Lyapunov rank of polyhedral positive operators

Michael Orlitzky

Background: Cones

Background: Cones

This story is set in a finite-dimensional real Hilbert space V. You can pretend that V is $\mathbb{R}^n$.

If W is another such space, then the set of all linear operators from V to W is $\mathcal{B}(V, W)$.

When $V = W$, we simply write $\mathcal{B}(V)$.

Background: Cones

Definition.

A nonempty subset K of V is a *cone* if $\lambda K = K$ for all $\lambda > 0$. A *closed convex cone* is a cone that is closed and convex as a subset of V.

You might also see this condition with $\lambda \geq 0$. They're the same thing for closed cones.

Background: Cones

Closed convex cones can contain subspaces or fail to have interior: $\mathbb{R}^2$ is a closed convex cone in $\mathbb{R}^3$.

Definition.

A full-dimensional closed convex cone that contains no subspaces is called *proper*.

Background: Cones

The *conic hull* of a nonempty subset X of V is

$$\operatorname{cone}(X) := \left\{ \sum_{i=1}^{m} \alpha_i x_i \;\middle|\; x_i \in X,\ \alpha_i \geq 0,\ m \in \mathbb{N} \right\}.$$

The conic hull is like a convex hull where we extend every point "up" as well as "in."

Background: Cones

Every proper cone K has a set of *extreme directions* $\operatorname{Ext}(K)$ such that

$$K = \operatorname{cone}(\operatorname{Ext}(K)).$$

$\operatorname{Ext}(K)$ is the smallest set with that property.

Definition.

If $\operatorname{Ext}(K)$ is finite, then K is *polyhedral*.

Background: Cones

Example.

The nonnegative orthant $\mathbb{R}^3_+$ in $\mathbb{R}^3$ has the standard basis as its extreme directions,

$$\mathbb{R}^3_+ = \operatorname{cone}\left(\{e_1, e_2, e_3\}\right).$$

$\operatorname{Ext}\left(\mathbb{R}^3_+\right)$ is finite, so $\mathbb{R}^3_+$ is polyhedral.

Example.

Any proper cone in $\mathbb{R}^2$ is polyhedral.

Start with two extreme directions in the plane and try to add a third. If it lies in the cone, it is redundant (not extreme). If it lies outside of the cone, then it renders one of the other two directions redundant.

Background: Cones

Example.

The ice-cream cone in $\mathbb{R}^3$ is not polyhedral.

Clearly, it is the conic hull of its boundary rays; however, if you attempt to remove any boundary ray from the conic hull, a part of the cone will become flat (no longer an ice-cream cone).

Background: Cones

Definition.

If K is a subset of V, then the *dual cone* of K is

$$K^{*} \coloneqq \left\{ y \in V \mid \forall x \in K,\ \langle x, y \rangle \geq 0 \right\}.$$

The dual K^{*} is always a closed convex cone. If K is a closed convex cone, then the duality is faithful and $\left(K^{*}\right)^{*} = K$.

Example.

The nonnegative orthant $\mathbb{R}^n_+$ is self-dual. Note that any element s in its dual must have $\langle s, e_i \rangle = s_i \geq 0$ for every basis vector e_i.

If the entries of s are all nonnegative, then $s \in \mathbb{R}^n_+$. The converse is easy.

Background: Cones

Example. $K = \operatorname{cone}\left(\left\{(3,2)^T, (3,-2)^T\right\}\right).$

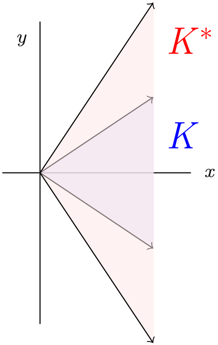

Background: Cones

Example. The ice-cream cone is self-dual.

This is a special case of the following result.

Theorem (Güler [4], 1996).

A cone is symmetric (self-dual and *homogeneous*) if and only if it is the cone of squares in some Euclidean Jordan Algebra.

Background: Cones

Proposition.

