Reducibility of convex cones

Michael Orlitzky

January 30, 2026

1 Convex cones

Convex cones are generalizations of real vector spaces. While the elements of a real vector space can be scaled in any direction (positively or negatively), the elements of a convex cone can only be scaled positively.

Convex cones arise naturally in optimization where typically you are trying to minimize of maximize some function within a *feasible* set, and (at least locally) there is a "good direction" you want to go in: it's OK to move positively in the good direction, and bad to move away from it. If your set is nice enough, the directions you're allowed to move in will also form (at least locally) a convex cone. Satisfying both at the same time leads to a third, smaller cone of directions that are both good and allowable.

In many instances, the feasible set will also involve a convex cone. The most familiar example is linear programming where you are optimizing over the nonnegative orthant $\mathbb{R}^n_+$, but second-order programming (over the ice-cream cone) and semidefinite programming over (over the real PSD cone) are also incredibly popular as well. In each case we start with a convex cone, take a slice of it, and then try to minimize or maximize some function over the resulting set. The (free) book *Convex Optimization* by Boyd and Vandenberghe covers this in depth [5].

This last sort is generically referred to as **conic programming**. If K_1 and K_2 are closed convex cones in some Euclidean space V, if $A : V \to V$ is linear, and if $b, c \in V$ are two points, then the "primal" conic programming problem is to

$$\begin{aligned} \text{minimize} \quad & \langle b, x \rangle \\ \text{subject to} \quad & A(x) \in K_2 + c \\ & x \in K_1 \end{aligned}$$

If you are willing to use weird-enough cones, conic programming can be used to solve NP-hard problems:

- Quadratic programs including the quadratic assignment problem
- Maximum stable set of a graph
- Chromatic number of a graph

The chapter "Copositive programming" by Samuel Burer in the *Handbook on Semidefinite, Conic and Polynomial Optimization* is a good overview [1].

While it may occasionally be overkill, we shall assume that we are working in a finite-dimensional real inner-product space. For all intents and purposes, this is $\mathbb{R}^n$. We call these *Euclidean spaces* from now on.

Definition 1 (convex cones)**.** A subset of a Euclidean space is *convex* if the line segment between any two points of the set is entirely contained in the set: if x and y belong to the set, then so does $\{\alpha x + (1-\alpha)\, y \mid \alpha \in [0,1]\}$. A *cone* is a set that is closed under positive scaling: if x belongs to the set, then so does αx when $\alpha > 0$. A *convex cone* is a cone that is convex. Another way to state this is to say that $x + y$ belongs to a convex cone whenever x and y do.

This definition of cone omits the origin, but we can remedy that by including its boundary. The result is a (topologically) closed set. There are a few other conic adjectives that will play important roles. They are usually desirable; the more adjectives we have, the nicer our cone will be[1].

Definition 2 (species of convex cone)**.** A *closed convex cone* is a convex cone that is topologically closed as a subset of the ambient Euclidean space. A convex cone is *pointed* if it contains no lines, and *solid* if it spans the ambient space. When all of these are true at once, the cone is called *proper*, the nicest adjective of all.

Closed convex cones are often thought of as being generated by some smaller set. If we can reconstruct a the cone K from another set G by taking all nonnegative linear combinations of the elements in G, then we say that K is generated by G, and write $K = \operatorname{cone}(G)$.

Formally,

$$\operatorname{cone}(G) \coloneqq \left\{ \sum_{i=1}^{m} \alpha_i g_i \;\middle|\; g_i \in G,\ \alpha_i \geq 0,\ m \in \mathbb{N} \right\},$$

and K is generated by G if $\operatorname{cone}(G) = K$. The terminology is a bit ambiguous, but we say that the elements of G are *generators* of K in this case.

Lemma 1. *If K is a convex cone in a nontrivial Euclidean space V, then the following are equivalent:*

1. *K is pointed*
2. *$K \cap (-K) \subseteq \{0\}$*
3. *$K \setminus \{0\}$ is strictly contained in some half-space of V*

Geometrically, the last item here means that a pointed cone K consists of all positive multiples of one of its cross-sections. If $q \in V$ defines the half-space, then $C \coloneqq \{x \in K \mid \langle x, q \rangle = \pm 1\}$ is one such cross-section (called a *base*).

[1] Differential geometers will disagree about "closed."

Lemma 2. *If K is a convex cone generated by a set G in the Euclidean space V, then the following are equivalent:*

1. *K is solid*
2. *K contains a basis for V*
3. *G contains a basis for V*

The first two items here are more or less the definition of "solid." The third item is however easy to see because whatever basis you obtain from the second item must be expressable in terms of the elements of G.

Example 1. In the space $\mathbb{R}^n$ with the standard basis $\{e_1, e_2, \ldots, e_n\}$, the non-negative orthant

$$\mathbb{R}^n_+ \coloneqq \left\{ (x_1, x_2, \ldots, x_n)^T \;\middle|\; x_i \geq 0 \text{ for all } i \right\}$$

is a closed, convex, pointed, and solid (i.e. proper) cone generated by the standard basis. That is, $\mathbb{R}^n_+ = \operatorname{cone}(\{e_1, e_2, \ldots, e_n\})$.

