

On the Absence of Identifiable Manifolds in Finite-Max Composite Optimization

Yifan Wang^{1*}, Jianhao Ma^{2*} and Salar Fattahi^{1*}

¹Department of Industrial and Operations Engineering, University of
Michigan, Ann Arbor, MI, USA.

²Department of Industrial Engineering, Tsinghua University, Beijing,
China.

Abstract

In nonsmooth optimization, identifiable sets describe the local region eventually reached by sequences converging to a prescribed critical point. When such a set is a $\mathcal{C}^2$ manifold on which the objective restricts to a $\mathcal{C}^2$ function, it is called an identifiable manifold. Their appeal lies in what they enable: many first-order methods identify these manifolds in finitely many iterations, after which the iterates enter a region in which the problem is effectively smooth. Consequently, many powerful tools and guarantees from smooth optimization transplant naturally to the nonsmooth setting. Owing to these properties, much existing work has focused on characterizing conditions that guarantee their existence. In this work, we study a complementary question: under what conditions is a critical point devoid of any identifiable manifold? We answer this by developing a deterministic branching criterion for a broad class of finite-max composite optimization problems, characterizing when a critical point admits no identifiable manifold. This criterion is surprisingly mild in certain classes of problems: it holds with high probability for overparameterized robust low-rank recovery and almost surely at common interpolators of random minimax regression, suggesting that the absence of identifiable manifolds may be the rule rather than the exception in modern optimization.

Keywords: identifiable sets, active manifolds, nonsmooth optimization, strict saddle points, overparameterized robust low-rank optimization

1 Introduction

A recurring theme in nonsmooth optimization is that apparent nonsmoothness often conceals a smoother underlying structure. The idea is to identify a local set $\mathcal{M}$ near a critical point $\bar{x}$ that captures all nearly critical behavior: whenever $x_\nu \rightarrow \bar{x}$ and $v_\nu \in \partial f(x_\nu)$ satisfies $v_\nu \rightarrow 0$, the points x_ν eventually lie in $\mathcal{M}$. The most powerful case is when $\mathcal{M}$ is a smooth manifold on which f restricts smoothly, so that the local analysis reduces to smooth analysis on $\mathcal{M}$. Moreover, once a first-order method identifies $\mathcal{M}$, its subsequent iterates evolve on the manifold and effectively optimize the smooth restricted objective, enabling guarantees such as convergence to minimizers and avoidance of strict saddle points [1, 2]. This perspective underlies a long line of work on partial smoothness [3–5], active-set methods [6–8], sensitivity theory [9–11], and, more recently, active-strict-saddle theory for weakly convex optimization [12, 13].

The interest in identifiable manifolds is driven by what they make possible at a critical point. Identifiable manifolds enable smooth local analysis at critical points. Because the restriction $f|_{\mathcal{M}}$ is C^2 , conditions such as positive-definiteness of the restricted Hessian can certify local minimality, and the restriction becomes the natural object for sensitivity and tilt-stability analysis [3, 9, 14]. Proximal-based and dual averaging methods can identify the manifold in finitely many iterations and subsequently inherit smooth convergence rates [5, 15–17], while manifold-restricted Newton updates attain fast local convergence [3, 15]. Identifiable manifolds also underlie the theory of active strict saddles [12]: critical points lying on a C^2 identifiable manifold along which the objective decreases smoothly and quadratically in some tangent direction. Such points are provably avoided by several algorithms, including stochastic subgradient [13] and proximal-type methods [12].

Owing to these benefits, the existing literature has focused almost entirely on sufficient conditions for the *existence* of identifiable manifolds. The strongest such result is a genericity theorem: for almost every linear perturbation of a weakly convex semi-algebraic function, the perturbed function admits an identifiable manifold at each of its critical points [18, Corollary 4.8]. A critical point with no identifiable manifold thus appears to be a measure-zero event, removable by a generic perturbation, leading most existing theory to exclude such points altogether. This view is mathematically sound, but it may understate the practical cost of perturbing the objective: perturbation does not merely alter the local geometry around a critical point but reshapes the geometry globally. As the following example illustrates, even an arbitrarily small linear perturbation can destroy useful global structure; a bounded problem with a benign landscape can become unbounded below.

Example 1.1 Let $r := \sqrt{x^2 + y^2}$ and consider

$$h(x, y) = \chi(r) \left(5 + (|x| + |y|)^2 - 2x^2 \right),$$

where χ is the C^2 cutoff

$$\chi(r) = \begin{cases} 1, & 0 \leq r \leq \frac{1}{2}, \\ \omega\left(\frac{2r-1}{3}\right), & \frac{1}{2} < r < 2, \\ 0, & r \geq 2, \end{cases} \quad \omega(t) := 1 - 10t^3 + 15t^4 - 6t^5.$$




Fig. 1: Global and local landscape of the function h in Example 1.1. *Left:* the benign landscape with the localized nonsmooth landscape at the origin; the red box marks the region enlarged in the right panel. *Right:* the corresponding local landscape near the origin, where the origin is an inactive strict saddle. The two black curves mark the coordinate axes, the nonsmooth locus of the local term $|xy|$.

The local minimizers of h are exactly the points satisfying $x^2 + y^2 \geq 4$, all of which attain the global minimum value 0, while every other critical point, including the origin, is a strict saddle; see Figure 1a. Specifically, near the origin, $h(x, y) = 5 + y^2 - x^2 + 2|xy|$, so $h(t, 0) = 5 - t^2$, and hence the origin is a strict saddle. For every sequence $(x_k, y_k) \rightarrow (0, 0)$, one can choose subgradients $v_k \in \partial h(x_k, y_k)$ with $v_k \rightarrow (0, 0)$. Consequently, any identifiable set at the origin must contain a full neighborhood of the origin. Since h is not C^2 on any such neighborhood, the origin admits no C^2 identifiable manifold and is therefore an inactive strict saddle; see Figure 1b. Any linear perturbation, no matter how small, makes the problem unbounded below. To see this, for any $a \neq 0$, define $h_a(z) := h(z) - \langle a, z \rangle$. Since $h(ta) = 0$ for all sufficiently large t ,

$$h_a(ta) = -t\|a\|^2 \longrightarrow -\infty, \quad \text{as } t \rightarrow +\infty.$$

The above example is not an isolated case. Figure 2 shows that the same tension arises—more sharply on the algorithmic side—in an important class of problems: robust low-rank matrix sensing with overparameterized rank, details of which we defer to Section 5. There, the subgradient method converges to the true low-rank solution, which we show is a critical point admitting no identifiable manifold, yet convergence fails once the problem is subject to a small random perturbation.

This motivates the main question studied in this paper:




Fig. 2: Behavior of the subgradient method on overparameterized low-rank matrix sensing with ℓ_1 -loss. The ground-truth factor U^* , satisfying $U^*U^{*\top} = X^*$, is a critical point admitting no identifiable manifold. From local random initializations, the unperturbed method converges sublinearly to X^* and retains effective rank near 1.1, whereas a small random perturbation prevents convergence and consistently increases the effective rank. The relative recovery error and factor effective rank are $\|U_t U_t^\top - X^*\|_F / \|X^*\|_F$ and $\|U_t\|_F^2 / \|U_t\|_{\text{op}}^2$, respectively. Thin curves show the 50 paired trials, and thick curves show their coordinatewise medians. See Appendix C for details.

Under what conditions does the original, unperturbed problem fail to admit a C^2 -identifiable manifold at a critical point?

Beyond a few special cases, this question has not been studied systematically; see Section 1.1. We address it for finite-max composite problems, a broad class encompassing robust estimation with ℓ_1 -loss [19, 20], ReLU and hinge losses, finite distributionally robust optimization [21], and exact-penalty reformulations of constrained problems [22]. For this class, we provide deterministic geometric conditions, called the *branching criterion*, which when satisfied at a critical point rule out the existence of any identifiable manifold, and with it the algorithmic guarantees described above. We summarize the contributions of this paper below.

1. *A deterministic obstruction.* For finite-max composite optimization, we provide a branching criterion that, when satisfied at a given critical point, rules out the existence of any C^2 identifiable manifold (Theorem 3.9). The key challenge is that establishing the nonexistence of a C^2 identifiable manifold is a priori challenging, as one must rule out all possible manifolds simultaneously. We address this via a contrapositive argument: if a C^2 identifiable manifold exists, then a minimal one exists and must be locally unique. We explicitly characterize this hypothetical minimal manifold and show that the branching criterion precisely excludes it, thereby certifying nonexistence. This criterion admits a particularly

geometric interpretation for stratifiable functions: it requires that two strata of the identifiable set, reachable along a common set of tangent directions, induce different quadratic forms on those directions. The existence of such incompatible directions precludes an identifiable manifold from existing.

2. *Applications.* We show that the branching criterion is far from being pathological: it is satisfied with high probability or almost surely for several important statistical estimation problems, including overparameterized nonnegative sparse regression (Theorem 5.1), symmetric matrix sensing (Theorem 5.2), asymmetric matrix sensing (Theorem 5.3), and finite-scenario minimax optimization (Theorem 4.1 and Corollary 4.2). These results point to a common structural phenomenon across overparameterized learning tasks: critical points without identifiable manifolds may be universal rather than exceptional.

1.1 Related Work

Identifiability and partial smoothness

The idea that nonsmooth optimization problems often carry a hidden smooth structure goes back at least to the 1990s, with the identifiable surfaces introduced by Wright [6] for constrained optimization and the active-set identification results of Burke and Moré [7] and Burke [8], which showed that the optimal active set is correctly identified once iterates are close enough to a solution. Lewis [3] unified these threads under the notion of *partial smoothness*, isolating the geometric conditions—smoothness along a manifold and sharpness normal to it—that make such identification possible; Hare and Lewis [4, 5] then tied this geometry to concrete algorithms. Two parallel perspectives refine the same picture: the $\mathcal{U}$ – $\mathcal{V}$ decomposition of Lemaréchal et al. [23], which splits the space into directions of smoothness and of nonsmoothness, and the cone-reducibility theory of Bonnans and Shapiro [9], developed for sensitivity analysis. Drusvyatskiy and Lewis [14, 24] recast the theory in terms of *identifiable sets*, which need not be assumed manifolds at the outset. We adopt this point of view in our paper.

Active-set and active-manifold algorithms

The terms *active manifold* (from partial-smoothness theory) and *identifiable manifold* (from the identifiable-set theory) refer to the same object at a critical pair: Drusvyatskiy and Lewis [24, Proposition 10.8] show that the active manifold of a partly smooth function is identifiable, and conversely a C^2 identifiable manifold is active in the partial-smoothness sense. A parallel line of work studies algorithms that identify this manifold in finite time and then exploit its smooth structure. Classical results include those of Burke and Moré [7] and Calamai and Moré [25] for projected gradient methods on polyhedral constraints. In the nonsmooth setting, Hare and Lewis [5] showed that the proximal point method identifies the active manifold of a partly smooth function, and Lewis and Wright [15] extended this to proximal-gradient methods for composite problems. Liang et al. [16] established analogous identification results for forward-backward and primal-dual splitting methods, and Lee et al. [17] obtained identification guarantees for regularized dual averaging in stochastic online learning. All of these results presuppose the existence of a smooth active manifold;

the present paper studies the obstruction to that presupposition. In the context of constrained optimization, both the second-order methods of Wright [26] and Wright [27] and the more recent first-order methods of Díaz et al. [28] guarantee convergence to degenerate solutions where strict complementarity fails. Although such degenerate solutions are often associated with the absence of an identifiable manifold, this connection is not causal: degeneracy alone does not guarantee absence of an identifiable manifold, as we discuss next.

Absence of identifiable manifolds for composite functions

For a composite function $h = g \circ F$, the absence of identifiable manifolds is often associated with degeneracy of the critical points, namely when the Jacobian DF loses rank at a critical point [29]. This degeneracy is closely related to *overparameterization* in modern statistical models, especially factored low-rank optimization and rank-overspecified robust recovery [19, 30–34]. Local-search algorithms often converge to such degenerate points [29, 32–37]. However, rank deficiency of DF alone does not preclude the existence of an identifiable manifold for h : it only means that F cannot serve as a local defining map for such a manifold [38, Definition 3.1]. For instance, the x -axis $\mathcal{M} = \{(x, y) : y = 0\}$ is a smooth manifold with valid local defining map $h_1(x, y) := y$ with a full-rank Jacobian, satisfying $\mathcal{M} = h_1^{-1}(0)$. The map $h_2(x, y) := y^2$ also satisfies $\mathcal{M} = h_2^{-1}(0)$ near the origin, but its Jacobian vanishes there, disqualifying it as a local defining map. This shows that the rank deficiency of one representation of $\mathcal{M}$ does not preclude the existence of a valid one.

Ruling out a C^2 identifiable manifold is therefore considerably more delicate: one must exclude every valid local defining map, not merely show that the natural candidate F fails. To the best of our knowledge, the only existing sufficient conditions are first-order. Specifically, Drusvyatskiy and Lewis [24, Proposition 10.12] rule out an identifiable manifold when zero does not lie in the relative interior of the Fréchet subdifferential, while Drusvyatskiy et al. [18, Theorem 4.11] do so when the critical cone is not a linear subspace. These conditions are equivalent [39, Proposition 2.1]. In our setting, however, both conditions are too restrictive to be applicable: Proposition 5.15 shows that they fail with high probability in robust low-rank recovery. Thus, first-order geometry alone does not detect inactivity in our setting. The obstruction instead appears at second order: distinct accessible branches induce incompatible quadratic expansions that cannot arise from a single C^2 restriction. Capturing this incompatibility is the purpose of our branching criterion.

1.2 Organization

Section 2 reviews the necessary background and introduces finite-max composite problems. Section 3 characterizes their locally minimal identifiable sets and establishes the branching criterion (Theorem 3.9). Section 4 applies the branching criterion to common interpolators of finite-scenario minimax problems. Section 5 specializes this criterion to overparameterized robust low-rank recovery, reducing it to a geometric cone-selection test (Theorem 5.11) and deriving model-specific guarantees for non-negative sparse regression, symmetric matrix sensing, and asymmetric matrix sensing (Theorems 5.1, 5.2, and 5.3). Concluding remarks are provided in Section 6.

2 Notation and Basic Constructions

All vector and matrix spaces are finite-dimensional Euclidean spaces over $\mathbb{R}$, equipped with their standard inner products and induced norms. On matrix spaces, these are the Frobenius inner product and norm. We write $\|A\|_{\text{op}}$ for the operator norm. The space of $n \times n$ symmetric matrices is denoted by $\mathbb{S}^n$. For $X \in \mathbb{S}^n$, the notation $X \succ 0$ and $X \succeq 0$ means that X is positive definite and positive semidefinite, respectively. We write I_n for the $n \times n$ identity matrix and e_i for the i th standard basis vector.

For a positive integer m , let $[m] := \{1, \dots, m\}$. For $S \subseteq [m]$, we write $S^c := [m] \setminus S$ and use x_S to denote the restriction of a vector $x \in \mathbb{R}^m$ to the coordinates in S . The sign function $\text{sign}(x) \in \{-1, 0, +1\}^m$ is applied componentwise. The Hadamard product is denoted by $\odot$, while $\circ$ denotes function composition.

