $b_1 = 1$, $b_m = d$ and

$$z_{\sigma(i)} = z_{\sigma(j)} \Leftrightarrow b_k \leq i, j < b_{k+1} \text{ for some } k.$$

Define the vectors $(\mathbf{h}(k) \in \mathbb{R}^d)_{k=1}^m$, by

$$h(k)_{\sigma(j)} = \begin{cases} 1 & \text{if } j \geq b_k \\ 0 & \text{otherwise.} \end{cases}$$

for $j = 1, \dots, d$. Note that for $r, j = 1, \dots, d$ and $k = 1, \dots, m$, we have

$$z_r > z_j \Rightarrow h(k)_r \geq h(k)_j, \quad \text{and} \quad z_r = z_j \Rightarrow h(k)_r = h(k)_j. \quad (39)$$

Let $P = \text{pattern}(\mathbf{z})$ and let $Q(k) = \text{pattern}(\mathbf{h}(k))$ for $k = 1, \dots, m$. Now suppose that $j \in P_i$, then

$$a_{ir} = -\infty \quad \forall r : z_r > z_j, \quad \text{and} \quad \exists s : z_s = z_j, \quad a_{is} = 0.$$

So from (39) we have that

$$a_{ir} = -\infty \quad \forall r : h(k)_r > h(k)_j, \quad \text{and} \quad \exists s : h(k)_s = h(k)_j, \quad a_{is} = 0,$$

and that $j \in Q(k)_i$, for all $k = 1, \dots, m$. Therefore $Q(k) \preceq P$ and $\mathbf{h}(k) \in \text{Cl}(X(P))$, so that

$$A \otimes \mathbf{h}(k) = A_P(\mathbf{h}(k)), \quad (40)$$

for all $k = 1, \dots, m$. The the local map in (40) is given by

$$A_P(\mathbf{x}) = L\mathbf{x} + \mathbf{y}_P,$$

where the vector \mathbf{y}_P is defined by $(\mathbf{y}_P)_i = a_{ip(i)}$, where p is the subpattern of P . Note that since A, y are a finite form and P is a feasible pattern and since A has only zero and minus infinity entries, we must have $\mathbf{y}_P = \underline{0} \in \mathbb{R}^n$ and therefore that A_P is a classically linear map.

Finally note that

$$\mathbf{z} = \sum_{k=2}^m \mathbf{h}(k) \alpha_k, \quad (41)$$

where the coefficients

$$\alpha_k = \left(z_{\sigma(b(k))} - z_{\sigma(b(k-1))} \right),$$

are strictly positive for $k = 2, \dots, m$. From (40), (41) and the fact that A_P is linear, we obtain

$$\langle A \otimes \mathbf{z}, \mathbf{y} \rangle = \sum_{k=2}^m \alpha_k \langle A \otimes \mathbf{h}(k), \mathbf{y} \rangle. \quad (42)$$

Therefore, using the result of the first part of the Lemma, if $z \in \mathbb{R}^d$ is a descent direction for $\underline{0} \in \mathbb{R}^d$ then $\langle A \otimes \mathbf{z}, \mathbf{y} \rangle > 0$ and from (42) there exists $2 \leq k \leq m$ such that $\langle A \otimes \mathbf{h}(k), \mathbf{y} \rangle > 0$, equivalently, such that $\mathbf{h}(k)$ is a descent direction for $\underline{0} \in \mathbb{R}^d$ with all entries in $\{0, 1\}$. \square

Theorem 2.12. Let $F = \{F_i : i = 1, \dots, m\}$ be a family of subsets of $\{1, \dots, n\}$ with $\cup_{i=1}^m F_i = \{1, \dots, n\}$ and let $1 < k < m$. There exists a subset $J \subset \{1, \dots, m\}$, such that $|J| \leq k$ and $\cup_{j \in J} F_j = \{1, \dots, n\}$, if and only if $\underline{0} \in \mathbb{R}_{\max}^m$ has a descent direction for the following $N \times m$ max-plus 2 norm regression problem, where $N = n + m + \frac{m(m-1)}{2} + 1$. Set

$$A = \begin{bmatrix} A_1 \\ A_2 \\ A_3 \\ A_4 \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \\ \mathbf{y}_4 \end{bmatrix},$$