Example 2. In the space $\mathbb{R}^n \cong \left(\mathbb{R}, \mathbb{R}^{n-1}\right)$, the ice-cream cone

$$\mathcal{L}^n_+ \coloneqq \left\{ (x_1, \bar{x})^T \;\middle|\; x_1 \geq \|\bar{x}\| \right\}$$

is a closed, convex, pointed, and solid (i.e. proper) cone. It is generated by the non-ice-cream part of the ice-cream cone.

Example 3. In the space $\mathcal{S}^n$ of real, symmetric $n \times n$ matrices, the set of all positive-semidefinite matrices

$$\mathcal{S}^n_+ \coloneqq \{A \in \mathcal{S}^n \mid \langle Ax, x \rangle \geq 0 \text{ for all } x \in \mathbb{R}^n\}$$

is a closed, convex, pointed, and solid (i.e. proper) cone generated by the rank-one matrices of the form xx^T for $x \in \mathbb{R}^n$. (Think spectral decomposition.)

Any generating set[2] can be scaled to have unit norm without affecting the cone. In Examples 2 and 3 we did not do this, so our generating sets contain "extra" elements.

Example 4. In the space $\mathbb{R}^n$, any vector subspace (including $\mathbb{R}^n$ itself) is a closed convex cone. The subspace $\mathbb{R}^n$ is solid and $\{0\}$ is pointed. How many vectors are needed for a generating set of $\mathbb{R}^2$?

Example 5. If K_1 and K_2 are both closed (pointed, solid) convex cones, then $K_1 \times K_2$ is a (pointed, solid) convex cone. If G_1 and G_2 generate K_1 and K_2 respectively, then $(G_1 \times \{0\}) \cup (\{0\} \times G_2)$ generates $K_1 \times K_2$.

[2] Ignoring generating sets that pointlessly contain the origin.

Example 6. If K is a (closed, pointed, solid) convex cone and if A is invertible, then $A(K)$ is a (closed, pointed, solid) convex cone. If G generates K, then $A(G)$ generates $A(K)$.

The last two examples hint at a generic way to construct big cones from little ones. Start with two cones K_1 and K_2 in spaces V_1 and V_2, respectively, and take their cartesian product. Then apply an invertible linear map

$$A : (V_1 \times V_2) \to (V_1 \times V_2) .$$

The result is a new convex cone $J \coloneqq A(K_1 \times K_2)$. But what if we we are *given* a cone like J? Can we reverse the process to recover K_1 and K_2?

2 Reducibility

A vector space V is a direct sum of subspaces V_1 and V_2 if $V = V_1 + V_2$ as sets and if $V_1 \cap V_2 = \{0\}$. In that case, we write $V = V_1 \oplus V_2$. Typically we want $V_1, V_2 \neq \{0\}$ and say that the direct sum is nontrivial in that case. This notion of direct sum easily extends to other non-subspace sets $S = S_1 \oplus S_2 \subseteq V$ by requiring that $S_1 \subseteq V_1$ and $S_2 \subseteq V_2$ for subspaces V_1, V_2 such that $V = V_1 \oplus V_2$.

Definition 3. A convex cone is *reducible* if it is the nontrivial direct sum of two convex cones, and *irreducible* if not. Formally, a convex cone $K \subseteq V$ is reducible if and only if there exist nontrivial subspaces $V_1, V_2 \subseteq V$ such that $V = V_1 \oplus V_2$, and convex cones $K_1 \subseteq V_1$ and $K_2 \subseteq V_2$ for which $K = K_1 + K_2$.

Reducibility is important in many contexts; classical examples can be found the famous book by Berman and Plemmons [3]. A few modern ones:

- The Jordan-automorphism group of a Euclidean Jordan algebra [18], acts transitively on the set of Jordan frames if and only if the cone of squares is irreducible or $\mathbb{R}^n_+$.
- If K is a proper polyhedral cone, then its Lyapunov rank is one if and only if it is irreducible [10].
- If K is a proper polyhedral cone in n dimensions and if $\mathbf{Z}(K)^*$ is the dual of the cone of $\mathbf{Z}$-operators on K, then K is irreducible if and only if $\dim\left(\mathbf{Z}(K)^*\right) \geq n^2 - 1$. This combines results of Orlitzky, Haynesworth, Fiedler, and Pták [17].

If you start with a reducible cone and reduce it repeatedly, what you wind up with is a direct sum of irreducible cones. The sum is unique (up to order) if the original cone was proper. This result has been proved many times in many different settings, to varying degrees of uniqueness. Taking Lemma 1 in to account, decompositions of compact convex sets can also be seen as decompositions of pointed convex cones.

1963 Proposition 1 in Chapter I §1 of Vinberg [25], for open, pointed, and solid convex cones.

1964 As Theorem 9.3 of Bleicher and Schneider, for pointed convex cones in finitely-generated modules over an ordered ring with the ascending chain condition on left ideals [4].

1966 Rothaus [21], for "regular" cones, but again in the context of constructing homogeneous cones.

1967 Kaneyuki [14], for affine-homogeneous Siegel domains (from which all homogeneous convex cones arise).

1968 Asano [2], for homogeneous cones in T-algebras.

1970 Gruber [11] for affine convex bodies (in German).

1975 Gale and Klee [9] in the more general context of topological commutive semigroups. (Euclidean spaces are commutative topological group with respect to addition.)

