A Collection of Convex Cones

Applications to Linear, Second-Order, Semidefinite, and Conic Programming

Printable Topics in Mathematics

Benjamin Grimmer

Johns Hopkins, AMS

Printable Topics in Mathematics Project #3 - First Posted 7/10/2022

What are Convex Cones? Shapes like Icecream Cones.

Formalizing Definitions for Convex Cones

We capture the idea of being like an icecream cone with the following two notions: cones and convexity. To make our definitions rigorous, we will talk about points (typically with names like p and q) inside a bigger space $\mathcal{E}$. You can safely think of this as being the three-dimensional world we live in, where each point p or q is really just three numbers giving the x, y, z coordinates of a point in space. We are concerned with describing shapes in this space, not just points, so we will consider some collection of these points, denoted by a set $\mathcal{K} \subseteq \mathcal{E}$.

What makes $\mathcal{K}$ a cone? Cones (colloquially) all have a tip that they extend outward from: the bottom of your icecream cone or the top of a traffic cone. Mathematically let's restrict to this point being the origin (the coordinates $(0, 0, 0)$). Then we say $\mathcal{K}$ is a *"cone"* if it is closed under positive rescalings. Written formally, we require for every $p \in \mathcal{K}$, the set contains the ray passing through it:

$$\lambda p \in \mathcal{K} \qquad \forall \lambda > 0 \ .$$

What makes $\mathcal{K}$ a convex? Convex shapes (colloquially) have no divots or dimples: a convex lens is rounded as opposed to concave lenses that bend inward. Icecream filling the whole interior of our cone. We say a set $\mathcal{K}$ is a *"convex"* if for every pair of points in $\mathcal{K}$, the line segment connecting them is also in the set (preventing any divots). Mathematically we examine all the points between p and q via a weighted averages of them $\lambda p + (1 - \lambda)q$, taking $\lambda \in [0, 1]$ fraction of p and filling the remainder with $1 - \lambda$ fraction with q. Algebraically then, convexity amounts to requiring for every $p, q \in \mathcal{K}$,

$$\lambda p + (1 - \lambda)q \in \mathcal{K} \qquad \forall \lambda \in [0, 1] \ .$$

Shapes with both of these properties are convex cones, extending radially outward with no holes or dimples.

Convex cones occur widely in mathematical modeling as a tool describing common problem structures. This captures problems dealing with nonnegative numbers (see Figure 1), with distances and norms (see Figure 2), and rather surprisingly, with spectral matrix properties (see Figures 3-4).

Numerous planning problems in operations research are formulated over convex cones: Minimizing costs while shipping nonnegative amounts of goods, maximizing profit constrained to only small changes in a system's design. For several of the cones described and illustrated below, the resulting optimization problems have grown their own areas of research: *"Linear Programming"* for nonnegative cones, *"Second-Order Cone Programming"* for Euclidean distances, and *"Semidefinite Programming"* for various problems concerning eigenvalues of matrices.

The Nonnegative Orthant

Perhaps the simplest convex cone is the set of nonnegative numbers living on the real line. It is immediate to check that this set is a cone (multiplying a nonnegative number by any positive λ is certainly still nonnegative) and is convex (the average of two nonnegative numbers always remains nonnegative). A mild generalization of this to vectors gives our first 3D convex cone: the set of vectors with all nonnegative entries.

The nonnegative orthant plays a fundamental role in linear programming. There one seeks a nonnegative vector x minimizing a linear function and satisfying a system of equations (for related algorithms, see the simplex and interior-point methods).

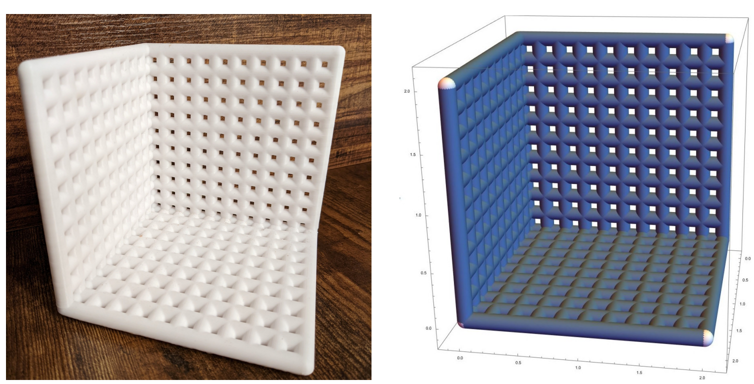

Figure 1. In three dimensions, this gives the nonnegative orthant, containing one-eighth of the whole space with three boundary planes where $x = 0$, where $y = 0$, and where $z = 0$.

3D Model Viewer and Source File Details

Disclaimer before you dive into these files: I am a mathematician professionally with only a self-taught/amateur background in three-dimensional printing and modeling.

- ❑ `.stl` files and 3D model viewer are available at

  `printables.com/model/239642`

- ❑ `.nb` Mathematica file is available at

  `github.com/profgrimmer/Cones`

- ❑ This `.pdf` is available at

  `ams.jhu.edu/~grimmer/Cones.pdf`

The Second-Order Cone

Norms provide a natural way to measure the size of a point $\|p\|$ and the distance between two points as $\|p-q\|$. The graph of any norm is always a cone (the proof amounts to noting every norm assigns the positive rescaling λp norm $\lambda\|p\|$). To make this cone convex, we include everything above the graph, giving the following set (called the *epigraph*)

$$\{(x,u)\in\mathcal{E}\times\mathbb{R} \mid u\geq\|x\|\}\ .$$

Considering the *Euclidean norm* $\|x\|_2=\sqrt{\sum_{i=1}^n x_i^2}$ gives the Second-Order Cone, capturing our typical idea of distance.

