



Terracini convexity

James Saunderson¹ · Venkat Chandrasekaran²

Received: 19 November 2020 / Accepted: 9 January 2022 / Published online: 5 March 2022
© The Author(s) 2022

Abstract

We present a generalization of the notion of neighborliness to non-polyhedral convex cones. Although a definition of neighborliness is available in the non-polyhedral case in the literature, it is fairly restrictive as it requires all the low-dimensional faces to be polyhedral. Our approach is more flexible and includes, for example, the cone of positive-semidefinite matrices as a special case (this cone is not neighborly in general). We term our generalization Terracini convexity due to its conceptual similarity with the conclusion of Terracini's lemma from algebraic geometry. Polyhedral cones are Terracini convex if and only if they are neighborly. More broadly, we derive many families of non-polyhedral Terracini convex cones based on neighborly cones, linear images of cones of positive-semidefinite matrices, and derivative relaxations of Terracini convex hyperbolicity cones. As a demonstration of the utility of our framework in the non-polyhedral case, we give a characterization based on Terracini convexity of the tightness of semidefinite relaxations for certain inverse problems.

Keywords Face lattice · Hyperbolic programming · Moments and nonnegative polynomials · Neighborly polytopes · Semidefinite programming

Mathematics Subject Classification 52A20 · 52A22 · 52B12 · 90C22

V. C. was supported in part by National Science Foundation grant CCF-1637598, in part by National Science Foundation grant DMS-2113724, and in part by AFOSR grant FA9550-20-1-0320. J. S. was supported in part by an Australian Research Council Discovery Early Career Researcher Award (project number DE210101056) funded by the Australian Government.

✉ James Saunderson
james.saunderson@monash.edu

Venkat Chandrasekaran
venkate@caltech.edu

¹ Department of Electrical and Computer Systems Engineering, Monash University, Clayton, VIC 3800, Australia

² Department of Computing and Mathematical Sciences and Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125, USA

1 Introduction

The combinatorial view of polytopes is a pillar of polyhedral theory which has played a prominent role both in deepening our understanding of the structure of polytopes as well as in illuminating those attributes of polytopes that are significant in the context of particular applications such as linear programming. A parallel perspective for non-polyhedral convex sets—even in the presence of additional structure—has generally been lacking. This limitation may be attributed to the fact that the central object of study in polyhedral combinatorics is the face lattice, and consequently, many of the key ideas and definitions in the field are face-centric. However, face-centric notions do not always carry over naturally to the non-polyhedral setting for a number of reasons; in particular, non-polyhedral closed convex sets consist of infinitely many faces, may contain non-exposed faces, may lack faces of all dimensions, may not be closed under linear images, and so forth. Motivated by this broad challenge of bridging the gap in our understanding between the polyhedral and non-polyhedral cases, we focus in this article on the question of obtaining a suitable generalization of *neighborliness* for non-polyhedral convex sets, with a less face-centric reformulation of neighborliness of polytopes playing a central role in our development.

A polyhedral cone that is pointed is called *k-neighborly* if the cone over any subset of up to k extreme rays forms a face [13].¹ Neighborliness arises in many contexts in geometry and polyhedral combinatorics, most notably in the characterization of various extremal classes of polytopes [13] and in conditions under which linear programming relaxations are tight for certain nonconvex inverse problems [9].

1.1 Motivation

We are aware that there is a definition available for non-polyhedral k -neighborly convex cones that are closed and pointed which parallels the polyhedral setting [14]—that is, the cone over any subset of up to k extreme rays forms an exposed face. However, this notion is too restrictive in the non-polyhedral case as it essentially requires that all the low-dimensional faces are polyhedral, and, in particular, are linearly isomorphic to orthants. This limitation restricts the utility of neighborliness in the non-polyhedral context in a number of ways.

As one example, the cone of positive semidefinite matrices is not k -neighborly for any $k > 1$ as all the faces other than the extreme rays are non-polyhedral, and as a consequence, neighborliness is not useful for characterizing tightness of semidefinite relaxations for nonconvex problems that are ubiquitous in many applications [8, 15], in contrast to the situation with linear programming. Concretely, Donoho and Tanner [9] used neighborliness of polytopes to characterize the exactness of linear programming relaxations for identifying nonnegative vectors with the smallest number of nonzeros in affine spaces. A similar characterization of the success of semidefinite relaxations for identifying low-rank positive semidefinite matrices in affine spaces—a problem

¹ Neighborliness is usually defined for convex polytopes but it is more convenient in this article to consider the polyhedral cones. One can recover equivalent notions for compact convex sets by taking bases of convex cones that are closed and pointed.

that arises in a range of applications such as factor analysis, collaborative filtering, and phase retrieval, and contains NP-hard problems as special cases—has been lacking. Thus, we seek a more flexible notion for non-polyhedral cones that specializes to the usual definition of neighborliness for polyhedral cones.

In a different vein, the utility of neighborliness lies in the fact that it provides a succinct characterization of the geometry of the ‘most singular’ pieces of the boundary of a polyhedral cone. It is of intrinsic interest to understand such geometry more generally for other families of structured cones. Hyperbolicity cones serve as an instructive case study in this regard. These are convex cones derived from hyperbolic polynomials, with the nonnegative orthant and the positive semidefinite matrices being prominent examples. Relaxations based on derivatives of hyperbolicity cones offer the prospect of computationally less expensive approaches for obtaining bounds on conic optimization problems with respect to hyperbolicity cones, and an intriguing feature of these relaxations is that they tend to preserve the low-dimensional faces of the original hyperbolicity cone. Formalizing and quantifying this assertion by leveraging the perspective of neighborliness would provide new insights into the facial geometry of a large class of structured convex cones.

In this paper, we describe a generalization of neighborliness for non-polyhedral cones that addresses the preceding objectives.

1.2 Towards a definition for non-polyhedral cones

In aiming at an appropriate generalization of neighborliness for non-polyhedral cones that overcomes the limitation of polyhedrality of the low-dimensional faces, a natural approach is to reformulate neighborliness via other geometric attributes that are less face-centric. As a first attempt, for a convex cone \mathcal{C} that is closed and pointed but not necessarily polyhedral, let $\mathcal{S}_{\mathcal{C}}(x)$ denote the linear span of the smallest exposed face of \mathcal{C} that contains x . Then one can check that if the extreme rays of \mathcal{C} are exposed, k -neighborliness of \mathcal{C} is equivalent to the following condition for any collection $x^{(1)}, \dots, x^{(k)}$ of generators of the extreme rays of \mathcal{C} :

$$\mathcal{S}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) = \sum_{i=1}^k \mathcal{S}_{\mathcal{C}}\left(x^{(i)}\right). \quad (1)$$

One can check that the left-hand-side of this equation always contains the right-hand-side, with the containment being strict in general and equality holding only for k -neighborly cones. It is instructive to consider the three cones in \mathbb{R}^3 that are shown in Fig. 1 from the perspective of the relation (1). The cone \mathcal{C}_1 is isomorphic to the orthant in \mathbb{R}^3 , which is 3-neighborly, and therefore the relation (1) holds for any subset of the generators of the three extreme rays. The cone \mathcal{C}_2 is not 2-neighborly as the cone over the generators $x^{(1)}, x^{(2)}$ is not a face of \mathcal{C}_2 ; accordingly, we note that $\mathcal{S}_{\mathcal{C}_2}(x^{(1)}+x^{(2)}) \supsetneq \mathcal{S}_{\mathcal{C}_2}(x^{(1)})+\mathcal{S}_{\mathcal{C}_2}(x^{(2)})$. Finally, the ice-cream cone \mathcal{C}_3 is evidently not 2-neighborly by considering the cone over the generators $x^{(1)}, x^{(2)}$; as expected, we again have the strict containment $\mathcal{S}_{\mathcal{C}_3}(x^{(1)}+x^{(2)}) \supsetneq \mathcal{S}_{\mathcal{C}_3}(x^{(1)})+\mathcal{S}_{\mathcal{C}_3}(x^{(2)})$. The cone \mathcal{C}_3 presents an interesting case study as it is also linearly isomorphic to the cone




Fig. 1 Illustration of neighborliness properties of three cones. \mathcal{C}_1 is neighborly while \mathcal{C}_2 is not. \mathcal{C}_3 is not neighborly but it serves as an instructive example for the definition of Terracini convexity

of 2×2 symmetric positive semidefinite matrices. As mentioned previously, developing a suitable generalization of neighborliness that encompasses the cone of positive semidefinite matrices is one of the motivations for this article, and we investigate next what precisely fails with the relation (1) for \mathcal{C}_3 .

For a polyhedral cone \mathcal{C} that is pointed, the map $\mathcal{S}_{\mathcal{C}}(x)$ represents a kind of “local linearization” of \mathcal{C} around the point x ; concretely, the set $\mathcal{S}_{\mathcal{C}}(x)$ is the largest subspace—also called the *lineality space*—in the cone of feasible directions from x into \mathcal{C} . However, the interpretation of $\mathcal{S}_{\mathcal{C}}(x)$ as a local linearization of \mathcal{C} at x no longer holds in general if \mathcal{C} is not polyhedral. For the cone \mathcal{C}_3 in Fig. 1, the set $\mathcal{S}_{\mathcal{C}_3}(x^{(1)})$ does not fully represent a local linearization of \mathcal{C}_3 around $x^{(1)}$ as it fails to account for the curvature of the boundary of \mathcal{C}_3 at $x^{(1)}$. Rather, the subspace $\mathcal{L}_{\mathcal{C}_3}(x^{(1)})$ in Fig. 1, akin to a tangent space at $x^{(1)}$ with respect to the boundary of \mathcal{C}_3 , provides a more accurate local linearization of \mathcal{C}_3 at $x^{(1)}$. Letting $\mathcal{L}_{\mathcal{C}_3}(x^{(2)})$ similarly denote an accurate local linearization of \mathcal{C}_3 at $x^{(2)}$, we observe that $\mathcal{L}_{\mathcal{C}_3}(x^{(1)}) + \mathcal{L}_{\mathcal{C}_3}(x^{(2)}) = \mathbb{R}^3$. As $x^{(1)} + x^{(2)}$ lies in the interior of \mathcal{C}_3 , a natural local linearization of \mathcal{C}_3 at $x^{(1)} + x^{(2)}$ is the full space \mathbb{R}^3 , i.e., $\mathcal{L}_{\mathcal{C}_3}(x^{(1)} + x^{(2)}) = \mathbb{R}^3$. Consequently, we have that the relation (1) holds for \mathcal{C}_3 with $k = 2$ if we substitute $\mathcal{S}_{\mathcal{C}_3}$ with $\mathcal{L}_{\mathcal{C}_3}$. Motivated by this discussion, our generalization of neighborliness to closed, convex, pointed cones is based on a criterion analogous to (1) with a more accurate notion of local linearization; as we discuss in the sequel, this criterion is satisfied by neighborly polyhedral cones, cones of positive semidefinite matrices, as well as many other families.

1.3 Terracini convex cones

We begin by giving a formal definition of the map $\mathcal{L}_{\mathcal{C}}(x)$. In the example with the cone \mathcal{C}_3 from Fig. 1, the set $\mathcal{L}_{\mathcal{C}_3}(x)$ corresponds to a tangent space. However, convex cones in general have both smooth and singular features in their boundary, and therefore we do not explicitly appeal to any differential notions. Our definition is stated in terms of the feasible directions $\mathcal{H}_{\mathcal{C}}(x)$ into a convex cone $\mathcal{C} \subset \mathbb{R}^d$ that is closed and pointed

from any $x \in \mathcal{C}$:

$$\mathcal{H}_{\mathcal{C}}(x) = \text{cone}\{z - x : z \in \mathcal{C}\}.$$

The closure of the cone of feasible directions $\overline{\mathcal{H}_{\mathcal{C}}(x)}$ is called the *tangent cone* of \mathcal{C} at x .

Definition 1 Let $\mathcal{C} \subset \mathbb{R}^d$ be a convex cone that is closed and pointed. For any $x \in \mathcal{C}$, the *convex tangent space* of \mathcal{C} at x is denoted by $\mathcal{L}_{\mathcal{C}}(x)$ and is defined as the lineality space of the tangent cone of \mathcal{C} at x :

$$\mathcal{L}_{\mathcal{C}}(x) = \overline{\mathcal{H}_{\mathcal{C}}(x)} \cap -\overline{\mathcal{H}_{\mathcal{C}}(x)}.$$

In some sense, the subspace $\mathcal{L}_{\mathcal{C}}(x)$ represents all those directions from x in which the cone \mathcal{C} is locally “flat”. For smooth convex cones \mathcal{C} that are closed and pointed, the convex tangent space $\mathcal{L}_{\mathcal{C}}(x)$ at a point $x (\neq 0)$ on the boundary is indeed the tangent space with respect to the boundary of \mathcal{C} at x . For polyhedral cones \mathcal{C} that are pointed, one can check that $\mathcal{L}_{\mathcal{C}}(x) = \mathcal{S}_{\mathcal{C}}(x)$. With this definition, we are in a position to present the main object of investigation of this article.

Definition 2 A convex cone $\mathcal{C} \subset \mathbb{R}^d$ that is closed and pointed is *k-Terracini convex* if the following condition holds for any collection $x^{(1)}, \dots, x^{(k)}$ of generators of extreme rays of \mathcal{C} :

$$\mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) = \sum_{i=1}^k \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right). \quad (2)$$

If \mathcal{C} is *k-Terracini convex* for all k , then we say that \mathcal{C} is *Terracini convex*.

One inclusion always holds as $\mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) \supseteq \sum_{i=1}^k \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right)$, and the relevant portion of this definition is the other inclusion. The reason for the terminology ‘Terracini convexity’ is that the stipulation in this definition mirrors the consequence of Terracini’s lemma in algebraic geometry [19], with convex tangent space playing the role in our context that a tangent space does in Terracini’s lemma.² We give next some preliminary examples of *k-Terracini convex* cones:

Example 1 To begin with, it is instructive to compare *k-Terracini convexity* to *k-neighborliness* for polyhedral cones. For a polyhedral cone \mathcal{C} that is pointed, we observed previously that $\mathcal{L}_{\mathcal{C}}(x) = \mathcal{S}_{\mathcal{C}}(x)$ for $x \in \mathcal{C}$. As \mathcal{C} has exposed extreme rays and as the relation (1) is equivalent to *k-neighborliness*, we have that *k-Terracini convexity* and *k-neighborliness* are equivalent for pointed polyhedral cones. We also prove this fact as a special case of a more general result (see Theorem 1 and Corollary 2).

² Consider a projective variety \mathcal{V} over an algebraically closed field, and let $\mathcal{V}^{(k)}$ be the k ’th secant variety of \mathcal{V} obtained by taking the closure of the set of spans of every collection of k points in \mathcal{V} . Informally, Terracini’s lemma states that for k generic points $X^{(1)}, \dots, X^{(k)} \in \mathcal{V}$, the tangent space at a generic point in the span of $\{X^{(1)}, \dots, X^{(k)}\}$ with respect to $\mathcal{V}^{(k)}$ is equal to the sum of the tangent spaces at each $X^{(i)}$ with respect to \mathcal{V} .

Example 2 All convex cones that are closed and pointed are trivially 1-Terracini convex. As a contrast, based on the generalization of [14] of neighborliness to non-polyhedral cones, a convex cone that is closed and pointed is 1-neighborly if and only if all its extreme rays are exposed.

Example 3 Let $\mathcal{C} \subset \mathbb{R}^d$ be a smooth convex cone that is closed and pointed. Then \mathcal{C} is Terracini convex. To see this, consider any collection $x^{(1)}, \dots, x^{(k)}$ of generators of extreme rays of \mathcal{C} . Due to the smoothness of \mathcal{C} , we have that $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}(x^{(i)}) = \text{span}(\mathcal{C})$ for $k \geq 2$, unless all the $x^{(i)}$'s generate the same extreme ray (in which case the Terracini convexity condition is trivially satisfied).

Example 4 As our next example, we consider the cone of positive semidefinite matrices \mathbb{S}_+^d in the space of $d \times d$ real symmetric matrices \mathbb{S}^d . This cone consists of both smooth and singular features in its boundary. For $X \in \mathbb{S}_+^d$, one can check that $\mathcal{L}_{\mathbb{S}_+^d}(X) = \{MX + XM : M \in \mathbb{S}^n\}$, from which it follows that \mathbb{S}_+^d is Terracini convex. We give an alternative proof of this fact via a dual perspective on Terracini convexity; see Example 5 after Proposition 1.

It is instructive to consider the definition of Terracini convexity from a dual perspective, as this leads to a characterization that is more easily verified in some cases. In preparation to state this dual criterion, we recall that the polar of a cone $\mathcal{S} \subset \mathbb{R}^d$ is the collection of linear functionals that are nonpositive on \mathcal{S} and is denoted \mathcal{S}° . With this notation, the normal cone to a convex cone $\mathcal{C} \subset \mathbb{R}^d$ at $x \in \mathcal{C}$ is denoted $\mathcal{N}_{\mathcal{C}}(x)$ and is the polar $\mathcal{H}_{\mathcal{C}}(x)^\circ$ of the cone of feasible directions from x into \mathcal{C} . As \mathcal{C} is a cone, one can check that the normal cone to \mathcal{C} at $x \in \mathcal{C}$ is given by:

$$\mathcal{N}_{\mathcal{C}}(x) = \mathcal{H}_{\mathcal{C}}(x)^\circ = \{\ell \in \mathcal{C}^\circ : \ell(x) = 0\}, \tag{3}$$

which is the set of linear functionals that are nonpositive on \mathcal{C} and vanish at x . We now establish an equivalent dual formulation of Terracini convexity.

Proposition 1 *A closed, pointed, convex cone $\mathcal{C} \subset \mathbb{R}^d$ is k -Terracini convex if and only if for any collection $x^{(1)}, \dots, x^{(k)}$ of generators of extreme rays of \mathcal{C} ,*

$$\text{span} \left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{C}}(x^{(i)}) \right) = \bigcap_{i=1}^k \text{span} \left(\mathcal{N}_{\mathcal{C}}(x^{(i)}) \right). \tag{4}$$

Remark 1 In the result above, one inclusion is trivial—we always have that the span of the intersection of the normal cones is contained inside the intersection of the spans of the normal cones. Terracini convexity corresponds to the reverse inclusion being true, and this is all we need to verify. This remark is dual to the assertion after Definition 2 about one inclusion always being true.

Proof The normal cone and the closure of the cone of feasible directions at a point $x \in \mathcal{C}$ are related via $\mathcal{N}_{\mathcal{C}}(x) = \mathcal{H}_{\mathcal{C}}(x)^\circ = \overline{\mathcal{H}_{\mathcal{C}}(x)}^\circ$, which implies that

$\mathcal{L}_{\mathcal{C}}(x)^\perp = \text{span}(\mathcal{N}_{\mathcal{C}}(x))$. Taking orthogonal complements in the definition of k -Terracini convexity, we see that \mathcal{C} is k -Terracini convex if and only if for any collection $x^{(1)}, \dots, x^{(k)}$ of generators of extreme rays of \mathcal{C} ,

$$\text{span} \left(\mathcal{N}_{\mathcal{C}} \left(\sum_{i=1}^k x^{(i)} \right) \right) = \bigcap_{i=1}^k \text{span} \left(\mathcal{N}_{\mathcal{C}}(x^{(i)}) \right). \tag{5}$$

Here we have used that the orthogonal complement of a sum of subspaces is the intersection of the orthogonal complements. To complete the proof, we note that $\mathcal{N}_{\mathcal{C}} \left(\sum_{i=1}^k x^{(i)} \right) = \bigcap_{i=1}^k \mathcal{N}_{\mathcal{C}}(x^{(i)})$ whenever $x^{(1)}, \dots, x^{(k)} \in \mathcal{C}$. For one inclusion, if $\ell \in \mathcal{C}^\circ$ and $\ell(x^{(i)}) = 0$ then $\ell \left(\sum_{i=1}^k x^{(i)} \right) = 0$. For the other inclusion, if $\ell \in \mathcal{C}^\circ$ and $\ell \left(\sum_{i=1}^k x^{(i)} \right) = \sum_{i=1}^k \ell(x^{(i)}) = 0$, then we have that $\ell(x^{(i)}) \leq 0$ for each i (as $\ell \in \mathcal{C}^\circ$) and therefore $\ell(x^{(i)}) = 0$ for each i (as $\sum_{i=1}^k \ell(x^{(i)}) = 0$). \square

To illustrate the utility of this dual formulation, we show that the positive semidefinite cone is Terracini convex.