1982 Meschiari [15], who cites for it the disserations of Gentili and O'Connor.

1994 As Proposition III.4.5 of Faraut and Korányi, for symmetric cones [7], where the it follows from the decomposition of a Euclidean Jordan algebra into simple components.

2002 As Theorem 4.3 of Hauser and Güler [13], for pointed convex cones in Euclidean spaces with uniqueness when the cone is solid. (The authors do mention Gruber's 1970 paper.)

2011 Theorem 1.13 of Rohrer [20] in the context of toric geometry. The uniqueness for a line is mentioned here.

2014 Theorem 4 of Bremner, Sikirić, Pasechnik, Rehn and Schürmann [6] along with its computability.

Perhaps the simplest example is the nonnegative orthant which can be expressed as the direct sum of one-dimensional rays.

Example 7. Recall the nonnegative orthant in $\mathbb{R}^n$,

$$\mathbb{R}^n_+ = \operatorname{cone}\left(\{e_1, e_2, \ldots, e_n\}\right).$$

If we write

$$\mathbb{R}^n = \bigoplus_{i=1}^{n} \operatorname{span}\left(\{e_i\}\right)$$

then the nonnegative orthant becomes a direct sum as well:

$$\mathbb{R}^n_+ = \bigoplus_{i=1}^{n} \operatorname{cone}\left(\{e_i\}\right).$$

Why bother? In most cases, it is easier to study two small objects than one big one.

If your cone is not pointed or solid, it can still be decomposed, though perhaps not in a unique manner. A cone that is not pointed cannot be uniquely decomposed unless it is a line because, for example, every basis for a two-dimensional subspace (a reducible convex cone) leads to a different decomposition of that subspace into a direct sum of two one-dimensional spaces (irreducible convex cones).

Example 8. The space $\mathbb{R}^2$ is a reducible convex cone in $\mathbb{R}^2$, yet it can be expressed as the sum of irreducible convex cones in an infinite number of ways:

$$\mathbb{R}^2 = \operatorname{span}\left(\left\{(1,0)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(0,1)^T\right\}\right),$$
$$\mathbb{R}^2 = \operatorname{span}\left(\left\{(1,1)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(1,-1)^T\right\}\right),$$

et cetera.

We note that every line is a sum of convex cones, since $\mathbb{R} = \mathbb{R}_+ + (-\mathbb{R}_+)$, but the sum is not direct because there is no (nontrivial) direct sum decomposition of a one-dimensional space.

A cone that is not solid can fail to decompose uniquely for similar reasons, though this is more of a technicality. Every decomposition of a convex cone corresponds to a decomposition of the ambient space. If your cone is one-dimensional and if the ambient space is three-dimensional, then again, you are free to choose any basis for the two-dimensional subspace that your cone doesn't touch. This leads to a farcical decomposition of the cone $\{0\}$ into $\{0\} \oplus \{0\}$, but in a way that is not unique, because the subspaces that are hidden by the notation are not unique.

Example 9. The set $\{0\}$ in $\mathbb{R}^2$ is a pointed closed convex cone, but it is not solid. And it can be decomposed into a direct sum but not in a unique way. Let,

$$W_1 \coloneqq \operatorname{span}\left(\left\{(1,0)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(0,1)^T\right\}\right), \text{ and}$$
$$W_2 \coloneqq \operatorname{span}\left(\left\{(1,1)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(1,-1)^T\right\}\right)$$

be two direct sum decompositions of $\mathbb{R}^2$, as in the previous example. Then

$$\{0\} = \{0\} \oplus \{0\} \subseteq W_1, \text{ and}$$
$$\{0\} = \{0\} \oplus \{0\} \subseteq W_2$$

are two irreducible decompositions of $\{0\}$.

Despite these counter-examples, there is a canonical decomposition for any convex cone. To work around the difficulties, we must deal with the lineal and trivial components of the cone separately.

Definition 4 (lineality space)**.** If K is a closed convex cone, it is convenient to define the *lineality space* of K to be the largest subspace contained in K, and denote it by $\text{linspace}(K)$.

The dimension of $\text{linspace}(K)$ measures "how not pointed" the cone K is. Similarly, the dimension of the subspace $\text{span}(K)^\perp$ measures "how not solid" it is. It is straightforward to verify that $\text{linspace}(K) = K \cap (-K)$, so by Lemma 1, K is pointed if and only if $\text{linspace}(K) = \{0\}$.

With that out of the way, the following is well-known.

Theorem 1 (Stoer and Witzgall, Theorem 2.10.5)**.** *Every closed convex cone K in a Euclidean space has an orthogonal direct sum decomposition,*

$$K = \text{linspace}(K) \oplus K \cap \text{linspace}(K)^\perp,$$

where $K \cap \text{linspace}(K)^\perp$ is pointed.