For Euclidean spaces $\mathcal{V}$ and $\mathcal{W}$ and a linear map $T : \mathcal{V} \rightarrow \mathcal{W}$, we write T^* , $\text{Im } T$, $\ker T$, and $\text{rank } T$ for its adjoint, range, nullspace, and rank, respectively. For a linear subspace $\mathcal{U} \subseteq \mathcal{V}$, we write $\mathcal{U}^\perp$ for its orthogonal complement and $P_{\mathcal{U}}$ for the orthogonal projection onto $\mathcal{U}$. For a matrix M , $\text{col } M$ denotes its column space, and $\text{span } \mathcal{S}$ denotes the linear span of a set $\mathcal{S} \subseteq \mathcal{V}$. For a set $\mathcal{X} \subseteq \mathbb{R}^n$, we write $\text{cl } \mathcal{X}$, $\text{relint } \mathcal{X}$, and $\text{conv } \mathcal{X}$ for its closure, relative interior, and convex hull, respectively. For a polytope $\mathcal{P}$, $\text{Vert}(\mathcal{P})$ denotes its set of vertices. Given a finite index set $\mathcal{I}$, the probability simplex on $\mathcal{I}$ is

$$\Delta_{\mathcal{I}} := \left\{ \lambda \in \mathbb{R}_{+}^{\mathcal{I}} : \sum_{i \in \mathcal{I}} \lambda_i = 1 \right\}, \quad \mathbb{R}_{+} := \{t \in \mathbb{R} : t \geq 0\}.$$

For $\lambda \in \mathbb{R}^{\mathcal{I}}$, write $\text{supp } \lambda := \{i \in \mathcal{I} : \lambda_i \neq 0\}$ and $\|\lambda\|_0 := |\text{supp } \lambda|$. For a function $h : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$, its domain, graph, and epigraph are

$$\begin{aligned} \text{dom } h &:= \{x \in \mathbb{R}^d : h(x) < +\infty\}, \\ \text{Gph } h &:= \{(x, h(x)) \in \mathbb{R}^d \times \mathbb{R} : x \in \text{dom } h\}, \\ \text{epi } h &:= \{(x, r) \in \mathbb{R}^d \times \mathbb{R} : r \geq h(x)\}. \end{aligned}$$

For a set-valued map $\Phi : \mathbb{R}^d \rightrightarrows \mathbb{R}^m$, we define $\text{Gph } \Phi := \{(x, v) \in \mathbb{R}^d \times \mathbb{R}^m : v \in \Phi(x)\}$.

We write $N(\mu, \Sigma)$ for the multivariate normal distribution with mean μ and covariance Σ , and $\text{GOE}(n)$ for the Gaussian Orthogonal Ensemble of $n \times n$ symmetric matrices with i.i.d. $N(0, 1/2)$ off-diagonal entries and i.i.d. $N(0, 1)$ diagonal entries. Probability is denoted by $\mathbb{P}$ and expectation by $\mathbb{E}$.

2.1 Manifolds and Definable Sets

We recall basic facts from differential geometry. For smooth manifolds, we follow Boumal [38] and Lee [40]; for o-minimal structures and stratifications, we follow Van den Dries [41] and Coste [42].

For a C^k map $F : \mathbb{R}^d \rightarrow \mathbb{R}^m$, we denote its derivative at x by a linear map $DF(x) : \mathbb{R}^d \rightarrow \mathbb{R}^m$. When $m = 1$, we identify $DF(x)$ with the gradient $\nabla F(x) \in \mathbb{R}^d$ and denote the Hessian by $\nabla^2 F(x)$. If F is C^2 , we write $D^2 F(x)[v, w]$ for its second

derivative evaluated along $v, w \in \mathbb{R}^d$ and $D^2F(x)[v, v]$ for the associated quadratic form along v .

A set $\mathcal{M} \subset \mathbb{R}^d$ is a C^k *embedded submanifold* (or simply a *manifold*) of dimension p around $\bar{x} \in \mathcal{M}$ if it is defined locally by a smooth nondegenerate equation: there is a neighborhood $\mathcal{U}$ of $\bar{x}$ and a C^k map $\Phi : \mathcal{U} \rightarrow \mathbb{R}^{d-p}$ with surjective derivative such that $\mathcal{M} \cap \mathcal{U} = \Phi^{-1}(0)$. Its *tangent* and *normal* spaces at $\bar{x}$ are then simply $T_{\bar{x}}\mathcal{M} := \ker D\Phi(\bar{x})$ and $N_{\bar{x}}\mathcal{M} := (T_{\bar{x}}\mathcal{M})^\perp$, both independent of the defining map Φ , and $P_{T_{\bar{x}}\mathcal{M}}$ denotes the orthogonal projection onto the tangent space.

Let h be C^1 on $\mathcal{M}$ near $\bar{x}$. Its *Riemannian gradient* is $\nabla_{\mathcal{M}}h(\bar{x}) := P_{T_{\bar{x}}\mathcal{M}}\nabla\tilde{h}(\bar{x})$, where $\tilde{h}$ is any C^1 extension of h ; this definition is independent of the extension. Thus, $\bar{x}$ is critical for $h|_{\mathcal{M}}$ when $\nabla_{\mathcal{M}}h(\bar{x}) = 0$. If $\mathcal{M}$ is C^2 , the second-order term of a local tangent-space parametrization captures its curvature, as formalized in the next lemma.

Lemma 2.1 (Second-order local parametrization, Chapter 5 of Boumal [38]) *Let $\mathcal{M}$ be a C^2 manifold around $\bar{x}$. There is a C^2 map γ from a neighborhood of 0 in $T_{\bar{x}}\mathcal{M}$ onto a neighborhood of $\bar{x}$ in $\mathcal{M}$, with $\gamma(0) = \bar{x}$, of the form*

$$\gamma(v) = \bar{x} + v + Z_{\mathcal{M}}(v) + o(\|v\|^2), \quad v \in T_{\bar{x}}\mathcal{M}, \quad (1)$$

where $Z_{\mathcal{M}} : T_{\bar{x}}\mathcal{M} \rightarrow N_{\bar{x}}\mathcal{M}$ is a homogeneous quadratic map.

Many sets in optimization are not manifolds, but their singularities often admit a structured decomposition into smooth manifolds, formalized by the notion of *stratification*.

Definition 2.2 (C^k stratification) A C^k *stratification* ($k \geq 1$) of a set $\mathcal{Q} \subset \mathbb{R}^d$ is a partition of $\mathcal{Q}$ into finitely many connected C^k manifolds, called *strata*, such that the closure of any stratum meets a second stratum only by containing it:

$$\text{for all strata } \mathcal{X}, \mathcal{Y} \subseteq \mathcal{Q} : \quad \mathcal{Y} \cap \text{cl } \mathcal{X} \neq \emptyset \implies \mathcal{Y} \subseteq \text{cl } \mathcal{X}.$$

The stratification is *compatible* with a family of sets if every stratum is either contained in or disjoint from each member of the family.

We focus on C^2 stratifications, henceforth called simply stratifications. To ensure their existence, we assume throughout that our objectives are definable in a fixed *o-minimal* structure. Such structures are closed under Boolean operations, products, and coordinate projections, while their definable subsets of $\mathbb{R}$ are precisely finite unions of points and intervals, ruling out pathologies such as infinite oscillation. They include the semialgebraic, globally subanalytic, and $\mathbb{R}_{\text{an}, \text{exp}}$ structures [41, 42]. A set is *definable* if it belongs to the chosen structure, and a map is definable if its graph is definable. All examples considered in this paper are semialgebraic and hence definable.

Theorem 2.3 (Theorem 1.3 and Corollary 1.11 of Loi [43]) *For every $k \geq 1$ and every finite family $\{\mathcal{Q}_1, \dots, \mathcal{Q}_p\}$ of definable subsets of $\mathbb{R}^d$, the union $\bigcup_{j=1}^p \mathcal{Q}_j$ admits a C^k stratification compatible with the family¹, in the sense of Definition 2.2.*

¹Loi [43, Theorem 1.3] shows the existence of C^k stratification for all of $\mathbb{R}^d$ compatible with the family $\{\mathcal{Q}_1, \dots, \mathcal{Q}_p\}$, which implies the existence of such stratification for the union $\bigcup_{j=1}^p \mathcal{Q}_j$.

We use the compatibility of this stratification in Section 3 to decompose an identifiable set into finitely many strata on which the objective is smooth.

2.2 Normal Cones and Subdifferentials

We next recall the variational objects used throughout; see [10, 11] for background.

Two normal cones to a closed set $\mathcal{Q} \subset \mathbb{R}^d$ at $\bar{x} \in \mathcal{Q}$ underlie the subdifferentials below. The *Fréchet normal cone* is defined as

$$\hat{N}_{\mathcal{Q}}(\bar{x}) := \left\{ v \in \mathbb{R}^d : \limsup_{x \rightarrow \bar{x}, x \in \mathcal{Q}} \frac{\langle v, x - \bar{x} \rangle}{\|x - \bar{x}\|} \leq 0 \right\}.$$

The *limiting normal cone* enlarges it by including the limits of nearby Fréchet normals:

$$N_{\mathcal{Q}}(\bar{x}) := \left\{ v : \exists x_\nu \rightarrow \bar{x}, v_\nu \rightarrow v, v_\nu \in \hat{N}_{\mathcal{Q}}(x_\nu) \right\}.$$

Indeed, $\hat{N}_{\mathcal{Q}}(\bar{x}) \subseteq N_{\mathcal{Q}}(\bar{x})$. Subgradients are defined through these cones via the epigraph. For lower semicontinuous $h : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ and $\bar{x} \in \text{dom } h$, the *Fréchet* and *limiting subdifferentials* are

$$\hat{\partial}h(\bar{x}) := \{v : (v, -1) \in \hat{N}_{\text{epi } h}(\bar{x}, h(\bar{x}))\}, \quad \partial h(\bar{x}) := \{v : (v, -1) \in N_{\text{epi } h}(\bar{x}, h(\bar{x}))\},$$

and the *horizon subdifferential* is $\partial^\infty h(\bar{x}) := \{v : (v, 0) \in N_{\text{epi } h}(\bar{x}, h(\bar{x}))\}$; we write $\text{Gph } \partial h$ for the graph of ∂h . For convex h both subdifferentials reduce to the convex subdifferential, and for C^1 functions to the gradient. The horizon subdifferential records recession directions of the subdifferential graph and vanishes whenever h is locally Lipschitz around $\bar{x}$.

Example 2.4 (Subdifferential of the ℓ_1 -norm) Let $h(z) = \|z\|_1 = \sum_{i=1}^m |z_i|$ on $\mathbb{R}^m$. Since the coordinates separate, the subdifferential is

$$\partial h(z) = \{y \in [-1, 1]^m : y_i = \text{sign}(z_i) \text{ for every } i \text{ with } z_i \neq 0\},$$

a face of the unit cube $[-1, 1]^m$ whose dimension equals the number of free coordinates of z . In particular, $\partial h(0) = [-1, 1]^m$ is the full cube, while $\partial h(z) = \{\text{sign}(z)\}$ is a single vertex when z has no zero entry. Since h is Lipschitz, $\partial^\infty h(z) = \{0\}$ for all z . This cube-face structure reappears in Section 5, where the multipliers at a critical point form a slice of the cube by an affine subspace.

A point $\bar{x} \in \text{dom } h$ is a *critical point* of h if $0 \in \partial h(\bar{x})$, and a *Fréchet-critical point* if $0 \in \hat{\partial}h(\bar{x})$; the function is *subdifferentially regular* at $\bar{x}$ if $\hat{\partial}h(\bar{x}) = \partial h(\bar{x})$. For locally Lipschitz functions, this is the standard Clarke regularity condition; in that case the Fréchet, limiting, and Clarke subdifferentials agree at $\bar{x}$. Regularity holds for every function class in this paper. A critical point $\bar{x}$ is a *strict saddle point* of h if, for some unit direction w , the standard second-order subderivative of h at $\bar{x}$ for the subgradient 0 satisfies $d^2h(\bar{x} | 0)(w) < 0$ [10, 12, 13].

2.3 Identifiable Sets and Identifiable Manifolds

We adopt the definitions of Drusvyatskiy and Lewis [14, 24] for identifiable sets and manifolds.

Definition 2.5 (Identifiable set) Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be locally Lipschitz, and let $(\bar{x}, \bar{v}) \in \text{Gph } \partial h$. A set $\mathcal{M} \subset \mathbb{R}^d$ containing $\bar{x}$ is identifiable for h at $\bar{x}$ for $\bar{v}$ if, for every sequence $x_\nu \rightarrow \bar{x}$ and every sequence $v_\nu \rightarrow \bar{v}$ with $v_\nu \in \partial h(x_\nu)$, all sufficiently large indices satisfy $x_\nu \in \mathcal{M}$.

Definition 2.6 (Locally minimal identifiable set) An identifiable set $\mathcal{M}$ for h at $\bar{x}$ for $\bar{v}$ is locally minimal if, for every other identifiable set $\mathcal{N}$ for the same pair, there exists a neighborhood $\mathcal{O}$ of $\bar{x}$ such that

$$\mathcal{M} \cap \mathcal{O} \subseteq \mathcal{N} \cap \mathcal{O}.$$

Definition 2.7 (C^2 identifiable manifold) Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be locally Lipschitz, and let $(\bar{x}, \bar{v}) \in \text{Gph } \partial h$. A set $\mathcal{M} \subset \mathbb{R}^d$ containing $\bar{x}$ is a C^2 identifiable manifold for h at $\bar{x}$ for $\bar{v}$ if it is identifiable for h at $\bar{x}$ for $\bar{v}$ and, in a neighborhood of $\bar{x}$, the set $\mathcal{M}$ is a C^2 manifold of $\mathbb{R}^d$ on which h restricts to a C^2 function.

Proposition 2.8 (Local minimality and local uniqueness, Proposition 8.2 of Drusvyatskiy and Lewis [14]) Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be lower-semicontinuous and subdifferentially regular at some critical point $\bar{x}$. If $\mathcal{M}$ is a C^2 -identifiable manifold for h at $\bar{x}$ for 0, then $\mathcal{M}$ is automatically a locally minimal identifiable set at $\bar{x}$ for 0.

This proposition is particularly useful in our analysis: any identifiable manifold, upon existence, must be determined *uniquely* by the locally minimal identifiable set.

Definition 2.9 (Inactive critical point) Let $0 \in \partial h(\bar{x})$. We say that $\bar{x}$ is inactive if no C^2 identifiable manifold exists for h at $\bar{x}$ for 0.

We next present a useful property of inactive critical points, the proof of which is presented in Appendix B.1.

Lemma 2.10 (Orthogonal invariance of criticality and identifiability) Let $\mathcal{X}$ be a finite-dimensional Euclidean space, let $h : \mathcal{X} \rightarrow \mathbb{R}$ be locally Lipschitz, and let $\mathcal{T} : \mathcal{X} \rightarrow \mathcal{X}$ be an orthogonal linear transformation satisfying $h \circ \mathcal{T} = h$. Then, for every $(x, v) \in \text{Gph } \partial h$, $\partial h(\mathcal{T}x) = \mathcal{T}(\partial h(x))$. Moreover, a set $\mathcal{M} \subset \mathcal{X}$ is identifiable for h at x for v if and only if $\mathcal{T}\mathcal{M}$ is identifiable for h at $\mathcal{T}x$ for $\mathcal{T}v$. The same equivalence holds for C^2 identifiable manifolds. Consequently, x is an inactive critical point if and only if $\mathcal{T}x$ is an inactive critical point.

We now specialize to finite-max composite objectives, a broad class for which the locally minimal identifiable set admits an explicit description.

2.4 Finite-Max Composite Functions

In this paper, we study *finite-max composite* objectives

$$f(x) = g(F(x)), \quad g(z) := \max_{i \in \mathcal{I}} g_i(z),$$

where $\mathcal{I}$ is a finite index set, $F : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is C^2 and definable, and each $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is C^2 and definable. We write $f_i := g_i \circ F$ and call each f_i a smooth *piece* of f and each g_i a smooth *piece* of g . The nonsmoothness of f comes entirely from the outer pointwise maximum. Fix a critical point $\bar{x}$ of f and set $\bar{z} := F(\bar{x})$. For z near $\bar{z}$, the *active set* is defined as

$$I(z) := \{i \in \mathcal{I} : g_i(z) = g(z)\}.$$

Since g is a finite maximum of C^2 functions, it is locally Lipschitz and subdifferentially regular, with

$$\partial g(\bar{z}) = \text{conv}\{\nabla g_i(\bar{z}) : i \in I(\bar{z})\} \quad \text{and} \quad \partial^\infty g(\bar{z}) = \{0\}.$$

The two results below characterize the locally minimal identifiable set for two building blocks of finite-max composite functions: finite-max of C^1 functions (Proposition 2.11) and composite functions (Proposition 2.13).