Example 5 (Positive semidefinite cone) Let $\mathcal{C} = \mathbb{S}_+^d$ be the cone of $d \times d$ positive semidefinite matrices. Given an extreme ray vv' for $v \in \mathbb{R}^d$, the corresponding normal cone from (3) is $\mathcal{N}_{\mathcal{C}}(vv') = \{Q \in -\mathbb{S}_+^d : v'Qv = 0\} = \{Q \in -\mathbb{S}_+^d : Qv = 0\}$. For any collection of generators of extreme rays $v^{(1)}v^{(1)'}, \dots, v^{(k)}v^{(k)'}$ of \mathcal{C} for $v^{(1)}, \dots, v^{(k)} \in \mathbb{R}^d$, we have that:

$$\begin{aligned} \text{span} \left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{C}} \left(v^{(i)}v^{(i)'} \right) \right) &= \{Q \in \mathbb{S}^d : Qv^{(i)} = 0, i = 1, \dots, k\} \\ &= \bigcap_{i=1}^k \{Q \in \mathbb{S}^d : Qv^{(i)} = 0\}. \end{aligned}$$

As $\text{span} \left(\mathcal{N}_{\mathcal{C}} \left(v^{(i)}v^{(i)'} \right) \right) = \{Q \in \mathbb{S}^d : Qv^{(i)} = 0\}$ and as k was arbitrary, it follows that \mathbb{S}_+^d is Terracini convex.

1.4 Outline of contributions

We initiate our study of Terracini convex cones by investigating the face structure of such cones. Specifically, in Sect. 2 we provide two conditions for a closed, pointed, convex cone to be Terracini convex based on order-theoretic properties of the faces of the cone. The first condition states that if a cone is k -Terracini convex for a sufficiently large k , which is a function of the height of the partially ordered set of faces, then the cone is Terracini convex. The second condition gives a necessary and sufficient characterization for a cone to be Terracini convex based on the collection of all convex tangent spaces of the cone inheriting some of the lattice structure of the subspace lattice.

From the examples in the previous subsection we see that Terracini convexity is equivalent to neighborliness for polyhedral cones, but there are many families of non-polyhedral cones that are also Terracini convex. Thus, a natural question is to clarify the distinction between Terracini convexity and neighborliness for non-polyhedral cones. In one direction, the cone of positive semidefinite matrices serves as an example that there are Terracini convex cones that are not neighborly. In the other direction, we prove in Sect. 3 that subject to a non-degeneracy condition that is of the form of a quadratic growth property, k -neighborly cones are k -Terracini convex. As a consequence of this result, we obtain that the cone over the (homogeneous) moment curve, which was studied by Kalai and Wigderson in [14], is Terracini convex; see Sect. 3.3 for more examples.

Next we demonstrate the utility of the notion of Terracini convexity in characterizing tightness of semidefinite relaxations for the problem of finding a positive semidefinite matrix of smallest rank in an affine space. A commonly employed heuristic to solve this problem is to compute the positive semidefinite matrix of smallest trace in the given affine space, which can be obtained via a tractable semidefinite program. In Sect. 4, we show that the success of this heuristic is closely tied to a certain cone being Terracini convex. Our result may be viewed as a generalization of Donoho and Tanner's result on using neighborliness to characterize the exactness of linear programming relaxations for identifying nonnegative vectors with the smallest number of nonzeros in affine spaces [9]. As a by-product of our result, we obtain that 'most' linear images of a cone of positive semidefinite matrices are k -Terracini convex, where the value of k depends on the dimension of the image of the linear map; see Theorem 4.

In Sect. 5, we investigate the Terracini convexity properties of derivative relaxations of hyperbolicity cones. We study conditions under which derivatives of Terracini convex hyperbolicity cones continue to be k -Terracini convex (for suitable k), and in particular the relationship between the number of derivatives and k . As a consequence, we obtain new examples of Terracini convex cones, and in particular ones that are basic semialgebraic; it is instructive to contrast these examples with the ones described in Sect. 4.3 of linear images of cones of positive semidefinite matrices, which are semialgebraic but not necessarily basic semialgebraic.

Sections 3, 4, and 5 illustrate the role that Terracini convexity plays in illuminating various aspects of the facial structure of convex cones. In each case, we obtain new examples of Terracini convex cones in the course of our discussion. We conclude in Sect. 6 with some open questions.

2 Order-theoretic conditions for Terracini convexity

In this section we discuss conditions under which a closed, pointed, convex cone is Terracini convex based on the order structure underlying the faces of a convex cone. Section 2.1 shows that a cone that is k -Terracini convex for sufficiently large k is Terracini convex, with the threshold value of k depending on the length of the longest chain of faces of the cone. In Sect. 2.2 we give a lattice-theoretic condition on the collection of lineality spaces that is necessary and sufficient for a cone to be Terracini convex.

In preparation for our discussion, we recall briefly a few relevant facts about the face structure of a convex cone. Let \mathcal{C} be a closed, pointed, convex cone. A subset $\mathcal{F} \subseteq \mathcal{C}$ is a *face* if $x, y \in \mathcal{C}$ and $x + y \in \mathcal{F}$ implies that $x, y \in \mathcal{F}$. A face $\mathcal{F} \subseteq \mathcal{C}$ is *exposed* if \mathcal{F} can be expressed as the intersection of \mathcal{C} and a hyperplane specified by a linear functional $\ell \in \mathcal{C}^\circ$, i.e., $\mathcal{F} = \{x \in \mathcal{C} : \ell(x) = 0\}$. By convention \mathcal{C} is itself an exposed face as one can take $\ell = 0$. The collection of (exposed) faces of \mathcal{C} form a partially ordered set (poset) by inclusion. For any subset $\mathcal{X} \subseteq \mathcal{C}$, let $\mathcal{F}_{\mathcal{C}}(\mathcal{X})$ (respectively, $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(\mathcal{X})$) denote the inclusion-wise minimal (exposed) face of \mathcal{C} containing \mathcal{X} . For any element $x \in \mathcal{C}$, one can check that the normal cone $\mathcal{N}_{\mathcal{C}}(x)$ depends only on $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$, which in turn depends only on $\mathcal{F}_{\mathcal{C}}(x)$; consequently, the convex tangent space $\mathcal{L}_{\mathcal{C}}(x)$ depends only on $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$ and in turn $\mathcal{F}_{\mathcal{C}}(x)$ [18]. Formally, for any $x^{(1)}, x^{(2)} \in \mathcal{C}$:

$$\begin{aligned} \mathcal{F}_{\mathcal{C}}(x^{(1)}) = \mathcal{F}_{\mathcal{C}}(x^{(2)}) &\Leftrightarrow \mathcal{F}_{\mathcal{C}}^{\text{exp}}(x^{(1)}) = \mathcal{F}_{\mathcal{C}}^{\text{exp}}(x^{(2)}) \\ &\Leftrightarrow \mathcal{N}_{\mathcal{C}}(x^{(1)}) = \mathcal{N}_{\mathcal{C}}(x^{(2)}) \Leftrightarrow \mathcal{L}_{\mathcal{C}}(x^{(1)}) = \mathcal{L}_{\mathcal{C}}(x^{(2)}). \end{aligned} \quad (6)$$

2.1 Terracini convexity and the height of the poset of faces

Given a closed, pointed, convex cone \mathcal{C} , consider a collection of points $x^{(1)}, \dots, x^{(k)} \in \mathcal{C}$. For large k , it is possible to replace the convex tangent space $\mathcal{L}_{\mathcal{C}}(\sum_{i=1}^k x^{(i)})$ by $\mathcal{L}_{\mathcal{C}}(\sum_{i \in I} x^{(i)})$ for a subset $I \subseteq \{1, \dots, k\}$ that is potentially much smaller than k , by appealing to the observation that the convex tangent space at a point depends only on the smallest face containing the point. This allows us to conclude that if \mathcal{C} is k -Terracini convex for sufficiently large k , then \mathcal{C} is Terracini convex.

We describe next the relevant terminology that we use in our result. A collection of faces $\mathcal{F}^{(i)}$, $i = 1, \dots, m$ of \mathcal{C} that satisfies $\mathcal{F}^{(1)} \subsetneq \dots \subsetneq \mathcal{F}^{(m)}$ is called a *chain* of faces. For a closed, pointed, convex cone \mathcal{C} , let $\mathcal{H}(\mathcal{C})$ denote the *height* of the poset of faces of \mathcal{C} , which is the length of the longest chain of faces of \mathcal{C} . As the dimension always increases strictly along chains of faces and as any maximal-length chain of faces begins with the zero-dimensional face³ $\{0\}$ and ends with \mathcal{C} , we have that $\mathcal{H}(\mathcal{C}) \leq \dim(\mathcal{C}) + 1$. We have next a result that allows us to replace the convex tangent space of a large sum of elements of \mathcal{C} by that of a smaller subset based on $\mathcal{H}(\mathcal{C})$:

Lemma 1 *Let \mathcal{C} be a closed, pointed, convex cone, and consider a collection of points $x^{(1)}, \dots, x^{(k)} \in \mathcal{C}$. There exists $I \subseteq \{1, \dots, k\}$ with $|I| \leq \mathcal{H}(\mathcal{C}) - 1$ such that $\mathcal{F}_{\mathcal{C}}(\sum_{i=1}^k x^{(i)}) = \mathcal{F}_{\mathcal{C}}(\sum_{i \in I} x^{(i)})$.*

Proof We explicitly construct a set I with $|I| \leq \mathcal{H}(\mathcal{C}) - 1$. Set $j = 0$, $I_0 = \emptyset$, $\mathcal{F}_{\mathcal{C}}^{(0)} = \{0\}$. Running sequentially through $i = 1, \dots, k$, if $x^{(i)} \notin \mathcal{F}_{\mathcal{C}}(I_j)$, then (a) increase j by one, (b) set $I_j = I_{j-1} \cup \{i\}$, and (c) set $\mathcal{F}_{\mathcal{C}}^{(j)} = \mathcal{F}_{\mathcal{C}}(\sum_{m \in I_j} x^{(m)})$.

The sequence of faces $\mathcal{F}_{\mathcal{C}}^{(0)}, \dots, \mathcal{F}_{\mathcal{C}}^{(j)}$ has the property that $\mathcal{F}_{\mathcal{C}}^{(0)} \subsetneq \dots \subsetneq \mathcal{F}_{\mathcal{C}}^{(j)} = \mathcal{F}_{\mathcal{C}}(\{x^{(1)}, \dots, x^{(k)}\})$, and therefore forms a chain of faces of \mathcal{C} of length at most

³ We do not consider the empty set to be a face of \mathcal{C} .

$\mathcal{H}(\mathcal{C})$. As $\mathcal{F}_{\mathcal{C}}(\{x^{(1)}, \dots, x^{(k)}\}) = \mathcal{F}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right)$ and as the index set I_j satisfies $|I_j| \leq \mathcal{H}(\mathcal{C}) - 1$, setting $I = I_j$ leads to the desired conclusion. \square

We are now in a position to state and prove the main result of this section.

Proposition 2 *Let \mathcal{C} be a closed, pointed, convex cone that is $(\mathcal{H}(\mathcal{C}) - 1)$ -Terracini convex. Then \mathcal{C} is Terracini convex.*

Proof Let $x^{(1)}, \dots, x^{(k)}$ be a collection of generators of extreme rays of \mathcal{C} . By Lemma 1, we know that there exists $I \subseteq \{1, \dots, k\}$ with $|I| \leq \mathcal{H}(\mathcal{C}) - 1$ such that $\mathcal{F}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) = \mathcal{F}_{\mathcal{C}}\left(\sum_{i \in I} x^{(i)}\right)$. From (6) we have that:

$$\mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) = \mathcal{L}_{\mathcal{C}}\left(\sum_{i \in I} x^{(i)}\right). \tag{7}$$

Since \mathcal{C} is $(\mathcal{H}(\mathcal{C}) - 1)$ -Terracini convex, it is $|I|$ -Terracini convex and therefore

$$\mathcal{L}_{\mathcal{C}}\left(\sum_{i \in I} x^{(i)}\right) = \sum_{i \in I} \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right). \tag{8}$$

Combining (7) and (8), and noting that $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right) \subseteq \mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right)$ as well as $\sum_{i \in I} \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right) \subseteq \sum_{i=1}^k \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right)$, we conclude that \mathcal{C} is k -Terracini convex. Since k was arbitrary, we have shown that \mathcal{C} is Terracini convex. \square

As a consequence of this result, we have the following corollary:

Corollary 1 *Let \mathcal{C} be a closed, pointed, convex cone that is $\dim(\mathcal{C})$ -Terracini convex. Then \mathcal{C} is Terracini convex.*

Proof This follows from the observation that $\mathcal{H}(\mathcal{C}) \leq \dim(\mathcal{C}) + 1$. \square

2.2 Terracini convexity and the lattice of subspaces

Motivated by the order-theoretic structure underlying the faces of a closed, pointed, convex cone $\mathcal{C} \subset \mathbb{R}^d$, we consider the order-theoretic aspects of the collection of convex tangent spaces associated to \mathcal{C} :

$$\mathfrak{L}(\mathcal{C}) = \{\mathcal{L}_{\mathcal{C}}(x) : x \in \mathcal{C}\}$$

As $\mathfrak{L}(\mathcal{C})$ is a subset of the collection of subspaces in \mathbb{R}^d , one may view $\mathfrak{L}(\mathcal{C})$ as a poset by inclusion. However, the collection of all subspaces in \mathbb{R}^d additionally forms a lattice (called the subspace lattice in \mathbb{R}^d) with the join of two subspaces given by their sum and the meet given by their intersection. In this section we relate Terracini convexity of \mathcal{C} to $\mathfrak{L}(\mathcal{C})$ inheriting some of the lattice structure of the collection of all subspaces in \mathbb{R}^d .

In preparation to present this result, we discuss next a link between the elements of $\mathfrak{L}(\mathcal{C})$ and the exposed faces of \mathcal{C} . As noted previously in (6), the convex tangent space at a point $x \in \mathcal{C}$ depends only on the smallest exposed face of \mathcal{C} containing x so that the elements of $\mathfrak{L}(\mathcal{C})$ are in one-to-one correspondence with the exposed faces of \mathcal{C} . The next result describes how one obtains an exposed face of \mathcal{C} given an element of $\mathfrak{L}(\mathcal{C})$:

Lemma 2 *Let \mathcal{C} be a closed, pointed, convex cone. For any $x \in \mathcal{C}$ we have that:*

$$\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x) = \mathcal{C} \cap \mathcal{L}_{\mathcal{C}}(x).$$

Proof One can check that $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x) \subseteq \mathcal{L}_{\mathcal{C}}(x)$, and therefore $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x) \subseteq \mathcal{C} \cap \mathcal{L}_{\mathcal{C}}(x)$. In the other direction, we begin by observing that any hyperplane supporting \mathcal{C} that contains $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$ must contain $\mathcal{L}_{\mathcal{C}}(x)$. Consider a hyperplane H supporting \mathcal{C} that exposes $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$, i.e., $\mathcal{C} \cap H = \mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$ (such a hyperplane must exist as $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$ is an exposed face). As $\mathcal{L}_{\mathcal{C}}(x) \subseteq H$, we have that $\mathcal{C} \cap \mathcal{L}_{\mathcal{C}}(x) \subseteq \mathcal{F}_{\mathcal{C}}^{\text{exp}}(x)$. This concludes the proof. \square

With this result in hand, we are now in a position to state and prove the following proposition:

Proposition 3 *Let $\mathcal{C} \subset \mathbb{R}^d$ be a closed, pointed, convex cone. The cone \mathcal{C} is Terracini convex if and only if $\mathfrak{L}(\mathcal{C})$ is a join sub-semilattice of the lattice of all subspaces in \mathbb{R}^d (i.e., the poset $\mathfrak{L}(\mathcal{C})$ has a join given by the sum of two subspaces).*

Proof Suppose first that \mathcal{C} is Terracini convex. Consider any pair $\mathcal{L}_{\mathcal{C}}(x), \mathcal{L}_{\mathcal{C}}(y) \in \mathfrak{L}(\mathcal{C})$ corresponding to $x, y \in \mathcal{C}$, and let $x = \sum_i x^{(i)}$ and $y = \sum_j y^{(j)}$ be decompositions in terms of generators of extreme rays of \mathcal{C} . As \mathcal{C} is Terracini convex, we have that:

$$\begin{aligned} \mathcal{L}_{\mathcal{C}}(x) + \mathcal{L}_{\mathcal{C}}(y) &= \sum_i \mathcal{L}_{\mathcal{C}}(x^{(i)}) + \sum_j \mathcal{L}_{\mathcal{C}}(y^{(j)}) \\ &= \mathcal{L}_{\mathcal{C}}\left(\sum_i x^{(i)} + \sum_j y^{(j)}\right) = \mathcal{L}_{\mathcal{C}}(x + y). \end{aligned}$$

Since $\mathcal{L}_{\mathcal{C}}(x + y) \in \mathfrak{L}(\mathcal{C})$, the poset $\mathfrak{L}(\mathcal{C})$ is a join sub-semilattice of the lattice of all subspaces in \mathbb{R}^d .

In the other direction, suppose that the poset $\mathfrak{L}(\mathcal{C})$ is a join sub-semilattice of the lattice of all subspaces in \mathbb{R}^d . Consider any collection $x^{(1)}, \dots, x^{(k)} \in \mathcal{C}$ of generators of extreme rays of \mathcal{C} . As the join is given by subspace sum, we have that $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}(x^{(i)}) \in \mathfrak{L}(\mathcal{C})$, which implies that $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}(x^{(i)})$ is the convex tangent space at some point $y \in \mathcal{C}$. Then, from Lemma 2 we see that $\mathcal{C} \cap \sum_{i=1}^k \mathcal{L}_{\mathcal{C}}(x^{(i)}) = \mathcal{F}_{\mathcal{C}}^{\text{exp}}(y)$, and in particular, $\sum_{i=1}^k x^{(i)} \in \mathcal{F}_{\mathcal{C}}^{\text{exp}}(y)$ as each $x^{(i)} \in \mathcal{L}_{\mathcal{C}}(x^{(i)})$. We also have that $\mathcal{C} \cap \mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right) = \mathcal{F}_{\mathcal{C}}^{\text{exp}}\left(\sum_{i=1}^k x^{(i)}\right)$. As $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}(x^{(i)}) \subseteq$

$\mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right)$, we conclude that $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(y) \subseteq \mathcal{F}_{\mathcal{C}}^{\text{exp}}\left(\sum_{i=1}^k x^{(i)}\right)$, which in turn implies that $\mathcal{F}_{\mathcal{C}}^{\text{exp}}(y) = \mathcal{F}_{\mathcal{C}}^{\text{exp}}\left(\sum_{i=1}^k x^{(i)}\right)$ because $\sum_{i=1}^k x^{(i)} \in \mathcal{F}_{\mathcal{C}}^{\text{exp}}(y)$. Appealing to (6), we can then conclude that $\sum_{i=1}^k \mathcal{L}_{\mathcal{C}}\left(x^{(i)}\right) = \mathcal{L}_{\mathcal{C}}\left(\sum_{i=1}^k x^{(i)}\right)$. \square

Therefore, Terracini convexity of a cone \mathcal{C} is linked to the poset $\mathfrak{L}(\mathcal{C})$ inheriting the join structure of the lattice of subspaces. In general, $\mathfrak{L}(\mathcal{C})$ does not inherit the meet structure of the lattice of subspaces as the intersection of the convex tangent spaces corresponding to two exposed faces does not usually yield a convex tangent space corresponding to an exposed face of \mathcal{C} (the positive semidefinite cone provides a counterexample); indeed, the preceding proposition makes no assumptions on the existence of a meet operation.