A closed convex cone is polyhedral if and only if its dual is polyhedral.

Intuition:

1. Only $\operatorname{Ext}(K)$ matters in the definition of K^*.

2. Every $x \in \operatorname{Ext}(K)$ defines a half-space.

3. K^* is the intersection of half-spaces.

Background: Cones

From now on, every cone will be both proper and polyhedral.

(That's the simplest possible case.)

We're interested in a quantity called the *Lyapunov rank* of a proper polyhedral cone. A few examples motivate its definition.

Part 1, Section 2

Background: Lyapunov rank

Background: Lyapunov rank

Lyapunov rank was introduced in *Bilinear optimality constraints for the cone of positive polynomials* by G. Rudolf, N. Noyan, D. Papp, and F. Alizadeh [6] (2011).

The authors intended to use it to show that the cone of positive polynomials was, in a sense, bad.

Oh, and they called it *bilinearity rank*.

Background: Lyapunov rank

The linear complementarity problem LCP (q, M).

Given: $q \in \mathbb{R}^n$ and $M \in \mathbb{R}^{n \times n}$.

Asked: find an $x \in \mathbb{R}^n$ such that

$$\begin{aligned} x &\geq 0 \\ q + Mx &\geq 0 \\ x^T (q + Mx) &= 0. \end{aligned}$$

Background: Lyapunov rank

If $s \coloneqq q + Mx$ and $K \coloneqq \mathbb{R}^n_+ = K^*$, then the LCP (q, M) asks for a pair (x, s) such that

$$x \in K, \quad s \in K^*, \quad \langle x, s \rangle = 0.$$

Note that $\langle x, s \rangle = 0$ is necessary for any solution.

Background: Lyapunov rank

The primal linear programming problem. Given:

- An objective function $\langle b, \cdot \rangle$ where $b \in \mathbb{R}^n$.
- Some linear constraints $L \in \mathbb{R}^{n \times n}$.
- A shift $c \in \mathbb{R}^n$ for those linear constraints.

We are asked to

$$\begin{aligned} \text{minimize} \ & \langle b, x \rangle \\ \text{subject to} \ & L(x) \geq c \\ & x \geq 0. \end{aligned}$$

Background: Lyapunov rank

The dual linear programming problem asks us to

$$\begin{aligned} \text{maximize} \quad & \langle c, s \rangle \\ \text{subject to } L^{*}(s) & \leq b \\ s & \geq 0. \end{aligned}$$

If $\bar{x}$ solves the primal problem and $\bar{s}$ solves the dual problem, then $\langle \bar{x}, \bar{s} \rangle = 0$. This is called "complementary slackness."

Background: Lyapunov rank

So, being able to solve the equation $\langle x, s \rangle = 0$ helps us solve optimization problems.

Over the cone $K = \mathbb{R}^n_+$, something nice happens.

If $x \in K$ and $s \in K^* = K$, then

$$\langle x, s \rangle = 0 \iff x_i s_i = 0 \text{ for all } i.$$

Background: Lyapunov rank

It's a lot easier to solve n equations $x_i s_i = 0$ than it is to solve the single equation $\langle x, s \rangle = 0$.

Can the same thing happen over other cones?

We can always write the identity operator in terms of some others, say, $\mathrm{id}_V = L_1 + L_2$.

Background: Lyapunov rank

Then,

$$\langle x, s \rangle = 0$$
$$\iff$$
$$\langle \operatorname{id}_V(x), s \rangle = 0$$
$$\iff$$
$$\langle L_1(x), s \rangle + \langle L_2(x), s \rangle = 0.$$

Background: Lyapunov rank

This won't split into two equations, but we can simply require that it does.

Definition.

$L \in \mathcal{B}(V)$ is *Lyapunov-like* on K if $\langle L(x), s \rangle = 0$ for all orthogonal $x \in K$ and $s \in K^{*}$.