And of course, if a closed convex cone K is not solid, then it is easy to decompose it into a direct sum $K = S \oplus \{0\}$ where S is solid within $\text{span}(K)$, and $\{0\}$ lives in the space orthogonal to K. By combining these two results we obtain a (unique) decomposition of any closed convex cone K into a sum,

$$\begin{aligned} K &= \text{linspace}(K) \oplus K \cap \text{linspace}(K)^\perp \oplus \{0\} \\ &\subseteq \\ V &= \text{linspace}(K) \oplus \Big(\text{span}(K) \cap \text{linspace}(K)^\perp\Big) \oplus \text{span}(K)^\perp. \end{aligned}$$

The middle term here is a proper cone and can be uniquely decomposed as a sum of irreducible proper cones. The remaining terms are a vector space (in itself) and the trivial cone (in the space orthogonal to K), which *are* unique, so long as we don't try to decompose them. This leads to the following "result."

Theorem 2. *If K is a closed convex cone in a finite dimensional real inner-product space V, then, up to the order of the factors, it has a unique decomposition*

$$K = L \oplus Z \oplus P$$

where

$$\begin{aligned} L &= \text{linspace}(K) \textit{ in itself}, \\ Z &= \{0\} \textit{ in } \text{span}(K)^\perp, \\ P &= \bigoplus_{i=1}^{m} P_i \textit{ in } \text{span}(K) \cap \text{linspace}(K)^\perp, \end{aligned}$$

and each P_i is proper and irreducible in $\text{span}(P_i)$.

If you are willing to abuse the notation for direct sum, the component $Z = \{0\}$ can of course be omitted, as if you were working with a solid cone in $\text{span}(K)$ all along.

Proposition 1 (Wets-Witzgall [26])**.** *If K is a closed convex cone in a Euclidean space V and if $K = \operatorname{cone}(G)$ for some $G \subseteq V$, then*

$$\operatorname{linspace}(K) = \operatorname{cone}(G \cap \operatorname{linspace}(K)).$$

In other words, the lineality space of K is generated by (only) those generators of K that lie in it.

Proof. As usual, one inclusion is obvious. Conic combinations are more restrictive than linear combinations, so they cannot be used to escape a vector space:

$$\operatorname{cone}(G \cap \operatorname{linspace}(K)) \subseteq \operatorname{span}(\operatorname{linspace}(K)) = \operatorname{linspace}(K).$$

For the other inclusion, suppose that $x \in \operatorname{linspace}(K) = K \cap (-K)$. Then $x \in \operatorname{cone}(G) \cap \operatorname{cone}(-G)$, and it can be expressed as two different conic combinations in terms of G and $-G$:

$$x = \sum_{i=1}^{\ell} \alpha_i g_i = \sum_{j=1}^{m} \beta_j h_j.$$

where $g_i \in G$, $h_j \in -G$, and each term of each sum is (without loss of generality) non-zero. If either sum is empty, then $x = 0$, which technically belongs to $\operatorname{cone}(G \cap \operatorname{linspace}(K))$, since $\{0\} = \operatorname{cone}(\emptyset)$. On the other hand, if both sums are nonempty, then,

$$\alpha_1 g_1 = \left(\sum_{j=1}^{m} \beta_j h_j \right) + \left(\sum_{i=2}^{\ell} \alpha_i (-g_i) \right).$$

We know that $\alpha_1 g_1 \in K$, and the right-hand side is in $\operatorname{cone}(-G) = -K$. Thus

$$g_1 \in K \cap (-K) = \operatorname{linspace}(K).$$

The choice of $\alpha_1 g_1$ here can of course be repeated with any other $\alpha_i g_i$. Doing so, we conclude that $x = \sum \alpha_i g_i$ belongs to $\operatorname{cone}(G \cap \operatorname{linspace}(K))$. □

Modulo the lineality space, a similar result holds for any direct sum decomposition of a pointed cone.

Proposition 2. *If $K = K_1 \oplus K_2$ is a pointed closed convex cone in $V = V_1 \oplus V_2$ and if $K = \operatorname{cone}(G)$ for some $G \subseteq V$, then $K_1 = \operatorname{cone}(G \cap V_1)$ and likewise for K_2.*

Proof. Clearly, if $x \in \operatorname{cone}(G \cap V_1)$, then $x \in K_1$.

Conversely, suppose $x \in K_1$ is nonzero. Expressing each $g \in G \subseteq K$ as $p + q$ with $p \in K_1$ and $q \in K_2$, we have $x = \sum_{i=1}^{m} \alpha_i g_i = \sum_{i=1}^{m} \alpha_i (p_i + q_i)$ with all $\alpha_i > 0$. By grouping these,

$$x = \underbrace{\left(\sum_{i=1}^{m} \alpha_i p_i \right)}_{\in V_1} + \underbrace{\left(\sum_{i=1}^{m} \alpha_i q_i \right)}_{\in V_2}.$$

Since $x \in V_1$, the second sum must be zero. But K_2 is pointed (because K is pointed), so the only way that can happen is if all q_i are zero (Lemma 1, strict half-space condition). Thus $g_i = p_i \in V_1$ for all g_i used to express x. □

Unfortunately we cannot just combine Propositions 1 and 2 to conclude that, upon decomposing *any* convex cone, its generators will lie in the subspaces involved in the direct sum. If $K = \text{cone}(G)$, the subset of G that lives outside of $\text{linspace}(K)$ will not necessarily generate the cone $K \cap \text{linspace}(K)^\perp$ in Theorem 1. The remaining generators can however be projected on to the subspace $\text{linspace}(K)^\perp$ to produce generators for the pointed cone; at that point Proposition 2 may be applied.