3 Neighborliness and Terracini convexity

Terracini convexity is one approach to extend neighborliness from polyhedral cones to non-polyhedral convex cones. As discussed in the introduction, there is already a previous notion of neighborliness available in the non-polyhedral case due to Kalai and Wigderson [14]. In this section we investigate the relationship between these two concepts, and in particular we show that k -neighborly convex cones (formally defined in Sect. 3.1) are k -Terracini convex subject to mild non-degeneracy conditions. Throughout this section we view \mathbb{R}^m as being equipped with an inner product (which varies based on context and is specified clearly in each case), and we define an associated set $\mathcal{S}^{m-1} \subset \mathbb{R}^m$ of unit-norm elements induced by the inner product. Doing so allows us to work with a distinguished set $\text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}$ of normalized extreme rays of a closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$.

3.1 k -Neighborly convex cones

In [14] Kalai and Wigderson extend the notion of a neighborly polytope to define a k -neighborly embedded smooth manifold. This concept serves as the point of departure for a definition of a k -neighborly convex cone that is expressed in convex-geometric terms with no reference to an underlying embedded manifold.

Definition 3 Let \mathcal{M} be a smooth manifold and let $\phi : \mathcal{M} \rightarrow \mathbb{R}^m$ be an embedding of \mathcal{M} in \mathbb{R}^m . The image $\phi(\mathcal{M})$ is a k -neighborly embedded manifold if for any collection $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ of elements of $\phi(\mathcal{M})$, there exists an affine function $\ell : \mathbb{R}^m \rightarrow \mathbb{R}$ such that $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$ and $\ell(x) > 0$ for all $x \in \phi(\mathcal{M}) \setminus \{x^{(1)}, x^{(2)}, \dots, x^{(k)}\}$.

This definition is a slight reformulation of that of Kalai and Wigderson and it is stated in a manner that is more convenient for our presentation. The neighborliness of $\phi(\mathcal{M})$ clearly only depends on the convex hull of $\phi(\mathcal{M})$, which suggests the following notion of a k -neighborly convex cone.

Definition 4 A closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$ is k -neighborly if for every collection $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ of normalized extreme rays of \mathcal{K} , there exists a linear functional $\ell : \mathbb{R}^m \rightarrow \mathbb{R}$ such that $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$ and $\ell(x) > 0$ for all $x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1} \setminus \{x^{(1)}, x^{(2)}, \dots, x^{(k)}\}$.

It is straightforward to check that if an embedded smooth manifold $\phi(\mathcal{M}) \subseteq \mathbb{R}^m$ is k -neighborly, then the cone over $\phi(\mathcal{M})$, i.e., $\text{cone}(\{1\} \times \phi(\mathcal{M})) \subseteq \mathbb{R}^{m+1}$, is a k -neighborly convex cone. A basic observation about k -neighborly convex cones is that all of their sufficiently low-dimensional faces are linearly isomorphic to a nonnegative orthant.

Proposition 4 Consider a closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$ that is k -neighborly, and suppose \mathcal{F} is a face of \mathcal{K} of dimension $d \leq k$. Then \mathcal{F} is linearly isomorphic to \mathbb{R}_+^d .

Proof As \mathcal{K} is a closed, pointed, convex cone, so is \mathcal{F} . Hence, \mathcal{F} is the conic hull of its extreme rays. Let $x^{(1)}, \dots, x^{(d)}$ be a choice of d linearly independent normalized extreme rays of \mathcal{F} (and hence of \mathcal{K}). Let ℓ be a linear functional satisfying $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, d$ and $\ell(x) > 0$ for all other normalized extreme rays of \mathcal{K} , whose existence is guaranteed due to the k -neighborliness of \mathcal{K} . Let $\tilde{\mathcal{F}} = \{x \in \mathcal{K} : \ell(x) = 0\}$ be the face of \mathcal{K} exposed by ℓ . Since every extreme ray of \mathcal{K} that belongs to $\tilde{\mathcal{F}}$ is also an extreme ray of \mathcal{F} , it follows from the definition of ℓ that $x^{(1)}, x^{(2)}, \dots, x^{(d)}$ are exactly the normalized extreme rays of $\tilde{\mathcal{F}}$. As such, $\tilde{\mathcal{F}}$ is a closed, pointed, convex cone with exactly d linearly independent extreme rays, and therefore it must be linearly isomorphic to \mathbb{R}_+^d . Finally, $\tilde{\mathcal{F}}$ and \mathcal{F} are both faces of \mathcal{K} such that their relative interiors have a point in common, so $\tilde{\mathcal{F}} = \mathcal{F}$ [18, Corollary 18.1.2]. \square

Proposition 4 makes it clear that k -Terracini convex cones are not necessarily k -neighborly. Indeed, we have seen that the positive semidefinite cone is Terracini convex, and yet its faces are not linearly isomorphic to nonnegative orthants in general. We describe next an example that serves as a running illustration throughout this section. This cone was considered by Kalai and Wigderson [14] in the language of neighborly manifolds.

Cone over the Veronese embedding The Veronese embedding $\phi_{n,2d} : \mathbb{R}^n \rightarrow \mathbb{R}^{\binom{n+2d-1}{2d}}$ is defined by the homogeneous moment map $\phi_{n,2d}(z) = (z^\alpha)_{\alpha \in \mathcal{A}_{n,2d}}$ where $\mathcal{A}_{n,2d} = \{\alpha \in \mathbb{N}^n : \sum_{i=1}^n \alpha_i = 2d\}$ and $z^\alpha := \prod_{i=1}^n z_i^{\alpha_i}$. We denote the cone over this embedding by

$$\mathcal{C}_{n,2d} := \text{cone}\{\phi_{n,2d}(z) : z \in \mathbb{R}^n\}.$$

When discussing this example, we let $m = \binom{n+2d-1}{2d}$ and equip \mathbb{R}^m with the inner product⁴ that satisfies

$$\langle \phi(y), \phi(z) \rangle_B := \langle y, z \rangle^{2d} \quad \text{for all } y, z \in \mathbb{R}^n$$

⁴ This inner product is variously referred to as the apolar, Bombieri, Weyl-Bombieri, Fisher, or Calderón inner product.

where the inner product on the right is the Euclidean inner product on \mathbb{R}^n . The norms associated with these inner products are denoted $\|\cdot\|_B$ and $\|\cdot\|$, respectively. Any linear functional $\ell : \mathbb{R}^m \rightarrow \mathbb{R}$ restricted to the extreme rays of the cone $\mathcal{C}_{n,2d}$ can be interpreted as a homogeneous polynomial of degree $2d$ in n variables, i.e.,

$$\ell(\phi_{n,2d}(z)) = \sum_{\alpha \in \mathcal{A}_{n,2d}} \ell_\alpha z^\alpha.$$

Under this interpretation, the dual cone $-\mathcal{C}_{n,2d}^\circ$ is the cone of (coefficients of) non-negative homogeneous polynomials of degree $2d$ in n variables.

Example 6 (Neighborliness of cones over Veronese embeddings [14]) The cone $\mathcal{C}_{n,2d}$ is a d -neighborly convex cone. To see this, consider a collection of up to d normalized extreme rays

$$\{\phi_{n,2d}(z^{(1)}), \dots, \phi_{n,2d}(z^{(d)})\} \subseteq \text{ext}(\mathcal{C}_{n,2d}) \cap \mathcal{S}^{m-1}$$

and define the linear functional

$$\ell(\phi_{n,2d}(z)) = \prod_{i=1}^d (\|z\|^2 \|z^{(i)}\|^2 - \langle z, z^{(i)} \rangle^2).$$

From the Cauchy-Schwarz inequality, we can see that this is a nonnegative polynomial in z . (In fact, it is a sum of squares.) As such, ℓ defines a linear functional that is nonnegative on the extreme rays of $\mathcal{C}_{n,2d}$, and hence on $\mathcal{C}_{n,2d}$ itself. Furthermore, the only normalized extreme rays at which ℓ vanishes are $\phi_{n,2d}(z^{(i)})$ for $i = 1, 2, \dots, d$.

3.2 Non-degeneracy and regularity of convex cones

Our approach to showing that a k -neighborly cone is k -Terracini convex is based on the dual characterization of k -Terracini convexity from Proposition 1. Specifically, for any collection of normalized extreme rays $x^{(1)}, \dots, x^{(k)}$ of a k -neighborly cone $\mathcal{K} \subseteq \mathbb{R}^m$, we wish to prove that $\bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)})) \subseteq \text{span}\left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})\right)$. Our strategy is to identify an $\ell \in -\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})$ such that

$$\ell + U \cap \left[\bigcap_{i=1}^k \text{span}\left(\mathcal{N}_{\mathcal{K}}(x^{(i)})\right) \right] \subseteq -\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)}) \tag{9}$$

for an open set $U \subseteq \mathbb{R}^m$ containing the origin. The linear functional that supports \mathcal{K} at the points $x^{(1)}, \dots, x^{(k)}$, which is available to us from the definition of k -neighborliness, serves as a natural candidate for ℓ . The key issue with executing this strategy is that we need to control the extent to which any $\Delta \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)}))$ perturbs ℓ . In particular, as $\Delta \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)}))$ may be decomposed as $\Delta =$

$\Delta_+^{(i)} - \Delta_-^{(i)}$ for each $i = 1, \dots, k$, (with $\Delta_+^{(i)}, \Delta_-^{(i)} \in -\mathcal{N}_{\mathcal{H}}(x^{(i)})$), we need to bound the amount that the ‘negative’ parts $\Delta_-^{(i)}$ perturb ℓ . We consider two conditions to address this point. The first one ensures that $\ell(x)$ grows sufficiently fast around $\{x^{(1)}, x^{(2)}, \dots, x^{(k)}\}$. The second one controls the growth of any linear functional in $-\mathcal{N}_{\mathcal{H}}(x)$ for any normalized extreme ray $x \in \mathcal{H}$. Under these conditions—with the second one applied to each $\Delta_-^{(i)}$ —we show that ℓ dominates $\Delta_-^{(i)}$; consequently, we prove that for each $\Delta \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{H}}(x^{(i)}))$ there exists $\gamma \neq 0$ such that $\ell + \gamma \Delta \in -\bigcap_{i=1}^k \mathcal{N}_{\mathcal{H}}(x^{(i)})$. The first condition is a requirement on k -neighborly cones and takes the form of a quadratic growth criterion, while the second one is a regularity property applicable to arbitrary closed, pointed, convex cones. Both of these conditions are mild; for example, we show that the cone over the Veronese embedding satisfies them. (That being said, we are unaware of a method to prove that a k -neighborly cone is k -Terracini convex without these two conditions.) We precisely describe the conditions next, and we prove in Sect. 3.3 that k -neighborly cones satisfying these conditions are k -Terracini convex.

3.2.1 Non-degenerate neighborliness

We present a non-degenerate extension of the notion k -neighborliness in which the linear functional exposing a subset of k extreme rays satisfies an additional growth condition when restricted to nearby extreme rays.

Definition 5 A closed, pointed, convex cone $\mathcal{H} \subseteq \mathbb{R}^m$ is *non-degenerate k -neighborly* if for every collection $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ of normalized extreme rays of \mathcal{H} , there exist $\epsilon > 0$, $\mu > 0$, and a linear functional $\ell : \mathbb{R}^m \rightarrow \mathbb{R}$, such that $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$, $\ell(x) > 0$ for all $x \in (\text{ext}(\mathcal{H}) \cap \mathcal{S}^{m-1}) \setminus \{x^{(1)}, \dots, x^{(k)}\}$, and

$$\ell(x) \geq \mu \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \quad \text{for all } x \in (\text{ext}(\mathcal{H}) \cap \mathcal{S}^{m-1}) \cap (\cup_{i=1}^k \mathcal{B}(x^{(i)}, \epsilon)). \quad (10)$$

The quadratic growth condition (10) is a mild restriction, and it is satisfied by the examples of k -neighborly convex cones we consider in this section.

Example 7 (k -neighborly polyhedral cones are non-degenerate k -neighborly) If $\mathcal{H} \subseteq \mathbb{R}^m$ is a k -neighborly polyhedral cone, then for any collection $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ of normalized extreme rays there is a linear functional ℓ such that $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$ and $\ell(x) > 0$ for all other normalized extreme rays of \mathcal{H} . As the set of normalized extreme rays is finite, one can choose ϵ smaller than half the minimum distance between normalized extreme rays and obtain that

$$(\text{ext}(\mathcal{H}) \cap \mathcal{S}^{m-1}) \cap (\cup_{i=1}^k \mathcal{B}(x^{(i)}, \epsilon)) = \{x^{(1)}, x^{(2)}, \dots, x^{(k)}\},$$

which implies that (10) is vacuously satisfied for any positive μ .

Example 8 (Cone $\mathcal{C}_{n,2d}$ over the Veronese embedding is non-degenerate d -neighborly)
 For $y, z \in \mathbb{R}^n$ with unit Euclidean norm so that $\|\phi(y)\|_B = \|\phi(z)\|_B = 1$ (this is the norm associated with the Bombieri inner product on \mathbb{R}^m), we have that

$$\begin{aligned} \frac{1}{2} \|\phi_{n,2d}(y) - \phi_{n,2d}(z)\|_B^2 &= 1 - \langle y, z \rangle^{2d} \\ &= (1 - \langle y, z \rangle^2)(1 + \langle y, z \rangle^2 + \dots + \langle y, z \rangle^{2d-2}) \\ &\leq d(1 - \langle y, z \rangle^2). \end{aligned}$$

Here, the inequality follows from the Cauchy-Schwarz inequality and the fact that y and z have unit Euclidean norm. For unit Euclidean norm $z^{(i)}$, $i = 1, 2, \dots, d$ and unit Euclidean norm $z \in \mathbb{R}^n$, the linear functional ℓ from Example 6 satisfies

$$\ell(\phi_{n,2d}(z)) = \prod_{i=1}^d (1 - \langle z, z^{(i)} \rangle^2) \geq \prod_{i=1}^d \frac{\|\phi_{n,2d}(z) - \phi_{n,2d}(z^{(i)})\|_B^2}{2d}.$$

Choosing $\epsilon = \frac{1}{2} \min_{i \neq j} \|\phi(z^{(i)}) - \phi(z^{(j)})\|_B > 0$, whenever $\phi_{n,2d}(z) \in \bigcup_{i=1}^k \mathcal{B}(\phi_{n,2d}(z^{(i)}), \epsilon)$ and $\|z\|^2 = 1$ we have that

$$\ell(\phi_{n,2d}(z)) \geq \frac{1}{2d} \left(\frac{\epsilon}{2d}\right)^{d-1} \min_i \|\phi_{n,2d}(z) - \phi_{n,2d}(z^{(i)})\|_B^2.$$

It follows that $\mathcal{C}_{n,2d}$ is non-degenerate d -neighborly.

Although the definition of being non-degenerate k -neighborly only requires quadratic growth locally around the set of minimizers, compactness of the sphere means that local quadratic growth implies global quadratic growth.

Lemma 3 *If a closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$ is non-degenerate k -neighborly then for every collection $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ of normalized extreme rays of \mathcal{K} , there exists $\mu_0 > 0$, and a linear functional ℓ , such that $\ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$ and*

$$\ell(x) \geq \mu_0 \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \text{ for all } x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}.$$

Proof Let $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ be a collection of normalized extreme rays of \mathcal{K} . Let ϵ and μ be the positive constants, and let ℓ be the linear functional, that exist because \mathcal{K} is non-degenerate k -neighborly. Let

$$\mathcal{W} = \{x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1} : \min_{i=1,2,\dots,k} \|x - x^{(i)}\| < \epsilon\}$$

and let $\mathcal{W}^c = \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1} \setminus \mathcal{W}$ be its complement in normalized extreme rays. By compactness of \mathcal{W}^c and the fact that $\ell(x) > 0$ on \mathcal{W}^c , there exists some $M > 0$ such that

$$\ell(x) \geq M \geq \frac{M}{4} \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \text{ for all } x \in \mathcal{W}^c$$

where the second inequality holds because $\|x - y\|^2 \leq 4$ whenever $x, y \in \mathcal{S}^{m-1}$. Since

$$\ell(x) \geq \mu \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \quad \text{for all } x \in \mathcal{W},$$

taking $\mu_0 = \min\{\mu, M/4\}$ completes the proof. \square

3.2.2 Regular cones

Our notion of regularity for a closed, pointed, convex cone requires that no linear functional in the dual cone grows too fast around its minimizer when restricted to extreme rays. This holds whenever the restriction of a linear functional to the extreme rays is smooth.

Definition 6 A closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$ is *regular* if for each $x_0 \in \text{ext}(\mathcal{K})$ and each $\ell \in -\mathcal{N}_{\mathcal{K}}(x_0)$, there exist $\delta > 0$ and $\nu > 0$ such that

$$\ell(x) \leq \nu \|x - x_0\|^2 \quad \text{for all } x \in (\text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}) \cap \mathcal{B}(x_0, \delta). \quad (11)$$

Example 9 (Polyhedral cones are regular) If $\mathcal{K} \subseteq \mathbb{R}^m$ is a proper polyhedral cone, then the set of normalized extreme rays is finite. Therefore, for sufficiently small δ , $(\text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}) \cap \mathcal{B}(x_0, \delta) = \{x_0\}$. If $\ell \in -\mathcal{N}_{\mathcal{K}}(x_0)$, then $\ell(x_0) = 0$ and so (11) is vacuously satisfied for any $\nu > 0$.

Example 10 (Cone over the Veronese embedding is regular) Suppose that $z_0 \in \mathcal{S}^{n-1}$ and $\ell(\phi_{n,2d}(z))$ is nonnegative and vanishes at z_0 . Consider the nonnegative homogeneous quadratic $\|z\|^2 \|z_0\|^2 - \langle z, z_0 \rangle^2$, which vanishes only on the line spanned by z_0 . Since both $\ell(\phi_{n,2d}(z))$ and its gradient vanish at $z = z_0$, there exists $M > 0$ such that $\ell(\phi_{n,2d}(z)) \leq M(\|z\|^2 \|z_0\|^2 - \langle z, z_0 \rangle^2)$ for all $z \in \mathcal{S}^{n-1}$. Then if $z \in \mathcal{S}^{n-1}$,

$$\ell(\phi_{n,2d}(z)) \leq M(1 - \langle z, z_0 \rangle^2) \leq M(1 - \langle z, z_0 \rangle^{2d}) = \frac{M}{2} \|\phi_{n,2d}(z) - \phi_{n,2d}(z_0)\|_B^2.$$

Since z_0 was arbitrary, it follows that $\mathcal{C}_{n,2d}$ is regular.