Proposition 2.11 (Identifiable set for a finite maximum, Theorem 7.3 of Lewis and Tian [44]) *Let $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ be C^1 for each $i \in \mathcal{I}$, where $\mathcal{I}$ is a finite index set, and define $g(z) := \max_{i \in \mathcal{I}} g_i(z)$. For $\bar{z} \in \text{dom } g$ and $\bar{v} \in \partial g(\bar{z})$, define the set of multipliers*

$$\Lambda_{\bar{v}} := \left\{ \lambda \in \Delta_{\mathcal{I}} : \bar{v} = \sum_{i \in I(\bar{z})} \lambda_i \nabla g_i(\bar{z}), \text{supp } \lambda \subseteq I(\bar{z}) \right\}.$$

Then

$$\mathcal{M}_{\bar{v}} := \bigcup_{\lambda \in \Lambda_{\bar{v}}} \{z \in \mathbb{R}^n : \text{supp } \lambda \subseteq I(z)\}$$

is the locally minimal identifiable set for g at $\bar{z}$ for $\bar{v}$.

Example 2.12 (Identifiable set of the ℓ_1 -norm) Take $g = \|\cdot\|_1$ on $\mathbb{R}^m$. Its pieces are the affine maps $g_\sigma(z) = \langle \sigma, z \rangle$ over $\sigma \in \{\pm 1\}^m =: \mathcal{I}$, with $\nabla g_\sigma = \sigma$. Moreover, σ is an element of the active set $I(z)$ exactly when $\sigma_i = \text{sign}(z_i)$ on every coordinate with $z_i \neq 0$. Fix a subgradient $\bar{v} \in \partial g(\bar{z})$. By Proposition 2.11, the locally minimal identifiable set for g at $\bar{z}$ for $\bar{v}$ is

$$\mathcal{M}_{\bar{v}} := \bigcup_{\lambda \in \Lambda_{\bar{v}}} \{z \in \mathbb{R}^m : \text{supp } \lambda \subseteq I(z)\}, \quad \Lambda_{\bar{v}} = \left\{ \lambda \in \Delta_{\mathcal{I}} : \bar{v} = \sum_{\sigma} \lambda_{\sigma} \sigma \right\}, \quad (2)$$

a union, over the representations of $\bar{v}$ as a convex combination of sign patterns, of the points z at which every pattern in $\text{supp } \lambda$ is simultaneously active.² In Appendix A, we show that this identifiable set reduces to:

$$\mathcal{S}_{\bar{v}} := \{z \in \mathbb{R}^m : z_i = 0 \text{ if } |\bar{v}_i| < 1, \ z_i \geq 0 \text{ if } \bar{v}_i = 1, \ z_i \leq 0 \text{ if } \bar{v}_i = -1\}, \quad (3)$$

²The support condition $\text{supp } \lambda \subseteq I(\bar{z})$ of Proposition 2.11 is automatic here: a sign pattern disagreeing with $\bar{v}$ on a coordinate where $|\bar{v}_i| = 1$ cannot appear in any representation of $\bar{v}$, while a coordinate where $|\bar{v}_i| < 1$ satisfies $\bar{z}_i = 0$ and therefore imposes no constraint on the support.

which we call the *sign cone* of $\bar{v}$ henceforth. Equivalently, $\mathcal{S}_{\bar{v}}$ is the normal cone $N_{[-1,1]^m}(\bar{v})$ to the cube $[-1, 1]^m$ at $\bar{v}$. Call a coordinate i *saturated* for $\bar{v}$ when $|\bar{v}_i| = 1$ and *free* when $|\bar{v}_i| < 1$. Then $\dim \mathcal{S}_{\bar{v}}$ is the number of saturated coordinates, complementary to the dimension of the cube face containing $\bar{v}$. For example, if $\bar{v}$ is a vertex of the cube, $\mathcal{S}_{\bar{v}}$ is a full-dimensional orthant; if $\bar{v}$ is the center, $\mathcal{S}_{\bar{v}} = \{0\}$. Figure 3 illustrates an example.




Fig. 3: The locally minimal identifiable set for $g = \|\cdot\|_1$ on $\mathbb{R}^3$. *Left:* two subgradients $\bar{v}, \bar{v}' \in [-1, 1]^3 = \partial g(0)$, where $\bar{v}$ lies on the front-right edge of the cube (two saturated coordinates) and $\bar{v}'$ lies on the top face (one saturated coordinate). *Right:* the corresponding identifiable sets $\mathcal{S}_{\bar{v}}$ and $\mathcal{S}_{\bar{v}'}$, which are the normal cones to the cube at $\bar{v}$ and $\bar{v}'$, respectively. The edge point $\bar{v}$ (orange) gives the two-dimensional cone $\{z \in \mathbb{R}^3 : z_1 \geq 0, z_2 \geq 0, z_3 = 0\}$; the face point $\bar{v}'$ (purple) gives the ray $\{z \in \mathbb{R}^3 : z_1 = z_2 = 0, z_3 \geq 0\}$.

Proposition 2.13 (Chain rule, Proposition 5.1 of Drusvyatskiy and Lewis [24]) *Consider a function $f = g \circ F$ defined on an open neighborhood $\mathcal{V} \subseteq \mathbb{R}^d$, where $F : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is C^1 and $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is lower-semicontinuous. Fix a critical point $\bar{x} \in \text{dom } f$ with $\bar{z} := F(\bar{x})$. Suppose the following qualification condition holds at $\bar{x}$*

$$\ker DF(\bar{x})^* \cap \partial^\infty g(F(\bar{x})) = \{0\},$$

and hence the inclusion $\partial f(\bar{x}) \subseteq DF(\bar{x})^(\partial g(\bar{z}))$ holds. Define the set of multipliers*

$$\Lambda := \{v \in \partial g(\bar{z}) : DF(\bar{x})^*(v) = 0\} = \partial g(\bar{z}) \cap \ker DF(\bar{x})^*.$$

Suppose

- *g is Clarke regular at all points in $\text{dom } g$ around $F(\bar{x})$;*
- *for each $v \in \Lambda$, there exists a locally minimal identifiable set $\mathcal{M}_v$ for g at $\bar{z}$ for v ;*
- *the collection $\{\mathcal{M}_v\}_{v \in \Lambda}$ is finite.*

Then, the set

$$\mathcal{M} := \bigcup_{v \in \Lambda} F^{-1}(\mathcal{M}_v)$$

is locally minimal identifiable for f at $\bar{x}$ for $0 \in \partial f(\bar{x})$.

Example 2.14 (Identifiable set of an ℓ_1 -composite) Take $g = \|\cdot\|_1$ and $f = \|F(\cdot)\|_1$ for a C^2 map $F : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with critical point $\bar{x}$ and $\bar{z} := F(\bar{x})$. The chain rule of Proposition 2.13

applies: the qualification condition holds because $\partial^\infty g(\bar{z}) = \{0\}$, the norm is convex and hence Clarke regular, and the identifiable set of g is the cone $\mathcal{S}_v$ defined in (3), taking finitely many values, one per sign pattern of v . Its relevant multipliers are

$$\Lambda = \partial g(\bar{z}) \cap \ker DF(\bar{x})^* = \{v \in [-1, 1]^m : v_i = \text{sign}(\bar{z}_i) \text{ for } \bar{z}_i \neq 0, DF(\bar{x})^*v = 0\}.$$

Since the preimage commutes with unions, the locally minimal identifiable set is

$$\mathcal{M} = \bigcup_{v \in \Lambda} F^{-1}(\mathcal{S}_v) = F^{-1}(\mathcal{E}), \quad \mathcal{E} := \bigcup_{v \in \Lambda} \mathcal{S}_v \subset \mathbb{R}^m, \quad (4)$$

where $\mathcal{S}_v$ is the sign cone of v defined in Example 2.12. The set $\mathcal{E}$ is a union of normal cones $\mathcal{S}_v = N_{[-1, 1]^m}(v)$ to the cube, taken over the multiplier slice $\Lambda = \partial g(\bar{z}) \cap \ker DF(\bar{x})^*$. Therefore, its geometry is determined by how the subspace $\ker DF(\bar{x})^*$ intersects the cube. Figure 4 illustrates an example: depending on the orientation of the slice Λ , the geometry of the set $\mathcal{E}$ can differ drastically. Whether its pullback via $F^{-1}(\mathcal{E})$ is a C^2 smooth manifold is precisely the question addressed by the branching criterion of Section 3.




Fig. 4: The identifiable set $\mathcal{E}$ for a tilted slice, with $m = 3$ and $\dim \Lambda = 2$. *Left:* the slice Λ is a hexagon (gray), with six vertices $\nu_1, \dots, \nu_6$ (colored). *Right:* the corresponding fan $\mathcal{E} = \bigcup_{\nu \in \text{Vert}(\Lambda)} \mathcal{S}_\nu$. The cone $\mathcal{S}_\nu = N_{[-1, 1]^3}(\nu)$ for each vertex ν is a two-dimensional sector in the matching color, meeting along the rays where adjacent vertices share a saturated coordinate. The set $\mathcal{E}$ is a nonconvex and nonsmooth polyhedral fan.

Together, these two propositions give a clean expression for the locally minimal identifiable set of the finite-max composite problem of Section 2.4. Proposition 2.11 considers the outer maximum: the identifiable set of g at a subgradient $\bar{v}$ is the union, over multipliers λ representing $\bar{v}$, of the regions where the pieces g_i of g in $\text{supp } \lambda$ stay active. Proposition 2.13 pulls this set back through the smooth map F , keeping only the multipliers compatible with criticality at $\bar{x}$. The result is that the locally minimal identifiable set of f is the F -preimage of a union of active-piece regions, which we derive explicitly in Section 3.1 as the starting point for the branching criterion.

3 Deterministic Branching Criterion

To derive the branching criterion for finite-max composite problems, we proceed as follows. In Section 3.1, we combine Propositions 2.11 and 2.13 to obtain an explicit description of the locally minimal identifiable set together with a compatible stratification. In Section 3.2, we work under the hypothesis that this set is a C^2 identifiable manifold and derive its second-order expansion. Section 3.3 proves the resulting compatibility condition across accessible strata (Proposition 3.8) and obtains the branching criterion by contraposition (Theorem 3.9).

3.1 Locally Minimal Identifiable Set for Finite-Max Composite Functions

The following proposition combines Propositions 2.11 and 2.13 to express the locally minimal identifiable set of the composite $f = g \circ F$ in closed form.

Proposition 3.1 (Locally minimal identifiable set for finite-max composites) *Consider the finite-max composite function $f = g \circ F$ of Section 2.4. Fix a critical point $\bar{x}$ with $\bar{z} := F(\bar{x})$ and active set $I(\bar{z})$. Define the critical multiplier set*

$$\Lambda_0 := \left\{ \lambda \in \Delta_{\mathcal{I}} : DF(\bar{x})^* \left(\sum_{i \in I(\bar{z})} \lambda_i \nabla g_i(\bar{z}) \right) = 0, \text{supp } \lambda \subseteq I(\bar{z}) \right\}, \quad (5)$$

the family of active-support sets $\mathcal{K} := \{\text{supp } \lambda : \lambda \in \Lambda_0\}$, and for nonempty $\mathcal{K} \subseteq I(\bar{z})$, the active locus

$$\mathcal{A}_{\mathcal{K}} := \{z : \mathcal{K} \subseteq I(z)\} = \bigcap_{i \in \mathcal{K}} \bigcap_{j \in \mathcal{I}} \{z : g_i(z) \geq g_j(z)\}.$$

Then the set

$$\mathcal{M} := F^{-1}(\mathcal{E}), \quad \mathcal{E} := \bigcup_{\mathcal{K} \in \mathcal{K}} \mathcal{A}_{\mathcal{K}},$$

is the locally minimal identifiable set for f at $\bar{x}$ for $0 \in \partial f(\bar{x})$.

Proof To apply Proposition 2.13, we first need to verify its conditions. Since g is a finite maximum of C^2 functions, it is locally Lipschitz near $\bar{z}$, so $\partial^\infty g(\bar{z}) = \{0\}$. Hence $\ker DF(\bar{x})^* \cap \partial^\infty g(\bar{z}) = \{0\}$ trivially, and the qualification condition holds. Moreover, g is the finite maximum of C^2 functions, and hence Clarke regular, at every point near $\bar{z}$. Moreover, by Proposition 2.11, the locally minimal identifiable set $\mathcal{M}_v$ for g at $\bar{z}$ for $v \in \partial g(\bar{z})$ is

$$\begin{aligned} \mathcal{M}_v &= \bigcup_{\lambda \in \Lambda_v} \{z : \text{supp } \lambda \subseteq I(z)\}, \\ \Lambda_v &= \left\{ \lambda \in \Delta_{\mathcal{I}} : v = \sum_{i \in I(\bar{z})} \lambda_i \nabla g_i(\bar{z}), \text{supp } \lambda \subseteq I(\bar{z}) \right\}. \end{aligned}$$

The set $\mathcal{M}_v$ depends on v only through the collection of supports $\{\text{supp } \lambda : \lambda \in \Lambda_v\}$, which is a family of subsets of the finite set $I(\bar{z})$. There are finitely many such families, so $\{\mathcal{M}_v\}_{v \in \Lambda}$ is finite. The conditions of Proposition 2.13 are thus satisfied, and the locally minimal identifiable set of f at $\bar{x}$ is

$$\mathcal{M} = \bigcup_{v \in \Lambda} F^{-1}(\mathcal{M}_v), \quad \Lambda = \partial g(\bar{z}) \cap \ker DF(\bar{x})^*.$$

Since $\partial g(\bar{z}) = \text{conv}\{\nabla g_i(\bar{z}) : i \in I(\bar{z})\}$, every $v \in \partial g(\bar{z})$ admits a representation $v = \sum_{i \in I(\bar{z})} \lambda_i \nabla g_i(\bar{z})$ for some $\lambda \in \Delta_I$ with $\text{supp } \lambda \subseteq I(\bar{z})$. The condition $v \in \ker DF(\bar{x})^*$ then reads $DF(\bar{x})^*(\sum_{i \in I(\bar{z})} \lambda_i \nabla g_i(\bar{z})) = 0$. Hence the weights representing multipliers in Λ are exactly the elements of Λ_0 defined in (5). Finally, we focus on simplifying the description of $\mathcal{M}$. Pulling F^{-1} outside the union and substituting the expression for $\mathcal{M}_v$ gives

$$\mathcal{M} = \bigcup_{v \in \Lambda} F^{-1}(\mathcal{M}_v) = F^{-1}\left(\bigcup_{v \in \Lambda} \mathcal{M}_v\right) = F^{-1}\left(\bigcup_{v \in \Lambda} \bigcup_{\lambda \in \Lambda_v} \{z : \text{supp } \lambda \subseteq I(z)\}\right).$$

As v ranges over Λ and λ over Λ_v , the multiplier λ ranges exactly over Λ_0 . Writing $\mathcal{K} = \text{supp } \lambda$ and recalling $\mathcal{A}_{\mathcal{K}} = \{z : \mathcal{K} \subseteq I(z)\}$ and $\mathcal{K} = \{\text{supp } \lambda : \lambda \in \Lambda_0\}$,

$$\bigcup_{v \in \Lambda} \bigcup_{\lambda \in \Lambda_v} \{z : \text{supp } \lambda \subseteq I(z)\} = \bigcup_{\lambda \in \Lambda_0} \mathcal{A}_{\text{supp } \lambda} = \bigcup_{\mathcal{K} \in \mathcal{K}} \mathcal{A}_{\mathcal{K}} = \mathcal{E}.$$

Therefore $\mathcal{M} = F^{-1}(\mathcal{E})$, as claimed. $\square$

The identity $\mathcal{E} = \bigcup_{\mathcal{K} \in \mathcal{K}} \mathcal{A}_{\mathcal{K}}$ expresses $\mathcal{E}$ through the supports of the critical multipliers, each locus $\mathcal{A}_{\mathcal{K}}$ collecting the points where the pieces g_i , $i \in \mathcal{K}$, stay active. We now refine it into a description of the stratification of $\mathcal{E}$. Being definable, $\mathcal{E}$ admits a C^2 stratification globally. Near $\bar{z}$, the active set is a subset of $I(\bar{z})$: an inactive piece with $g_i(\bar{z}) < g(\bar{z})$ stays below g around $\bar{z}$, so $I(z) \subseteq I(\bar{z})$ on a neighborhood of $\bar{z}$. On such a neighborhood, the geometry of $\mathcal{E}$ is governed by which subset of $I(\bar{z})$ stays active, and the following proposition shows this subset to be locally constant: on a neighborhood of $\bar{z}$, the strata of $\mathcal{E}$ each carry a fixed active subset of $I(\bar{z})$. For $g = \|\cdot\|_1$, Example 2.14 makes the decomposition concrete: $\mathcal{E} = \bigcup_{v \in \Lambda} \mathcal{S}_v$ is a polyhedral fan whose maximal cones are the top-dimensional strata, meeting along the lower-dimensional cones where sign patterns coincide (Figure 4).