where

1. $A_1 \in \mathbb{R}_{\max}^{n \times m}$ is defined by $(A_1)_{ij} = 0$, if $i \in F_j$ and $(A_1)_{ij} = -\infty$, otherwise, for $i = 1, \dots, n$, $j = 1, \dots, m$. $\mathbf{y}_1 \in \mathbb{R}_{\max}^n$ is defined by $(\mathbf{y}_1)_i = a$ for all $i = 1, \dots, n$.
2. $A_2 \in \mathbb{R}_{\max}^{m \times m}$ is the $m \times m$ max-plus identity matrix defined by $(A_2)_{ij} = 0$, if $i = j$ and $(A_2)_{ij} = -\infty$, otherwise, for $i, j = 1, \dots, m$. $\mathbf{y}_2 \in \mathbb{R}_{\max}^m$ is defined by $(\mathbf{y}_2)_i = b$ for all $i = 1, \dots, m$.
3. $A_3 \in \mathbb{R}_{\max}^{m(m-1)/2 \times m}$ is defined by $(A_3)_{ij} = 0$, for $j \in p(i)$ and $(A_3)_{ij} = -\infty$, otherwise, for $i = 1, \dots, m(m-1)/2$, $j = 1, \dots, m$, where $p(1), p(2), \dots, p(m(m-1)/2)$ is a list of all the unordered pairs of elements of $\{1, \dots, m\}$. $\mathbf{y}_3 \in \mathbb{R}_{\max}^{m(m-1)/2}$ is defined by $(\mathbf{y}_3)_i = c$, for all $i = 1, \dots, m(m-1)/2$.
4. $A_4 \in \mathbb{R}_{\max}^{1 \times m}$ is defined by $(A_4)_{1j} = 0$ for $j = 1, \dots, m$. $\mathbf{y}_4 \in \mathbb{R}_{\max}$ is given by

$$\mathbf{y}_4 = -na - mb - \frac{m(m-1)}{2}c.$$

Where the coefficients are given by

$$a = m(k+1), \quad b = m - k - 3/2, \quad c = -2.$$

Proof. First note that

$$\sum_{i=1}^N y_i = 0,$$

and that A, \mathbf{y} are a finite form so the results of Lemma 2.11 apply. Now suppose that $\mathbf{z} \in \{0, 1\}^m$ and let $J = \{j : z_j = 1\}$, $x = |J|$ and $u = |\cup_{j \in J} F_j|$. Then we have

$$\begin{aligned} \langle A \otimes \mathbf{z}, \mathbf{y} \rangle &= ua + xb + \left(x(m-x) + \frac{x(x-1)}{2}\right)c - na - mb - \frac{m(m-1)}{2}c \\ &= (x-m)(x - (k+1/2)) + a(u-n), \end{aligned}$$

unless $x = 0$, in which case we have $\langle A \otimes \mathbf{z}, \mathbf{y} \rangle = 0$. Therefore since $0 \leq x \leq m$ and $0 \leq u \leq n$, we have

$$\langle A \otimes \mathbf{z}, \mathbf{y} \rangle > 0 \Leftrightarrow u = n, x \leq k. \quad (43)$$

Now suppose that the set covering problem has a solution $J \subset \{1, \dots, m\}$. Then $|J| \leq k$ and $|\cup_{j \in J} F_j| = n$. Define $\mathbf{z} \in \{0, 1\}^m$, by $z_j = 1$, if and only if $j \in J$, then from (43) we have

$\langle A \otimes \mathbf{z}, \mathbf{y} \rangle > 0$, which from the first part of Lemma 2.11 is equivalent to \mathbf{z} being a descent direction for $\underline{0}$. Conversely suppose that $\mathbf{z} \in \mathbb{R}^m$ is a descent direction for $\underline{0}$, then from the second part of Lemma 2.11 we have that there exists a descent direction $\mathbf{z}' \in \{0, 1\}^m$ and from the first part of Lemma 2.11 we have that $\langle A \otimes \mathbf{z}', \mathbf{y} \rangle > 0$. Now let $J = \{j : z'_j = 1\}$, then from (43) we must have that $|J| \leq k$ and $|\cup_{j \in J} F_j| = n$, i.e. that J is a solution to the set covering problem. \square

Corollary 2.13. *Problem 2.9 is NP-hard.*

Proof. Theorem 2.12 shows that any instance of the set-covering problem can be reduced to an instance of Problem 2.9 with size polynomial in that of the original problem. From [10] we have that the set-covering problem is NP-hard and therefore Problem 2.9 is also NP-hard. \square

2.5 Brute force method

An exhaustive approach to solving the max-plus 2-norm regression problem exactly is to search through all of the feasible patterns, computing the normal projections and checking their admissibility for each one in turn. We then select the closest admissible normal projection for our solution.