Although the definition of a cone being regular only bounds the growth of a linear functional on normalized extreme rays locally around its minimizer, such a local bound can be extended to a global bound.

Lemma 4 If a closed, pointed, convex cone $\mathcal{K} \subseteq \mathbb{R}^m$ is regular then for each $x_0 \in \text{ext}(\mathcal{K})$ and each $\ell \in -\mathcal{N}_{\mathcal{K}}(x_0)$ there exists $\nu_0 > 0$ such that $\ell(x) \leq \nu_0 \|x - x_0\|^2$ for all $x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}$.

Proof If $x_0 \in \text{ext}(\mathcal{K})$, the cone \mathcal{K} is regular, and $\ell \in -\mathcal{N}_{\mathcal{K}}(x_0)$, then there exist $\delta > 0$ and $\nu \geq 0$ such that $x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}$ and $\|x - x_0\| < \delta$ implies $\ell(x) \leq \nu \|x - x_0\|^2$. If, on the other hand, $\frac{\|x_0 - x\|}{\delta} \geq 1$ and $L = \max_{x \in \mathcal{S}^{m-1}} \ell(x)$ then

$$\ell(x) = \ell(x - x_0) \leq L\delta \left(\frac{\|x - x_0\|}{\delta} \right) \leq L\delta \left(\frac{\|x - x_0\|}{\delta} \right)^2 = \frac{L}{\delta} \|x - x_0\|^2.$$

Choosing $\nu_0 = \max\{\nu, L/\delta\}$ completes the proof. \square

3.3 Terracini convexity of neighborly cones

We are now in a position to state and prove the main result of this section.

Theorem 1 *If a closed, pointed, convex cone is non-degenerate k -neighborly and regular, then it is k -Terracini.*

Proof Let $x^{(1)}, x^{(2)}, \dots, x^{(k)}$ be a collection of normalized extreme rays of a closed, pointed, non-degenerate k -neighborly convex cone \mathcal{K} . To establish that \mathcal{K} is k -Terracini convex, by Remark 1 it suffices to show that $\bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)})) \subseteq \text{span}\left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})\right)$. As such, let $\Delta \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)}))$ be arbitrary.

Let ℓ be a linear functional from the definition of non-degenerate k -neighborliness of \mathcal{K} . Since this functional is nonnegative on \mathcal{K} and vanishes on $x^{(i)}$ for $i = 1, 2, \dots, k$, it follows that $\ell \in -\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})$. Further, from Lemma 3 there exists $\mu_0 > 0$ such that

$$\ell(x) \geq \mu_0 \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \quad \text{for all } x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}. \tag{12}$$

Since $\Delta \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{K}}(x^{(i)}))$, for each i we have a decomposition of Δ as $\Delta = \Delta_+^{(i)} - \Delta_-^{(i)}$ where $\Delta_+^{(i)}, \Delta_-^{(i)} \in -\mathcal{N}_{\mathcal{K}}(x^{(i)})$. As \mathcal{K} is regular, for each $i = 1, 2, \dots, k$, there exists $v_0^{(i)} > 0$ such that $\Delta_-^{(i)} \leq v_0^{(i)} \|x - x^{(i)}\|^2$ for all $x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}$ from Lemma 4. Setting $v_0 = \max_i \{v_0^{(i)}\}$ we have that

$$\Delta(x) \geq -v_0 \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \quad \text{for all } x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}. \tag{13}$$

If we choose $0 < \gamma < \mu_0/v_0$ it follows from (12) and (13) that

$$(\ell + \gamma\Delta)(x) \geq (\mu_0 - \gamma v_0) \min_{i=1,2,\dots,k} \|x - x^{(i)}\|^2 \quad \text{for all } x \in \text{ext}(\mathcal{K}) \cap \mathcal{S}^{m-1}. \tag{14}$$

Using the fact that $\Delta(x^{(i)}) = \ell(x^{(i)}) = 0$ for $i = 1, 2, \dots, k$, we can conclude that $\ell + \gamma\Delta \in -\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})$. Since $\gamma \neq 0$, it follows that $\Delta \in \text{span}\left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{K}}(x^{(i)})\right)$, and so \mathcal{K} is k -Terracini convex. \square

This theorem yields two immediate corollaries based on the examples in Sects. 3.2.1 and 3.2.2.

Corollary 2 *A pointed k -neighborly polyhedral cone is k -Terracini convex.*

Proof This follows immediately from Theorem 1 and Examples 7 and 9. \square

Corollary 3 *The cone $\mathcal{C}_{n,2d}$ over the Veronese embedding is d -Terracini convex.*

Proof This follows immediately from Theorem 1 and Examples 8 and 10. \square

While Corollary 3 holds for general cones over Veronese embeddings, for the special case of the cone over the moment curve, i.e., the case $n = 2$, a stronger conclusion is possible.

Corollary 4 *The cone $\mathcal{C}_{2,2d}$ over the homogeneous moment curve is Terracini convex, i.e., is k -Terracini convex for all k .*

Proof Let $x^{(1)}, \dots, x^{(k)}$ generate distinct extreme rays of $\mathcal{C}_{2,2d}$. Then there exist points $z^{(1)}, \dots, z^{(k)} \in \mathbb{R}^2$ such that $\phi_{2,2d}(z^{(i)}) = x^{(i)}$ for $i = 1, 2, \dots, k$ and $z_1^{(i)}z_2^{(j)} - z_1^{(j)}z_2^{(i)} \neq 0$ whenever $i \neq j$. In other words, the $z^{(i)}$ represent distinct elements of the real projective line. Since $\mathcal{C}_{2,2d}$ is d -Terracini convex, to conclude that $\mathcal{C}_{2,2d}$ is Terracini convex it suffices to show that if $k \geq d + 1$ then

$$\bigcap_{i=1}^k \text{span} \left(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}) \right) = \{0\} \subseteq \text{span} \left(\bigcap_{i=1}^k \mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}) \right).$$

Elements $\ell \in \mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)})$ are exactly the linear functionals with the property that $\ell(\phi_{2,2d}(z))$ is a bivariate homogeneous polynomial of degree $2d$ that is non-positive and vanishes at $z^{(i)}$. As such $\ell \in \bigcap_{i=1}^d \mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)})$ if and only if $\ell(\phi_{2,2d}(z))$ is a non-negative multiple of $p(z) = -\prod_{i=1}^d (z_1z_2^{(i)} - z_2z_1^{(i)})^2$. From d -Terracini convexity of $\mathcal{C}_{2,2d}$, it follows that $\ell \in \bigcap_{i=1}^d \text{span}(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}))$ if and only if $\ell(\phi_{2,2d}(z))$ is a scalar multiple of $p(z)$.

Consider any $\tilde{\ell} \in \bigcap_{i=1}^k \text{span}(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}))$ for $k \geq d + 1$ and let $q(z) = \tilde{\ell}(\phi_{2,2d}(z))$. Then $q(z) = \alpha p(z)$ for some scalar α since

$$\tilde{\ell} \in \bigcap_{i=1}^d \text{span} \left(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}) \right) \subseteq \bigcap_{i=1}^k \text{span} \left(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(i)}) \right).$$

Furthermore, $q(z^{(d+1)}) = 0$ since $\tilde{\ell} \in \text{span}(\mathcal{N}_{\mathcal{C}_{2,2d}}(x^{(d+1)}))$. Since $z_1^{(i)}z_2^{(j)} - z_1^{(j)}z_2^{(i)} \neq 0$ whenever $i \neq j$, this is only possible if $\alpha = 0$ and hence $\tilde{\ell} = 0$. \square

A natural question at this stage is whether cones $\mathcal{C}_{n,2d}$ over Veronese embeddings for $n > 2$ are also Terracini convex, rather than merely being d -Terracini convex. For the case of $n = 3$, this question is open, and (to the best of our knowledge) cannot be resolved given the current understanding of the structure of $\mathcal{C}_{3,2d}$. For the case of $n = 4$, the following example shows that $\mathcal{C}_{4,4}$ is not Terracini convex based on Blekherman's study of dimensional differences between faces of nonnegative polynomials and sums of squares [5].

Example 11 ([5, Section 2.2]) Consider the cone $\mathcal{C}_{4,4}$, which can be viewed as dual to nonnegative quartic forms in four variables. Let $S = \{(1, 1, 0, 0), (1, 0, 1, 0), (1, 0, 0, 1), (0, 1, 1, 0), (0, 1, 0, 1), (0, 0, 1, 1), (1, 1, 1, 1)\}$. Blekherman shows that the face of nonnegative quartic forms in four variables that vanish on S has dimension 6, i.e.,

$$\dim \operatorname{span} \left(\bigcap_{z \in S} \mathcal{N}_{\mathcal{C}_{4,4}}(\phi_{4,4}(z)) \right) = 6. \quad (15)$$

Furthermore, each of the subspaces $\operatorname{span}(\mathcal{N}_{\mathcal{C}_{4,4}}(\phi_{4,4}(z)))$ for $z \in S$ has codimension 4 in the 35-dimensional space of quartic forms in four variables. The intersection (over $z \in S$) of these subspaces are exactly the forms that double vanish on S . Consequently

$$\dim \left(\bigcap_{z \in S} \operatorname{span}(\mathcal{N}_{\mathcal{C}_{4,4}}(\phi_{4,4}(z))) \right) \geq 35 - 4|S| = 7. \quad (16)$$

In Blekherman's language, the set S is not 2-independent. It follows from (15) and (16) that $\mathcal{C}_{4,4}$ is not 7-Terracini convex, and hence not Terracini convex.

4 Preservation of Terracini convexity under linear images

In this section, we consider the Terracini convexity properties of linear images of Terracini convex cones such as the nonnegative orthant and the positive semidefinite matrices. We carry out our investigation by analyzing the performance of convex relaxations for nonconvex inverse problems. Specifically, we consider the problem of finding the componentwise nonnegative vector with the smallest number of nonzero entries (i.e., nonnegative sparse vectors) in an affine space, and that of finding the smallest rank positive semidefinite matrix in an affine space. Both of these problems arise commonly in many applications and they have been widely studied in the literature. In Sect. 4.1 we consider sparse vector recovery and we reprove a result of Donoho and Tanner that a natural linear programming relaxation succeeds in recovering nonnegative sparse vectors in an affine space if and only if a particular linear image of the nonnegative orthant is k -Terracini convex for an appropriate k [9]. Donoho and Tanner's original proof was given in the language of neighborly polytopes. We provide an alternate proof in Sect. 4.1 by appealing to the dual relation Proposition 1 as it is instructive in our subsequent analysis on recovering low-rank matrices in affine spaces. In Sect. 4.2 we prove that the success of a semidefinite programming relaxation in recovering positive semidefinite low-rank matrices implies k -Terracini convexity of a particular linear image of the cone of positive semidefinite matrices for a suitable k ; in the reverse direction, we show that a 'robust' analog of k -Terracini convexity implies success of the semidefinite relaxation. The results in Sect. 4.2 lead to a new family of non-polyhedral Terracini convex cones, which we describe in Sect. 4.3. Thus, this section supplies new examples of Terracini convex cones, and our results also highlight the utility of our definition of Terracini convexity in generalizing neighborly polyhedral cones, as the usual notion of neighborliness for non-polyhedral cones is not the right one for characterizing the performance of semidefinite relaxations for low-rank matrix recovery.

4.1 Linear images of the nonnegative orthant

In applications ranging from feature selection in machine learning to recovering signals and images from a limited number of measurements, a frequently encountered question is that of finding vectors with the smallest number of nonzero entries in a given affine space. Consider the following model problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^d} \quad & |\text{support}(x)| \\ \text{s.t.} \quad & Ax = b, \quad x \geq 0. \end{aligned} \tag{P0}$$

Here $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is a linear map, $b \in \mathbb{R}^n$, $x \geq 0$ denotes componentwise nonnegativity of x , and $|\text{support}(x)|$ denotes the number of nonzero entries of x . As solving (P0) is NP-hard in general, the following tractable linear programming relaxation is the method of choice that is employed in most contexts:

$$\begin{aligned} LP(A, b) = \arg \min_{x \in \mathbb{R}^d} \quad & \langle 1, x \rangle \\ \text{s.t.} \quad & Ax = b, \quad x \geq 0. \end{aligned} \tag{P1}$$

In assessing the performance of the relaxation (P1), the usual mode of analysis is to suppose that there exists a nonnegative vector $x^* \in \mathbb{R}^d$ with a small number of nonzeros such that $b = Ax^*$, and to then ask whether x^* is the unique optimal solution of (P1), i.e., whether $LP(A, Ax^*) = \{x^*\}$. The main result of Donoho and Tanner [9] relates the success of (P1) to neighborliness properties of images of the d -simplex $\Delta^d = \{x \in \mathbb{R}^d : x \geq 0, \langle 1, x \rangle = 1\}$ under the map A .

In Theorem 2, to follow, we state a conic analog of the result in [9], and we reprove it in two stages. The proof we give offers a template for our generalization in Sect. 4.2 on relating the performance of semidefinite relaxations for low-rank matrix recovery to Terracini convexity of linear images of the cone of positive semidefinite matrices. Our analysis relies on relating the following three properties; each of these is stated with respect to a positive integer k , which will be clear from context.

- A linear map $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$ satisfies the *exact recovery property* if, for each $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$, the unique optimal solution of the linear programming relaxation (P1) is $LP(A, Ax^*) = \{x^*\}$.
- Consider a linear map $B : \mathbb{R}^d \rightarrow \mathbb{R}^N$. The cone $B(\mathbb{R}_+^d)$ satisfies the *unique preimage property* if, for each $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$, the point Bx^* has a unique preimage in \mathbb{R}_+^d .
- Consider a linear map $B : \mathbb{R}^d \rightarrow \mathbb{R}^N$. The cone $B(\mathbb{R}_+^d)$ satisfies the *Terracini convexity property* if it is pointed, it has d extreme rays, and it is k -Terracini convex.

Given these notions we state next the result of Donoho and Tanner in conic form:

Theorem 2 Consider a linear map $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$ that is surjective and define the linear map $B : \mathbb{R}^d \rightarrow \mathbb{R}^{n+1}$ as $Bx = \begin{pmatrix} Ax \\ \langle 1, x \rangle \end{pmatrix}$. Suppose that $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$.

Fix a positive integer $k < d$. The map A satisfies the exact recovery property if and only if the cone $B(\mathbb{R}_+^d)$ satisfies the Terracini convexity property.

Remark 2 This result is a conic analog of those in [9]. The assumption that A is surjective is to ensure a cleaner argument; if this condition is not satisfied, the proof can be adapted by restricting to the image of A . Finally, the results in [9] do not require the condition $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$, and they are described in terms of a property termed ‘outward neighborliness’. However, the particular restriction on which we focus suffices for our purposes and leads to a simpler exposition.

This result leads to two types of consequences in [9]. In one direction, Donoho and Tanner leveraged results on constructions of neighborly polytopes to obtain new families of linear maps A for which the linear program (P1) succeeds in sparse recovery. Conversely, by building on results in the sparse recovery literature, they constructed new families of neighborly polytopes.

Our proof proceeds in two steps and is based on the following intermediate results.

Lemma 5 Consider a linear map $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$ and define the linear map $B : \mathbb{R}^d \rightarrow \mathbb{R}^{n+1}$ as $Bx = \begin{pmatrix} Ax \\ \langle 1, x \rangle \end{pmatrix}$. Suppose that $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$. Fix a positive integer $k < d$. The map A satisfies the exact recovery property if and only if the cone $B(\mathbb{R}_+^d)$ satisfies the unique preimage property.

Proof For the case $x^* = 0$, one can check that $LP(A, 0) = \{0\}$ and that the unique preimage of $0 \in \mathbb{R}^{n+1}$ under the map B in \mathbb{R}_+^d is also $\{0\}$. For nonzero x^* , in considering the exact recovery property and the unique preimage property, we may assume without loss of generality that $\langle 1, x^* \rangle = 1$. The reason for this that $LP(A, \alpha b) = \alpha LP(A, b)$ for any $\alpha > 0$; the unique preimage property is similarly unaffected by such scaling. With this normalization, the exact recovery property is equivalent to the fact that for any $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$, the point Ax^* has a unique preimage in the solid simplex $\Delta_0^d = \{x \in \mathbb{R}^d : \langle 1, x \rangle \leq 1, x \geq 0\}$.

Consider the implication that the exact recovery property implies the unique preimage property. Assume that the unique preimage property does not hold. Then there exists $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$ and $\tilde{x} \in \mathbb{R}_+^d$ such that $B\tilde{x} = Bx^*$, $\tilde{x} \neq x^*$. Based on the description of B , we can conclude that $\langle 1, \tilde{x} \rangle = 1$ and therefore $\tilde{x} \in \Delta^d$. This violates the property that Ax^* has a unique preimage in Δ_0^d ; hence the exact recovery property does not hold.

Conversely, consider the implication that the unique preimage property implies the exact recovery property. Assume for the sake of a contradiction that there exists $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$ and $\tilde{x} \in \Delta_0^d$ such that $A\tilde{x} = Ax^*$, $\tilde{x} \neq x^*$. As $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$, there exists $x^0 \in \Delta^d$ with $|\text{support}(x^0)| = d$ such that $Ax^0 = 0$. The point $x' = (1 - \langle 1, \tilde{x} \rangle)x^0 + \tilde{x}$ has the property that $Bx' = Bx^*$. Consequently, we have that $x^* = x' = (1 - \langle 1, \tilde{x} \rangle)x^0 + \tilde{x}$, which in turn implies that x^0 and \tilde{x} belong to the smallest face of \mathbb{R}_+^d containing x^* , i.e., $\text{support}(x^0) \subseteq \text{support}(x^*)$ and $\text{support}(\tilde{x}) \subseteq \text{support}(x^*)$. However, as $|\text{support}(x^0)| = d$ but $|\text{support}(x^*)| \leq k < d$, we have the desired contradiction. \square

Our next result relates the unique preimage property to the Terracini convexity property:

Proposition 5 Consider a linear map $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$ and define the linear map $B : \mathbb{R}^d \rightarrow \mathbb{R}^{n+1}$ as $Bx = \begin{pmatrix} Ax \\ (1, x) \end{pmatrix}$. Suppose the map B is surjective. Fix a positive integer k . The cone $B(\mathbb{R}_+^d)$ satisfies the unique preimage property if and only if it satisfies the Terracini convexity property.

Proof First, we give a dual reformulation of the unique preimage property. For each $x^* \in \mathbb{R}_+^d$ with $|\text{support}(x^*)| \leq k$, the property that Bx^* has a unique preimage in \mathbb{R}_+^d is equivalent to the transverse intersection condition $\text{null}(B) \cap \mathcal{N}_{\mathbb{R}_+^d}(x^*) = \{0\}$. The cone $\mathcal{N}_{\mathbb{R}_+^d}(x^*)$ is closed and therefore one can check that this transverse intersection condition is equivalent to $\text{null}(B)^\perp \cap \text{ri}(\mathcal{N}_{\mathbb{R}_+^d}(x^*)) \neq \emptyset$. As the nonnegative orthant is a self-dual cone, the normal cone $\mathcal{N}_{\mathbb{R}_+^d}(x^*)$ is given by a face of \mathbb{R}_+^d of co-dimension at most k . In summary, the unique preimage property states that for any face Ω of \mathbb{R}_+^d of co-dimension at most k , we have that $\text{null}(B)^\perp \cap \text{ri}(\Omega) \neq \emptyset$.