Background: Lyapunov rank

If L_1 and L_2 are Lyapunov-like, then that's exactly what we need to split

$$\langle L_1(x), s \rangle + \langle L_2(x), s \rangle = 0$$

into

$$\begin{aligned} \langle L_1(x), s \rangle &= 0 \\ \langle L_2(x), s \rangle &= 0. \end{aligned}$$

Background: Lyapunov rank

But *how many* equations can we get?

The set of all Lyapunov-like operators on K turns out to be a vector space $\mathbf{LL}(K)$ whose dimension is the number of equations we can obtain.

Definition. The *Lyapunov rank* of K is

$$\beta(K) \coloneqq \dim(\mathbf{LL}(K)).$$

(Mnemonic: "beta" is for "bilinearity.")

Background: Lyapunov rank

Example.

The Lyapunov rank of $\mathbb{R}^n_+$ is n because we can get n equations from $\langle x, s \rangle = 0$ when $x, s \in \mathbb{R}^n_+$:

$$\begin{aligned} x_1 s_1 &= 0 \\ x_2 s_2 &= 0 \\ &\vdots \\ x_n s_n &= 0. \end{aligned}$$

Background: Lyapunov rank

Example (Gowda and Tao [3], 2013).

The Lyapunov rank of the ice-cream cone in $\mathbb{R}^n$ is $\left(n^2 - n + 2\right)/2$, much larger than n.

Background: Lyapunov rank

Example (Gowda and Tao [3], 2013).

The cone $\mathcal{S}^n_+$ of symmetric positive semidefinite $n \times n$ matrices has Lyapunov rank n^2.

Note: the elements of $\mathcal{S}^n_+$ live in a space of dimension $\left(n^2 + n\right)/2$ which is less than n^2.

Background: Lyapunov rank

Example.

The positive operators on a proper polyhedral cone K, denoted by $\pi(K)$, have Lyapunov rank

$$\beta(\pi(K)) = \beta(K)^2.$$

Just kidding, I'm going to prove this.

Background: Lyapunov rank

Theorem (Rudolf et al. [6], 2011).

The Lyapunov rank of a proper cone is,

- invariant under invertible linear operators
- additive on cartesian products
- the same as the Lyapunov rank of its dual.

Background: Lyapunov rank

The first two items show that

$$\beta\left(K \oplus H\right) = \beta\left(K\right) + \beta\left(H\right)$$

for proper cones K and H.

This follows since any direct sum can be sent to a cartesian product by an invertible linear operator.

Background: Lyapunov rank

Definition.

A proper cone is *(ir)reducible* if it is (not) a nontrivial direct sum of proper cones.

Theorem (Gowda and Tao [3], 2013).

The Lyapunov rank of any irreducible proper polyhedral cone is one.

$\pi(K)$: *Definition*

$\pi(K)$: Definition

Every closed convex cone K orders its ambient vector space V by

$$x \succcurlyeq y \iff x - y \in K.$$

If K is proper, then this ordering is "nice," it respects the linear structure of V.

$\pi(K)$: Definition

In any ordered vector space $(V, \succcurlyeq)$, an element $x \in V$ is called a *positive element* if $x \succcurlyeq 0$.

A *positive operator* on V is an $L \in \mathcal{B}(V)$ such that $L(x) \succcurlyeq 0$ for all $x \succcurlyeq 0$.

Positive operators preserve positivity.

(The term *positive* is wrong, but standard.)

$\pi(K)$: Definition

Notice that with a proper cone ordering,

- x is a positive element $\iff x \in K$.
- L is a positive operator $\iff L(K) \subseteq K$.

By example, we define *positive operators on* K,

$$\pi(K) \coloneqq \{L \in \mathcal{B}(V) \mid L(K) \subseteq K\}.$$

$\pi(K)$: Definition

Example (Perron-Frobenius).

Let $K = \mathbb{R}^n_+$, the nonnegative orthant in $\mathbb{R}^n$. Then the positive operators on K are the real $n \times n$ matrices having nonnegative elements.

Let $L \in \pi(K)$ and $\rho(L)$ be its spectral radius. The Perron-Frobenius theorem states that $L(x) = \rho(L)\,x$ for some $x \succcurlyeq 0$.