To search efficiently through the set of all feasible patterns we consider the tree T , with vertices V_0, \dots, V_n at depths $0, 1, \dots, n$ respectively. The depth k vertices V_k are ordered k -tuples of the form $v = (P_1, \dots, P_k)$, with $P_i \subset \{1, \dots, d\}$, for $i = 1, \dots, k$. The depth k vertices can be interpreted as patterns of support for the max-plus regression problem formed from the first k rows of A and \mathbf{y} . A vertex $v \in V_k$ is parent to $v' \in V_{k+1}$ if and only if $v' = (v, P_{k+1})$ for some $P_{k+1} \subset \{1, \dots, d\}$.

In analogy to the feasibility matrix for a pattern (15), define the feasibility matrix for a vertex $v \in V_k$, by $F_v \in \mathbb{R}_{\max}^{d \times d}$, with

$$f_{jk} = \begin{cases} 0, & \text{for } j = k, \\ \max \{ -\infty, \max \{ a_{ik} - a_{ij} : j \in P_i, i = 1, \dots, k \} \}, & \text{otherwise.} \end{cases} \quad (44)$$

We say that $v \in V_k$ is feasible if $\lambda(F_v^*) = 0$. Note that the set of feasible leaf vertices is identical to the set of feasible patterns of support. It is easy to show that if $v \in V_k$ is feasible then all of its ancestors are feasible and at least one of its children is feasible. Also if $v \in V_k$ is not feasible then none of its children are feasible.

Next we define an order \leq_L on the vertices of T by taking an arbitrary ordering \leq on the subsets of $\{1, \dots, d\}$ and extending this order lexicographically to T . We start at the vertex $v_0 = ()$ and proceed to search through T , in order of \leq_L . At each vertex $v \in V_k$, we check for feasibility by computing $\lambda(F_v^*)$, with worst case cost $\Theta(d^3)$. If $\lambda(F_v^*) > 0$, then v is non-feasible and we skip all of its decedents. Whenever we reach a feasible leaf vertex $P \in V_n$, we compute $\Phi(P, \mathbf{y})$ and $\Psi(P, \mathbf{y}, \underline{-\infty})$ with cost $\Theta(n)$. We check for admissibility by computing $F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty})$ with cost $\Theta(d^2)$. If $F_P \otimes \Psi(P, \mathbf{y}, \underline{-\infty}) = \Psi(P, \mathbf{y}, \underline{-\infty})$ then we compute $\|\mathbf{y} - \Phi(P, \mathbf{y})\|_2$ and if this residual is the best that we have seen so far, then we save $\Psi(P, \mathbf{y}, \underline{-\infty})$ as the interim optimal solution. The algorithm terminates when we have either checked or skipped all vertices in T . See Algorithm 1.

Applied to an $n \times d$ problem, Algorithm 1 must check the feasibility of c_v vertices and compute normal projections and check admissibility for c_l leaf vertices. From (14) we have

$$c_l \leq \sum_{k=1}^n \frac{(n + d - k - 1)!}{(n - k)! \cdot (d - k)! \cdot (k - 1)!}$$

Algorithm 1 Returns an optimal solution to Problem 1.1 with $p = 2$.

```

1: set  $v = ()$ 
2: set  $r_{\min} = \infty$ 
3: while not all of  $T$  explored or skipped do
4:   if  $\lambda(F_v^*) = 0 \Leftrightarrow v$  is feasible then
5:     if  $v = P$  is a leaf vertex then
6:       compute  $\Phi(P, \mathbf{y})$  and  $\Psi(P, \mathbf{y}, -\infty)$ 
7:       if  $F_P \otimes \Psi(P, \mathbf{y}, -\infty) = \Psi(P, \mathbf{y}, -\infty) \Leftrightarrow \Phi(P, \mathbf{y})$  is admissible then
8:         if  $\|\Phi(P, \mathbf{y}) - \mathbf{y}\|_2 < r_{\min}$  then
9:           update  $r_{\min} \leftarrow \|\Phi(P, \mathbf{y}) - \mathbf{y}\|_2$  and  $\mathbf{x} \leftarrow \Psi(P, \mathbf{y}, -\infty)$ 
10:        end if
11:       end if
12:     end if
13:   else
14:     skip decedents of  $v$ 
15:   end if
16:   go to next  $v$ 
17: end while
18: return solution  $\mathbf{x}$ 