Second, we note that the cone $B(\mathbb{R}_+^d)$ is pointed by construction. As the linear map B is surjective, elements of the normal cone $\mathcal{N}_{B(\mathbb{R}_+^d)}(Bx)$ for any $x \in \mathbb{R}_+^d$ are in one-to-one correspondence with $\text{null}(B)^\perp \cap \mathcal{N}_{\mathbb{R}_+^d}(x)$. Consequently, by appealing to Proposition 1, the cone $B(\mathbb{R}_+^d)$ being k -Terracini convex is equivalent to the condition that for any face Ω of \mathbb{R}_+^d of co-dimension at most k , we have that $\text{span}(\text{null}(B)^\perp \cap \Omega) = \text{null}(B)^\perp \cap \text{span}(\Omega)$.

With these two reformulations of the unique preimage property and the Terracini convexity property in hand, we proceed to establish the desired result.

Consider the implication that the unique preimage property implies the Terracini convexity property. Based on the unique preimage property applied to elements of \mathbb{R}_+^d with one nonzero entry, we conclude that $B(\mathbb{R}_+^d)$ has d extreme rays. Let $v \in \text{null}(B)^\perp \cap \text{ri}(\Omega)$. Letting U be an open set in \mathbb{R}^d containing the origin, we have that $v + \epsilon[U \cap \text{null}(B)^\perp \cap \text{span}(\Omega)] \subset \text{null}(B)^\perp \cap \text{ri}(\Omega)$ for a sufficiently small $\epsilon > 0$. Consequently, we can conclude that $\text{span}(\text{null}(B)^\perp \cap \Omega) = \text{null}(B)^\perp \cap \text{span}(\Omega)$, which is equivalent to $B(\mathbb{R}_+^d)$ being k -Terracini convex.

Next, consider the implication that the Terracini convexity property implies the unique preimage property. We prove this by induction on k . For the base case $k = 1$, as the cone $B(\mathbb{R}_+^d)$ has d extreme rays, we have that the unique preimage property holds for $k = 1$. For $k > 1$, suppose for the sake of a contradiction that $\text{null}(B)^\perp \cap \text{ri}(\Omega) = \emptyset$. Thus, there exists a face $\hat{\Omega}$ of \mathbb{R}_+^d contained strictly in Ω , i.e., $\hat{\Omega} \subsetneq \Omega$ such that $\text{null}(B)^\perp \cap \hat{\Omega} = \text{null}(B)^\perp \cap \Omega$. We have the following sequence of containment relations:

$$\begin{aligned} \text{null}(B)^\perp \cap \text{span}(\hat{\Omega}) &\subseteq \text{null}(B)^\perp \cap \text{span}(\Omega) \\ &= \text{span}(\text{null}(B)^\perp \cap \Omega) \\ &= \text{span}(\text{null}(B)^\perp \cap \hat{\Omega}) \\ &\subseteq \text{null}(B)^\perp \cap \text{span}(\hat{\Omega}). \end{aligned}$$

The first relation follows from $\hat{\Omega} \subseteq \Omega$, the second one follows from the Terracini convexity property, the third one follows from $\text{null}(B)^\perp \cap \hat{\Omega} = \text{null}(B)^\perp \cap \Omega$, and the final one follows from the fact that the span of the intersection of two sets is contained inside the intersection of the spans of the sets. In conclusion, all the containments are satisfied with equality and we have that $\text{null}(B)^\perp \cap \text{span}(\hat{\Omega}) = \text{null}(B)^\perp \cap \text{span}(\Omega)$, or equivalently that:

$$\text{null}(B) + \text{span}(\hat{\Omega})^\perp = \text{null}(B) + \text{span}(\Omega)^\perp. \tag{17}$$

As \mathbb{R}_+^d is a polyhedral cone, we note that $\text{span}(\hat{\Omega})^\perp$ and $\text{span}(\Omega)^\perp$ are themselves spans of faces of \mathbb{R}_+^d . In particular, let $\mathcal{F}, \hat{\mathcal{F}}$ be faces of \mathbb{R}_+^d such that $\mathcal{F} \subsetneq \hat{\mathcal{F}}$, and $\text{span}(\mathcal{F}) = \text{span}(\Omega)^\perp, \text{span}(\hat{\mathcal{F}}) = \text{span}(\hat{\Omega})^\perp$. The relationship (17) implies that there exists a generator \hat{x} of an extreme ray of \mathbb{R}_+^d in $\hat{\mathcal{F}} \setminus \mathcal{F}$ such that $\hat{x} = (x^{(+)} - x^{(-)}) + v$ for $x^{(+)}, x^{(-)} \in \mathcal{F}$ with disjoint supports and $v \in \text{null}(B)$. Hence, we have that $B(\hat{x} + x^{(-)}) = Bx^{(+)}$. As $\dim(\mathcal{F}) = k$, the sum of the sizes of the supports of $\hat{x} + x^{(-)}$ and of $x^{(+)}$ is at most $k + 1$. If $x^{(+)} \neq 0$ we have a contradiction due to the inductive hypothesis. If $x^{(+)} = 0$ we have $\langle 1, (\hat{x} + x^{(-)}) \rangle = 0$, which implies that $\hat{x} + x^{(-)} = 0$ and in turn that $\hat{x} = 0$, also a contradiction. \square

Based on these two results, we are now in a position to prove Theorem 2.

Proof of Theorem 2 As $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$ and $k < d$ by assumption, we can apply Lemma 5. Specifically, the exact recovery property for A is equivalent to the unique preimage property for $B(\mathbb{R}_+^d)$.

Next, in preparation to apply Proposition 5, we need to verify that the linear map B is surjective. The surjectivity of B is equivalent to A being surjective and $1 \notin \text{null}(A)^\perp$. The former condition holds by assumption and the latter condition is in turn equivalent to $\text{null}(A) \not\subseteq \text{span}(1)^\perp$. The assumption $\text{null}(A) \cap \mathbb{R}_{++}^d \neq \emptyset$ implies that $\text{null}(A) \not\subseteq \text{span}(1)^\perp$. Thus, we are in a position to apply Proposition 5 and obtain that the unique preimage property of the cone $B(\mathbb{R}_+^d)$ is equivalent to $B(\mathbb{R}_+^d)$ satisfying the Terracini convexity property. This concludes the proof. \square

4.2 Linear images of the positive semidefinite matrices

The development of convex relaxations for obtaining low-rank matrices in affine spaces largely paralleled and built upon the literature on sparse recovery. Notable examples of such problems include factor analysis and collaborative filtering. Concretely, given an affine space in \mathbb{S}^d of the form $\{X \in \mathbb{S}^d : \mathcal{A}(X) = b\}$ where $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ is a linear map and $b \in \mathbb{R}^n$, consider the following optimization problem for identifying a positive-semidefinite low-rank matrix in this space:

$$\begin{aligned} \min_{X \in \mathbb{S}^d} \quad & \text{rank}(X) & \tag{R0} \\ \text{s.t.} \quad & \mathcal{A}(X) = b, X \succeq 0. \end{aligned}$$

As with the problem (P0), the program (R0) is also NP-hard to solve in general. Consequently, the following semidefinite relaxation is widely employed in practice:

$$SDP(\mathcal{A}, b) = \arg \min_{X \in \mathbb{S}^d} \text{tr}(X) \tag{R1}$$

$$\text{s.t. } \mathcal{A}(X) = b, X \succeq 0.$$

By analogy with the analysis of the performance of (P1), we are interested in obtaining conditions under which the unique optimal solution of (R1) with $b = \mathcal{A}(X^*)$ for a low-rank matrix $X^* \in \mathbb{S}_+^d$ is equal to X^* , i.e., whether $SDP(\mathcal{A}, \mathcal{A}(X^*)) = \{X^*\}$. Our objective in the remainder of this section is to relate such exact recovery to Terracini convexity of an appropriate linear image of \mathbb{S}_+^d .

As with the previous subsection, our analysis is organized in terms of three properties:

- A linear map $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ satisfies the *exact recovery property* if for any $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$, the unique optimal solution of the semidefinite programming relaxation (R1) is $SDP(\mathcal{A}, \mathcal{A}(X^*)) = \{X^*\}$.
- Consider a linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^N$. The cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the *unique preimage property* if for any $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$, the point $\mathcal{B}(X^*)$ has a unique preimage in \mathbb{S}_+^d .
- Consider a linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^N$. The cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the *Terracini convex property* if it is closed and pointed, its extreme rays are in one-to-one correspondence with those of \mathbb{S}_+^d , and it is k -Terracini.

In what follows, let $\mathcal{O}^d = \{X \in \mathbb{S}^d : \text{tr}(X) = 1, X \succeq 0\}$ be the spectraplex. This plays the same role as the simplex Δ^d did in Sect. 4.1. We are now in a position to state the main new result of this section.

Theorem 3 Consider a linear map $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ and fix a positive integer $k < d$. Then the following two statements hold:

1. Suppose that \mathcal{A} is surjective and $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$. Consider the linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ defined as $\mathcal{B}(X) = \begin{pmatrix} \mathcal{A}(X) \\ \text{tr}(X) \end{pmatrix}$. If the map \mathcal{A} satisfies the exact recovery property, then the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the Terracini convexity property.
2. Assume that $n > \binom{d+1}{2} - \binom{d-k+1}{2}$. Suppose there exists an open set \mathcal{G} in the space of linear maps from \mathbb{S}^d to \mathbb{R}^n with the following properties:
 - $\mathcal{A} \in \mathcal{G}$
 - For each $\tilde{\mathcal{A}} \in \mathcal{G}$, the map $\tilde{\mathcal{A}}$ is surjective and satisfies $\text{null}(\tilde{\mathcal{A}}) \cap \mathbb{S}_{++}^d \neq \emptyset$.
 - For each $\tilde{\mathcal{A}} \in \mathcal{G}$ with associated $\tilde{\mathcal{B}} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ defined as $\tilde{\mathcal{B}}(X) = \begin{pmatrix} \tilde{\mathcal{A}}(X) \\ \text{tr}(X) \end{pmatrix}$, the cone $\tilde{\mathcal{B}}(\mathbb{S}_+^d)$ satisfies the Terracini convexity property.

Then the map \mathcal{A} satisfies the exact recovery property.

The proof in the direction from the exact recovery property to the Terracini convexity property largely follows the same sequence of steps as the proof of the analogous

direction of Theorem 2, although technical care is required due to the fact that the cone of feasible directions into the cone of positive-semidefinite matrices is not closed. In the direction from the Terracini convexity property to the exact recovery property, we require a robust analog of the Terracini convexity property. This condition is in the same spirit as constraint qualification type assumptions that are required in the semidefinite programming literature in order to guarantee strict complementarity [16].

Inspired by the two-stage proof in Sect. 4.1, we begin with the following result that parallels Lemma 5:

Lemma 6 Consider a linear map $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ and define the linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ as $\mathcal{B}(X) = \begin{pmatrix} \mathcal{A}(X) \\ \text{tr}(X) \end{pmatrix}$. Suppose that $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$. Fix a positive integer $k < d$. The map \mathcal{A} satisfies the exact recovery property if and only if the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the unique preimage property.

Proof As with the proof of Lemma 5, in considering the exact recovery property and the unique recovery property, we assume without loss of generality that $\text{tr}(X^*) = 1$. With this normalization, the exact recovery property is equivalent to the fact that for any $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$, the point $\mathcal{A}(X^*)$ has a unique preimage in the solid spectraplex $\mathcal{O}_0^d = \{X \in \mathbb{S}^d : \text{tr}(X) \leq 1, X \geq 0\}$.

Consider the implication that the exact recovery property implies the unique preimage property. Assume that the unique preimage property does not hold. Then there exists $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$ and $\tilde{X} \in \mathbb{S}_+^d$ such that $\mathcal{B}(\tilde{X}) = \mathcal{B}(X^*)$, $\tilde{X} \neq X^*$. Based on the description of \mathcal{B} , we can conclude that $\text{tr}(\tilde{X}) = 1$ and therefore $\tilde{X} \in \mathcal{O}_0^d$. This violates the property that $\mathcal{A}(X^*)$ has a unique preimage in \mathcal{O}_0^d ; hence the exact recovery property does not hold.

Conversely, consider the implication that the unique preimage property implies the exact recovery property. Assume for the sake of a contradiction that there exists $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$ and $\tilde{X} \in \mathcal{O}_0^d$ such that $\mathcal{A}(\tilde{X}) = \mathcal{A}(X^*)$, $\tilde{X} \neq X^*$. As $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$, there exists $X^0 \in \mathcal{O}^d$ with $\text{rank}(X^0) = d$ such that $\mathcal{A}(X^0) = 0$. The point $X' = (1 - \text{tr}(\tilde{X}))X^0 + \tilde{X}$ has the property that $\mathcal{B}(X') = \mathcal{B}(X^*)$. Consequently, we have that $X^* = X' = (1 - \text{tr}(\tilde{X}))X^0 + \tilde{X}$, which in turn implies that X^0 and \tilde{X} belong to the smallest face of \mathbb{S}_+^d containing X^* . However, as $\text{rank}(X^0) = d$ but $\text{rank}(X^*) \leq k < d$, we have the desired contradiction. \square

The next proposition represents the main new component of the proof of Theorem 3:

Proposition 6 Consider a linear map $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ and define the linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ as $\mathcal{B}(X) = \begin{pmatrix} \mathcal{A}(X) \\ \text{tr}(X) \end{pmatrix}$. Fix a positive integer k . Then we have the following two results:

1. Suppose the map \mathcal{B} is surjective. If the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the unique preimage property, then it satisfies the Terracini convexity property.
2. Assume that $n > \binom{d+1}{2} - \binom{d-k+1}{2}$. Suppose there exists an open set \mathcal{G} in the space of linear maps from \mathbb{S}^d to \mathbb{R}^n satisfying the following conditions:

- $\mathcal{A} \in \mathfrak{S}$
- For each $\tilde{\mathcal{A}} \in \mathfrak{S}$, the associated linear map $\tilde{\mathcal{B}} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ defined as $\tilde{\mathcal{B}}(X) = \begin{pmatrix} \tilde{\mathcal{A}}(X) \\ \text{tr}(X) \end{pmatrix}$ is surjective and the cone $\tilde{\mathcal{B}}(\mathbb{S}_+^d)$ satisfies the Terracini convexity property.

Then the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the unique preimage property.

Remarks: In the direction from the Terracini convexity property to the unique preimage property, the fact that \mathbb{S}_+^d is not polyhedral, unlike \mathbb{R}_+^d , complicates matters in comparison to the proof of Proposition 5. Specifically, translated to the context of the present theorem, the reasoning up to (17) in Proposition 5 continues to hold, but the sentence immediately after (17) is no longer true. As stated previously, the nature of this difficulty is akin to the lack of strict complementarity in semidefinite programs (in contrast to linear programs), thus necessitating some type of constraint qualification assumption. The ‘robust Terracini’ form of the assumption in the second part of this result is similar in spirit to assumptions discussed in [16] to ensure strong duality in conic programs.

Proof We begin by presenting a dual reformulation of the unique preimage property. For each $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$, the property that $\mathcal{B}(X^*) \in \mathcal{C}$ has a unique preimage in \mathbb{S}_+^d is equivalent to the transverse intersection condition $\text{null}(\mathcal{B}) \cap \mathcal{K}_{\mathbb{S}_+^d}(X^*) = \{0\}$. Unlike the situation with Proposition 5, the cone of feasible directions $\mathcal{K}_{\mathbb{S}_+^d}(X^*)$ is not closed, which presents additional complications. We prove next that we must have $\text{null}(\mathcal{B}) \cap \overline{\mathcal{K}_{\mathbb{S}_+^d}(X^*)} = \{0\}$ by reasoning that if there exists a nonzero $M \in \text{null}(\mathcal{B}) \cap \overline{\mathcal{K}_{\mathbb{S}_+^d}(X^*)}$ then there is a low-rank matrix near X^* for which the unique preimage property does not hold.

Concretely, suppose for the sake of a contradiction that $M \in \text{null}(\mathcal{B}) \cap \overline{\mathcal{K}_{\mathbb{S}_+^d}(X^*)}$ with $M \neq 0$. Without loss of generality, we assume that X^* has rank $r \in \{1, \dots, k\}$ with the row/column space equal to the span of the first r standard basis vectors. For such an X^* , the closure of the cone of feasible directions $\overline{\mathcal{K}_{\mathbb{S}_+^d}(X^*)}$ takes on a convenient block-diagonal form, so that $M \in \overline{\mathcal{K}_{\mathbb{S}_+^d}(X^*)}$ may be viewed as follows:

$$M = \begin{pmatrix} P & V' \\ V & Q \end{pmatrix},$$

with $P \in \mathbb{S}^r$, $V \in \mathbb{R}^{(n-r) \times r}$, $Q \in \mathbb{S}_+^{(n-r)}$. We now construct a rank- r matrix for which the unique preimage property does not hold, thus violating the given assumption. Choose any matrix $W \in \mathbb{S}^r$ such that W and $W + P$ are strictly positive definite. We have that the matrix $\begin{pmatrix} W & -V' \\ -V & VW^{-1}V' \end{pmatrix}$ belongs to \mathbb{S}_+^d and has rank equal to r . Further, we also have that the matrix $\begin{pmatrix} W + P & 0 \\ 0 & Q + VW^{-1}V' \end{pmatrix}$ lies in \mathbb{S}_+^d . Consequently, we

have that the matrix:

$$\begin{pmatrix} W + P & 0 \\ 0 & Q + VW^{-1}V' \end{pmatrix} - \begin{pmatrix} W & -V' \\ -V & VW^{-1}V' \end{pmatrix} = \begin{pmatrix} P & V' \\ V & Q \end{pmatrix}$$

lies in the cone of feasible directions from $\begin{pmatrix} W & -V' \\ -V & VW^{-1}V' \end{pmatrix}$ into \mathbb{S}_+^d . Since $M \in \text{null}(\mathcal{B})$,

$$\mathcal{B} \begin{pmatrix} W & -V' \\ -V & VW^{-1}V' \end{pmatrix} = \mathcal{B} \begin{pmatrix} W + P & 0 \\ 0 & Q + VW^{-1}V' \end{pmatrix}$$

and so the image of the rank- r matrix $\begin{pmatrix} W & -V' \\ -V & VW^{-1}V' \end{pmatrix}$ under the map \mathcal{B} does not have a unique preimage in \mathbb{S}_+^d , which gives us the desired contradiction. In summary, we have for each $X^* \in \mathbb{S}_+^d$ with $\text{rank}(X^*) \leq k$ that $\text{null}(\mathcal{B}) \cap \overline{\mathcal{N}_{\mathbb{S}_+^d}(X^*)} = \{0\}$, which in turn is equivalent to $\text{null}(\mathcal{B})^\perp \cap \text{ri}(\mathcal{N}_{\mathbb{S}_+^d}(X^*)) \neq \emptyset$. In analogy to the case of the nonnegative orthant, the positive-semidefinite cone \mathbb{S}_+^d is self-dual and the normal cone $\mathcal{N}_{\mathbb{S}_+^d}(X^*)$ is given by a face of \mathbb{S}_+^d of dimension at least $\binom{d-k+1}{2}$ (corresponding to positive-semidefinite matrices with row/column space orthogonal to those of X^*). Thus, the unique preimage property states that for any face Ω of \mathbb{S}_+^d of dimension at least $\binom{d-k+1}{2}$, we have that $\text{null}(\mathcal{B})^\perp \cap \text{ri}(\Omega) \neq \emptyset$.