Proposition 3.2 (Stratification compatible with constant active set) *Let $\mathcal{E} = \bigcup_{\mathcal{K} \in \mathcal{K}} \mathcal{A}_{\mathcal{K}}$ be as in Proposition 3.1. For each $\mathcal{J} \subseteq I(\bar{z})$, define the active-set region*

$$\mathcal{R}_{\mathcal{J}} := \{z : I(z) = \mathcal{J}\}.$$

In a neighborhood of $\bar{z}$, the set $\mathcal{E}$ admits a finite C^2 stratification $\mathcal{Z}$ compatible with the family $\{\mathcal{R}_{\mathcal{J}} : \mathcal{J} \subseteq I(\bar{z})\}$ in the sense of Definition 2.2. Equivalently, for every stratum $\mathcal{Z} \in \mathcal{Z}$ there is an active set $\mathcal{I}_{\mathcal{Z}} \subseteq I(\bar{z})$ such that

$$I(z) = \mathcal{I}_{\mathcal{Z}} \quad \text{for all } z \in \mathcal{Z}.$$

In particular, g agrees with each active piece g_i , $i \in \mathcal{I}_{\mathcal{Z}}$, on $\mathcal{Z}$, so $g|_{\mathcal{Z}}$ is the restriction of a C^2 function.

Proof Choose a definable open neighborhood $\mathcal{U}$ of $\bar{z}$ with $I(z) \subseteq I(\bar{z})$ for all $z \in \mathcal{U}$. We work within $\mathcal{U}$ throughout, intersecting every set with $\mathcal{U}$ without further mention. At each point of $\mathcal{U}$ the maximum defining g is attained, so the active set is a nonempty subset of $I(\bar{z})$, and the family $\{\mathcal{R}_{\mathcal{J}} : \mathcal{J} \subseteq I(\bar{z}), \mathcal{J} \neq \emptyset\}$ partitions $\mathcal{U}$. Each $\mathcal{R}_{\mathcal{J}}$ is definable, being cut out by the equalities $g_i = g$ for $i \in \mathcal{J}$ together with the strict inequalities $g_i < g$ for $i \in I(\bar{z}) \setminus \mathcal{J}$.

We first rewrite $\mathcal{E}$ in terms of the regions $\mathcal{R}_{\mathcal{J}}$. By definition, $\mathcal{A}_{\mathcal{K}} = \{z : \mathcal{K} \subseteq I(z)\}$, and a point lies in $\mathcal{A}_{\mathcal{K}}$ exactly when its active set is one of the subsets of $I(\bar{z})$ containing $\mathcal{K}$. Therefore

$$\mathcal{A}_{\mathcal{K}} = \bigcup_{\mathcal{K} \subseteq \mathcal{J} \subseteq I(\bar{z})} \mathcal{R}_{\mathcal{J}}, \quad \text{and hence} \quad \mathcal{E} = \bigcup_{\mathcal{K} \in \mathcal{K}} \mathcal{A}_{\mathcal{K}} = \bigcup_{\mathcal{J} \in \mathcal{K}^\uparrow} \mathcal{R}_{\mathcal{J}},$$

where $\mathcal{K}^\uparrow := \{\mathcal{J} \subseteq I(\bar{z}) : \mathcal{K} \subseteq \mathcal{J} \text{ for some } \mathcal{K} \in \mathcal{K}\}$ is the collection of active sets that contain some member of $\mathcal{K}$. Thus $\mathcal{E}$ is a union of regions, and any region $\mathcal{R}_{\mathcal{J}}$ with $\mathcal{J} \notin \mathcal{K}^\uparrow$ is disjoint from $\mathcal{E}$. Now apply Theorem 2.3 to the finite family $\{\mathcal{R}_{\mathcal{J}} : \mathcal{J} \subseteq I(\bar{z}), \mathcal{J} \neq \emptyset\}$, whose union is $\mathcal{U}$. The resulting C^2 stratification of $\mathcal{U}$ is compatible with this family, so each stratum is contained in a single $\mathcal{R}_{\mathcal{J}}$. Since $\mathcal{E}$ is a union of the sets $\mathcal{R}_{\mathcal{J}}$, every stratum is either contained in or disjoint from $\mathcal{E}$, and we let $\mathcal{Z}$ be those contained in $\mathcal{E}$. Recalling that $\mathcal{E} = \bigcup_{\mathcal{J} \in \mathcal{K}^\uparrow} \mathcal{R}_{\mathcal{J}}$, the members of $\mathcal{Z}$ are exactly the strata contained in some $\mathcal{R}_{\mathcal{J}}$ with $\mathcal{J} \in \mathcal{K}^\uparrow$, and they partition $\mathcal{E}$.

Fix a stratum $\mathcal{Z} \in \mathcal{Z}$ and let $\mathcal{J}$ be the unique index with $\mathcal{Z} \subseteq \mathcal{R}_{\mathcal{J}}$. The active set is then constant on $\mathcal{Z}$, equal to $\mathcal{I}_{\mathcal{Z}} := \mathcal{J}$, so that $g_i = g$ on $\mathcal{Z}$ for every $i \in \mathcal{I}_{\mathcal{Z}}$. It follows that $g|_{\mathcal{Z}}$ coincides with the restriction of the C^2 function g_i for any such i , and is in particular C^2 . $\square$

Each active locus is a union of the sets $\mathcal{R}_{\mathcal{J}}$, namely $\mathcal{A}_{\mathcal{K}} = \bigcup_{\mathcal{K} \subseteq \mathcal{J}} \mathcal{R}_{\mathcal{J}}$, and each $\mathcal{R}_{\mathcal{J}}$ is a union of strata of $\mathcal{Z}$. Consequently $\mathcal{E}$ is a union of strata, and on each stratum a single active set, hence a single piece g_i , agrees with g . Example 3.10 below, and the robust low-rank recovery of Section 5, identify the strata explicitly and compare the pieces g_i across them.

Propositions 3.1 and 3.2 together concretize the question of whether a critical point $\bar{x}$ admits a C^2 identifiable manifold. Such a manifold, if it exists, must coincide near $\bar{x}$ with $\mathcal{M} = F^{-1}(\mathcal{E})$, on which f restricts to a C^2 function with a single second-order expansion at $\bar{x}$. The stratification of $\mathcal{E}$ turns this into a condition on the individual strata: each contributes its own second-order behavior to f along the manifold. The next subsection computes the second-order contribution of each stratum, and Section 3.3 derives the compatibility that these contributions must satisfy. The branching criterion is then obtained by contraposition.

3.2 Second-Order Expansion Under the Activity Hypothesis

We derive a necessary condition for activity under the following standing hypothesis.

Assumption 3.3 (Activity hypothesis) The locally minimal identifiable set $\mathcal{M} = F^{-1}(\mathcal{E})$ is a C^2 identifiable manifold for f at $\bar{x}$ for $0 \in \partial f(\bar{x})$.

This hypothesis entails no loss in deriving a necessary condition: if any C^2 identifiable manifold exists, Proposition 2.8 and the local uniqueness of the locally minimal identifiable set force it to agree with $\mathcal{M}$ near $\bar{x}$. In particular, under the activity hypothesis, $\mathcal{M}$ is a C^2 manifold near $\bar{x}$ and $f|_{\mathcal{M}}$ is C^2 . All tangent spaces, local charts, and chart-dependent quantities introduced below are understood under this hypothesis until it is discharged in Theorem 3.9.

Each stratum of $\mathcal{E}$ determines a quadratic form, the leading second-order term contributed by that stratum. The activity hypothesis requires these forms to agree

whenever the corresponding strata are reached along a common family of tangent directions, because the C^2 restriction $f|_{\mathcal{M}}$ has a unique second-order expansion at $\bar{x}$. We first compute the contribution of each active piece and then identify which strata are accessible from $\mathcal{M}$.

Writing $T_{\bar{x}}\mathcal{M}$ for the tangent space, Lemma 2.1 expresses points of $\mathcal{M}$ near $\bar{x}$ as

$$\gamma(v) = \bar{x} + v + Z_{\mathcal{M}}(v) + o(\|v\|^2), \quad v \in T_{\bar{x}}\mathcal{M}, \quad (6)$$

where $Z_{\mathcal{M}} : T_{\bar{x}}\mathcal{M} \rightarrow N_{\bar{x}}\mathcal{M}$ is the homogeneous quadratic second-order term. Define

$$\eta_{\mathcal{M}}(v) := DF(\bar{x})[Z_{\mathcal{M}}(v)] + \frac{1}{2}D^2F(\bar{x})[v, v], \quad (7)$$

and, for each active piece $i \in I(\bar{z})$,

$$\zeta_i(v) := \langle \nabla g_i(\bar{z}), \eta_{\mathcal{M}}(v) \rangle + \frac{1}{2} \langle DF(\bar{x})[v], \nabla^2 g_i(\bar{z}) DF(\bar{x})[v] \rangle. \quad (8)$$

The next proposition shows that $\eta_{\mathcal{M}}(v)$ and $\zeta_i(v)$ capture the second-order deviations of $\gamma(v)$ when pushed forward through F and $g_i \circ F$, respectively.

Proposition 3.4 (Tangential first-order and second-order expansions) *Let $\bar{x}$ be a critical point of f , and define $\bar{z} = F(\bar{x})$. Suppose the activity hypothesis (Assumption 3.3) holds. Then:*

1. *for every active index $i \in I(\bar{z})$,*

$$P_{T_{\bar{x}}\mathcal{M}} \nabla(g_i \circ F)(\bar{x}) = \nabla_{\mathcal{M}} f(\bar{x}) = 0;$$

2. *for all sufficiently small $v \in T_{\bar{x}}\mathcal{M}$,*

$$F(\gamma(v)) = \bar{z} + DF(\bar{x})[v] + \eta_{\mathcal{M}}(v) + o(\|v\|^2);$$

3. *for all sufficiently small $v \in T_{\bar{x}}\mathcal{M}$,*

$$f(\gamma(v)) = g(\bar{z}) + \max_{i \in I(\bar{z})} \zeta_i(v) + o(\|v\|^2).$$

Proof Write $f_i := g_i \circ F$, so that $f = \max_{i \in I(\bar{z})} f_i$ on a neighborhood of $\bar{x}$, where each f_i is C^2 with $\nabla f_i(\bar{x}) = DF(\bar{x})^* \nabla g_i(\bar{z})$.

Proof of (i). Fix $v \in T_{\bar{x}}\mathcal{M}$ and let $t \mapsto \gamma(t)$ be a C^1 curve in $\mathcal{M}$ with $\gamma(0) = \bar{x}$ and $\dot{\gamma}(0) = v$. Since each f_i is differentiable and, on a neighborhood of $\bar{x}$, the active indices $I(\bar{z})$ are the only ones attaining the maximum, the one-sided derivative of f along γ at $t = 0$ is

$$\lim_{t \downarrow 0} \frac{f(\gamma(t)) - f(\bar{x})}{t} = \max_{i \in I(\bar{z})} \langle \nabla f_i(\bar{x}), v \rangle.$$

Since $f|_{\mathcal{M}}$ is differentiable at $\bar{x}$, the left-hand side equals $\langle \nabla_{\mathcal{M}} f(\bar{x}), v \rangle$, so

$$\langle \nabla_{\mathcal{M}} f(\bar{x}), v \rangle = \max_{i \in I_{\star}} \langle \nabla f_i(\bar{x}), v \rangle \quad \text{for all } v \in T_{\bar{x}}\mathcal{M}. \quad (9)$$

The left-hand side of (9) is linear in v . Noting that $-v \in T_{\bar{x}}\mathcal{M}$ and replacing v by $-v$ gives

$$\max_{i \in I(\bar{z})} \langle \nabla f_i(\bar{x}), v \rangle = - \max_{i \in I(\bar{z})} \langle \nabla f_i(\bar{x}), -v \rangle = \min_{i \in I(\bar{z})} \langle \nabla f_i(\bar{x}), v \rangle.$$

Thus the maximum and minimum over $i \in I(\bar{z})$ agree for every $v \in T_{\bar{x}}\mathcal{M}$, so all $\langle \nabla f_i(\bar{x}), v \rangle$ coincide and, by (9), equal $\langle \nabla_{\mathcal{M}} f(\bar{x}), v \rangle$. Equivalently, $P_{T_{\bar{x}}\mathcal{M}} \nabla f_i(\bar{x}) = \nabla_{\mathcal{M}} f(\bar{x})$ for every $i \in I(\bar{z})$. Since $0 \in \partial f(\bar{x}) = \text{conv}\{\nabla f_i(\bar{x}) : i \in I(\bar{z})\}$, projecting onto $T_{\bar{x}}\mathcal{M}$ gives $\nabla_{\mathcal{M}} f(\bar{x}) = 0$.

Proof of (ii). Using Taylor expansion of F along the chart (6) and grouping by order gives

$$\begin{aligned} F(\gamma(v)) &= \bar{z} + DF(\bar{x})[v] + (DF(\bar{x})[Z_{\mathcal{M}}(v)] + \tfrac{1}{2}D^2F(\bar{x})[v, v]) + o(\|v\|^2) \\ &= \bar{z} + DF(\bar{x})[v] + \eta_{\mathcal{M}}(v) + o(\|v\|^2), \end{aligned}$$

where the second line uses the definition of $\eta_{\mathcal{M}}(v)$ in (7).

Proof of (iii). Fix $i \in I(\bar{z})$. Using Taylor expansion of $g_i \circ F$ along the chart (6), and using the conclusion (ii), we have

$$\begin{aligned} g_i(F(\gamma(v))) &= g_i(\bar{z}) + \langle \nabla g_i(\bar{z}), DF(\bar{x})[v] \rangle + \langle \nabla g_i(\bar{z}), \eta_{\mathcal{M}}(v) \rangle + \tfrac{1}{2} \langle DF(\bar{x})[v], \nabla^2 g_i(\bar{z}) DF(\bar{x})[v] \rangle + o(\|v\|^2) \\ &= g_i(\bar{z}) + \langle \nabla g_i(\bar{z}), DF(\bar{x})[v] \rangle + \zeta_i(v) + o(\|v\|^2) \end{aligned}$$

where the third line uses the definition of ζ_i in (8). By conclusion (i) the linear term satisfies $\langle \nabla g_i(\bar{z}), DF(\bar{x})[v] \rangle = \langle \nabla f_i(\bar{x}), v \rangle = 0$ for all $v \in T_{\bar{x}}\mathcal{M}$, and $g_i(\bar{z}) = g(\bar{z})$ since i is active, so

$$g_i(F(\gamma(v))) = g(\bar{z}) + \zeta_i(v) + o(\|v\|^2). \quad (10)$$

By continuity, each inactive piece stays strictly below g on a neighborhood of $\bar{z}$, and I is finite. Therefore, for small enough v , the active piece of f is attained on $I(\bar{z})$. This implies

$$f(\gamma(v)) = \max_{i \in I(\bar{z})} \{g_i(F(\gamma(v)))\} = g(\bar{z}) + \max_{i \in I(\bar{z})} \zeta_i(v) + o(\|v\|^2).$$

This completes the proof. $\square$

Under the activity hypothesis, conclusion (i) says that every active piece of f has the same first-order variation along $T_{\bar{x}}\mathcal{M}$, equal to that of $f|_{\mathcal{M}}$, and this common value is zero. Crucially, this does *not* require $DF(\bar{x})[v]$ itself to vanish: the inner displacement may be nonzero, provided it is orthogonal to every active outer gradient $\nabla g_i(\bar{z})$. Tangent directions are thus constrained only through the active gradients, and first-order information alone cannot pin down the manifold. This is why the second-order coefficients ζ_i of conclusion (iii) are needed.