```

with equality in the generic case. The total cost of Algorithm 1 is $\Theta(c_v d^3 + c_l d^2)$. We do not currently have any convenient expression for c_v but the contribution to the cost from the leaf vertices alone is greater than polynomial. Although such exhaustive search algorithms will never be suitable for applying to very large problems, they can still be extremely valuable for use on smaller data sets and as a way to benchmark the performance of faster approximate algorithms.

2.6 Newton's method

The results of the previous section and Theorem 2.12 suggest that computing exact local minima for the max-plus $p = 2$ regression problem might not be computationally feasible for larger problems. Instead we propose using the following technique, which consists of Newton's method with an undershooting parameter.

Recall that the squared residual $R : \mathbb{R}_{\max}^d \mapsto \mathbb{R}_+$, given by $R(\mathbf{x}) = \|A \otimes \mathbf{x} - \mathbf{y}\|_2^2 / 2$ is piecewise quadratic and that for $\mathbf{x} \in \text{Cl}(X(P))$, we have $R(\mathbf{x}) = R_P(\mathbf{x})$. Newton's method minimizes a function by iteratively mapping to the minimum of a local quadratic approximation to that function. In the case of the squared residual R this means iteratively mapping to the minimum of the locally quadratic piece. There are several options when implementing Newton's method, for example when \mathbf{x} is contained in the closure of more than one domain, which pattern do we choose? Also, how do we choose between non-unique minima? The method we set out below is chosen primarily for its simplicity.

For $\mathbf{x} \in \mathbb{R}_{\max}^d$, define the *subpattern* $p(\mathbf{x})$ to be the subpattern of $\text{pattern}(\mathbf{x})$

$$p(\mathbf{x})_i = \min(\text{pattern}(\mathbf{x})_i). \quad (45)$$

Then $p(\mathbf{x}) \preceq \text{pattern}(\mathbf{x})$ and $\mathbf{x} \in \text{Cl}(X(p(\mathbf{x})))$. Define the *Newton update map* $\mathcal{N} : \mathbb{R}_{\max}^d \mapsto \mathbb{R}_{\max}^d$,

by

$$\mathcal{N}(\mathbf{x}) = \Psi(p(\mathbf{x}), \mathbf{y}, \mathbf{x}). \quad (46)$$

The map (46) is set to always chooses a pattern whose domain is of the maximum possible dimension. In the case where the minima is non-unique, it returns the one that is closest to the current point.

A difficulty for Newton’s method is that the non-differentiability of R means that the iteration needn’t converge to a local minima and can instead get caught in a periodic orbit. This makes choosing a stopping condition difficult. We use the rule that if the residual has not decreased in some fixed number of steps then we terminate the algorithm and return the best solution from the iterations orbit. We also include a shooting parameter $\mu \in (0, 1)$ and make the update $\mathbf{x} \leftarrow (1 - \mu)\mathbf{x} + \mu\mathcal{N}(\mathbf{x})$. Choosing $\mu < 1$ causes the method to undershoot and so avoid being caught in the periodic orbits mentioned previously. If Algorithm 2 iterates k times then it has cost $\Theta(knd)$. In the Numerical examples that follow we randomly sample ten different initial conditions then apply Algorithm 2 once with $\mu = 1$ then once more with $\mu = 0.05$, each time using $t = 5$, then pick the best approximate solution. Optimizing the choice of parameters and random starting conditions is an important topic for future research.