Next, we note that for each $\tilde{\mathcal{A}} \in \mathfrak{S}$, the associated linear map $\tilde{\mathcal{B}}$ is such that the cone $\tilde{\mathcal{B}}(\mathbb{S}_+^d)$ is closed and pointed by construction. Further, each $\tilde{\mathcal{B}}$ is surjective by assumption. Thus, elements of the normal cone $\mathcal{N}_{\tilde{\mathcal{B}}(\mathbb{S}_+^d)}(\mathcal{B}(X))$ are in one-to-one correspondence with those of $\text{null}(\tilde{\mathcal{B}})^\perp \cap \mathcal{N}_{\mathbb{S}_+^d}(X)$ for each $X \in \mathbb{S}_+^d$. Hence, by appealing to Proposition 1, the Terracini convexity property states that for any face Ω of \mathbb{S}_+^d of dimension at least $\binom{d-k+1}{2}$, we have that $\text{span}(\text{null}(\tilde{\mathcal{B}})^\perp \cap \Omega) = \text{null}(\tilde{\mathcal{B}})^\perp \cap \text{span}(\Omega)$.

With these reformulations of the unique preimage property and the Terracini convexity property, we now proceed to establish the result.

Proof of Statement 1 To prove the first result, we begin by noting that unique preimage property applied to rank-one elements of \mathbb{S}_+^d implies that the cone $\mathcal{B}(\mathbb{S}_+^d)$ has extreme rays in one-to-one correspondence with those of \mathbb{S}_+^d . Next, let $M \in \text{null}(\mathcal{B})^\perp \cap \text{ri}(\Omega)$. Letting U be an open set in \mathbb{S}^d containing the origin, we have that $M + \epsilon[U \cap \text{null}(\mathcal{B})^\perp \cap \text{span}(\Omega)] \subset \text{null}(\mathcal{B})^\perp \cap \text{ri}(\Omega)$ for a sufficiently small $\epsilon > 0$. Consequently, we can conclude that $\text{span}(\text{null}(\mathcal{B})^\perp \cap \Omega) = \text{null}(\mathcal{B})^\perp \cap \text{span}(\Omega)$, which is equivalent to the Terracini convexity condition.

Proof of Statement 2 Next we consider the second statement. Fix a face Ω of \mathbb{S}_+^d of co-dimension at least $\binom{d-k+1}{2}$. Suppose for the sake of a contradiction that $\text{span}(\text{null}(\mathcal{B})^\perp \cap \text{ri}(\Omega)) = \emptyset$. As $n > \binom{d+1}{2} - \binom{d-k+1}{2}$, we have that $\text{null}(\mathcal{B})^\perp \cap \text{span}(\Omega)$ is a subspace of positive dimension in \mathbb{S}^d . By the Terracini convexity property applied to the cone $\mathcal{B}(\mathbb{S}_+^d)$, we have that $\text{null}(\mathcal{B})^\perp \cap \Omega \neq \emptyset$. Hence, there

exists a proper face $\hat{\Omega}$ of \mathbb{S}_+^d such that $\hat{\Omega} \subsetneq \Omega$, $\text{null}(\mathcal{B})^\perp \cap \Omega = \text{null}(\mathcal{B})^\perp \cap \hat{\Omega}$, and $\text{null}(\mathcal{B})^\perp \cap \text{ri}(\hat{\Omega}) \neq \emptyset$. As a consequence, there also exists an element $W \in [\Omega \cap \text{span}(\hat{\Omega}^\perp)] \setminus \{0\}$.

We use the W available to us to construct a linear map $\tilde{\mathcal{A}} \in \mathfrak{S}$. Specifically, there exists $\epsilon > 0$ such that:

$$\text{null}(\tilde{\mathcal{A}})^\perp = \{M - \epsilon \|M\| W : M \in \text{null}(\mathcal{A})^\perp\}$$

for some $\tilde{\mathcal{A}} \in \mathfrak{S}$. Associated to this $\tilde{\mathcal{A}}$ is the linear map $\tilde{\mathcal{B}}$. We show next that $\text{null}(\tilde{\mathcal{B}})^\perp \cap \Omega = \{0\}$. As $\tilde{\mathcal{B}}$ is surjective, we may consider the direct sum decomposition $\text{null}(\tilde{\mathcal{B}})^\perp = \text{null}(\tilde{\mathcal{A}})^\perp \oplus \text{span}(I)$. Thus, for any $Y \in \text{null}(\tilde{\mathcal{B}})^\perp$, we have the decomposition $Y = M - \epsilon \|M\| W + cI$ for some $M \in \text{null}(\mathcal{A})^\perp$, $c \in \mathbb{R}$. If $Y \in \Omega$ then one can check that $Y + \epsilon \|M\| W \in \Omega$, and in particular, that $Y + \epsilon \|M\| W \notin \hat{\Omega}$ based on the construction of W , unless $M = 0$. But we also have that $Y + \epsilon \|M\| W = M + cI$ and $M + cI \in \hat{\Omega}$, which implies that $M = 0$ and in turn that $c = 0$. In summary, we obtain that $\text{null}(\tilde{\mathcal{B}})^\perp \cap \Omega = \{0\}$.

Next, we prove that $\text{null}(\tilde{\mathcal{B}})^\perp \cap \text{span}(\Omega)$ is a subspace of positive dimension by constructing a nonzero element in this subspace. Recall that $\text{null}(\mathcal{B})^\perp \cap \text{ri}(\hat{\Omega}) \neq \emptyset$ and that $\hat{\Omega} \subset \Omega$. Consider any $Z \in [\text{null}(\mathcal{B})^\perp \cap \text{ri}(\hat{\Omega})]$, which by construction is nonzero. We have the expression $Z = M + cI$ with $M \in \text{null}(\mathcal{A})^\perp \setminus \{0\}$ and $c \in \mathbb{R}$ based on the surjectivity of \mathcal{B} and that $I \notin \hat{\Omega}$. It follows that $Z - \epsilon \|M\| W \in \text{span}(\Omega) \setminus \{0\}$ as $Z \in \hat{\Omega} \setminus \{0\}$ and $W \in [\Omega \cap \text{span}(\hat{\Omega}^\perp)] \setminus \{0\}$. Further, we also have that $Z - \epsilon \|M\| W = M - \epsilon \|M\| W + cI \in \text{null}(\tilde{\mathcal{B}})^\perp$, as $M - \epsilon \|M\| W \in \text{null}(\tilde{\mathcal{A}})^\perp$ and $\text{null}(\tilde{\mathcal{B}})^\perp = \text{null}(\tilde{\mathcal{A}})^\perp \oplus \text{span}(I)$. As a result, we have that $Z - \epsilon \|M\| W \in \text{null}(\tilde{\mathcal{B}})^\perp \cap \text{span}(\Omega) \setminus \{0\}$.

Finally, we consider the preceding two paragraphs together in the context of the Terracini convex property of the cone $\tilde{\mathcal{B}}(\mathbb{S}_+^d)$. Specifically, we have that $\text{null}(\tilde{\mathcal{B}})^\perp \cap \Omega = \{0\}$ and that $\text{null}(\tilde{\mathcal{B}})^\perp \cap \text{span}(\Omega)$ is a subspace of positive dimension. This violates the reformulation of Terracini convexity of $\tilde{\mathcal{B}}(\mathbb{S}_+^d)$ that $\text{span}(\text{null}(\tilde{\mathcal{B}})^\perp \cap \Omega) = \text{null}(\tilde{\mathcal{B}})^\perp \cap \text{span}(\Omega)$. This gives us the desired contradiction. \square

Given the preceding two results, we now prove Theorem 3:

Proof of Theorem 3 For the first statement, we are given that $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$. Hence, we can apply Lemma 6 and obtain that the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the unique preimage property. Next, in preparation to apply the first part of Proposition 6, we need to check that the linear map \mathcal{B} is surjective, which is equivalent to \mathcal{A} being surjective and $I \notin \text{null}(\mathcal{A})^\perp$. The former condition holds by assumption and the latter condition is in turn equivalent to $\text{null}(\mathcal{A}) \not\subseteq \text{span}(I)^\perp$. The assumption $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$ implies that $\text{null}(\mathcal{A}) \not\subseteq \text{span}(I)^\perp$. Thus, we are in a position to apply Proposition 6 and obtain that the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the Terracini convexity property.

For the second statement, we can apply the second part of Proposition 6 to conclude that the cone $\mathcal{B}(\mathbb{S}_+^d)$ satisfies the unique preimage property. Applying Lemma 6, we conclude that the map \mathcal{A} satisfies the exact recovery property. \square

4.3 New families of Terracini convex cones

The results from the preceding section lead naturally to new families of Terracini convex cones. Specifically, from the literature on the semidefinite relaxation (R1) we have that the exact recovery property is satisfied with high probability by random linear maps \mathcal{A} of suitable dimension [7, 15]. Combined with the first part of Theorem 3, we obtain Terracini convex cones that are specified as linear images of the cone of the positive-semidefinite matrices.

Theorem 4 *Let $A_1, \dots, A_n \in \mathbb{R}^{d \times d}$ be a collection of independent random matrices in which each A_i is a Gaussian random matrix with i.i.d entries that have zero-mean and variance $\frac{1}{n}$, and suppose $n \leq (1/2 - \epsilon) \binom{d+1}{2}$ for some $\epsilon \in (0, 1/2)$. Consider the*

linear map $\mathcal{B} : \mathbb{S}^d \rightarrow \mathbb{R}^{n+1}$ defined as $\mathcal{B}(X) = \begin{pmatrix} \text{tr}(A_1 X) \\ \vdots \\ \text{tr}(A_n X) \\ \text{tr}(X) \end{pmatrix}$. There exist constants

$c_1, c_2 > 0$ and $c_3(\epsilon) > 0$ (depending on ϵ), such that for $k = \lfloor \frac{c_1 n}{d} \rfloor$, the cone $\mathcal{B}(\mathbb{S}^d) \subset \mathbb{R}^{n+1}$ is k -Terracini convex with probability greater than $1 - 2e^{-c_2 n} - e^{-c_3(\epsilon)n}$.

Proof We begin with a geometric reformulation of the exact recovery property of Sect. 4.2 based on the argument presented in Lemma 6. Specifically, for a given linear map $\mathcal{A} : \mathbb{S}^d \rightarrow \mathbb{R}^n$ and a positive integer k , the exact recovery property of Sect. 4.2 is equivalent to the condition that for any $X^* \in \mathbb{S}^d$ with $\text{rank}(X^*) \leq k$ and $\text{tr}(X^*) = 1$, we have that $\mathcal{A}(X^*)$ has a unique preimage in the solid spectraplex $\mathcal{O}_0^d = \{X \in \mathbb{S}^d : \text{tr}(X) \leq 1, X \geq 0\}$.

The results in [7, 15] concern a more general geometric criterion which can be specialized to our context. These results are stated in terms of the matrix nuclear norm $\|\cdot\|_* = \sum_i \sigma_i(\cdot)$ (i.e., the sum of the singular values). Consider the linear map $\hat{\mathcal{A}} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^n$ defined in terms of the Gaussian random matrices A_1, \dots, A_n as
$$\hat{\mathcal{A}}(M) = \begin{pmatrix} \text{tr}(A_1 M) \\ \vdots \\ \text{tr}(A_n M) \end{pmatrix}.$$
 There exist constants $c_1, c_2 > 0$ such that if $k = \lfloor \frac{c_1 n}{d} \rfloor$,

then with probability at least $1 - 2e^{-c_2 n}$, for every $M^* \in \mathbb{R}^{d \times d}$ with $\text{rank}(M^*) \leq k$ and $\|M^*\|_* = 1$, the point $\hat{\mathcal{A}}(M^*)$ has a unique preimage in the nuclear norm ball $\{M \in \mathbb{R}^{d \times d} : \|M\|_* \leq 1\}$ [7]. Note that the solid spectraplex $\mathcal{O}_0^d \subset \mathbb{S}^d \subset \mathbb{R}^{d \times d}$ is a subset of the nuclear norm unit ball. Thus, with the same value of $k = \lfloor \frac{c_1 n}{d} \rfloor$, one can conclude that the linear map \mathcal{A} defined by the restriction of $\hat{\mathcal{A}}$ to the domain \mathbb{S}^d satisfies the exact recovery property of Sect. 4.2 for $k = \lfloor \frac{c_1 n}{d} \rfloor$ with probability greater than $1 - 2e^{-c_2 n}$.

Further, we have that $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$ with probability at least $1 - e^{-c_3(\epsilon)n}$. This follows from the observation that the probability that $\text{null}(\mathcal{A}) \cap \mathbb{S}_{++}^d \neq \emptyset$ is the same as the probability that $\text{null}(\mathcal{A}) \cap \mathbb{S}_+^d \neq \{0\}$. This latter quantity can be estimated using the results from [1, 8, 12], using the fact that the positive semidefinite cone is self-dual.

Therefore, by a union bound, the assumptions of the first part of Theorem 3 are satisfied, and hence the cone $\mathcal{B}(\mathbb{S}^d)$ is k -Terracini convex, with probability at least $1 - 2e^{-c_2 n} - e^{-c_3(\epsilon)n}$. \square

Thus, in some sense ‘most’ linear images of the cone of positive semidefinite matrices are k -Terracini convex for a suitable k depending on the dimension of the image of the linear map. This result offers a semidefinite analog of the result of Donoho and Tanner [9] on neighborliness of linear images the nonnegative orthant. Linear images of the positive semidefinite cone are semialgebraic but are generally not basic semialgebraic (as this property is not preserved under linear projections). In the next section, we describe an approach to obtaining basic semialgebraic Terracini convex cones from the positive semidefinite cone via an different construction based on the viewpoint of hyperbolic programming.

5 Terracini convexity and derivative relaxations of hyperbolicity cones

In this section, we study Terracini convexity from a more algebraic perspective by focusing on a class of convex cones that are obtained from hyperbolic polynomials, which are multivariate polynomials possessing certain real-rootedness properties. The associated cones are called hyperbolicity cones, and among the prototypical examples of such cones are the nonnegative orthant and the positive semidefinite cone. Where the previous section demonstrated that generic linear images of the nonnegative orthant and of the positive semidefinite cone are k -Terracini convex for suitable k , here we show that the (algebraically defined) operation of taking derivative relaxations of the nonnegative orthant and of the positive semidefinite cone lead to hyperbolicity cones with non-trivial Terracini convexity properties. As hyperbolicity cones are basic semialgebraic, i.e., they are defined by finitely many polynomial inequalities, a remarkable fact about the k -Terracini convex cones we construct in this section is that they are all basic semialgebraic. In contrast, the k -Terracini convex cones constructed in Sect. 4 by taking projections of the positive semidefinite cone are, in general, not basic semialgebraic.

The rest of the section is organized in the following way. In Sect. 5.1, we briefly state basic definitions and terminology related to hyperbolic polynomials, hyperbolicity cones, and their derivative relaxations, as well as reviewing properties of the boundary and extreme rays of hyperbolicity cones. In Sect. 5.2 we study tangent cones of hyperbolicity cones and how these interact with derivative relaxations. In particular, we show that the tangent cone to a hyperbolicity cone at a point is the hyperbolicity cone associated with the localization of the associated hyperbolic polynomial at that point. This gives us an algebraic handle on the objects arising in the definition of Terracini convexity. Sect. 5.3 is focused on establishing the main result on Terracini convexity properties of derivative relaxations of a class of hyperbolicity cones that includes the orthant, the positive semidefinite cone, and the cone of positive semidefinite Hankel matrices.

5.1 Hyperbolicity cones and their derivative relaxations

Hyperbolic polynomials Let p be a polynomial with real coefficients that is homogeneous of degree d in n variables, and let $e \in \mathbb{R}^n$. We say that p is *hyperbolic with respect to e* if $p(e) > 0$ and, for each $x \in \mathbb{R}^n$, the univariate polynomial $t \mapsto p(te - x)$ has only real roots. Given $x \in \mathbb{R}^n$ let $\lambda_{\max}^{p,e}(x) = \lambda_1^{p,e}(x) \geq \lambda_2^{p,e}(x) \geq \dots \geq \lambda_d^{p,e}(x) = \lambda_{\min}^{p,e}(x)$ denote the roots of $t \mapsto p(te - x)$, or *hyperbolic eigenvalues* of x with respect to p and e . If p and e are clear from the context, we write $\lambda_1(x), \dots, \lambda_d(x)$. The *rank* of $x \in \mathbb{R}^n$, denoted $\text{rank}_p(x)$, is the number of non-zero hyperbolic eigenvalues of x with respect to p and e . The *multiplicity* of x is $\text{mult}_p(x) = \deg(p) - \text{rank}_p(x)$, the number of zero hyperbolic eigenvalues of x .

Hyperbolicity cones Associated with a hyperbolic polynomial p and direction of hyperbolicity e is the *closed hyperbolicity cone* $\Lambda_+(p, e) = \{x \in \mathbb{R}^n : \lambda_{\min}^{p,e}(x) \geq 0\}$. This is a convex cone, a result due to Gårding [11]. We denote the interior of this cone by $\Lambda_{++}(p, e)$. If $\tilde{e} \in \Lambda_{++}(p, e)$, then p is hyperbolic with respect to \tilde{e} and $\Lambda_+(p, e) = \Lambda_+(p, \tilde{e})$ [11]. If p and e are clear from the context, we write Λ_+ instead of $\Lambda_+(p, e)$ for brevity of notation.

Although the hyperbolic eigenvalues of x with respect to p depend on the choice of e , the multiplicity, $\text{mult}_p(x)$, and rank, $\text{rank}_p(x)$, are independent of the choice of direction of hyperbolicity [17, Proposition 22]. The lineality space of the hyperbolicity cone Λ_+ is exactly the set of points with multiplicity $\deg(p)$ (or rank zero), i.e.,

$$\Lambda_+ \cap (-\Lambda_+) = \{x \in \mathbb{R}^n : \text{mult}_p(x) = \deg(p)\}. \tag{18}$$

(see, e.g., [17, Proposition 11]). If we expand $p(x + te)$ in powers of t as

$$p(x + te) = a_0 t^d + a_1(x) t^{d-1} + \dots + a_{d-2}(x) t^2 + a_{d-1}(x) t + a_d(x), \tag{19}$$

then Descartes’ rule of signs gives an equivalent description of the hyperbolicity cone as

$$\Lambda_+(p, e) = \{x \in \mathbb{R}^n : a_d(x) \geq 0, a_{d-1}(x) \geq 0, a_{d-2}(x) \geq 0, \dots, a_1(x) \geq 0\}.$$

This shows that any hyperbolicity cone is a *basic semialgebraic* set, i.e., it can be expressed via finitely many polynomial inequalities.

Derivative relaxations If p is hyperbolic with respect to e and $\tilde{e} \in \Lambda_{++}(p, e)$, then the directional derivative

$$D_{\tilde{e}} p(x) := \left. \frac{d}{dt} p(x + t\tilde{e}) \right|_{t=0}$$

is again hyperbolic with respect to e (by Rolle’s theorem). The hyperbolicity cone $\Lambda_+(D_{\tilde{e}} p, e)$ satisfies $\Lambda_+(D_{\tilde{e}} p, e) \supseteq \Lambda_+(p, e)$. As such, it is often referred to as a *derivative relaxation* of $\Lambda_+(p, e)$. When p, e , and \tilde{e} are clear from the context we abuse notation and write $\Lambda'_+ := \Lambda_+(D_{\tilde{e}} p, e)$ for brevity.