Lemma 3. *If $K = \text{cone}(G)$, then the generators of $K \cap \text{linspace}(K)^\perp$ are the orthogonal projections of G onto the subspace $\text{linspace}(K)^\perp$.*

Proof. If $x \in K \cap \text{linspace}(K)^\perp$, then $x = \sum_{i=1}^m \alpha_i g_i$ for $g_i \in G$ because x is in K. Project x onto $\text{linspace}(K)^\perp$, which is linear and does nothing to x, to express x as a conic combination of projections of the g_i. □

In the next theorem we refer to $\mathbb{A}$, the field of algebraic real numbers[3]. Basically these are the rational numbers, but with all non-imaginary square, cube (etc.), roots adjoined. The reason for the switch is that we cannot honestly claim to have an algorithm that works over $\mathbb{R}$, but neither do we want to introduce the nightmare called floating point arithmetic. The rational numbers would be perfect, but we cannot perform orthogonal projections over the rationals because we cannot normalize vectors. (Try to normalize $(1,1)^T \in \mathbb{Q}^2$.) The algebraic reals, on the other hand, are a field—so linear algebra works—and have the roots that we need to normalize.

Remark 1. Though the decimal expansion of $\sqrt{2}$ cannot be stored on a computer, we can identify $\sqrt{2}$ with a polynomial $X^2 - 2$ that has integer coefficients. With a bit of cleverness and a lot of abstract algebra, this idea can be used to perform exact computations involving $\sqrt{2}$. This is implemented to various extents by the FLINT library and the SageMath system [23, 24].

Theorem 3. *If $K = \text{cone}(G)$ and if $G \subseteq \mathbb{A}^n$ is finite, then there is an algorithm to compute the decomposition in Theorem 2.*

Proof. The component $L = \text{linspace}(K)$ is easy to compute—Wets and Witzgall tell you how to do it [26]. Afterwards, $\text{linspace}(K)^\perp$ is easy to find, and the subset of G that does *not* live in $\text{linspace}(K)$ can be projected on to it using basic linear algebra. Thus we obtain a pointed cone $P = \text{cone}(H)$ to which we may apply the algorithm [6] from Theorem 4 of Bremner et al.

H contains a basis for $\text{span}(P)$ by Lemma 2. Place the generators $h \in H$ in a matrix as the columns, and use Gaussian elimination to reduce it to row-echelon

[3] https://en.wikipedia.org/wiki/Algebraic_number

form. The pivot columns will then indicate which of the original columns formed a linearly-independent set[4], i.e. a basis $\{h_{i_1}, h_{i_2}, \ldots, h_{i_m}\} \subseteq H$ for $\operatorname{span}(P)$.

At this point, pause momentarily and compute $\operatorname{span}(K)^{\perp}$, the orthogonal complement of $L \oplus \operatorname{span}(P)$. This is the subspace that the component $Z = \{0\}$ will live in.

Continuing now, it remains to find the components P_i such that $P = \bigoplus P_i$. Express each $h_i \in H$ in terms of the basis that we computed a moment ago:

$$h_i = \alpha_1(h_i)\, h_{i_1} + \alpha_2(h_i)\, h_{i_2} + \cdots + \alpha_m(h_i)\, h_{i_m}.$$

Construct a graph (V, E) with vertex set $V \coloneqq H$ and edges

$$E \coloneqq \{(h_i, h_{i_k}) \mid \alpha_k(h_i) \neq 0\}.$$

In this graph, the edges indicate whether or not each generator h_i of P has a nonzero coefficient in terms of the basis element h_{i_k}, which also happens to be a generator of P. From Proposition 2, there is a path in this graph from h_i to h_{i_k} if and only if h_i and h_{i_k} are generators of the same irreducible component of P. All that remains is to compute the vertex sets H_i in each connected component of the graph (V, E). This is a standard depth-first search algorithm in graph theory. We then set $P_i \coloneqq \operatorname{cone}(H_i)$ and are done. □

This algorithm has been implemented in SageMath [24]. Perhaps the simplest example we can give is of decomposing the nonnegative orthant in $\mathbb{Q}^3$ into a direct sum of three rays. Here there is no need to project because the cone is proper, so we can get away with using the faster (and conceptually simpler) rational arithmetic.