3.3 Accessible Strata and Branching Criterion

We now identify which strata contribute to the local expansion under the activity hypothesis. For a comparison between two strata to be meaningful, the manifold must reach both strata, and do so along an open set of tangent directions rather than a single curve. We introduce the following definition to isolate this condition.

Definition 3.5 (Accessible stratum) Under the activity hypothesis (Assumption 3.3), let $\mathcal{W} \subseteq T_{\bar{x}}\mathcal{M}$ be a linear subspace. A stratum $\mathcal{Z} \in \mathcal{Z}$ is *accessible* from $\mathcal{M}$ at $\bar{x}$ along $\mathcal{W}$ if there is a nonempty relatively open cone $\mathcal{C} \subset \mathcal{W} \setminus \{0\}$ such that, for every unit vector $v \in \mathcal{C}$, there is $\varepsilon(v) > 0$ with

$$F(\gamma(tv)) \in \mathcal{Z}, \quad 0 < t < \varepsilon(v).$$

Note that the open cone is essential in this definition. Reaching a stratum along a single direction would describe f only along that one direction, whereas reaching it

along an open cone describes it on a full solid angle of directions in $\mathcal{W}$. Intuitively, $\mathcal{Z}$ is accessible when the curve $t \mapsto F(\gamma(tv))$ enters $\mathcal{Z}$ for a non-zero measure of tangent directions v in $\mathcal{W}$.

Lemma 3.6 (Quadratic forms agreeing on an open cone) *Let q_1, q_2 be quadratic forms on a finite-dimensional vector space $\mathcal{W}$. If $q_1 = q_2$ on a nonempty relatively open cone in $\mathcal{W}$, then $q_1 = q_2$ on $\mathcal{W}$.*

Proof The difference $q_1 - q_2$ is a homogeneous polynomial of degree two. A polynomial that vanishes on a nonempty open subset of its domain vanishes identically. $\square$

On an accessible stratum, all active pieces share a single quadratic form, making the following definition unambiguous.

Lemma 3.7 (Branch quadratic of a stratum) *Under the activity hypothesis (Assumption 3.3), let $\mathcal{Z} \in \mathcal{Z}$ be accessible from $\mathcal{M}$ at $\bar{x}$ along a linear subspace $\mathcal{W} \subseteq T_{\bar{x}}\mathcal{M}$. Then the restrictions $\zeta_i|_{\mathcal{W}}$, with ζ_i as in (8), coincide for all $i \in \mathcal{I}_{\mathcal{Z}}$. We call this common quadratic form the branch quadratic of $\mathcal{Z}$ along $\mathcal{W}$ and denote it $\zeta_{\mathcal{Z}, \mathcal{W}}$.*

Proof Let $\mathcal{C} \subset \mathcal{W} \setminus \{0\}$ be an accessibility cone for $\mathcal{Z}$, and take $i, i' \in \mathcal{I}_{\mathcal{Z}}$. For every unit $v \in \mathcal{C}$ and all sufficiently small $t > 0$, the point $F(\gamma(tv))$ lies in $\mathcal{Z}$, and since the active set is constant on $\mathcal{Z}$,

$$g_i(F(\gamma(tv))) = g_{i'}(F(\gamma(tv))).$$

By the per-branch expansion (10), we have $g_i(F(\gamma(tv))) = g(\bar{z}) + t^2 \zeta_i(v) + o(t^2)$ for each active i , so

$$0 = g_i(F(\gamma(tv))) - g_{i'}(F(\gamma(tv))) = t^2 (\zeta_i(v) - \zeta_{i'}(v)) + o(t^2).$$

Dividing by t^2 and letting $t \downarrow 0$ gives $\zeta_i(v) = \zeta_{i'}(v)$ for every unit $v \in \mathcal{C}$, and hence on all of $\mathcal{C}$ by homogeneity. Lemma 3.6 then gives $\zeta_i|_{\mathcal{W}} \equiv \zeta_{i'}|_{\mathcal{W}}$, so $\zeta_{\mathcal{Z}, \mathcal{W}}$ is well-defined. $\square$

The activity hypothesis imposes a single quadratic expansion across all accessible strata. This is the key necessary condition behind the branching criterion.

Proposition 3.8 (Compatibility of accessible branch quadratics) *Under the activity hypothesis, let $\mathcal{W} \subseteq T_{\bar{x}}\mathcal{M}$ be a linear subspace. If $\mathcal{Z}_1, \mathcal{Z}_2 \in \mathcal{Z}$ are accessible from $\mathcal{M}$ at $\bar{x}$ along $\mathcal{W}$, then $\zeta_{\mathcal{Z}_1, \mathcal{W}} = \zeta_{\mathcal{Z}_2, \mathcal{W}}$ on $\mathcal{W}$.*

Proof By Proposition 3.4, $\nabla_{\mathcal{M}} f(\bar{x}) = 0$. Since $f|_{\mathcal{M}}$ is C^2 under the activity hypothesis, there is a unique quadratic form q on $T_{\bar{x}}\mathcal{M}$ such that

$$f(\gamma(v)) = f(\bar{x}) + q(v) + o(\|v\|^2).$$

Fix $j \in \{1, 2\}$ and let $\mathcal{C}_j \subset \mathcal{W} \setminus \{0\}$ be an accessibility cone for $\mathcal{Z}_j$. For $v \in \mathcal{C}_j$ and all sufficiently small $t > 0$, the point $F(\gamma(tv))$ lies in $\mathcal{Z}_j$. Hence, for any $i \in \mathcal{I}_{\mathcal{Z}_j}$, the per-branch expansion (10) and Lemma 3.7 give

$$f(\gamma(tv)) = g(\bar{z}) + t^2 \zeta_{\mathcal{Z}_j, \mathcal{W}}(v) + o(t^2).$$

Comparing the two expansions shows that $q(v) = \zeta_{\mathcal{Z}_j, \mathcal{W}}(v)$ on $\mathcal{C}_j$. Lemma 3.6 therefore yields $q|_{\mathcal{W}} \equiv \zeta_{\mathcal{Z}_j, \mathcal{W}}$. Applying this identity for $j = 1, 2$ proves the result. $\square$

We now discharge the activity hypothesis. The following contrapositive is the main result of this section.

Theorem 3.9 (Branching criterion) *Let $\bar{x}$ be a critical point of the finite-max composite $f = g \circ F$, with locally minimal identifiable set $\mathcal{M} = F^{-1}(\mathcal{E})$ from Proposition 3.1 and compatible stratification $\mathcal{Z}$ from Proposition 3.2. Suppose that the activity hypothesis implies the existence of a linear subspace $\mathcal{W} \subseteq T_{\bar{x}}\mathcal{M}$ and two strata $\mathcal{Z}_1, \mathcal{Z}_2 \in \mathcal{Z}$ such that*

- (a) *both $\mathcal{Z}_1$ and $\mathcal{Z}_2$ are accessible from $\mathcal{M}$ at $\bar{x}$ along $\mathcal{W}$;*
- (b) *their branch quadratics disagree: $\zeta_{\mathcal{Z}_1, \mathcal{W}} \neq \zeta_{\mathcal{Z}_2, \mathcal{W}}$.*

Then f admits no C^2 identifiable manifold at $\bar{x}$ for $0 \in \partial f(\bar{x})$; that is, $\bar{x}$ is inactive.

Proof Suppose, toward a contradiction, that a C^2 identifiable manifold exists at $\bar{x}$ for $0 \in \partial f(\bar{x})$. By Proposition 2.8 and the local uniqueness of the locally minimal identifiable set, this manifold agrees near $\bar{x}$ with $\mathcal{M} = F^{-1}(\mathcal{E})$. Thus the activity hypothesis holds. The premise of the theorem then provides $\mathcal{W}$, $\mathcal{Z}_1$, and $\mathcal{Z}_2$ satisfying (a)–(b), whereas Proposition 3.8 forces $\zeta_{\mathcal{Z}_1, \mathcal{W}} \equiv \zeta_{\mathcal{Z}_2, \mathcal{W}}$, a contradiction. $\square$

We now illustrate the branching criterion of Theorem 3.9 with an example from [12], explicitly verifying each of its conditions.

Example 3.10 (A two-piece inactive critical point) Consider

$$f(x_1, x_2) = (|x_1| + |x_2|)^2 - 2x_1^2.$$

Davis and Drusvyatskiy [12, Figure 1a] visually illustrate that the origin is an inactive critical point of f (in fact, a strict saddle). We now confirm this rigorously by directly applying our branching criterion. This function can be rewritten as

$$f(x_1, x_2) = x_2^2 - x_1^2 + 2|x_1x_2| = \max\{x_2^2 - x_1^2 + 2x_1x_2, x_2^2 - x_1^2 - 2x_1x_2\}.$$

Therefore, f admits a finite-max composite form $f = g \circ F$, where $F : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the identity map, $g = \max\{g_1, g_2\}$, and $g_1(z_1, z_2) := z_2^2 - z_1^2 + 2z_1z_2$ and $g_2(z_1, z_2) := z_2^2 - z_1^2 - 2z_1z_2$, both C^2 . To apply the branching criterion, we first obtain the locally minimal identifiable set $\mathcal{M}$ at the critical point $\bar{x} = (0, 0)$ via Proposition 3.1. Both pieces and their gradients vanish at $\bar{x}$, so the active set $I(\bar{z})$ at $\bar{z} = (0, 0)$, the critical multiplier set Λ_0 , and the family of active support sets $\mathcal{K}$ reduce to

$$I(\bar{z}) = \{1, 2\}, \quad \Lambda_0 = \Delta_{\mathcal{I}}, \quad \mathcal{K} = \{\{1\}, \{2\}, \{1, 2\}\}.$$

Since $g_1(z) - g_2(z) = 4z_1z_2$, the active loci are $A_{\{1\}} = \{z_1z_2 \geq 0\}$, $A_{\{2\}} = \{z_1z_2 \leq 0\}$, and $A_{\{1, 2\}} = \{z_1z_2 = 0\}$. Together they cover $\mathbb{R}^2$, so $\mathcal{E} = \mathbb{R}^2$ and $\mathcal{M} = F^{-1}(\mathcal{E}) = \mathbb{R}^2$.

We now identify the stratification of Proposition 3.2. The regions on which the active set is constant are

$$\mathcal{R}_{\{1\}} = \{z_1z_2 > 0\}, \quad \mathcal{R}_{\{2\}} = \{z_1z_2 < 0\}, \quad \mathcal{R}_{\{1, 2\}} = \{z_1z_2 = 0\},$$

on which the active piece is g_1 , g_2 , or both, respectively. Each is a union of strata: $\mathcal{R}_{\{1\}}$ and $\mathcal{R}_{\{2\}}$ split into their connected components, the open quadrants, while $\mathcal{R}_{\{1, 2\}}$ is the union

of the two coordinate axes, which stratify into the four open half-axes and the origin. The resulting strata are the four open quadrants

$\mathcal{Z}_1 = \{z_1 > 0, z_2 > 0\}$, $\mathcal{Z}_2 = \{z_1 < 0, z_2 > 0\}$, $\mathcal{Z}_3 = \{z_1 < 0, z_2 < 0\}$, $\mathcal{Z}_4 = \{z_1 > 0, z_2 < 0\}$, together with the four open half-axes, and the origin. The constant active sets are: $\{1\}$ on $\mathcal{Z}_1$ and $\mathcal{Z}_3$, $\{2\}$ on $\mathcal{Z}_2$ and $\mathcal{Z}_4$, and $\{1, 2\}$ on the open half axes and the origin. We compare the two quadrants $\mathcal{Z}_1$ and $\mathcal{Z}_4$, with $\mathcal{I}_{\mathcal{Z}_1} = \{1\}$ and $\mathcal{I}_{\mathcal{Z}_4} = \{2\}$.

To verify the implication required by Theorem 3.9, suppose the activity hypothesis holds. Since $\mathcal{M} = \mathbb{R}^2$, the chart of Lemma 2.1 is $\gamma(v) = \bar{x} + v$, and we take $\mathcal{W} = \mathbb{R}^2$. Both quadrants are accessible from $\mathcal{M}$ at $\bar{x}$ along $\mathcal{W}$: the cones $\mathcal{C}_1 = \{v_1 > 0, v_2 > 0\}$ and $\mathcal{C}_4 = \{v_1 > 0, v_2 < 0\}$ are nonempty and relatively open, and for every unit vector $v \in \mathcal{C}_1$ (resp. $v \in \mathcal{C}_4$) and sufficiently small $t > 0$ one has $F(\gamma(tv)) = tv \in \mathcal{Z}_1$ (resp. $\mathcal{Z}_4$), which is the accessibility condition of Definition 3.5. With $\gamma(v) = \bar{x} + v$, each piece is a homogeneous quadratic, so $\zeta_i(v) = g_i(v)$, and Lemma 3.7 identifies the branch quadratics

$$\zeta_{\mathcal{Z}_1, \mathcal{W}}(v) = v_2^2 - v_1^2 + 2v_1v_2, \quad \zeta_{\mathcal{Z}_4, \mathcal{W}}(v) = v_2^2 - v_1^2 - 2v_1v_2,$$

which disagree on $\mathcal{W}$. Thus the activity hypothesis implies the incompatible accessible branches required by Theorem 3.9, and $\bar{x}$ is inactive.

The example above is deliberately simple: inactivity can be verified by inspection, without invoking the branching criterion. The remainder of the paper demonstrates the criterion in two practically important problem classes: finite-scenario min-max optimization and robust low-rank recovery. The former models worst-case estimation across finitely many scenarios and provides a transparent first application: at a common interpolator, all scenario gradients vanish, so disagreement among their quadratic growth directly triggers the branching criterion. In the random overparameterized two-group model, every common interpolator is consequently an inactive global minimizer almost surely. Robust low-rank recovery, by contrast, seeks to recover a low-rank object from corrupted measurements and requires a substantially more involved reduction to a geometric cone-selection test and its probabilistic verification.

4 Inactive Global Minimizers in Finite-Scenario Minimax Optimization

We consider the application of the branching criterion to finite-scenario minimax optimization [45], a standard instance of group distributionally robust optimization [46]. Suppose the data are partitioned into S prespecified groups, and let $R_s : \mathbb{R}^d \rightarrow \mathbb{R}^{m_s}$ denote the residual map for group s , which we assume to be C^2 and definable. We consider

$$f_{\text{GDRO}}(x) := \max_{s \in [S]} \ell_s(x), \quad \ell_s(x) := \frac{1}{2} \|R_s(x)\|^2. \quad (11)$$

Equivalently, $f_{\text{GDRO}}(x) = \max_{p \in \Delta_{[S]}} \sum_{s=1}^S p_s \ell_s(x)$, matching the loss function of the group distributionally robust problem, wherein an adversary may choose any mixture of the empirical group distributions [46]. It has the finite-max composite form of Section 2.4: take $F(x) = (R_1(x), \dots, R_S(x))$ and, for $z = (z_1, \dots, z_S)$, define $g_s(z) := \frac{1}{2} \|z_s\|^2$ and $g := \max_{s \in [S]} g_s$.

We call $\bar{x}$ a *common interpolator* if $R_s(\bar{x}) = 0$ for every $s \in [S]$. Such a point is a global minimizer of f_{GDRO} , and every smooth piece has zero gradient there. Nevertheless, the pieces may grow according to different quadratic forms around $\bar{x}$. Such common interpolators arise naturally in modern overparameterized learning, where models often fit every training sample—and hence every prespecified group—simultaneously; see [46–48].