One of the most interesting aspects of derivative relaxations is that boundary points (of high enough multiplicity) of Λ_+ remain boundary points of Λ'_+ .

Theorem 5 (Renegar [17, Theorem 12]) *Let p be hyperbolic with respect to e with hyperbolicity cone Λ_+ , and for any $\tilde{e} \in \Lambda_{++}(p, e)$ let the associated derivative relaxation be Λ'_+ . If $m \geq 3$ then*

$$\{x \in \Lambda_+ : \text{mult}_p(x) = m\} = \{x \in \Lambda'_+ : \text{mult}_{D_{\tilde{e}}p}(x) = m - 1\}.$$

As a straightforward corollary, we obtain a relationship between the lineality spaces of a hyperbolicity cone and its derivative relaxation.

Corollary 5 *Under the same hypotheses as Theorem 5, if $\deg(p) \geq 3$ then $\Lambda_+ \cap (-\Lambda_+) = \Lambda'_+ \cap (-\Lambda'_+)$.*

Proof This follows from Theorem 5 by noting that the lineality space of Λ_+ is exactly the set of x with $\text{mult}_p(x) = \deg(p)$ and the lineality space of Λ'_+ is exactly the set of x with $\text{mult}_{D_{\tilde{e}}p}(x) = \deg(p) - 1$. \square

One consequence of Corollary 5 is that if $\deg(p) \geq 3$ then Λ_+ being a pointed cone implies that any derivative relaxation Λ'_+ is also pointed. Building on Corollary 5, we can understand how the extreme rays of the derivative cone and the original cone relate to each other. In particular, the extreme rays of derivative relaxations are either extreme rays of the original cone or extreme rays of multiplicity one.

Corollary 6 *Assume that Λ_+ is pointed and $\deg(p) \geq 3$, and let Λ'_+ be the derivative relaxation associated to any $\tilde{e} \in \Lambda_{++}(p, e)$. If x generates an extreme ray of Λ'_+ then either $\text{mult}_{D_{\tilde{e}}p}(x) = 1$ or x generates an extreme ray of Λ_+ and $\text{mult}_p(x) \geq 3$.*

Proof As Λ_+ is pointed and $\deg(p) \geq 3$ it follows from Corollary 5 that Λ'_+ is pointed. If x generates an extreme ray of Λ'_+ and $\text{mult}_{D_{\tilde{e}}p}(x) \geq 2$ then, by Theorem 5, we can conclude that $\text{mult}_p(x) \geq 3$ and $x \in \Lambda_+$. Since $x \in \Lambda_+ \supseteq \Lambda'_+$ and x generates an extreme ray of Λ'_+ , it follows that x generates an extreme ray of Λ_+ . \square

5.2 Tangent cones and derivative relaxations

In this section we study tangent cones of hyperbolicity cones, and in particular how tangent cones change when we take derivative relaxations. We first show that the tangent cone of a hyperbolicity cone $\Lambda_+(p, e)$ at a point x is again a hyperbolicity cone (Theorem 6) and that the corresponding hyperbolic polynomial is the localization of p at x (Definition 7). The main result of the section (Theorem 7) is that the tangent cone to Λ'_+ at a boundary point x is the corresponding derivative relaxation of the tangent cone to Λ_+ at that same point x . This is the key technical result that enables us to understand how k -Terracini convexity is affected by taking derivative relaxations (see Sect. 5.3).

Definition 7 If p is a hyperbolic polynomial with respect to e and with associated hyperbolicity cone Λ_+ , then the *localization of p at $x \in \Lambda_+$* is the polynomial of

degree $\text{mult}_p(x)$ defined by

$$\text{Loc}_x(p)(y) = \lim_{\lambda \rightarrow \infty} \lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}y) = \lim_{\lambda \rightarrow \infty} \lambda^{-\text{rank}_p(x)} p(\lambda x + y).$$

Example 12 Let $p(X) = \det(X)$ where X is a $d \times d$ symmetric matrix of indeterminates, and let $e = I$. The corresponding hyperbolicity cone is the cone of $d \times d$ positive semidefinite matrices. Suppose that $X = \begin{bmatrix} Z & 0 \\ 0 & 0 \end{bmatrix}$ where Z is $k \times k$ and positive definite. Then, by the formula for the determinant of a block matrix in terms of the Schur complement,

$$\begin{aligned} \text{Loc}_X(p) \left(\begin{bmatrix} Y_{11} & Y_{12} \\ Y_{12}^T & Y_{22} \end{bmatrix} \right) &= \lim_{\lambda \rightarrow \infty} \lambda^{d-k} \det \left(\begin{bmatrix} Z + \lambda^{-1}Y_{11} & \lambda^{-1}Y_{12} \\ \lambda^{-1}Y_{12}^T & \lambda^{-1}Y_{22} \end{bmatrix} \right) \\ &= \lim_{\lambda \rightarrow \infty} \lambda^{d-k} \det(\lambda^{-1}Y_{22}) \det(Z + \lambda^{-1}Y_{11} - \lambda^{-1}Y_{12}Y_{22}^{-1}Y_{12}^T) \\ &= \det(Z) \det(Y_{22}). \end{aligned}$$

There is an alternative formulation of $\text{Loc}_x(p)$ in terms of directional derivatives of p in the x direction. This alternative formulation is particularly useful in understanding how derivative relaxations interact with localization. In the forthcoming discussion, we refer on several occasions to higher-order directional derivatives of a hyperbolic polynomial, which we denote as a composition of first-order directional derivatives as $D_{y^{(k)}} \cdots D_{y^{(1)}} p$; if the directions $y^{(1)}, \dots, y^{(k)}$ are the same, we denote the associated higher-order directional derivative in a more compact manner as $D_y^k p$.

Lemma 7 *If p is a hyperbolic polynomial with respect to e then*

$$\text{Loc}_x(p)(y) = \frac{1}{\text{mult}_p(x)!} D_y^{\text{mult}_p(x)} p(x).$$

Proof By a Taylor expansion,

$$p(x + \lambda^{-1}y) = \sum_{k=0}^{\text{deg}(p)} \frac{\lambda^{-k}}{k!} D_y^k p(x).$$

Since p vanishes to order $\text{mult}_p(x)$ as $\lambda \rightarrow \infty$, and $\text{mult}_p(x)$ is independent of the choice of e in the interior of the hyperbolicity cone, it follows that $D_y^k p(x) = 0$ whenever $y \in \text{int}(\Lambda_+(p, e))$ and $0 \leq k < \text{mult}_p(x)$. As such, if $k < \text{mult}_p(x)$ then $y \mapsto D_y^k p(x)$ is a polynomial that vanishes on the interior of the (full-dimensional) hyperbolicity cone, so it must be identically zero. Hence

$$\lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}y) = \sum_{k=\text{mult}_p(x)}^{\text{deg}(p)} \frac{\lambda^{\text{mult}_p(x)-k}}{k!} D_y^k p(x).$$

Taking the limit as $\lambda \rightarrow \infty$ we obtain the stated result. □

We now consider localization of a hyperbolic polynomial from the point of view of its zeros. To do so, we use the following basic fact about how hyperbolic eigenvalues change along different directions.

Lemma 8 ([3, Lemma 3.27]) *Suppose p is hyperbolic with respect to e . If $x, u \in \mathbb{R}^n$ then*

$$p(x - te + su) = p(e) \prod_{i=1}^{\deg(p)} (t_i(s; x, u) - t)$$

where the functions $s \mapsto t_i(s; x, u)$ are real analytic functions of s . Furthermore, if $u \in \Lambda_+(p, e)$ then $t'_i(s; x, u) := \frac{d}{ds}t_i(s; x, u) \geq 0$ for all s .

The roots of the polynomial $t \mapsto p(x - te + su)$ are the eigenvalues of $x + su$, and therefore each $t_i(s; x, u)$ in the above lemma is an eigenvalue of $x + su$. The assertion that the functions $s \mapsto t_i(s; x, u)$ are real analytic functions of s corresponds to the eigenvalues of $x + su$ being analytic functions of s , and the nonnegativity of each of the derivatives $t'_i(s; x, u)$ (when $u \in \Lambda_+(p, e)$) corresponds to each of the eigenvalues of $x + su$ being non-decreasing functions of s . This result is useful because it allows us to understand localization from the point of view of eigenvalues.

Lemma 9 *Suppose p is hyperbolic with respect to e and fix some $x \in \Lambda_+(p, e)$. Letting $m = \text{mult}_p(x)$ we have that*

$$\text{Loc}_x(p)(y - te) = p(e) \prod_{i=1}^{\deg(p)-m} \lambda_i(x) \prod_{j=\deg(p)-m+1}^{\deg(p)} (t'_j(0; x, y) - t). \tag{20}$$

Proof If $x \in \Lambda_+(p, e)$ has multiplicity $m := \text{mult}_p(x)$ then the functions t_i in the factorization of Lemma 8 have the property that $t_1(0; x, u) = \dots = t_{\deg(p)-m}(0; x, u) > 0$ and $t_{\deg(p)-m+1}(0; x, u), \dots, t_{\deg(p)}(0; x, u) = 0$ by virtue of $t_i(0; x, u)$ being eigenvalues of x . Using the factorization of Lemma 8, we see that

$$\begin{aligned} &\lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}(y - te)) \\ &= p(e) \prod_{i=1}^{\deg(p)-m} (t_i(\lambda^{-1}; x, y) - \lambda^{-1}t) \prod_{j=\deg(p)-m+1}^{\deg(p)} (\lambda t_j(\lambda^{-1}; x, y) - t). \tag{21} \end{aligned}$$

Expanding $t_i(\lambda^{-1}; x, y)$ about $t_i(0; x, y)$, gives $t_i(\lambda^{-1}; x, y) = \lambda_i(x) + \lambda^{-1}t'_i(0; x, y) + O(\lambda^{-2})$. We obtain (20) by taking the limit as $\lambda \rightarrow \infty$. □

We are interested in the localization of a hyperbolic polynomial at a point because it turns out to be the algebraic analogue of the geometric operation of taking the tangent cone to a hyperbolicity cone at a point. Although this is probably well-known, we have included a proof because we had difficulty finding an explicit statement of this type in the literature.

We now show that localization at x is the algebraic analog of the tangent cone to the hyperbolicity cone at x .

Theorem 6 *If p is hyperbolic with respect to e and $x \in \Lambda_+(p, e)$ then*

1. $\text{Loc}_x(p)$ is hyperbolic with respect to e ; and
2. $\Lambda_+(\text{Loc}_x(p), e) = \overline{\mathcal{K}_{\Lambda_+(p,e)}(x)}$ is the tangent cone of $\Lambda_+(p, e)$ at x .

Proof The fact that the localization is hyperbolic with respect to e is exactly [3, Lemma 3.42], and also follows immediately from Lemma 9 and the fact that the $t'_i(0; x, y)$ are always real.

For the second part, we first show that the hyperbolicity cone of the localization at x is contained in the tangent cone of $\Lambda_+(p, e)$ at x . Let $z \in \Lambda_{++}(\text{Loc}_x(p), e)$ be in the interior of the hyperbolicity cone of the localization at x . Then, from (20) we know that $t'_j(0; x, z) > 0$ for $j = \text{deg}(p) - \text{mult}_p(x) + 1, \dots, \text{deg}(p)$. Furthermore, since $x \in \Lambda_+(p, e)$ we know that $t_i(0; x, z) = \lambda_i(x) > 0$ for $i = 1, \dots, \text{mult}_p(x)$. From (21) we know that the roots of $t \mapsto \lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}(z - te))$ are $\lambda t_i(\lambda^{-1}; x, z) = \lambda \lambda_i(x) + t'_i(0; x, z) + O(\lambda^{-1})$ for $i = 1, \dots, \text{deg}(p)$. As such, there exists a sufficiently large positive λ_0 such that if $\lambda \geq \lambda_0$ then all of these roots are positive. Hence we have that $x + z/\lambda_0 \in \Lambda_{++}(p, e)$ and therefore $z \in \mathcal{K}_{\Lambda_+(p,e)}(x)$, the cone of feasible directions with respect to $\Lambda_+(p, e)$ at x . We have shown that $\Lambda_{++}(\text{Loc}_x(p), e) \subseteq \mathcal{K}_{\Lambda_+(p,e)}(x)$. Taking closures shows that $\Lambda_+(\text{Loc}_x(p), e)$ is contained in the tangent cone of $\Lambda_+(p, e)$ at x .

For the reverse inclusion, suppose that $z \in \mathcal{K}_{\Lambda_+(p,e)}(x)$. In other words, there exists a sufficiently large positive λ_0 such that $x + \lambda^{-1}z \in \Lambda_+(p, e)$ for all $\lambda \geq \lambda_0$. Then

$$t \mapsto \lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}(z + te))$$

has nonnegative coefficients for all $\lambda \geq \lambda_0$. By continuity of the coefficients as functions of λ , it follows that

$$t \mapsto \lim_{\lambda \rightarrow \infty} \lambda^{\text{mult}_p(x)} p(x + \lambda^{-1}(z + te)) = \text{Loc}_x(p)(z + te)$$

has non-negative coefficients. Consequently, we have that $z \in \Lambda_+(\text{Loc}_x(p), e)$ and so the cone of feasible directions is contained in $\Lambda_+(\text{Loc}_x(p), e)$. Taking the closure shows that the tangent cone is contained in $\Lambda_+(\text{Loc}_x(p), e)$, completing the proof. \square

Example 13 (Example 12 continued) Suppose that $p(X) = \det(X)$ where X is a $d \times d$ symmetric matrix of indeterminates, $e = I$, and $Z = \begin{bmatrix} Z & 0 \\ 0 & 0 \end{bmatrix}$ where Z is $k \times k$ and positive definite. The hyperbolicity cone of $\text{Loc}_X(p)$ is

$$\Lambda_+(\text{Loc}_X(p), I) = \left\{ \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{12}^T & Y_{22} \end{bmatrix} : Y_{22} \geq 0 \right\}$$

which coincides with the tangent cone to the positive semidefinite cone at X .

The main technical result of this section, a fairly immediate corollary of Lemma 7, is that localization and taking derivatives commute.

Theorem 7 *If p is hyperbolic with respect to e and $x \in \Lambda_+(p, e)$ with $\text{mult}_p(x) \geq 1$, then $\Lambda_+(D_{\tilde{e}}\text{Loc}_x(p), e) = \Lambda_+(\text{Loc}_x(D_{\tilde{e}}p), e)$ for any $\tilde{e} \in \Lambda_{++}(p, e)$.*

Proof Let $m = \text{mult}_p(x) \geq 1$. On the one hand $\text{Loc}_x(p)(y) = \frac{1}{m!} D_y^m p(x)$. Differentiating in the direction \tilde{e} gives

$$\begin{aligned} D_{\tilde{e}} \text{Loc}_x(p)(y) &= \left. \frac{d}{dt} \frac{1}{m!} D_{y+t\tilde{e}}^m p(x) \right|_{t=0} \\ &= \left. \frac{d}{dt} \frac{1}{m!} \left[\sum_{i=0}^m \binom{m}{i} t^i D_{\tilde{e}}^i D_y^{m-i} p(x) \right] \right|_{t=0} \\ &= \frac{1}{(m-1)!} D_{\tilde{e}} D_y^{m-1} p(x). \end{aligned}$$

where we have used the fact that $D_{y^{(1)}} \cdots D_{y^{(m)}} p(x)$ is invariant under permutations of $y^{(1)}, \dots, y^{(m)}$. On the other hand x has multiplicity $m-1 \geq 0$ with respect to $D_{\tilde{e}} p$. As such

$$\text{Loc}_x(D_{\tilde{e}} p)(y) = \frac{1}{(m-1)!} D_y^{m-1} D_{\tilde{e}} p(x).$$

We have shown that $D_{\tilde{e}} \text{Loc}_x(p) = \text{Loc}_x(D_{\tilde{e}} p)$, from which the result directly follows. \square

The fact that localization and taking derivatives commute tells us that the convex tangent space of a hyperbolicity cone is exactly the same as the convex tangent space of its derivative relaxation at points of high enough multiplicity.

Corollary 7 *If p is hyperbolic with respect to e , $x \in \Lambda_+$, and $\text{mult}_p(x) \geq 3$, then $\mathcal{L}_{\Lambda'_+}(x) = \mathcal{L}_{\Lambda_+}(x)$ for any derivative relaxation $\Lambda'_+ = \Lambda_+(D_{\tilde{e}} p, e)$ for $\tilde{e} \in \Lambda_{++}(p, e)$.*

Proof From Theorem 6, the convex tangent space of Λ_+ at x is the lineality space of $\Lambda_+(\text{Loc}_x(p), e)$. Similarly, the convex tangent space of Λ'_+ at x is the lineality space of $\Lambda_+(\text{Loc}_x(D_{\tilde{e}} p), e)$, which is the lineality space of $\Lambda_+(D_{\tilde{e}} \text{Loc}_x(p), e)$ from Theorem 7. Since $\deg(\text{Loc}_x(p)) = \text{mult}_p(x) \geq 3$ Corollary 5 tells us that the lineality space of $\Lambda_+(\text{Loc}_x(p), e)$ is equal to the lineality space of $\Lambda_+(D_{\tilde{e}} \text{Loc}_x(p), e)$, completing the proof. \square

5.3 Derivative relaxations of Terracini convex hyperbolicity cones

In this section we state and prove two results related to Terracini convexity properties of derivative relaxations of hyperbolicity cones. The first, Proposition 7, gives a sufficient condition under which any derivative relaxation of a hyperbolicity cone that is k -Terracini convex also has non-trivial Terracini convexity properties. It is, *a priori*, unclear whether the hypotheses of Proposition 7 hold for any interesting examples. In the main result of this section (Theorem 8), we show that if Λ_+ is a hyperbolicity cone that is Terracini convex and for which all of its extreme rays have hyperbolic rank one, then repeatedly taking derivative relaxations produces new examples of hyperbolicity cones with Terracini convexity properties. Examples of hyperbolicity cones to which

Theorem 8 applies are the nonnegative orthant and the positive semidefinite cone, as well as other examples such as the cone of $d \times d$ positive semidefinite Hankel matrices.

Proposition 7 *Suppose that p is hyperbolic with respect to e , the degree $\deg(p) \geq 3$, and the associated hyperbolicity cone Λ_+ is pointed and k -Terracini convex. If each collection $x^{(1)}, \dots, x^{(k')}$ of k' extreme rays of Λ_+ satisfies one of the following conditions:*

- there exists j such that $\text{mult}_p(x^{(j)}) \leq 2$, or
- $\text{mult}_p\left(\sum_{i=1}^{k'} x^{(i)}\right) \geq 3$,

then any derivative relaxation $\Lambda'_+ = \Lambda_+(D_{\tilde{e}}p, e)$ for $\tilde{e} \in \Lambda_{++}(p, e)$ is $\min\{k, k'\}$ -Terracini convex.

Remarks: The case $\deg(p) = 1$ is vacuous as Λ_+ is a halfspace and Terracini convexity requires a cone to be pointed. For similar reasons, the case $\deg(p) = 2$ is not interesting as $\deg(D_{\tilde{e}}p) = 1$ and Λ'_+ is a halfspace.

Proof Let $\ell = \min\{k, k'\}$ and let $x^{(1)}, \dots, x^{(\ell)}$ be extreme rays of Λ'_+ (note that Λ'_+ is pointed as Λ_+ is pointed and $\deg(p) \geq 3$). We consider next two cases based on the multiplicities of the $x^{(j)}$'s with respect to the derivative polynomial $D_{\tilde{e}}p$.