Algorithm 2 (Newton’s method) Given an initial guess \mathbf{x} , returns an approximate solution to Problem 1.1 with $p = 2$. Parameters are $t \in \mathbb{N}$ the number of iterations for stopping condition and $\mu \in (0, 1)$ the undershooting parameter, which may be allowed to vary during the computation.

```

1: set  $r_{\min} = \infty$ 
2: while not terminated do
3:   update  $\mathbf{x} \leftarrow (1 - \mu)\mathbf{x} + \mu\mathcal{N}(\mathbf{x})$ 
4:   if  $R(\mathbf{x}) < r_{\min}$  then  $r_{\min} = R(\mathbf{x})$ ,  $\hat{\mathbf{x}} = \mathbf{x}$ ,
5:   end if
6:   if  $r_{\min}$  not decreased for  $t$  iterations then terminate
7:   end if
8: end while
9: return approximate solution  $\hat{\mathbf{x}}$ 

```

Applied to an $n \times d$ problem Algorithm 2 has cost $\Theta(nd)$ per iteration.

3 System identification

Consider the d -dimensional stochastic max-plus linear dynamical system

$$\mathbf{x}(n+1) = M \otimes \mathbf{x}(n) + \zeta(n), \quad (47)$$

where $M \in \mathbb{R}_{\max}^{d \times d}$ and $\zeta(0), \zeta(1), \dots \in \mathbb{R}^d$ are i.i.d Gaussians with mean zero and covariance matrix $\sigma^2 I$. Suppose that we do not know M , but that we have observed an orbit $\mathbf{x}(0), \mathbf{x}(1), \dots, \mathbf{x}(N)$ and want to estimate M from this data. The maximum likelihood estimate for this inference problem is given by

$$\max_{A \in \mathbb{R}_{\max}^{d \times d}} \mathbb{P}\{\mathbf{x}(0), \mathbf{x}(1), \dots, \mathbf{x}(N) \mid \mathbf{x}(k+1) = A \otimes \mathbf{x}(k) + \zeta(k), k = 0, 1, \dots, N-1\}. \quad (48)$$

This problem can be expressed as d independent regression problems as follows. Expanding (48) yields

$$\mathbb{P}\{\mathbf{x}(0), \dots, \mathbf{x}(N) \mid A\} = \prod_{n=0}^{N-1} \mathbb{P}\{\zeta(n) = \mathbf{x}(n+1) - A \otimes \mathbf{x}(n)\} \quad (49)$$

$$= \prod_{n=0}^{N-1} \prod_{k=1}^d \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\left(\mathbf{x}(n+1) - A \otimes \mathbf{x}(n)\right)_k^2}{2\sigma^2}\right). \quad (50)$$

The negative log likelihood is therefore given by

$$-\log(\mathbb{P}\{\mathbf{x}(0), \dots, \mathbf{x}(N) \mid A\}) = \frac{Nd}{2} \log(2\pi\sigma^2) + \frac{1}{2\sigma^2} \|A \otimes X(:, 1:N) - X(:, 2:N+1)\|_F^2, \quad (51)$$

where $X \in \mathbb{R}_{\max}^{d \times (N+1)}$ is the matrix whose columns are the time series observations $\mathbf{x}(0), \dots, \mathbf{x}(N)$ and where we use the Matlab style notation $X(\mathcal{I}, \mathcal{J})$ to indicate the submatrix of formed from the intersection of the \mathcal{I} rows and \mathcal{J} columns of X and use the symbol $:$ alone to denote the full range of row/cols. Next note that

$$\|A \otimes X(:, 1:N) - X(:, 2:N+1)\|_F^2 = \sum_{k=1}^d \|A(k, :) \otimes X(:, 1:N) - X(k, 2:N+1)\|_2^2. \quad (52)$$

Maximizing (48) is therefore equivalent to minimizing each of the terms summed over in (52). The k th of these terms measures our model's ability to predict the value of the k th variable at the next time step. To minimize this error we choose the k th row of A by

$$A(k, :) = \arg \min_{\mathbf{x} \in \mathbb{R}_{\max}^{1 \times d}} \|\mathbf{x} \otimes X(:, 1:N) - X(k, 2:N+1)\|_2, \quad (53)$$

which requires us to solve an $n \times d$ max-plus 2-norm regression problem. We can therefore solve (48) by solving d such regression problems.