$\pi(K)$: Definition

In fact, we can extend the definition of a positive operator to two cones $K \subseteq V$ and $H \subseteq W$,

$$\pi(K, H) := \{L \in \mathcal{B}(V, W) \mid L(K) \subseteq H\}.$$

We will need the general version to prove our result for the simpler $\pi(K)$ case.

$\pi(K)$: Lyapunov rank

$\pi(K)$: Lyapunov rank

Goal: compute the Lyapunov rank $\beta(\pi(K))$.

Note: this goal makes sense.

Proposition (Schneider and Vidyasagar [7], 1970).

If K and H are proper polyhedral cones, then $\pi(K, H)$ is too.

$\pi(K)$: Lyapunov rank

What we'd like to do:

1. Decompose $\pi(K, H)$ into a direct sum of irreducible cones.
2. Use the fact that Lyapunov rank is additive on a direct sum.
3. Conclude that $\beta(\pi(K, H)) = \beta(K)\,\beta(H)$ is one in the base case.
4. Hand-wave induction.

$\pi(K)$: Lyapunov rank

Here's what was known towards that goal.

Proposition (Barker and Loewy [1], 1975).

K is reducible if and only if $\pi(K)$ is reducible.

Proposition (Haynsworth, Fiedler, and Pták [5], 1976).

If K or H is reducible, then $\pi(K, H)$ is reducible.

$\pi(K)$: Lyapunov rank

And here's what's missing:

Theorem.

$\pi(K, H)$ is reducible if and only if either K or H is reducible.

(The converse of Haynsworth, Fiedler, and Pták.)

$\pi(K)$: Lyapunov rank

Proof.

Copy the proof of Barker and Loewy, who proved the result for $H = K$, line-for-line. Then change K^* to H^* everywhere. $\square$

Now, when K and H are irreducible, we know that $\pi(K, H)$ is too.

$\pi(K)$: Lyapunov rank

Recall: the Lyapunov rank of a proper polyhedral irreducible cone is one. So suppose that K and H are irreducible. Then,

$$\beta(K)\,\beta(H) = 1.$$

For the same reason, $\beta(\pi(K, H)) = 1$. Thus

$$\beta(\pi(K, H)) = \beta(K)\,\beta(H)$$

when K and H are irreducible.

$\pi(K)$: Lyapunov rank

For the general case, suppose $K = K_1 \oplus K_2$ and $H = H_1 \oplus H_2$ are direct sums of irreducible cones. Lyapunov rank is additive on a direct sum, so

$$\begin{aligned}\beta(K)\,\beta(H) &= \beta(K_1)\,\beta(H_1) \\ &+ \beta(K_1)\,\beta(H_2) \\ &+ \beta(K_2)\,\beta(H_1) \\ &+ \beta(K_2)\,\beta(H_2) \\ &= 4.\end{aligned}$$

$\pi(K)$: Lyapunov rank

What about $\pi(K, H)$ in this case?

There exist invertible linear A and B such that

$$A(K) = K_1 \times K_2$$
$$B(H) = H_1 \times H_2.$$

Lyapunov rank is invariant under invertible linear operators, so the extra A, B won't matter.

$\pi(K)$: Lyapunov rank

It turns out that

$$\pi(A(K), B(H)) = B \circ \pi(K, H) \circ A^{-1}.$$

But, $X \mapsto BXA^{-1}$ is an invertible linear operator, so that won't matter either.

$\pi(K)$: Lyapunov rank

Since our maps A and B won't matter, throw them away for simplicity, and pretend that

$$K = K_1 \times K_2$$
$$H = H_1 \times H_2.$$

Now what is $\pi(K, H)$?