Case 1: Assume that there exists j such that $\text{mult}_{D_{\tilde{e}}p}(x^{(j)}) = 1$. In this case, the localization of $D_{\tilde{e}}p$ at $x^{(j)}$ has degree one, which implies that $\Lambda_+(\text{Loc}_x(D_{\tilde{e}}p), e)$ is a halfspace; therefore, from the second part of Theorem 6, the convex tangent space of Λ'_+ at $x^{(j)}$ is a subspace of codimension one. If all of $x^{(1)}, \dots, x^{(\ell)}$ generate the same extreme ray, then so does $\sum_{i=1}^{\ell} x^{(i)}$. This means that all of the convex tangent spaces of Λ'_+ at these points are the same, so certainly $\mathcal{L}_{\Lambda'_+}\left(\sum_{i=1}^{\ell} x^{(i)}\right) = \sum_{i=1}^{\ell} \mathcal{L}_{\Lambda'_+}(x^{(i)})$. Otherwise there is some $x^{(j')}$ that generates an extreme ray that is distinct from $x^{(j)}$. Since $\mathcal{L}_{\Lambda'_+}(x^{(j)}) \cap \Lambda'_+$ exposes the extreme ray generated by $x^{(j)}$ (from Lemma 2, as hyperbolicity cones are facially exposed [17, Theorem 23]), it follows that $x^{(j')} \notin \mathcal{L}_{\Lambda'_+}(x^{(j)})$. Since the convex tangent space of Λ'_+ at $x^{(j)}$ has codimension one and does not contain $x^{(j')}$,

$$\begin{aligned} \sum_{i=1}^{\ell} \mathcal{L}_{\Lambda'_+}(x^{(i)}) &\supseteq \mathcal{L}_{\Lambda'_+}(x^{(j')}) + \mathcal{L}_{\Lambda'_+}(x^{(j)}) \\ &\supseteq \text{span}(x^{(j')}) + \mathcal{L}_{\Lambda'_+}(x^{(j)}) = \mathbb{R}^n \supseteq \mathcal{L}_{\Lambda'_+}\left(\sum_{i=1}^{\ell} x^{(i)}\right). \end{aligned}$$

Case 2: Assume that $\text{mult}_{D_{\tilde{e}}p}(x^{(i)}) \geq 2$ for all $i = 1, 2, \dots, \ell$. From Corollary 6, it follows that $\text{mult}_p(x^{(i)}) \geq 3$ for all $i = 1, 2, \dots, \ell$ and that the $x^{(i)}$ all generate extreme rays of Λ_+ . As $\ell \leq k'$ and by our assumption on the extreme rays of Λ_+ , it follows that $\text{mult}_p\left(\sum_{i=1}^{\ell} x^{(i)}\right) \geq 3$. Then

$$\mathcal{L}_{\Lambda'_+} \left(\sum_{i=1}^{\ell} x^{(i)} \right) = \mathcal{L}_{\Lambda_+} \left(\sum_{i=1}^{\ell} x^{(i)} \right) = \sum_{i=1}^{\ell} \mathcal{L}_{\Lambda_+} (x^{(i)}) = \sum_{i=1}^{\ell} \mathcal{L}_{\Lambda'_+} (x^{(i)}). \quad (22)$$

The first and third equalities in (22) follow from Corollary 7 together with the fact that $x^{(i)}$ (for each i) and $\sum_{i=1}^{\ell} x^{(i)}$ have multiplicity at least three with respect to p . The second equality in (22) follows from the fact that Λ_+ is k -Terracini convex and $\ell \leq k$. \square

While Proposition 7 may appear rather technical, it is useful because it applies when we repeatedly take derivative relaxations. Indeed, we have as an immediate consequence that for a hyperbolic polynomial p with $\deg(p) = 3$, if Λ_+ is Terracini convex then so is any derivative relaxation Λ'_+ ; this follows from the observation that multiplicity of any generator of an extreme ray of Λ_+ is at most two. For higher-degree hyperbolic polynomials, we present next the main result of this section which shows that for Terracini convex hyperbolicity cones with all the extreme rays having hyperbolic rank one, the derivative relaxations yield new hyperbolicity cones with non-trivial Terracini convexity properties. Recall that the rank of a point with respect to a hyperbolic polynomial is the number of non-zero eigenvalues, and that a cone is Terracini convex if it is k -Terracini convex for all k .

Theorem 8 *Let p be hyperbolic with respect to e and let $d = \deg(p)$ with $d > 3$. Suppose that $\Lambda_+(p, e)$ is pointed and Terracini convex and that whenever x generates an extreme ray of $\Lambda_+(p, e)$ then $\text{rank}_p(x) = 1$. If $1 \leq \ell \leq d - 3$ is a positive integer and $e^{(1)}, \dots, e^{(\ell)} \in \Lambda_{++}(p, e)$ then $\Lambda_+(D_{e^{(\ell)}} D_{e^{(\ell-1)}} \cdots D_{e^{(1)}} p, e)$ is $(d - \ell - 2)$ -Terracini convex.*

Proof For brevity of notation, we write $p^{(i)} = D_{e^{(i)}} \cdots D_{e^{(1)}} p$ and $\Lambda_+^{(i)} = \Lambda_+(p^{(i)}, e)$ for $i = 1, \dots, \ell$.

It is helpful in our proof to use the observation that any x that generates an extreme ray of $\Lambda^{(\ell)}$ either generates an extreme ray of $\Lambda_+ := \Lambda_+(p, e)$ or satisfies $\text{mult}_{p^{(\ell)}}(x) = 1$. We show both this secondary result as well as the primary result via induction.

For the base case of the secondary result, note that if x generates an extreme ray of $\Lambda_+^{(1)}$ then by Corollary 6 either $\text{mult}_{p^{(1)}}(x) = 1$ or x is an extreme ray of Λ_+ (with $\text{mult}_p(x) \geq 3$). For the primary result, note that Λ_+ is Terracini convex. If $x^{(1)}, \dots, x^{(d-3)}$ are extreme rays of Λ_+ , then their sum has rank at most $d - 3$ (since the hyperbolic rank function is subadditive [2]) and hence has multiplicity at least three with respect to p . It follows from Proposition 7 that $\Lambda_+^{(1)}$ is $(d - 3)$ -Terracini convex.

For the inductive hypothesis of the secondary result, assume that if x generates an extreme ray of $\Lambda_+^{(\ell-1)}$ then either x generates an extreme ray of Λ_+ or $\text{mult}_{p^{(\ell-1)}}(x) = 1$. For the primary result, assume that $\Lambda_+^{(\ell-1)}$ is $(d - \ell - 1)$ -Terracini convex.

We now establish the inductive step for the secondary result. If x generates an extreme ray of $\Lambda_+^{(\ell)}$ then by Corollary 6 either $\text{mult}_{p^{(\ell)}}(x) = 1$ or x generates an extreme ray of $\Lambda_+^{(\ell-1)}$ with $\text{mult}_{p^{(\ell-1)}}(x) \geq 3$, and so by the inductive hypothesis is an extreme ray of Λ_+ .

Finally we establish the inductive step of the primary result by applying Proposition 7. Let $x^{(1)}, \dots, x^{(d-\ell-2)}$ be extreme rays of $\Lambda_+^{(\ell-1)}$. Assume that each $x^{(i)}$ has multiplicity at least three with respect to $p^{(\ell-1)}$ (otherwise we are done). Based on the inductive hypothesis, each $x^{(i)}$ must be an extreme ray of Λ_+ and so must have rank one with respect to p by assumption. Then $\sum_{i=1}^{d-\ell-2} x^{(i)}$ has rank at most $d - \ell - 2$, and hence multiplicity at least $\ell + 2$, with respect to p . By applying Theorem 5 $\ell - 1$ times, and noting that $p, p^{(1)}, \dots, p^{(\ell-1)}$ all have degree at least three, we see that $\sum_{i=1}^{d-\ell-2} x^{(i)}$ has multiplicity at least $\ell + 2 - (\ell - 1) = 3$ with respect to $p^{(\ell-1)}$. Then, by Proposition 7, $\Lambda_+^{(\ell)}$ is $(d - \ell - 2)$ -Terracini convex. \square

We conclude by discussing three concrete special cases of Theorem 8.

Example 14 (Hyperbolicity cones associated with permanents) If $p(x) = \prod_{i=1}^d x_i$ and e is the vector of all ones, then the corresponding hyperbolicity cone is the nonnegative orthant. This is Terracini convex and all of its extreme rays have rank one. In this case, if $e^{(1)}, \dots, e^{(\ell)} \in \mathbb{R}_{++}^d$, then $D_{e^{(\ell)}} \cdots D_{e^{(1)}} p(x)$ is the *permanent* of the $d \times d$ matrix with columns $e^{(1)}, \dots, e^{(\ell)}$ and $d - \ell$ copies of x . Theorem 8 then tells us that the hyperbolicity cone associated with this permanent is $(d - \ell - 2)$ -Terracini convex as long as $1 \leq \ell \leq d - 3$.

Example 15 [Hyperbolicity cones associated with mixed discriminants] If $p(X) = \det(X)$ and e is the identity matrix, then the corresponding hyperbolicity cone is the positive semidefinite cone. This is Terracini convex and all of its extreme rays have rank one. In this case, if $E^{(1)}, \dots, E^{(\ell)}$ are positive definite matrices then the quantity $D_{E^{(\ell)}} \cdots D_{E^{(1)}} p(X)$ is known as the *mixed discriminant* of the d -tuple of matrices $(E^{(1)}, \dots, E^{(\ell)}, X, \dots, X)$. Theorem 8 then tells us that the hyperbolicity cone associated with this mixed discriminant is $(d - \ell - 2)$ -Terracini convex as long as $1 \leq \ell \leq d - 3$.

Example 16 (Hyperbolicity cones associated with mixed discriminants of Hankel matrices) Consider the cone \mathcal{H}_{d+1} of $(d + 1) \times (d + 1)$ symmetric positive semidefinite Hankel matrices. This can be viewed as the hyperbolicity cone associated with the determinant restricted to the $2d + 1$ -dimensional subspace of Hankel matrices. Its extreme rays have the form

$$\phi_{2,d}(x, y)\phi_{2,d}(x, y)' = \begin{bmatrix} x^d \\ x^{d-1}y \\ \vdots \\ xy^{d-1} \\ y^d \end{bmatrix} \begin{bmatrix} x^d & x^{d-1}y & \dots & xy^{d-1} & y^d \end{bmatrix}$$

and are rank one as symmetric matrices, and therefore have rank one with respect to the determinant polynomial. The cone \mathcal{H}_{d+1} is also linearly isomorphic to the cone $\mathcal{C}_{2,2d}$ over the homogeneous moment curve of degree $2d$, which is Terracini convex from Corollary 4. As such, if we choose $E^{(1)}, \dots, E^{(\ell)}$ to be positive definite $(d + 1) \times (d + 1)$ Hankel matrices and if $1 \leq \ell \leq d - 2$, then the mixed discriminant

of $(E^{(1)}, \dots, E^{(\ell)}, X, \dots, X)$ restricted to Hankel matrices X yields an associated hyperbolicity cone that is $(d - 1 - \ell)$ -Terracini convex.

6 Discussion

In this paper we introduced the notion of Terracini convex cones, generalizing the notion of neighborly polyhedral cones to the non-polyhedral setting in a way that includes examples such as the positive semidefinite cone and the cone over the moment curve. This suggests the pursuit of a broader program that seeks to extend key notions from polyhedral combinatorics to more general convex cones.

Explicit constructions A significant feature of the literature on neighborly polytopes—arguably, a principle reason for considering such polytopes in the first place—is that they offer examples of various extremal polyhedral constructions. Obtaining similar constructions with non-polyhedral Terracini convex cones would offer an interesting point of comparison with the polyhedral case. For example, we are not aware whether the non-degeneracy and regularity conditions of Sect. 3 are necessary to conclude that k -neighborly cones are k -Terracini convex, and identifying potential counterexamples would provide an interesting extremal class of convex cones. In a different direction, explicit constructions for linear images of the positive-semidefinite cone that are Terracini-convex would immediately yield explicit (non-random) families of linear maps for which the associated low-rank inverse problems considered in Sect. 4.2 may be solved exactly via semidefinite programming; despite significant attention devoted to this question, we are not aware of any such families of linear maps.

Beyond generalizing neighborliness A simplicial polytope is one in which every proper face is a simplex. In the spirit of this paper, a natural analogue in the non-polyhedral conic setting would be a closed pointed convex cone for which every proper face is Terracini convex. Let us call such convex cones *boundary Terracini convex*. Clearly the cone over any simplicial polytope is boundary Terracini convex, but boundary Terracini convex cones are a much richer class. One interesting example is the epigraph of the nuclear norm, i.e., $\{(X, t) \in \mathbb{R}^{m \times m} \times \mathbb{R} : \|X\|_* \leq t\}$. One can check that all of the proper faces of this convex cone are linearly isomorphic to positive semidefinite cones. Moreover, we can deduce from [17, Corollary 17] that if Λ_+ is a hyperbolicity cone that is boundary Terracini convex, then so are derivative relaxations of Λ_+ (as long as they are pointed). It would be interesting to study such boundary Terracini convex cones in more detail.

Weaker notions of Terracini convexity The key condition (2) in the definition of k -Terracini convexity is required to hold for every subset of at most k extreme rays. It is natural to consider weaker notions of k -Terracini convexity that only require (2) to hold for ‘many’ subsets of at most k extreme rays. By ruling out certain explicit configurations of k extreme rays such a definition would generalize important existing variations on neighborliness, such as k -neighborly centrally symmetric polytopes (in which subsets of k extreme points containing an antipodal pair are excluded). Another approach would be to require that (2) hold for suitably generic subsets of at most k extreme rays. Seeking and studying examples of convex cones that are generically

k -Terracini convex but not k -Terracini convex, would lead to a deeper understanding of Terracini convexity and its variants.

Possible further constructions of k -Terracini convex cones We have seen that the positive semidefinite cone and the cone over the moment curve (or equivalently the cone of Hankel positive semidefinite matrices) are Terracini convex. These are both examples of spectrahedral cones (intersections of a positive semidefinite cone with a subspace) with all extreme rays having rank one. This very special class of spectrahedral cones were classified by Blehkerman, Sinn, and Velasco [6] and are closely connected to questions about the relationship between nonnegative polynomials and sums of squares. It would be interesting to investigate the Terracini convexity properties of spectrahedral cones with only rank one extreme rays. Going one step further, one could similarly investigate the Terracini convexity properties of hyperbolicity cones with only (hyperbolic) rank one extreme rays. Unlike the spectrahedral setting, we are not aware of any nontrivial characterization of this class of convex cones.

Theorem 4 shows that with high probability, Gaussian random linear images of the positive semidefinite cone are k -Terracini convex, for a suitable k . The specific properties of the positive semidefinite cone are only used in isolated places in the argument, and do not seem to be essential. It is plausible that there an analogue of Theorem 4 where the positive semidefinite cone is replaced with any Terracini convex hyperbolicity cone, or perhaps even any Terracini convex cone. If this were the case, it would be a substantial further generalization of the fact that Gaussian random linear images of the simplex are k -neighborly polytopes, for suitable k [9]. It would also suggest the broader applicability of the notion of Terracini convexity for understanding convex relaxations of inverse problems.

Obstructions to lifts of convex sets Another setting in which neighborliness is useful, and Terracini convexity may find applications, is in the study of lifted representations of convex sets. Given a convex set \mathcal{C} and a closed convex cone \mathcal{H} , we say that \mathcal{C} has a \mathcal{H} -lift if we can express \mathcal{C} as the linear projection of an affine slice of \mathcal{H} . Such representations of convex sets are of importance when convex optimization problems are expressed in conic form. In particular, they play a prominent role in the study of the expressive power of linear, second-order cone, and semidefinite programming of a given size (see, e.g., [10] for a recent survey). It turns out if a convex set satisfies a notion of k -neighborliness that is somewhat weaker than that studied in this paper, then it cannot have a \mathcal{H} -lift where \mathcal{H} is a product of finitely many $k \times k$ positive semidefinite cones [4]. A natural extension of this line of inquiry would be to investigate whether Terracini convexity properties also provide obstructions to the existence of certain lifted representations of convex sets.

Acknowledgements The authors would like to thank Rainer Sinn for helpful conversations, and the anonymous referees for their useful comments and suggestions.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included

in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Amelunxen, D., Lotz, M., McCoy, M., Tropp, J.: Living on the edge: phase transitions in convex programs with random data. *Inf. Inference J. IMA* **3**(3), 224–294 (2014)
2. Amini, N., Brändén, P.: Non-representable hyperbolic matroids. *Adv. Math.* **334**, 417–449 (2018)
3. Atiyah, M.F., Bott, R., Gårding, L.: Lacunas for hyperbolic differential operators with constant coefficients I. *Acta Math.* **124**(1), 109–189 (1970)
4. Averkov, G.: Optimal size of linear matrix inequalities in semidefinite approaches to polynomial optimization. *SIAM J. Appl. Algebra Geom.* **3**(1), 128–151 (2019)
5. Blekherman, G.: Dimensional differences between nonnegative polynomials and sums of squares (2009). arXiv preprint [arXiv:0907.1339](https://arxiv.org/abs/0907.1339)
6. Blekherman, G., Sinn, R., Velasco, M.: Do sums of squares dream of free resolutions? *SIAM J. Appl. Algebra Geom.* **1**(1), 175–199 (2017)
7. Candès, E., Plan, Y.: Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inf. Theory* **57**(4), 2342–2359 (2011)
8. Chandrasekaran, V., Recht, B., Parrilo, P.A., Willsky, A.S.: The convex geometry of linear inverse problems. *Found. Comput. Math.* **12**(6), 805–849 (2012)
9. Donoho, D.L., Tanner, J.: Sparse nonnegative solution of underdetermined linear equations by linear programming. *Proc. Natl. Acad. Sci.* **102**(27), 9446–9451 (2005)
10. Fawzi, H., Gouveia, J., Parrilo, P.A., Saunderson, J., Thomas, R.R.: Lifting for simplicity: Concise descriptions of convex sets (2020). arXiv preprint [arXiv:2002.09788](https://arxiv.org/abs/2002.09788)
11. Gårding, L.: An inequality for hyperbolic polynomials. *J. Math. Mech.* **8**, 957–965 (1959)
12. Gordon, Y.: On Milman's inequality and random subspaces which escape through a mesh in \mathbb{R}^n . *Geom. Asp. Funct. Anal. Isr. Semin.* **1986–1987**(1317), 84–106 (1988)
13. Grunbaum, B.: *Convex Polytopes*. Springer, Berlin (2003)
14. Kalai, G., Wigderson, A.: Neighborly embedded manifolds. *Discrete Comput. Geom.* **40**(3), 319–324 (2008)
15. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.* **52**(3), 471–501 (2010)
16. Renegar, J.: *A Mathematical View of Interior-Point Methods in Convex Optimization*. MOS-SIAM Series on Optimization, Lausanne (2001)
17. Renegar, J.: Hyperbolic programs, and their derivative relaxations. *Found. Comput. Math.* **6**(1), 59–79 (2006)
18. Rockafellar, R.T.: *Convex Analysis*. Princeton University Press, Princeton (2015)
19. Terracini, A.: Sulle v_k per cui la varietà degli S_h ($h + 1$)-seganti ha dimensione minore dell'ordinario. *Rendiconti del Circolo Matematico di Palermo* **31**, 392–396 (1911)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.