Example 3.1. Consider the matrix

$$M = \begin{bmatrix} 7 & 15 & 10 & -\infty \\ 14 & -\infty & 11 & 11 \\ 14 & -\infty & -\infty & -\infty \\ 15 & 8 & 7 & 9 \end{bmatrix}.$$

From the initial condition $\mathbf{x}(0) = [0, 0, 0, 0]^\top$ we generate two orbits of length $n = 200$ by iterating (47). One with a low noise level, $\sigma = 1$ and one with a high noise level, $\sigma = 5$. Next we compute the maximum likelihood estimate for M from the time series, by applying Algorithm 2 to each of the row problems (53), for $k = 1, \dots, d$. Our estimates are given by

$$A(\sigma = 1) = \begin{bmatrix} 2.65 & 14.9 & 10.6 & 10 \\ 13.8 & -30.4 & -53.2 & 10.9 \\ 14 & 7.06 & 8.5 & 5.7 \\ 15 & 9.4 & -50.3 & 8.28 \end{bmatrix}, \quad A(\sigma = 5) = \begin{bmatrix} 8.24 & 14.1 & 11 & 1.67 \\ 13.8 & -\infty & 9.28 & 11.1 \\ 13.8 & -\infty & -\infty & -\infty \\ 14.3 & 9.21 & 7.49 & 7.62 \end{bmatrix}.$$

Table 1 displays the Frobenius error term (52) for each of these estimates. Note that both of these estimates fit the data better than the true system matrix M , which indicates that Algorithm 2 is able to find close to optimal solutions to the regression problem.

Comparing our estimates to M , we see that in both cases we have inferred values that are roughly correct for the larger entries in the matrix but that the minus infinities are poorly approximated in both cases and that some of the smaller finite entries are poorly approximated in the low noise case. For each orbit we record the matrix $S \in \mathbb{N}^{4 \times 4}$, with

$$s_{ij} = |\{0 \leq n < N - 1 : (A \otimes x(n))_i = a_{ij} + x(n)_j\}|,$$

which records how often variable j attains the maximum in determining variable i at the next time step, for $i, j = 1, \dots, 4$. Therefore s_{ij} can be thought of as a measure of how much evidence we have to infer the parameter a_{ij} from the orbit. These matrices are given by

$$S(\sigma = 1) = \begin{bmatrix} 0 & 201 & 0 & 0 \\ 167 & 0 & 5 & 29 \\ 201 & 0 & 0 & 0 \\ 197 & 0 & 0 & 4 \end{bmatrix}, \quad S(\sigma = 5) = \begin{bmatrix} 30 & 137 & 34 & 0 \\ 108 & 0 & 41 & 52 \\ 201 & 0 & 0 & 0 \\ 133 & 30 & 12 & 26 \end{bmatrix}.$$

Comparing the results it is clear that our inferences are more accurate for entries with more evidence. In the low noise case the evidence is all contained on a small number of entries, as under the nearly deterministic behavior of this regime only a few positions are ever able to attain the maximum. In the high noise case the more random behavior means that more entries are able to attain the maximum and therefore the evidence is more uniformly distributed, except onto the minus infinity entries, which can never attain the maximum.

Inferring the values of entries that do not play a role in the dynamics or only play a very small role is therefore an ill posed problem and consequently we obtain MLE matrices A that do a good job of fitting the data but which are not close to the true system matrix M . There are two common strategies for coping with such ill posed inverse problems. The first is to choose a prior distribution for the inferred parameters, then compute a maximum a posteriori estimate which minimizes the likelihood times the prior probability. The second approach is to add a regularization penalty to the targeted residual. Regularization is typically used to improved the well-posedness of inverse problems and to promote solutions which are in some way simpler. Typical choices for conventional linear regression problems are the 1-norm or 2-norm of the solution. For max-plus linear 2-norm regression we propose the following regularization penalty, which is chosen to promote solutions $\mathbf{x} \in \mathbb{R}_{\max}^d$ with smaller entries and with more entries equal to $-\infty$.

Problem 3.2. For $A \in \mathbb{R}_{\max}^{n \times d}$, $\mathbf{y} \in \mathbb{R}_{\max}^n$ and $\lambda \geq 0$, we seek

$$\min_{\mathbf{x} \in \mathbb{R}_{\max}^d} \left(\|(A \otimes \mathbf{x}) - \mathbf{y}\|_2^2 + \lambda \sum_{j=1}^d \mathbf{x}_j \right), \quad (54)$$

where we use the convention that in the case when the residual is $+\infty$ and the regularization term is $-\infty$ the total is $+\infty$. When solutions with $-\infty$ totals are possible we order them first by the number of components in \mathbf{x} equal to $-\infty$ and then by the total in (54) not including those $-\infty$ terms.