[4]Starting with a matrix in row-echelon form, the pivot columns must form a linearly-independent set, because each one has a nonzero coordinate in a position where the others don't. And at each step, the elementary row operations are invertible, and do not alter the linear-independence of a set of linearly-independent columns. Thus the original matrix, being obtainable by this procedure (row-echelon in reverse), must have the same set of linearly-independent columns.

```
sage: e1 = vector(QQ, (1,0,0))
sage: e2 = vector(QQ, (0,1,0))
sage: e3 = vector(QQ, (0,0,1))
sage: K = Cone([e1,e2,e3])
sage: K.irreducible_factors()
{1-d cone in 3-d lattice N,
 1-d cone in 3-d lattice N,
 1-d cone in 3-d lattice N}
sage: r1 = Cone([e1])
sage: r2 = Cone([e2])
sage: r3 = Cone([e3])
sage: K.irreducible_factors() == {r1,r2,r3}
True
```

This makes it easy to test whether or not a proper cone is irreducible. If we decompose it and there's only one cone in the decomposition, then that cone is the one we started with, and it's irreducible. For example, the ℓ_1-norm cone is irreducible.

```
sage: g1 = ( 1,  0, 1)
sage: g2 = ( 0,  1, 1)
sage: g3 = (-1,  0, 1)
sage: g4 = ( 0, -1, 1)
sage: K = Cone([g1, g2, g3, g4])
sage: K.irreducible_factors()
{3-d cone in 3-d lattice N}
sage: K.is_reducible()
False
```

3 Automorphisms

If K is a convex cone in a Euclidean space V, then an automorphism of K is any invertible linear map $A : V \to V$ such that $A(K) = K$. Unless your set carries some additional structure (like a group or manifold), this is simply the definition of an automorphism of a set in a vector space, applied to a set that happens to be a convex cone.

The set of all automorphisms of K form a group that is denoted by $\mathrm{Aut}(K)$. Knowing the automorphism group of your cone can be a big help, but it is usually

not very easy to find. Thankfully for polyhedral cones there is *some* hope, and it just so happens to provide us with an application for our decomposition.

Definition 5. A convex cone K is *polyhedral* if $K = \operatorname{cone}(G)$ for some finite set G. Alternatively, K is polyhedral if and only if it is the intersection of a finite number of half-spaces [19].

These two methods of describing a polyhedral convex cone are known as the $\mathcal{V}$- and $\mathcal{H}$-representations, for vertex ($\mathcal{V}$) and half-space ($\mathcal{H}$). Often one is more useful than the other. Or you can use both: that would be the famous double-description method of Motzkin et al. from 1953 [8].

Note that, per the half-space characterization, all polyhedral convex cones are closed. Polyhedrality comes in to play because when your cone has a finite number of generators (as in Theorem 3), it becomes possible to enumerate them. Many properties of closed convex cones need only be checked on a generating set as opposed to at every point of the cone, the latter of course being impossible unless the cone is trivial (consists of only the origin). But if some generating set is finite, we can check that instead.

Proposition 3 (Güler and Tunçel [12], Proposition 2.2)**.** *If K is a pointed polyhedral cone and if G is a* ***minimal*** *generating set of K, then up to positive scaling, every $A \in \operatorname{Aut}(K)$ permutes the elements of G. That is, for all $g \in G$, there exists a $\tilde{g} \in G$ such that $A(\operatorname{cone}(\{g\})) = \operatorname{cone}(\{\tilde{g}\})$.*

The "minimal" here is important, because there is nothing preventing you from adding extra junk elements to a generating set.

Example 10. Let $G_1 \coloneqq \{e_1, e_2, e_3\}$ be the standard basis in $\mathbb{R}^3$. Then as we have already seen, $\mathbb{R}^3_+ = \operatorname{cone}(G_1)$. But suppose we define

$$G_2 \coloneqq G_1 \cup \{e_1 + e_2\}.$$

Then $\mathbb{R}^3_+ = \operatorname{cone}(G_2)$ as well, but G_2 is a strict superset of G_1.

If we are checking some property on the generators of $\mathbb{R}^3_+$, then it is a waste of time to check it on $e_1 + e_2$, so in a practical sense it is better to have a minimal set of generators. But apropos of Proposition 3, define $A : \mathbb{R}^3 \to \mathbb{R}^3$ by

$$\begin{aligned} A(e_1) &= e_2, \\ A(e_2) &= e_3, \\ A(e_3) &= e_1. \end{aligned}$$

This is an automorphism of $\mathbb{R}^3_+$, and it permutes the minimal generating set G_1, but not G_2 because it sends $(e_1 + e_2)$ to $(e_2 + e_3)$, which isn't a positive multiple of anything in G_2.

The "minimal generating set" for a pointed cone hints at an important concept, the extreme ray. Formally, an extreme ray is a simply a one-dimensional face, but we wish to define neither of these terms. Intuitively, an extreme ray of

a pointed convex cone is one whose elements are not generated by the other rays of the cone. So a minimal generating set must contain the extreme rays, and conversely, the extreme rays generate the entire cone. The latter is a famous theorem of Minkowski (1911) and Steinitz (1916), though nowadays the general case is known as the Krein-Milman[5] theorem [22]. This only works for pointed cones because the theorem is stated for compact convex sets—not cones—so to use it, we must be able to take a compact convex slice of the cone (cf. the paragraph following Lemma 1) that scales to the whole thing.

The moral of the story is,

1. Up to positive scalings, automorphisms of a pointed convex cone must preserve a minimal generating set.

2. For polyhedral cones, minimal generating sets are finite.

3. The set of all permutations of a finite set is again finite.

The set of *scalings* is not finite, but with a little extra work and some Theorem 3, enumerating the permutations of a minimal generating set will produce the full automorphism group [6]. In the interest of brevity, however, we will focus only on the automorphisms that permute without scaling. Bremner et al. call this the **linear symmetry group** of K corresponding to the vectors $v \in G$.