Optimization over this cone arises in stochastic linear programs, modeling Gaussian noise, and in control theory, bounding the magnitude of adjustments and deviations from targets.

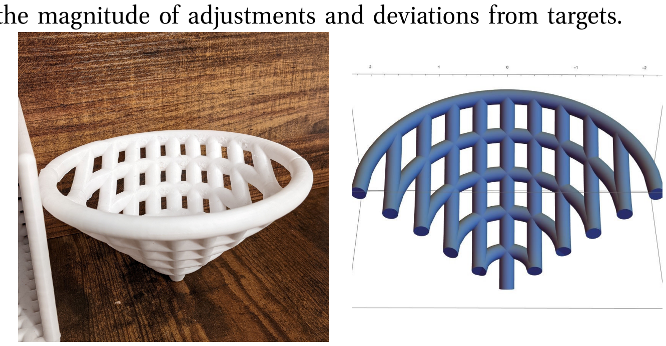

Figure 2. Taking the Second-Order Cone for two-dimensional vectors gives a graph/epigraph living in three dimensions. The model slices the cone, showing only the cone's back-half.

The Positive Semidefinite Matrix Cone

Beyond considering vectors, matrices give rise to very interesting and useful convex cones. One natural generalization of nonnegative vectors to $n\times n$ symmetric matrices is given by *positive semidefiniteness* ("psd" for short). A psd matrix P is one with

$$x^\top Px := \sum P_{ij}x_ix_j \geq 0 \quad \text{for every vector } x.$$

The set of psd matrices form a cone as $x^\top(\lambda P)x=\lambda x^\top Px\geq 0$.

Surprisingly, the psd cone strictly generalizes both the Nonnegative Orthant and Second-Order Cone. A diagonal matrix is psd if and only if its diagonal entries form a nonnegative vector. Only slightly more involved, considering Schur complements gives a reduction for the above Second-Order Cone.

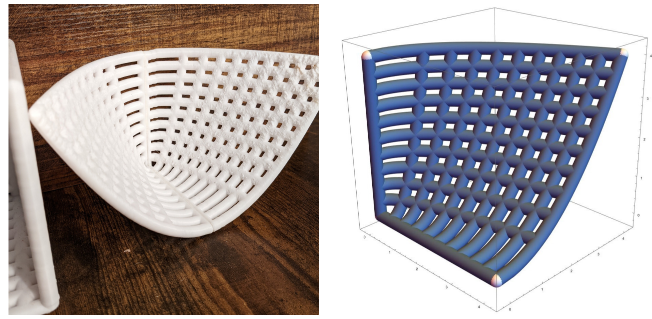

Figure 3. All 2×2 symmetric matrices P are given by some

$$(x,y,z)\mapsto\begin{bmatrix} x & y\\ y & z\end{bmatrix}\ .$$

The printed cone is all (x,y,z) mapping to a psd matrix. The model slices the cone, showing only the cone's right-half.

The Copositive Matrix Cone

A slightly enlarged family of matrices is given by relaxing the psd definition to only require nonnegativity when nonnegative vectors x are applied. A matrix P is *copositive* if

$$x^\top Px\geq 0 \quad \text{for every nonnegative vector } x.$$

Although this change seems mild, it fundamentally changes the nature of the resulting cone. To date, no efficient algorithms for checking if a matrix is copositive are known (in fact, this is an NP-Complete task whereas checking psd-ness is relatively easy).

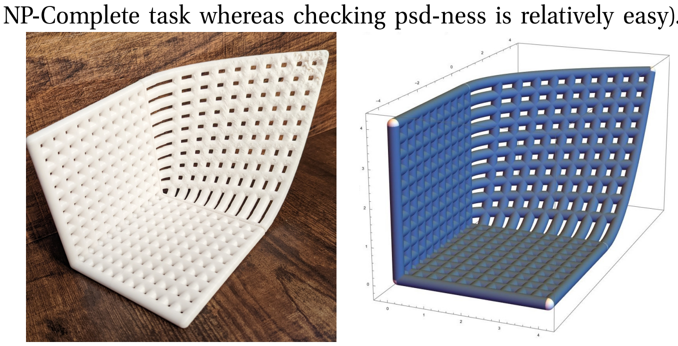

Figure 4. The enlarged set of 2×2 copositive matrices with the same embedding as the positive semidefinite matrices in Figure 3.

Finally, a Trivial Example: Every Subspace

Lastly, note that every subspace is a cone since subspaces are closed under scalar multiplication (not just positive scalings as cones require) and are closed under addition (not just weighted averages as convexity requires).

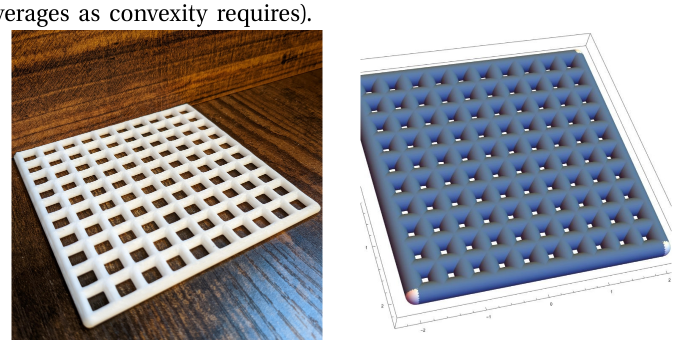

Figure 5. The $z=0$ two-dimensional plane lying inside three dimensions gives an admittedly rather degenerate convex cone.