Rather than searching for a global optima for Problem 3.2, which would contain the largest number of $-\infty$ entries possible, we instead start from a good (ideally optimal) solution to the unregularized problem and search for a nearby solution to Problem 3.2. In this way the regularization term provides a kind of downwards pull on the entries with very little evidence but not enough to disrupt the entries for which there is plenty of evidence.

We solve Problem 3.2 by solving a sequence of max-plus 2-norm regression problems by an approach which is inspired by the iteratively reweighed least squares method for solving conventional 1-norm regularized 2-norm regression problems [11]. Let $I \in \mathbb{R}_{\max}^{d \times d}$ be the max-plus identity matrix with zeros on the diagonal and minus infinities off of the diagonal and consider the residual

$$\left\| \begin{bmatrix} A \\ I \end{bmatrix} \otimes \mathbf{x} - \begin{bmatrix} \mathbf{y} \\ \mathbf{x}' - \lambda/2 \end{bmatrix} \right\|_2^2 = \|A \otimes \mathbf{x} - \mathbf{y}\|_2^2 + \sum_{j=1}^d (\lambda/2 + \mathbf{x} - \mathbf{x}')^2 \quad (55)$$

$$= \|A \otimes \mathbf{x} - \mathbf{y}\|_2^2 + \lambda \sum_{j=1}^d (\mathbf{x} - \mathbf{x}')_j + \mathcal{O}((\mathbf{x} - \mathbf{x}')^2). \quad (56)$$

Therefore, for \mathbf{x} close to $\mathbf{x}' \in \mathbb{R}^d$, (55) only differs from (54) by a constant factor. Algorithm 3 computes a sequence of approximate solutions, each time shifting the additional target variables in (55) to match the gradient of the residual in (54). If a component of the solution appears to be diverging to minus infinity, then we set it to equal this limit.

Algorithm 3 (Iteratively reshifted least squares) Given an initial guess \mathbf{x} , returns an approximate solution to Problem 3.2.

- 1: **while** not converged **do**
- 2: apply a max-plus 2-norm regression solver to compute

$$\mathbf{x} \leftarrow \arg \min_{\mathbf{x}' \in \mathbb{R}_{\max}^d} \left\| \begin{bmatrix} A \\ I \end{bmatrix} \otimes \mathbf{x}' - \begin{bmatrix} \mathbf{y} \\ \mathbf{x}(k-1) - \lambda/2 \end{bmatrix} \right\|_2^2$$

- 3: **end while**
 - 4: return approximate solution \mathbf{x}
-

Example 3.3. *Returning to the problem of Example 3.1. We repeat our analysis of the time-series data only this time we include a regularization term when computing each row via (53). The results of our regularized inference are as follows*

$$A(\sigma = 1, \lambda = 10) = \begin{bmatrix} -\infty & 15 & -\infty & -\infty \\ 13.9 & -\infty & -\infty & 11 \\ 14 & -\infty & -\infty & -\infty \\ 14.9 & -\infty & -\infty & -\infty \end{bmatrix}, \quad A(\sigma = 5, \lambda = 10) = \begin{bmatrix} 8.05 & 14 & 10.5 & -\infty \\ 13.7 & -\infty & 9.03 & 10.9 \\ 13.8 & -\infty & -\infty & -\infty \\ 14.1 & 8.97 & 7.26 & 7.4 \end{bmatrix}.$$

Note that any entry with little of no evidence is set to minus infinity and that the remaining entries are all fairly accurate approximations of the entries in the true system matrix M . Table 1 shows that applying the regularization penalty with $\lambda = 10$ only results in a tiny degradation in the solutions fit to the data.

Table 1: For the inverse problem of Example 3.1. Squared Frobenius norm residual $\|A \otimes X(:, 1 : n) - X(:, 2 : n + 1)\|_F^2$, for low and high noise orbits, with and without regularization penalty. All numeric values given to two decimal places.

	M	$A(\sigma, \lambda = 0)$	$A(\sigma, \lambda = 10)$
$\sigma = 1$	233.78	227.41	251.86
$\sigma = 5$	5308.58	5267.86	5275.12

4 Discussion

In this paper we presented theory and algorithms for max-plus 2-norm regression and then demonstrated how it could be applied to system identification of max-plus linear systems with Gaussian noise.