In the context of automorphisms, if the polyhedral cone K is not pointed or not solid, it's not a big deal. We already know the automorphisms of $\text{linspace}(K)$ and of $\{0\} \subseteq \text{span}(K)^\perp$, they're the invertible linear maps on $\text{linspace}(K)$ and $\text{span}(K)^\perp$ respectively. So without loss of generality, we now assume that our cone is proper. As in Theorem 3 we work over $\mathbb{A}^n$ because, if the original cone was not pointed, the generators of the proper cone that we obtain via projection will live in $\mathbb{A}^n$, even if the original generators had rational coordinates.

Theorem 4. *If K is a proper polyhedral cone and if $G \subseteq \mathbb{A}^n$ is a finite generating set for K, then there is an algorithm to find the linear symmetry group of K corresponding to G.*

Proof. Redundant generators in G can be removed: with an $\mathcal{H}$-representation it's easy, but if all else fails, checking if one point is in the conic hull of the others is a standard LP feasibility problem. So with G minimal, we now consult Proposition 3.

Though K is n-dimensional, G can contain more than n elements, and no fewer than n of them will fully determine a linear transformation on $\text{span}(K)$. We therefore consider every pair of n-element subsets S_1 and S_2 of G. If there is a solution A to the linear system $AS_1 = S_2$ and if the resulting A also permutes the remaining elements of G, then we keep it; otherwise, we move on to the next pair of n-element subsets (S_1, S_2). Note that if A permutes G, then A must be invertible because G contains a basis $\text{span}(K)$. Every A that passes the test is therefore an automorphism of K. □

[5] https://en.wikipedia.org/wiki/Krein-Milman_theorem

We rather conspicuously did not reduce K to a direct sum before computing its automorphisms. To avoid a lengthy digression, we will not delve into the complexity of Theorem 4; we will however note that most of its run-time comes from the combinatorial step, enumerating the pairs of subsets (S_1, S_2). If G contains N elements, then there are $2\binom{N}{n}$ ways to choose (S_1, S_2).

For an example, consider the cone $K = \operatorname{cone}(G)$ in $\mathbb{A}^3$ formed by raising the regular octagon in $\mathbb{A}^2$ up to height one and by taking nonnegative multiples of the resulting set. For this K we have $N = 8$ and $n = 3$, so there are $2\binom{8}{3} = 112$ ways to choose S_1 and S_2. Now let's glue together two copies of K to form $J \coloneqq K \times K = \operatorname{cone}(H)$, where

$$H \coloneqq \operatorname{cone}\left((G \times \{0\}) \cup (\{0\} \times G)\right).$$

is minimal after we remove $(0, 0)$ from it. We then have $N = 16$ and $n = 6$, leading to $2\binom{16}{6} = 16,016$ pairs of subsets. Clearly, it is better to enumerate the 112 possibilities for K twice than it is to enumerate the $16,016$ possibilities for J. This is true even when you consider that performing the decomposition requires approximately $N\left(n^2 + n\right) = 672$ operations at the outset [6].

There is one final caveat: when we are treating the two copies of K separately as subsets $K \times \{0\}$ and $\{0\} \times K$, we must also enumerate the automorphisms that swap them. If you approach this naively and are unlucky, it can require an additional 112 attempts to map the subsets $S_1 \subseteq (G \times \{0\})$ to $S_2 \subseteq (\{0\} \times G)$. In any case, the numbers remain heavily in favor of reduction-first.

This tradeoff won't always be beneficial (take $\mathbb{R}^n_+$ for example), but efficiency aside, the irreducible factors are needed to compute the *full* automorphism group [6], so it is pointless to worry about.

4 Conclusion

We have tried to motivate the reduction of a closed convex cone to its constituent parts, but we have also left a lot out:

1. All of Section 3 can be reformulated in terms of isomorphisms (two cones) rather than automorphisms (one cone).

2. While discussing the decomposition of $J = K \times K$ in $\mathbb{A}^6$, we counted the operations as if the linear algebra all took place in $\mathbb{A}^6$. In reality, it is better to perform a change of basis $\phi : K \times \{0\} \to K \subseteq \mathbb{A}^3$, at which point all of the $n = 6$ in the complexity calculation can be replaced by $n = 3$. At the end the change-of-basis must be undone, but ideally the bargain will be in your favor.

3. In practice, enumerating subsets is not the smartest way to compute automorphisms. The face lattice of a pointed polyhedral cone forms a poset, and isomorphic cones have isomorphic (as posets) face lattices [6]. As with the connected components of a graph, this is a great opportunity to punt to someone else's algorithm.

4. We chose the linear symmetry computation as an example because its combinatorial nature highlights the benefits of reduction. However there are many others. Take for example the computation of the positive operators between two proper cones K and H. It is known that

$$\pi(K, H) \cong \bigtimes_{i=1}^{m} \bigtimes_{j=1}^{n} \pi(K_i, H_j),$$

and the algorithm that computes it involves a generating set and taking duals. It is much more efficient on simpler cones in smaller spaces [16, 17].