The next result shows that whenever two different pieces dominate along different sets of directions, the common interpolator admits no C^2 identifiable manifold. For $\mathcal{J}_s[v] := DR_s(\bar{x})[v]$, define the scenario quadratic

$$q_s(v) := \frac{1}{2} \|\mathcal{J}_s[v]\|^2, \quad v \in \mathbb{R}^d, \quad (12)$$

and, for a linear subspace $\mathcal{W} \subseteq \mathbb{R}^d$, the strict dominance cone

$$\mathcal{D}_s(\mathcal{W}) := \{v \in \mathcal{W} \setminus \{0\} : q_s(v) > q_j(v) \text{ for every } j \neq s\}.$$

Theorem 4.1 (Inactive minimax interpolators) *Let $\bar{x}$ be a common interpolator of f_{GDRO} . Suppose there are a linear subspace $\mathcal{W} \subseteq \mathbb{R}^d$ and distinct indices $s_1, s_2 \in [S]$ such that both $\mathcal{D}_{s_1}(\mathcal{W})$ and $\mathcal{D}_{s_2}(\mathcal{W})$ are nonempty. Then $\bar{x}$ is an inactive global minimizer of f_{GDRO} .*

Proof Since every loss is nonnegative and vanishes at $\bar{x}$, the point $\bar{x}$ is a global minimizer. Moreover,

$$\nabla \ell_s(\bar{x}) = \mathcal{J}_s^*[R_s(\bar{x})] = 0 \quad \text{for every } s \in [S].$$

Since f_{GDRO} is a finite maximum of C^2 functions, it is locally Lipschitz and subdifferentially regular near $\bar{x}$, and the subdifferential formula of Section 2.4 gives $\partial f_{\text{GDRO}}(\bar{x}) = \text{conv}\{\nabla \ell_s(\bar{x}) : s \in [S]\} = \{0\}$; in particular, $\bar{x}$ is critical.

We next identify the locally minimal identifiable set. At $\bar{z} = F(\bar{x}) = 0$, every outer piece is active and $\nabla g_s(\bar{z}) = 0$, so the critical multiplier set (5) is the entire simplex: $\Lambda_0 = \Delta_{[S]}$. In particular, every singleton $\{s\}$ occurs as the support of a critical multiplier, and the corresponding active loci $\mathcal{A}_{\{s\}}$ already cover the image space, because every z has at least one active piece. Proposition 3.1 therefore gives

$$\mathcal{E} = \bigcup_{s \in [S]} \mathcal{A}_{\{s\}} = \prod_{s=1}^S \mathbb{R}^{m_s}, \quad \mathcal{M} = F^{-1}(\mathcal{E}) = \mathbb{R}^d,$$

To invoke Theorem 3.9, we verify its accessibility and incompatibility conditions under the activity hypothesis. For each $s \in [S]$, let $\mathcal{R}_{\{s\}} := \{z : I(z) = \{s\}\}$ be the open region on which g_s is uniquely active. Choose the compatible stratification in Proposition 3.2 so that every connected component of each nonempty $\mathcal{R}_{\{s\}}$ is a stratum; this is possible by taking these open components as the top-dimensional strata and stratifying their definable complement compatibly with their closures.

Work under the activity hypothesis. Since $\mathcal{M} = \mathbb{R}^d$, we have $T_{\bar{x}}\mathcal{M} = \mathbb{R}^d$ and may take the local chart to be $\gamma(v) = \bar{x} + v$. Fix $s \in \{s_1, s_2\}$ and choose a unit vector $v_s \in \mathcal{D}_s(\mathcal{W})$. Because $\mathcal{D}_s(\mathcal{W})$ is relatively open, there is a connected relatively open neighborhood $\mathcal{V}_s$ of v_s in the unit sphere of $\mathcal{W}$ such that $\text{cl } \mathcal{V}_s \subset \mathcal{D}_s(\mathcal{W})$. Uniformly for $v \in \text{cl } \mathcal{V}_s$,

$$g_s(F(\bar{x} + tv)) - g_j(F(\bar{x} + tv)) = t^2(q_s(v) - q_j(v)) + o(t^2), \quad j \neq s.$$

The strict inequalities on the compact set $\text{cl } \mathcal{V}_s$ therefore imply that, for some $\varepsilon_s > 0$,

$$F(\bar{x} + tv) \in \mathcal{R}_{\{s\}} \quad \text{for all } v \in \mathcal{V}_s \text{ and } 0 < t < \varepsilon_s.$$

The image of the connected set $\mathcal{V}_s \times (0, \varepsilon_s)$ under $(v, t) \mapsto F(\bar{x} + tv)$ is connected and hence lies in a single connected component $\mathcal{Z}_s$ of $\mathcal{R}_{\{s\}}$. Therefore, $\mathcal{Z}_s$ is accessible from $\mathcal{M}$ at $\bar{x}$ along $\mathcal{W}$, with accessibility cone $\mathcal{C}_s := \{\alpha v : \alpha > 0, v \in \mathcal{V}_s\}$.

It remains to compare the branch quadratics. Since $\nabla g_s(0) = 0$ and $\nabla^2 g_s(0)$ is the identity on the s -th group and zero on all other groups, definition (8) gives

$$\zeta_{\mathcal{Z}_s, \mathcal{W}}(v) = \frac{1}{2} \|\mathcal{J}_s[v]\|^2 = q_s(v) \quad \text{for every } v \in \mathcal{W}.$$

The two accessible strata $\mathcal{Z}_{s_1}$ and $\mathcal{Z}_{s_2}$ have incompatible branch quadratics: any $v \in \mathcal{D}_{s_1}(\mathcal{W})$ satisfies $q_{s_1}(v) > q_{s_2}(v)$, and hence $q_{s_1}|_{\mathcal{W}} \not\equiv q_{s_2}|_{\mathcal{W}}$. Thus, under the activity hypothesis, conditions (a)–(b) of Theorem 3.9 hold. That theorem rules out any C^2 identifiable manifold at $\bar{x}$, proving that $\bar{x}$ is inactive. $\square$

For two scenarios, the condition above reduces to a familiar matrix test. Indeed, $\mathcal{D}_1(\mathcal{W})$ and $\mathcal{D}_2(\mathcal{W})$ are both nonempty exactly when $\mathcal{J}_1^* \mathcal{J}_1 - \mathcal{J}_2^* \mathcal{J}_2$ is indefinite on $\mathcal{W}$. A useful sufficient condition is that the two kernels are incomparable: there exist $v_1 \in \ker \mathcal{J}_2 \setminus \ker \mathcal{J}_1$ and $v_2 \in \ker \mathcal{J}_1 \setminus \ker \mathcal{J}_2$. Then $q_1(v_1) > q_2(v_1)$ and $q_2(v_2) > q_1(v_2)$.

We next show that this kernel geometry is generic in overparameterized two-group linear regression. Given $A_s \in \mathbb{R}^{m_s \times d}$ and $b_s \in \mathbb{R}^{m_s}$, consider

$$f_{2G}(x) := \max_{s \in [2]} \frac{1}{2m_s} \|A_s x - b_s\|^2. \quad (13)$$

Corollary 4.2 (Gaussian two-group minimax regression) *Let $A_1 \in \mathbb{R}^{m_1 \times d}$ and $A_2 \in \mathbb{R}^{m_2 \times d}$ have i.i.d. standard normal entries, and suppose $m_1, m_2 \geq 1$ and $m_1 + m_2 \leq d$. Then, with probability one, for every $b_1 \in \mathbb{R}^{m_1}$ and $b_2 \in \mathbb{R}^{m_2}$, the common interpolators $\{\bar{x} \in \mathbb{R}^d : A_1 \bar{x} = b_1, A_2 \bar{x} = b_2\}$ form a nonempty affine subspace of dimension $d - m_1 - m_2$ and every one of them is an inactive global minimizer of f_{2G} .*

Proof The stacked matrix $[A_1^\top \ A_2^\top]^\top$ has i.i.d. standard normal entries and $m_1 + m_2$ rows. Since $m_1 + m_2 \leq d$, it has full row rank almost surely. Hence the two systems $A_1 \bar{x} = b_1$ and $A_2 \bar{x} = b_2$ admit a common interpolator for every (b_1, b_2) , and its solution set is an affine subspace of dimension $d - m_1 - m_2$. Full row rank also implies that neither $\ker A_1$ nor $\ker A_2$ contains the other: either inclusion would force the rows of one nonempty block to lie in the row span of the other, contradicting full row rank of the stack. Consequently, there exist

$$v_1 \in \ker A_2 \setminus \ker A_1, \quad v_2 \in \ker A_1 \setminus \ker A_2.$$

At any common interpolator, the scenario Jacobians in (12) are $\mathcal{J}_s = A_s / \sqrt{m_s}$. Thus $q_1(v_1) > q_2(v_1)$ and $q_2(v_2) > q_1(v_2)$, and Theorem 4.1 applies with $\mathcal{W} = \mathbb{R}^d$. $\square$

The mechanism here is particularly transparent. At a common interpolator, all scenario gradients vanish, so the locally minimal identifiable set fills the ambient space. Nevertheless, different scenarios dominate on different open cones and induce incompatible quadratic expansions, precluding a single C^2 local model. The inactive points identified here are global minimizers of a natural finite-max problem. By contrast, the next section turns to robust low-rank recovery, where the identified inactive critical points can be strict saddles rather than global/local minimizers.

5 Inactive Criticality in Robust Low-Rank Recovery

We next consider the class of *robust low-rank recovery* problems

$$\min_{X \in \mathbb{S}^n} \|\mathcal{A}(\mathcal{L}(X)) - y\|_1 \quad \text{s.t.} \quad X \succeq 0, \quad \text{rank}(X) \leq k, \quad (14)$$

for which the branching criterion takes a particularly geometric form. Here, $\mathcal{A} : \mathbb{S}^n \rightarrow \mathbb{R}^m$ is the linear *measurement operator* defined by $[\mathcal{A}(Z)]_i = \langle A_i, Z \rangle$ for $i = 1, \dots, m$, where $\{A_i\}_{i=1}^m \subset \mathbb{S}^n$ are known measurement matrices. The linear *design operator* $\mathcal{L} : \mathbb{S}^n \rightarrow \mathbb{S}^n$ is determined by the underlying generative model, such as matrix sensing or nonnegative sparse regression. The observations satisfy $y = \mathcal{A}(\mathcal{L}(X^*)) + \varepsilon$, where $X^* \succeq 0$ is the unknown ground-truth matrix with $\text{rank}(X^*) = r \leq k$, and $\varepsilon \in \mathbb{R}^m$ is a sparse vector of outliers. The ℓ_1 loss provides robustness against these outliers.

Instances of (14) arise in a wide range of robust statistical recovery applications, including signal processing, recommendation systems, and control; see the surveys by Chi et al. [30], Wright and Ma [49].

Problem (14) is often solved via the *Burer–Monteiro* factorization [50], which replaces the constraint set $\{X : X \succeq 0, \text{rank}(X) \leq k\}$ with $\{UU^\top : U \in \mathbb{R}^{n \times k}\}$, reducing (14) to

$$\min_{U \in \mathbb{R}^{n \times k}} \|\mathcal{A}(\mathcal{L}(UU^\top)) - y\|_1, \quad (15)$$

which is in the finite-max composite form considered in this paper. Despite its non-convex and nonsmooth nature, local-search algorithms are known to converge to the ground truth $X^* = U^*U^{*\top}$ when applied directly to (15). For instance, the subgradient method converges to X^* for matrix sensing even when $\varepsilon \neq 0$ and the model is severely *overparameterized*, i.e., the search rank dominates the true rank ($k \gg r$) [29, 33]. Analogous results hold for nonnegative sparse regression [29, 51].

Ma and Fattahi [52] showed that the (balanced) ground-truth solutions from $\{U^* : U^*U^{*\top} = X^*\}$ in matrix sensing emerge not as local or global minima, but as strict saddles. The convergence of the subgradient method to these points (or equivalently, its inability to escape them) therefore cannot be attributed to local optimality, leading the authors to conjecture that these strict saddles are inactive. In this work, we confirm this conjecture and prove a considerably stronger result:

Under appropriate randomness in the measurements and noise, and with overparameterization, every (balanced) critical true solution in $\{U^ : U^*U^{*\top} = X^*\}$ satisfies the branching criterion (Theorem 3.9) with high probability. Consequently, these critical true solutions are inactive, admitting no C^2 identifiable manifold; in regimes where they are strict saddles, they become inactive strict saddles.*

We formally present our results for two representative classes of robust low-rank recovery: nonnegative sparse regression and matrix sensing.

Nonnegative sparse regression: The goal of nonnegative sparse regression is to recover a nonnegative vector $x^* \in \mathbb{R}_+^n$ with at most $s \ll n$ nonzero entries from linear measurements $y_i = \langle a_i, x^* \rangle + \varepsilon_i$, where $\varepsilon \in \mathbb{R}^m$ is a sparse outlier vector. A common approach is to introduce the parametrization $x = u \odot u$, where $\odot$ denotes the

elementwise product, and apply the subgradient method directly to

$$f_{\text{sp}}(u) := \sum_{i=1}^m |\langle a_i, u \odot u \rangle - y_i|. \quad (16)$$

Nonnegative sparse regression is a special case of (15): take $k = 1$, set $A_i = \text{diag}(a_i)$, where $\text{diag}(a_i)$ denotes the diagonal matrix with diagonal a_i , and let $\mathcal{L}$ be the projection onto the diagonal matrices. Applying our branching criterion under Gaussian measurements yields the following high-probability guarantee.

Theorem 5.1 (Nonnegative sparse regression) *Consider the nonnegative sparse regression problem (16). Suppose that the measurement vectors $a_1, \dots, a_m$ have i.i.d. standard normal entries, and that the outlier vector ε is independent of the measurement vectors and satisfies $\|\varepsilon\|_0 \leq pm$, where $0 \leq p < 1/2$. Let $x^* \in \mathbb{R}_+^n$ satisfy $\|x^*\|_0 = s$ and $n \geq s + 2$. Then there exist constants $c_1, c_2 > 0$, depending only on p , such that, whenever $m \geq c_1 s$, with probability at least*

$$1 - e^{-c_2 s} - 2^{-(n-s-1)},$$

every ground-truth factor u^ satisfying $u^* \odot u^* = x^*$ is an inactive critical point of f_{sp} .*

Two features of the theorem are worth emphasizing. First, the assumption $n - s \geq 2$ ensures the existence of at least two coordinates outside the support of x^* . These off-support coordinates provide the directions used by the branching criterion to construct incompatible branch quadratics with high probability. The presence of off-support coordinates is not merely a technical requirement; it is, in fact, necessary. Indeed, when $n = s$, an argument analogous to that of Aravkin et al. [53, Theorem V.2] shows that f_{sp} is locally sharp away from the singleton $\mathcal{M} = \{u^*\}$.³ Together with the weak convexity of f_{sp} , this sharpness implies that $\mathcal{M}$ is a locally minimal identifiable set [14, Proposition 8.4]. Since a singleton is a C^2 manifold, u^* is therefore an active critical point when $n = s$. Second, the sample-size requirement $m \geq c_1 s$ is information-theoretically optimal up to constant factors. Thus, the order of measurements required to make recovery information-theoretically possible is already sufficient to ensure that the ground-truth factors are inactive critical points with high probability.