$\pi(K)$: Lyapunov rank

If $V_i := \operatorname{span}(K_i)$ and $W_i := \operatorname{span}(H_i)$,

$$\pi(K, H) \subseteq \left\{ \begin{bmatrix} A & B \\ C & D \end{bmatrix} \;\middle|\; \begin{array}{l} A \in \mathcal{B}(V_1, W_1) \\ B \in \mathcal{B}(V_2, W_1) \\ C \in \mathcal{B}(V_1, W_2) \\ D \in \mathcal{B}(V_2, W_2) \end{array} \right\}.$$

$\pi(K)$: Lyapunov rank

It's easy to check that for $\pi(K, H)$,

$$
\begin{aligned}
A &\in \pi(K_1, H_1) \\
B &\in \pi(K_2, H_1) \\
C &\in \pi(K_1, H_2) \\
D &\in \pi(K_2, H_2).
\end{aligned}
$$

If any of those fail, the same counterexample shows that the whole thing isn't in $\pi(K, H)$.

$\pi(K)$: Lyapunov rank

For example, the space of 2×2 real matrices is isomorphic to $\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}$. Likewise,

$$
\begin{aligned}
\pi(K, H) \cong\; & \pi(K_1, H_1) \\
& \times \pi(K_2, H_1) \\
& \times \pi(K_1, H_2) \\
& \times \pi(K_2, H_2).
\end{aligned}
$$

$\pi(K)$: Lyapunov rank

Each factor $\pi(K_j, H_i)$ is irreducible, because K_j and H_i are. The additivity of Lyapunov rank therefore gives,

$$\beta(\pi(K, H)) = 1 + 1 + 1 + 1 = \beta(K)\,\beta(H)\,.$$

$\pi(K)$: Lyapunov rank

If it works with two factors, it works for more:

The number of terms in $\beta(K)\,\beta(H)$ is equal to the number of blocks possessed by a block-form operator in $\pi(K, H)$.

Each term/block contributes one to the Lyapunov rank.

$\pi(K)$: Lyapunov rank

Theorem.

If K and H are proper polyhedral cones, then $\beta(\pi(K, H)) = \beta(K)\,\beta(H)$.

Corollary.

When $H = K$, we have $\beta(\pi(K)) = \beta(K)^2$.

Part 2, Section 5

$\pi(K)$: Lyapunov-like operators

$\pi(K)$: Lyapunov-like operators

Definition.

If $x, s \in V$, we define $s \otimes x$ to be the linear map $t \mapsto \langle x, t \rangle s$. That is,

$$(s \otimes x)(t) := \langle x, t \rangle s.$$

In finite dimensions, $s \otimes x$ can be thought of as the matrix sx^T.

$\pi(K)$: Lyapunov-like operators

For subsets $X, S \subseteq V$ we will write

$$S \otimes X := \{s \otimes x \mid s \in S, x \in X\}.$$

This is simply Minkowski notation.

It is known that $\dim(S \otimes X) = \dim(S)\dim(X)$.

$\pi(K)$: Lyapunov-like operators

Proposition (Berman and Gaiha [2], 1972).

If K and H are proper polyhedral cones, then,

$$\pi(K, H)^* = \operatorname{cone}(H^* \otimes K).$$

For polyhedral cones, it follows that

$$\operatorname{Ext}(\pi(K, H)^*) = \operatorname{Ext}(H^*) \otimes \operatorname{Ext}(K).$$

$\pi(K)$: Lyapunov-like operators

Recall that the Lyapunov rank of a cone's dual is the same as that of the original cone. Thus,

$$\beta(\pi(K, H)^*) = \beta(K)\,\beta(H)\,.$$

We're going to conjure up some Lyapunov-like operators on $\pi(K, H)^*$, and this equation tells us when to quit.

$\pi(K)$: Lyapunov-like operators

Theorem (Gowda and Tao [3], 2013).

If K is a proper polyhedral cone, then L is Lyapunov-like on K if and only if every element of $\operatorname{Ext}(K)$ is an eigenvector of L.

Since we know $\operatorname{Ext}(\pi(K, H)^{*})$, its Lyapunov-like operators are now within reach.

$\pi(K)$: Lyapunov-like operators

The elements of $\operatorname{Ext}(\pi(K, H)^{*})$ look like $s \otimes x$ where $x \in \operatorname{Ext}(K)$ and $s \in \operatorname{Ext}(H^{*})$.