We have shown how the geometry of max-plus linear spaces give rise to non-convex optimization problems. We have also proven that the exponentially many different patterns of support in a max-plus regression problem means that even the simple problem of determining whether a point is a local minimum is NP-hard.

In spite of these difficulties we did find that Algorithm 2 worked well enough in our example problem. However, developing efficient algorithms that are able to provide some better performance guarantees would be very desirable. Theorem 2.12 seems to stand somewhat in the way of this goal, but note that the theorem relates to a specific point for a highly structured (i.e. degenerate) problem. So it may still be possible to developing an efficient residual descending algorithm.

As noted in the introduction, most applications of max-plus linear dynamical systems use petri-net models, which result in highly structured iteration matrices and noise processes. Further work is needed to adapt the approach used in Section 3 to this setting. Similarly to include a control input as in [14]. Framing these more general inverse problems explicitly in terms of linear regression problems might inspire new techniques, possibly by trying to develop further max-plus analogues of classical linear systems theory.

Acknowledgement

We thank Henning Makhholm for answering a question on Math Stack Exchange, which helped in the formulation of Theorem 2.12.

References

- [1] F. Baccelli and D. Hong. TCP is max-plus linear and what it tells us on its throughput. *SIGCOMM Comput. Commun. Rev.*, 30(4):219–230, 2000.
- [2] A. Björck. *Numerical Methods for Least Squares Problems*. Society for Industrial and Applied Mathematics, 1996. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611971484>.
- [3] C. A. Brackley, D. S. Broomhead, M. C. Romano, and M. Thiel. A max-plus model of ribosome dynamics during mRNA translation. *Journal of Theoretical Biology*, 303(Supplement C):128 – 140, 2012.
- [4] P. Butkovič. *Max-Linear Systems: Theory and Algorithms*. Springer, 2010.

- [5] M. Develin and B. Sturmfels. Tropical convexity. *Documenta mathematica*, 9:1–27, 2004.
- [6] S. S. Farahani, T. van den Boom, and B. De Schutter. Exact and approximate approaches to the identification of stochastic max-plus-linear systems. *Discrete Event Dynamic Systems*, 24(4):447–471, 2014.
- [7] F. Gallot, J. L. Boimond, and L. Hardouin. Identification of simple elements in max-algebra: Application to siso discrete event systems modelisation. In *1997 European Control Conference (ECC)*, pages 1866–1871, 1997.
- [8] B. Heidergott, G. J. Olsder, and J. Woude. *Max Plus at Work: Modeling and Analysis of Synchronized Systems: A Course on Max-Plus Algebra and Its Applications*. Princeton University Press, 2006.
- [9] J. Hook. Critical path statistics of max-plus linear systems with gaussian noise. *Journal of Applied Probability*, (3):654–670, 09 2013.
- [10] R. M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. Springer US, Boston, MA, 1972. ISBN 978-1-4684-2001-2. URL https://doi.org/10.1007/978-1-4684-2001-2_9.
- [11] Q. Lyu, Z. Lin, Y. She, and C. Zhang. A comparison of typical ℓ_p minimization algorithms. *Neurocomput.*, 119:413–424, Nov. 2013. ISSN 0925-2312. URL <http://dx.doi.org/10.1016/j.neucom.2013.03.017>.
- [12] J. Richter-Gebert, B. Sturmfels, and T. Theobald. First steps in tropical geometry. *arXiv:math/0306366*, 2003.
- [13] B. D. Schutter and T. van den Boom. Model predictive control for max-plus-linear discrete event systems. *Automatica*, 37(7):1049 – 1056, 2001.
- [14] B. D. Schutter, T. J. J. van den Boom, and V. Verdult. State space identification of max-plus-linear discrete event systems from input-output data. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 4, pages 4024–4029, 2002.
- [15] S. Sergeev, H. Schneider, and P. Butkovič. On visualization scaling, subeigenvectors and Kleene stars in max algebra. *Linear Algebra and its Applications*, 431:2395–2406, 2009.
- [16] T. J. J. van den Boom and B. De Schutter. Modeling and control of switching max-plus-linear systems with random and deterministic switching. *Discrete Event Dynamic Systems*, 22(3): 293–332, 2012.
- [17] T. J. J. van den Boom, B. D. Schutter, and V. Verdult. Identification of stochastic max-plus-linear systems. In *2003 European Control Conference (ECC)*, pages 618–623, 2003.