Matrix sensing: Matrix sensing is the matrix analogue of nonnegative sparse regression, with low rank playing the role of sparsity in the spectral domain. The goal is to recover a rank- r matrix $X^* \in \mathbb{R}^{n_1 \times n_2}$ from linear measurements $y_i = \langle \Gamma_i, X^* \rangle + \varepsilon_i$, where $\Gamma_i \in \mathbb{R}^{n_1 \times n_2}$ have i.i.d. standard normal entries, and $\varepsilon \in \mathbb{R}^m$ is a sparse outlier vector. A common approach is to introduce a low-rank factorization and apply the subgradient method directly to one of the factorized objectives

$$f_{\text{sym}}(U) := \sum_{i=1}^m |\langle \Gamma_i, UU^\top \rangle - y_i|, \quad U \in \mathbb{R}^{n \times k}, \quad (17)$$

³That is, there exist $\mu > 0$ and a neighborhood $\mathcal{U}$ of u^* such that $f_{\text{sp}}(u) \geq f_{\text{sp}}(u^*) + \mu \text{dist}(u, \mathcal{M})$ for every $u \in \mathcal{U}$.

$$f_{\text{asym}}(U_1, U_2) := \sum_{i=1}^m |\langle \Gamma_i, U_1 U_2^\top \rangle - y_i|, \quad (U_1, U_2) \in \mathbb{R}^{n_1 \times k} \times \mathbb{R}^{n_2 \times k}, \quad (18)$$

where $k \geq r$. The symmetric formulation applies when $n_1 = n_2 = n$ and X^* is symmetric and positive semidefinite, whereas the asymmetric formulation accommodates general, possibly rectangular, matrices. Both formulations are special cases of (15). In the symmetric setting, it suffices to take $A_i = \frac{1}{2}(\Gamma_i + \Gamma_i^\top)$ and let $\mathcal{L}$ be the identity map. For the asymmetric setting, set $n := n_1 + n_2$ and introduce the stacked factor $U := [U_1^\top \ U_2^\top]^\top \in \mathbb{R}^{n \times k}$. Define

$$A_i := \frac{1}{2} \begin{bmatrix} 0 & \Gamma_i \\ \Gamma_i^\top & 0 \end{bmatrix}, \quad \mathcal{L}(X) := \begin{bmatrix} 0 & X_{12} \\ X_{12}^\top & 0 \end{bmatrix},$$

where $X_{12} \in \mathbb{R}^{n_1 \times n_2}$ denotes the upper-right off-diagonal block of $X \in \mathbb{S}^n$. Since the upper-right off-diagonal block of $UU^\top$ is $U_1 U_2^\top$, we have $\langle A_i, \mathcal{L}(UU^\top) \rangle = \langle \Gamma_i, U_1 U_2^\top \rangle$, which establishes the claimed reduction. Applying our branching criterion to the symmetric and asymmetric formulations, again under Gaussian measurements, yields the following high-probability guarantees.

Theorem 5.2 (Symmetric matrix sensing) *Consider the symmetric matrix sensing problem (17). Suppose that the measurement matrices $\Gamma_1, \dots, \Gamma_m$ have i.i.d. standard normal entries, and that the outlier vector ε is independent of the measurement matrices and satisfies $\|\varepsilon\|_0 \leq pm$ for some $0 \leq p < 1/2$. Let $X^* \succeq 0$ have rank r , and suppose that $k \geq r + 1$ and $n \geq r + 2$. Then there exist constants $c_1, c_2 > 0$, depending only on p , such that, whenever $m \geq c_1 nr$, with probability at least*

$$1 - e^{-c_2 nr} - 2^{-(n-r-1)},$$

every ground-truth factor U^ satisfying $U^* U^{*\top} = X^*$ is an inactive critical point of f_{sym} .*

Theorem 5.3 (Asymmetric matrix sensing) *Consider the asymmetric matrix sensing problem (18). Suppose that the measurement matrices $\Gamma_1, \dots, \Gamma_m$ have i.i.d. standard normal entries, and that the outlier vector ε is independent of the measurement matrices and satisfies $\|\varepsilon\|_0 \leq pm$ for some $0 \leq p < 1/2$. Let $X^* \in \mathbb{R}^{n_1 \times n_2}$ have rank r , and suppose that $k \geq r + 1$, $n_1 \geq r + 1$, and $n_2 \geq r + 1$. Then there exist constants $c_1, c_2 > 0$, depending only on p , such that, whenever $m \geq c_1(n_1 + n_2)r$, with probability at least*

$$1 - e^{-c_2(n_1 + n_2)r},$$

every balanced ground-truth factor pair (U_1^, U_2^*) satisfying $U_1^* U_2^{*\top} = X^*$ and $U_1^{*\top} U_1^* = U_2^{*\top} U_2^*$ is an inactive critical point of f_{asym} .*

Several observations regarding Theorems 5.2 and 5.3 are in order. First, Ma and Fattahi [52, Corollary 1] show that, with high probability, every ground-truth factor U^* satisfying $U^* U^{*\top} = X^*$ is a strict saddle point of f_{sym} , provided that $cnr \leq m \leq c'nk$ for suitable constants $c, c' > 0$. Theorem 5.2 complements this result by establishing that these ground-truth factors are also inactive. Consequently, within this sample-size

regime, they are inactive strict saddles with high probability. This identifies inactivity as the geometric mechanism underlying the failure of the subgradient method to escape these strict saddles. An analogous conclusion for asymmetric matrix sensing follows by combining Ma and Fattahi [52, Corollary 2] with Theorem 5.3. Second, as in nonnegative sparse regression, overparameterization—that is, $k > r$ —is necessary for the ground-truth factors to be inactive. Indeed, when $k = r$, Charisopoulos et al. [19, Propositions 4.6 and 4.7] establish that f_{sym} and f_{asym} grow sharply away from their respective ground-truth solution manifolds. Since these manifolds are C^2 smooth, this sharpness implies that the corresponding ground-truth factors are active critical points. Finally, the sample-size requirements in both theorems match the corresponding information-theoretic limits up to constant factors.

Structure of the proof and geometric intuition: The crux of our analysis is to show how the branching criterion manifests itself in robust low-rank recovery. In Section 5.1, we reduce the criterion to a deterministic cone-selection test involving the normal cones at the vertices of the clean multiplier polytope. Geometrically, each tangent direction in a hypothetical identifiable manifold generates a projected second-order residual, and the normal cone containing this residual determines the branch that is accessed. If two relatively open sets of tangent directions select normal cones corresponding to distinct vertices, then they access branches with incompatible quadratic expansions, forcing the ground-truth factor to be inactive. In Section 5.2, we show that, after an appropriate whitening, Gaussian measurements make the directions of these projected residuals uniformly distributed on the relevant sphere. Since each pointed normal cone occupies at most half of the sphere, the projected residuals are unlikely to all select the same branch, yielding the desired high-probability guarantees.

5.1 Branching Criterion and Cone Selection

Problem (15) takes the finite-max composite form of Section 2.4, $f_{\text{LR}} = g \circ F$, where

$$F(U) = \mathcal{A}(\mathcal{L}(UU^\top)) - y = \mathcal{A}(\mathcal{L}(UU^\top - X^\star)) - \varepsilon, \quad g(z) = \|z\|_1.$$

Here $F : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^m$ is C^2 and definable, and $g(z) = \max_{\sigma \in \{\pm 1\}^m} g_\sigma(z)$ with each piece $g_\sigma(z) := \langle \sigma, z \rangle$ being linear. Let $\Omega_C := \{i : \varepsilon_i = 0\}$ and $\Omega_N := \{i : \varepsilon_i \neq 0\}$ denote the index sets of clean and noisy measurements, with $m_C := |\Omega_C|$ and $m_N := |\Omega_N|$. We write $\mathcal{A}_C : \mathbb{S}^n \rightarrow \mathbb{R}^{m_C}$ and $F_C : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{m_C}$ for the restrictions of $\mathcal{A}$ and F to the clean measurements, and $\mathcal{A}_N, F_N$ for the corresponding restrictions to the noisy measurements. Analogously, we write $\varepsilon_N \in \mathbb{R}^{m_N}$ to denote the restriction of ε to the noisy measurements. The clean and noisy Jacobians at a ground-truth factor $U^\star$ are defined as

$$\mathcal{J}_C(U^\star)[V] := DF_C(U^\star)[V] = \mathcal{A}_C(\mathcal{L}(U^\star V^\top + V U^{\star\top})), \quad (19)$$

$$\mathcal{J}_N(U^\star)[V] := DF_N(U^\star)[V] = \mathcal{A}_N(\mathcal{L}(U^\star V^\top + V U^{\star\top})). \quad (20)$$

When the base point U^* is clear from context, we suppress it from the notation and write $\mathcal{J}_C[V]$ and $\mathcal{J}_N[V]$ in place of $\mathcal{J}_C(U^*)[V]$ and $\mathcal{J}_N(U^*)[V]$, respectively. Likewise, $\mathcal{J}_C^*[v]$ and $\mathcal{J}_N^*[v]$ denote the actions of the corresponding adjoint linear maps on v .

The following proposition gives the inclusion-minimal object that any C^2 identifiable manifold at a critical point U^* would have to equal locally.

Proposition 5.4 (Clean reduction of the identifiable set) *Fix a critical point U^* of f_{LR} satisfying $U^*U^{*\top} = X^*$. Define the clean multiplier polytope*

$$\Lambda_C := \{v \in [-1, 1]^{m_C} : \mathcal{J}_C^*[v] = \mathcal{J}_N^*[\text{sign}(\varepsilon_N)]\}, \quad (21)$$

Then, letting $\mathcal{S}_v$ denote the sign cone (3) in $\mathbb{R}^{m_C}$, the set

$$\mathcal{M} = F_C^{-1}(\mathcal{E}_C), \quad \mathcal{E}_C := \bigcup_{v \in \Lambda_C} \mathcal{S}_v \subset \mathbb{R}^{m_C} \quad (22)$$

is the locally minimal identifiable set for f_{LR} at U^ for $0 \in \partial f_{\text{LR}}(U^*)$.*

The proof of Proposition 5.4 is presented in Appendix B.2. It shows that the branching criterion can be studied on $F_C^{-1}(\mathcal{E}_C)$, where $\mathcal{E}_C$ is the identifiable set in the clean coordinates before pullback by F_C ; the corrupted coordinates contribute only a smooth, sign-fixed term to the local expansion of the objective. Consequently, the nonsmoothness of f sits entirely on the clean coordinates. To obtain a more tractable description of $\mathcal{E}_C$, we introduce

$$\mathcal{R}_C := \text{Im } \mathcal{J}_C, \quad \mathcal{N}_C := \ker \mathcal{J}_C^* = \mathcal{R}_C^\perp. \quad (23)$$

in terms of which Λ_C in (21) can be written as

$$\Lambda_C = (\bar{v} + \mathcal{N}_C) \cap [-1, 1]^{m_C},$$

where $\bar{v} \in \Lambda_C$ is arbitrary. Recall from Example 2.12 that a coordinate i of $\nu \in [-1, 1]^{m_C}$ is saturated when $|\nu_i| = 1$ and free when $|\nu_i| < 1$. Write $\mathcal{I}_<(\nu) := \{i : |\nu_i| < 1\}$ for the free coordinates of ν . Next, we impose two geometric conditions on the clean multiplier polytope Λ_C . Together they make Λ_C a full-dimensional affine-cube section whose vertices index the full-dimensional cones of $\mathcal{E}_C$.

Assumption 5.5 (Affine-cube geometry of the multiplier polytope) The clean multiplier polytope Λ_C satisfies:

- (I) (*Interior condition*) Λ_C meets the interior of the cube, i.e., $\Lambda_C \cap (-1, 1)^{m_C} \neq \emptyset$;
- (N) (*Nondegeneracy condition*) For every vertex ν of Λ_C , the restriction $\mathcal{J}_{C, \mathcal{I}_<(\nu)} : \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{|\mathcal{I}_<(\nu)|}$ of $\mathcal{J}_C$ to the free coordinates of ν has full rank, i.e., $\text{rank } \mathcal{J}_{C, \mathcal{I}_<(\nu)} = \text{rank } \mathcal{J}_C$.

In Section 5.2, we show that under appropriate randomness in the model, Assumption 5.5 holds with high probability. Under these assumptions, our next proposition shows that the set $\mathcal{E}_C$ decomposes into the sign cones $\mathcal{S}_\nu$ over the vertices of Λ_C , which are equivalently the normal cones of $[-1, 1]^{m_C}$ at the vertices of Λ_C , where each normal cone has dimension $\dim \mathcal{N}_C$.

Proposition 5.6 (Structure of $\mathcal{E}_C$) *The set $\mathcal{E}_C$ of (22) is a polyhedral fan:*

$$\mathcal{E}_C = \bigcup_{\nu \in \text{Vert}(\Lambda_C)} \mathcal{S}_\nu, \quad (24)$$

with distinct vertices ν yielding distinct sign cones $\mathcal{S}_\nu$. Moreover:

- Under condition (I), Λ_C is a full-dimensional polytope of dimension $\dim \mathcal{N}_C$, and $N_{\Lambda_C}(v) = \mathcal{R}_C + \mathcal{S}_v$ for every $v \in \Lambda_C$.
- Under conditions (I) and (N), every cone $\mathcal{S}_\nu$ for $\nu \in \text{Vert}(\Lambda_C)$ is full-dimensional with $\dim \mathcal{S}_\nu = \dim \mathcal{N}_C$, and the union (24) is irredundant: no $\mathcal{S}_\nu$ is contained in another.

Proof Recall the construction $\mathcal{E}_C = \bigcup_{v \in \Lambda_C} \mathcal{S}_v$ from Proposition 5.4, where $\mathcal{S}_v$ is the sign cone of v , or equivalently the normal cone to the cube at v ,

$$\mathcal{S}_v := \{z \in \mathbb{R}^{m_C} : z_i = 0 \text{ if } |v_i| < 1, z_i \geq 0 \text{ if } v_i = 1, z_i \leq 0 \text{ if } v_i = -1\}.$$

For two vertices ν, ν' , the intersection $\mathcal{S}_\nu \cap \mathcal{S}_{\nu'}$ is the set of z supported on the coordinates $\{i : |\nu_i| = |\nu'_i| = 1 \text{ and } \nu_i = \nu'_i\}$. This is a common face of the two cones, and every face of $\mathcal{S}_\nu$ arises in the same way. The cones $\mathcal{S}_\nu$ together with their faces are therefore a polyhedral fan. To establish (24), take $v \in \Lambda_C$ and let ν be a vertex of the smallest face of Λ_C containing v . Passing from v to ν only saturates further coordinates, so the saturated set of v lies in that of ν and $\mathcal{S}_v \subseteq \mathcal{S}_\nu$, which gives (24). To see distinct vertices give distinct cones, note that if $\mathcal{S}_\nu = \mathcal{S}_{\nu'}$, ν and ν' have the same saturated coordinates and signs. A vertex has no nonzero direction of $\mathcal{N}_C$ along its free coordinates, since such a direction would place it inside a segment of Λ , so its saturated coordinates determine it and $\nu = \nu'$.

Now assume condition (I), and fix $\bar{v} \in \Lambda_C \cap (-1, 1)^{m_C}$. As $\bar{v}$ is interior to the cube and $\Lambda_C = (\bar{v} + \mathcal{N}_C) \cap [-1, 1]^{m_C}$, Λ_C contains a relative neighborhood of $\bar{v}$ in $\bar{v} + \mathcal{N}_C$. Hence, it is full-dimensional there, of dimension $\dim \mathcal{N}_C$. For the normal cone formula, let $v \in \Lambda_C$ be arbitrary. Since $\bar{v} + \mathcal{N}_C$ meets the interior of the cube, the normal-cone sum rule [54, Corollary 23.8.1] applies at v : the normal cone to $\bar{v} + \mathcal{N}_C$ is $\mathcal{N}_C^\perp = \mathcal{R}_C$ and the normal cone to the cube at v is the sign cone $\mathcal{S}_v$, so $N_{\Lambda_C}(v) = N_{\bar{v} + \mathcal{N}_C}(v) + N_{[-1, 1]^{m_C}}(v) = \mathcal{R}_C + \mathcal{S}_v$.

Now assume conditions (I) and (N), and let ν be a vertex of Λ_C . By condition (N), the restriction $\mathcal{J}_{C, \mathcal{I}_<(\nu)}$ has rank equal to $\text{rank } \mathcal{J}_C$, which requires $|\mathcal{I}_<(\nu)| \geq \text{rank } \mathcal{J}_C$. For the reverse inequality, we claim ν has at least $m_C - \text{rank } \mathcal{J}_C$ saturated coordinates, i.e., $|\mathcal{I}_<(\nu)| \leq \text{rank } \mathcal{J}_C$. Indeed, if ν had fewer than $m_C - \text{rank } \mathcal{J}_C$ saturated coordinates, there would exist a nonzero direction $h \in \mathcal{N}_C$ supported on the free coordinates of ν , placing ν in the interior of a segment of Λ_C and contradicting that ν is a vertex. Hence $|\mathcal{I}_<(\nu)| = \text{rank } \mathcal{J}_C$, and therefore

$$\dim \mathcal{S}_\nu = m_C - |\mathcal{I}_<(\nu)| = m_C - \text{rank } \mathcal{J}_C = \dim \mathcal{N}_C.$$

Finally, if $\mathcal{S}_\nu \subseteq \mathcal{S}_{\nu'}$ for two vertices ν, ν' , then $\mathcal{S}_\nu$ is a face of $\mathcal{S}_{\nu'}$ of equal dimension, hence equals it, giving $\nu = \nu'$. The union (24) is therefore irredundant. $\square$

In light of Proposition 5.6, the set $\mathcal{E}_C = \bigcup_{\nu \in \text{Vert}(\Lambda_C)} \mathcal{S}_\nu$ forms a polyhedral fan and admits a natural stratification: the top-dimensional strata are the relative interiors $\text{relint } \mathcal{S}_\nu$ of the maximal cones, one per vertex ν of Λ_C , and lower-dimensional strata appear along their common faces. The branching criterion then compares the second-order behavior across these top-dimensional strata. Figure 5 illustrates the geometry of Λ_C , $\mathcal{S}_\nu$, and $\mathcal{E}_C$ for $\dim \mathcal{N}_C = 2$ with $m_C = 3$.