Consider the following operator on such a thing:

$$[M \odot L](s \otimes x) \coloneqq M(s) \otimes L(x) \cong (Ms)(Lx)^{T}.$$

This is the *Kronecker product* of M and L.

$\pi(K)$: Lyapunov-like operators

The Kronecker product is another type of tensor product, but the symbol $\otimes$ is worn out.

However, $\dim(\mathbf{M} \odot \mathbf{L}) = \dim(\mathbf{M}) \dim(\mathbf{L})$, since that was true of sets of tensor products.

$\pi(K)$: Lyapunov-like operators

Proposition.

Let K and H be proper polyhedral cones.

If L is Lyapunov-like on K and M is Lyapunov-like on H^*, then $M \odot L$ is Lyapunov-like on $\pi(K, H)^*$.

Proof.

Let $s \otimes x \in \operatorname{Ext}\left(\pi\left(K, H\right)^{*}\right)$ be arbitrary, and show that it's an eigenvector of $M \odot L$.

We have $x \in \operatorname{Ext}\left(K\right)$ and $s \in \operatorname{Ext}\left(H^{*}\right)$, so x is an eigenvector of L and s is an eigenvector of M. Thus,

$$M\left(s\right) \otimes L\left(x\right) = \lambda_{1}\lambda_{2}\left(s \otimes x\right). \qquad \square$$

$\pi(K)$: Lyapunov-like operators

Now consider the space of all such operators,

$$\operatorname{span}\left(\mathbf{LL}\left(H^{*}\right) \odot \mathbf{LL}\left(K\right)\right).$$

This has dimension $\beta(K)\,\beta(H)$, which we now know to be the Lyapunov rank of $\pi(K, H)^{*}$. And, they're all Lyapunov-like on $\pi(K, H)^{*}$.

It follows that the two spaces are equal.

$\pi(K)$: Lyapunov-like operators

We're almost there, we need one more result.

Proposition (Rudolf et al. [6], 2011).

L is Lyapunov-like on K if and only if its adjoint L^{*} is Lyapunov-like on the dual K^{*}.

$\pi(K)$: Lyapunov-like operators

Theorem.

If K and H are proper polyhedral cones, then

$$\mathbf{LL}(\pi(K, H)) = \operatorname{span}(\mathbf{LL}(H) \odot \mathbf{LL}(K^*)).$$

Proof.

Use the result for $\pi(K, H)^*$ and take duals/adjoints on both sides. $\square$

$\pi(K)$: Lyapunov-like operators

Corollary.

If K is a proper polyhedral cone, then

$$\mathbf{LL}\left(\pi\left(K\right)\right) = \operatorname{span}\left(\mathbf{LL}\left(K\right) \odot \mathbf{LL}\left(K^{*}\right)\right).$$

References

[1] G. P. Barker and R. Loewy.
The structure of cones of matrices.
Linear Algebra and its Applications, 12:87–94, 1975.

[2] A. Berman and P. Gaiha.
A generalization of irreducible monotonicity.
Linear Algebra and its Applications, 5:29–38, 1972.

[3] M. S. Gowda and J. Tao.
On the bilinearity rank of a proper cone and Lyapunov-like transformations.
Mathematical Programming, 147:155–170, 2014.

[4] O. Güler.
Barrier functions in interior point methods.
Mathematics of Operations Research, 21:860–885, 1996.

References

[5] E. Haynsworth, M. Fiedler, and V. Pták.
Extreme operators on polyhedral cones.
Linear Algebra and its Applications, 13:163–172, 1976.

[6] G. Rudolf, N. Noyan, D. Papp, and F. Alizadeh.
Bilinear optimality constraints for the cone of positive polynomials.
Mathematical Programming, 129:5–31, 2011.

[7] H. Schneider and M. Vidyasagar.
Cross-positive matrices.
SIAM Journal on Numerical Analysis, 7:508–519, 1970.