5. Though we have swept the details under the rug, the algorithm to compute the full automorphism group is itself an intriguing application of reducibility [6].

References

[1] Miguel F. Anjos and Jean B. Lasserre. *Handbook on Semidefinite, Conic and Polynomial Optimization*, vol. 166 of *International Series in Operations Research & Management Science*. Springer, New York, 2012. ISBN 9781461407683, doi:10.1007/978-1-4614-0769-0.

[2] Hiroshi Asano. *On the irreducibility of homogeneous convex cones*. Journal of the Faculty of Science, University of Tokyo, 15(2):201–208, 1968, doi:10.15083/00039854.

[3] Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. SIAM, Philadelphia, 1994. ISBN 9780898713213, doi:10.1137/1.9781611971262.

[4] Michael N. Bleicher and Hans Schneider. *The decomposition of cones in modules over ordered rings*. Journal of Algebra, 1(3):233–258, 1964, doi:10.1016/0021-8693(64)90021-3.

[5] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2009. ISBN 9780521833783.

[6] David Bremner, Mathieu Dutour Sikirić, Dmitrii V. Pasechnik, Thomas Rehn, and Achill Schürmann. *Computing symmetry groups of polyhedra*. LMS Journal of Computation and Mathematics, 17(1):568–581, 2014, doi:10.1112/S1461157014000400.

[7] Jacques Faraut and Adam Korányi. *Analysis on Symmetric Cones*. Clarendon Press, Oxford, 1994. ISBN 9780198534778.

[8] Komei Fukuda and Alain Prodon. *Double description method revisited*. In Michel Deza, Reinhardt Euler, and Ioannis Manoussakis (eds.), *Combinatorics and Computer Science*, 91–111. Springer, Berlin, Heidelberg, 1996. ISBN 9783540706274, doi:10.1007/3-540-61576-8_77.

[9] David Gale and Victor Klee. *Unique reducibility of subsets of commutative topological groups and semigroups*. Mathematica Scandinavica, 36(1):174–198, 1975, doi:10.7146/math.scand.a-11570.

[10] Muddappa Seetharama Gowda and Jiyuan Tao. *On the bilinearity rank of a proper cone and Lyapunov-like transformations*. Mathematical Programming, Series A, 147:155–170, 2014, doi:10.1007/s10107-013-0715-3.

[11] Peter M. Gruber. *Zur charakterisierung konvexer körper. über einen satz von Rogers und Shephard. II.* Mathematische Annalen, 184(1):79–105, 1970, doi:10.1007/BF01350311.

[12] Osman Güler and Levent Tunçel. *Characterization of the barrier parameter of homogenous convex cones*. Mathematical Programming, 81(1):55–76, 1998, doi:10.1007/BF01584844.

[13] Raphael A. Hauser and Osman Güler. *Self-scaled barrier functions on symmetric cones and their classification*. Foundations of Computational Mathematics, 2(2):121–143, 2002, doi:10.1007/s102080010022.

[14] Soji Kaneyuki. *On the automorphism groups of homogeneuous bounded domains*. Journal of the Faculty of Science, University of Tokyo, 14(1):89–130, 1967, doi:10.15083/00039871.

[15] Mauro Meschiari. *The group of isometries of a homogeneous convex irreducible self adjoint cone*. Annali di Matematica Pura ed Applicata, 132(1):47–68, 1982, doi:10.1007/BF01760976.

[16] Michael Joseph Orlitzky. *Lyapunov rank of polyhedral positive operators*. Linear and Multilinear Algebra, 66(5):992–1000, 2018, doi:10.1080/03081087.2017.1331998.

[17] Michael Joseph Orlitzky. *Positive and Z-operators on closed convex cones*. Electronic Journal of Linear Algebra, 34:444–458, 2018, doi:10.13001/1081-3810.3782.

[18] Michael Joseph Orlitzky. *Proscribed normal decompositions of Euclidean Jordan algebras*. Journal of Convex Analysis, 29(3):755–766, 2022.

[19] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970. ISBN 9780691015866.

[20] Fred Rohrer. *Completions of fans*. Journal of Geometry, 100(1):147–169, 2011, doi:10.1007/s00022-011-0073-3.

[21] Oscar S. Rothaus. *The construction of homogeneous convex cones*. Annals of Mathematics, Second Series, 83(2):358–376, 1966, doi:10.2307/1970436.

[22] Walter Rudin. *Functional Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, Inc., New York, second ed., 1991. ISBN 0070542368.

[23] The FLINT team. *FLINT: Fast Library for Number Theory*, 2025. URL https://flintlib.org/.

[24] The Sage Developers. *SageMath, the Sage Mathematics Software System*, 2026. URL https://www.sagemath.org/.

[25] Ernest Borisovich Vinberg. *The theory of convex homogeneous cones*. Transactions of the Moscow Mathematical Society, 12:340–403, 1963.

[26] Roger J.-B. Wets and Christoph Witzgall. *Algorithms for frames and lineality spaces of cones*. Journal of Research of the National Bureau of Standards, Section B: Mathematics and Mathematical Physics, 71B(1):1–7, 1967, doi:10.6028/jres.071B.001.