To apply the branching criterion (Theorem 3.9), we work under the activity hypothesis and show that it forces two incompatible accessible branches. Specifically, we




Fig. 5: The case $m_C = 3$, $\dim \mathcal{N}_C = 2$. *Left:* the clean multiplier polytope Λ_C is a hexagon obtained by intersecting $\bar{v} + \mathcal{N}_C$ with $[-1, 1]^3$, satisfying Assumption 5.5. Three of its six vertices are ν_1 (purple), ν_2 (orange), and ν_3 (cyan). *Right:* by Proposition 5.6, $\mathcal{E}_C = \bigcup_{\nu} \mathcal{S}_{\nu}$ over the six vertices; the cones of ν_1, ν_2, ν_3 are shown in matching colors, each two-dimensional ($\dim \mathcal{S}_{\nu} = \dim \mathcal{N}_C = 2$). Adjacent cones share a lower-dimensional face: for a point ι (green) on the open edge $\nu_2\nu_3$, the cone $\mathcal{S}_{\iota} = \mathcal{S}_{\nu_2} \cap \mathcal{S}_{\nu_3}$ (green) is their common boundary ray.

must determine (a) which cones $\mathcal{S}_{\nu}$ are accessible from $\mathcal{M}$ at U^* along directions in some subspace $\mathcal{W} \subseteq T_{U^*}\mathcal{M}$, and (b) whether their branch quadratics are incompatible. Condition (b) turns out to be easy: any two distinct strata $\text{relint } \mathcal{S}_{\nu_1}$ and $\text{relint } \mathcal{S}_{\nu_2}$ yield incompatible branch quadratics. Condition (a) is more delicate because it requires characterizing how curves on $\mathcal{M}$ through U^* behave after being pushed forward through F_C . As we show below, accessibility reduces to a cone-selection test. The following lemma characterizes $T_{U^*}\mathcal{M}$ under the activity hypothesis.

Lemma 5.7 (Tangent space of the clean identifiable set) *Fix a critical point U^* of f_{LR} satisfying $U^*U^{*\top} = X^*$. Under Assumption 5.5, assume $\mathcal{M} = F_C^{-1}(\mathcal{E}_C)$ is a C^2 identifiable manifold for f_{LR} at U^* for 0. Then,*

$$T_{U^*}\mathcal{M} = \ker \mathcal{J}_C.$$

We defer the proof of this lemma to Appendix B.3 as it requires tools we shall develop later.⁴ Under the activity hypothesis, Section 3.2 locally represents $\mathcal{M}$, along each tangent direction $V \in T_{U^*}\mathcal{M}$, by a curve starting from U^* with velocity V . Consider such a curve

$$\gamma_V(t) = U^* + tV + t^2Z + o(t^2) \quad (25)$$

with some second-order correction Z (cf. Lemma 2.1). Defining the clean second-order residual

$$\mathcal{Q}_C(V) := \frac{1}{2}D^2F_C(U^*)[V, V] = \mathcal{A}_C(\mathcal{L}(VV^\top)) \quad (26)$$

⁴One might expect this to follow directly from the standard characterization of the tangent space as the kernel of a local defining map [40, Proposition 5.38]. However, we cannot assume $\mathcal{J}_C$ to be a local defining map for $\mathcal{M}$, so this standard result does not apply directly.

and passing the curve $\gamma_V(t)$ through F_C , we obtain

$$\begin{aligned} F_C(\gamma_V(t)) &= F_C(U^*) + t\mathcal{J}_C[V] + t^2(\mathcal{Q}_C(V) + \mathcal{J}_C[Z]) + o(t^2) \\ &= t^2(\mathcal{Q}_C(V) + \mathcal{J}_C[Z]) + o(t^2) \in \mathcal{E}_C, \end{aligned} \quad (27)$$

where we use $F_C(U^*) = 0$ and $\mathcal{J}_C[V] = 0$, the latter following from $V \in \ker \mathcal{J}_C$ by Lemma 5.7. Since $\mathcal{J}_C[Z] \in \mathcal{R}_C$, the leading coefficient of $F_C(\gamma_V(t))$ lies in the affine subspace $\mathcal{Q}_C(V) + \mathcal{R}_C$ and in $\mathcal{E}_C$. The curvature of $\mathcal{M}$ along γ enters only through $\mathcal{J}_C[Z] \in \mathcal{R}_C$. The next proposition shows that this curvature term is immaterial: the affine subspace $\mathcal{Q}_C(V) + \mathcal{R}_C$ meets $\mathcal{E}_C$ in a single point, determined by $\mathcal{Q}_C(V)$ alone.

Proposition 5.8 (Second-order cone selection) *Under Assumption 5.5, for every $a \in \mathbb{R}^{m_C}$, the affine slice $(a + \mathcal{R}_C) \cap \mathcal{E}_C$ is a single point, denoted $\mathcal{S}_{\Lambda_C}(a)$:*

$$\{\mathcal{S}_{\Lambda_C}(a)\} := (a + \mathcal{R}_C) \cap \mathcal{E}_C. \quad (28)$$

The selection map $\mathcal{S}_{\Lambda_C}(a)$ depends only on the projection $P_{\mathcal{N}_C}a$, so it descends to a map $\mathcal{N}_C \rightarrow \mathcal{E}_C$. This descended map is a two-sided inverse of the projection $P_{\mathcal{N}_C}$ restricted to $\mathcal{E}_C$, so $P_{\mathcal{N}_C}$ restricts to a bijection $\mathcal{E}_C \rightarrow \mathcal{N}_C$. Moreover $\mathcal{S}_{\Lambda_C}$ is positively homogeneous: $\mathcal{S}_{\Lambda_C}(\lambda a) = \lambda \mathcal{S}_{\Lambda_C}(a)$ for all $\lambda \geq 0$.

The proof of Proposition 5.8 is presented in Appendix B.4. It implies that $\mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V))$ depends only on the projection $P_{\mathcal{N}_C}\mathcal{Q}_C(V)$. Geometrically, this projection flattens the polyhedral fan $\mathcal{E}_C$ along the first-order range $\mathcal{R}_C = \text{Im } \mathcal{J}_C$. By Proposition 5.8, $P_{\mathcal{N}_C}$ restricts to a bijection $\mathcal{E}_C \rightarrow \mathcal{N}_C$ with inverse $\mathcal{S}_{\Lambda_C}$, so no information is lost under projection. Each cone $\mathcal{S}_\nu$ maps to its image $P_{\mathcal{N}_C}\mathcal{S}_\nu$, one per vertex ν of Λ_C , and these images partition $\mathcal{N}_C$, meeting only along their boundaries. Since Λ_C lies in the affine subspace $\bar{v} + \mathcal{N}_C$, let $N_{\Lambda_C}^{\mathcal{N}_C}(\nu) \subset \mathcal{N}_C$ denote the normal cone of Λ_C at ν relative to $\bar{v} + \mathcal{N}_C$. The normal cone formula of Proposition 5.6 gives

$$P_{\mathcal{N}_C}\mathcal{S}_\nu = N_{\Lambda_C}^{\mathcal{N}_C}(\nu),$$

so flattening identifies each maximal cone of $\mathcal{E}_C$ with a normal cone of Λ_C in $\mathcal{N}_C$, as illustrated in Figure 6. Consequently, determining which cone $\text{relint } \mathcal{S}_\nu$ contains $\mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V))$ is equivalent to identifying which normal cone $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ contains $P_{\mathcal{N}_C}\mathcal{Q}_C(V)$, as formally established in the next lemma.

Lemma 5.9 (Projected cone selection) *Under Assumption 5.5, let ν be a vertex of Λ_C and $a \in \mathbb{R}^{m_C}$. Then*

$$\mathcal{S}_{\Lambda_C}(a) \in \text{relint } \mathcal{S}_\nu \quad \Longleftrightarrow \quad P_{\mathcal{N}_C}a \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu).$$

Proof By Proposition 5.6, $N_{\Lambda}(\nu) = \mathcal{R}_C + \mathcal{S}_\nu$. Therefore, $P_{\mathcal{N}_C}\mathcal{S}_\nu = N_{\Lambda_C}(\nu) \cap \mathcal{N}_C = N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$. By the linearity of $P_{\mathcal{N}_C}$, we have

$$P_{\mathcal{N}_C} \text{span } \mathcal{S}_\nu = \text{span } P_{\mathcal{N}_C}\mathcal{S}_\nu = \text{span } N_{\Lambda_C}^{\mathcal{N}_C}(\nu),$$




Fig. 6: Cone selection via flattening, continuing the hexagonal example of Figure 4. *Left:* the polyhedral fan $\mathcal{E}_C \subset \mathbb{R}^{m_C}$ with maximal cones $\mathcal{S}_{\nu_1}, \dots, \mathcal{S}_{\nu_6}$, one per vertex ν_i of Λ_C , colored as in Figure 4. Two tangent directions $V_1, V_2 \in \ker \mathcal{J}_C$ generate curves $\gamma_{V_1}, \gamma_{V_2}$ of the form (25); by Proposition 5.8, their images $F_C(\gamma_{V_1})$ (blue) and $F_C(\gamma_{V_2})$ (red) leave the origin along the rays spanned by the selected points $\mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V_1)) \in \text{relint } \mathcal{S}_{\nu_3}$ and $\mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V_2)) \in \text{relint } \mathcal{S}_{\nu_6}$. *Right:* the image of $\mathcal{E}_C$ under the projection $P_{\mathcal{N}_C}$ (arrow between the panels), inside $\mathcal{N}_C$, with Λ_C overlaid (dark green, translated from $\bar{v} + \mathcal{N}_C$ to $\mathcal{N}_C$). Each cone $\mathcal{S}_{\nu_i}$ flattens to the normal cone $N_{\Lambda_C}^{\mathcal{N}_C}(\nu_i)$ at the corresponding vertex, drawn in the matching color, and these cover $\mathcal{N}_C$. By Lemma 5.9, the cone selected along V_j is the one whose normal cone contains the projected direction $P_{\mathcal{N}_C} \mathcal{Q}_C(V_j)$: the blue and red arrows depict $P_{\mathcal{N}_C} \mathcal{Q}_C(V_1) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_3)$ and $P_{\mathcal{N}_C} \mathcal{Q}_C(V_2) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_6)$, respectively, reducing cone selection from $\mathbb{R}^{m_C}$ to the lower-dimensional space $\mathcal{N}_C$.

implying that $P_{\mathcal{N}_C}$ is surjective. Additionally, note that $P_{\mathcal{N}_C}$ is injective on $\mathcal{E}_C$, and hence on $\mathcal{S}_{\nu} \subseteq \mathcal{E}_C$, due to Proposition 5.8. Since $\mathcal{S}_{\nu}$ is a convex cone, $\text{span } \mathcal{S}_{\nu} = \mathcal{S}_{\nu} - \mathcal{S}_{\nu}$. If $h = z_1 - z_2 \in \text{span } \mathcal{S}_{\nu}$ satisfies $P_{\mathcal{N}_C} h = 0$, then $P_{\mathcal{N}_C} z_1 = P_{\mathcal{N}_C} z_2$, and the injectivity of $P_{\mathcal{N}_C}$ on $\mathcal{S}_{\nu}$ gives $z_1 = z_2$, hence $h = 0$. So $P_{\mathcal{N}_C}$ is injective on $\text{span } \mathcal{S}_{\nu}$, and therefore a linear isomorphism onto $\text{span } N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$. Thus, it maps the relative interior of $\mathcal{S}_{\nu}$ onto the relative interior of $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$. That is, for any $z \in \mathcal{E}_C$,

$$z \in \text{relint } \mathcal{S}_{\nu} \iff P_{\mathcal{N}_C} z \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu).$$

Now take $z = \mathcal{S}_{\Lambda}(a)$. Since $\mathcal{S}_{\Lambda_C}(a) \in a + \mathcal{R}_C$, $P_{\mathcal{N}_C} \mathcal{S}_{\Lambda_C}(a) = P_{\mathcal{N}_C} a$. Substituting this identity into the preceding equivalence concludes the proof. $\square$

Armed with these results, we now show that, under the activity hypothesis, the two conditions in Theorem 3.9 follow from cone selection: accessibility of the strata (condition (a)) and incompatibility of their branch quadratics (condition (b)). We begin with condition (a). The following lemma shows that $P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ for V ranging over a relatively open cone in $\ker \mathcal{J}_C$ is sufficient to guarantee accessibility of the stratum $\text{relint } \mathcal{S}_{\nu}$.

Lemma 5.10 (Accessibility of a vertex cone) *Fix a critical point $U^\star$ of f_{LR} satisfying $U^\star U^{\star\top} = X^\star$, and suppose that $\mathcal{M} = F_C^{-1}(\mathcal{E}_C)$ is a C^2 identifiable manifold for f_{LR} at $U^\star$ for 0. Under Assumption 5.5, suppose there exists a vertex ν of Λ_C , a linear subspace $\mathcal{W} \subseteq \ker \mathcal{J}_C$, and a nonempty relatively open cone $\mathcal{C} \subseteq \mathcal{W}$ such that*

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu), \quad \forall V \in \mathcal{C}.$$

Then, $\text{relint } \mathcal{S}_\nu$ is accessible from $\mathcal{M}$ at $U^\star$ along $\mathcal{W}$.

Proof Let γ be the chart in (6), and fix a unit vector $V \in \mathcal{C}$. By Lemma 5.7, $V \in T_{U^\star} \mathcal{M} = \ker \mathcal{J}_C$. Hence, using the homogeneity of $Z_{\mathcal{M}}$ and the definition (26),

$$\begin{aligned} F_C(\gamma(tV)) &= t^2(\mathcal{Q}_C(V) + \mathcal{J}_C[Z_{\mathcal{M}}(V)]) + o(t^2) \\ &=: t^2 \eta_C(V) + o(t^2). \end{aligned}$$

Here $\eta_C(V) := (\eta_{\mathcal{M}}(V))_{i \in \Omega_C}$ is the clean block of the coefficient in (7). For $t > 0$, set $z_t := t^{-2} F_C(\gamma(tV))$. Since $\gamma(tV) \in \mathcal{M} = F_C^{-1}(\mathcal{E}_C)$ and $\mathcal{E}_C$ is a cone, we have $z_t \in \mathcal{E}_C$. Moreover, $z_t \rightarrow \eta_C(V)$ as $t \downarrow 0$. Proposition 5.6 expresses $\mathcal{E}_C$ as a finite union of closed polyhedral cones, so $\mathcal{E}_C$ is closed and therefore $\eta_C(V) \in \mathcal{E}_C$. On the other hand,

$$\eta_C(V) - \mathcal{Q}_C(V) = \mathcal{J}_C[Z_{\mathcal{M}}(V)] \in \mathcal{R}_C.$$

Consequently,

$$\eta_C(V) \in (\mathcal{Q}_C(V) + \mathcal{R}_C) \cap \mathcal{E}_C.$$

Proposition 5.8 now yields

$$\eta_C(V) = \mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V)). \quad (29)$$

By the cone-selection hypothesis and Lemma 5.9,

$$\eta_C(V) \in \text{relint } \mathcal{S}_\nu.$$

Because the maximal cones of $\mathcal{E}_C$ form a finite polyhedral fan, every point of $\mathcal{E}_C$ sufficiently close to $\eta_C(V)$ also lies in $\text{relint } \mathcal{S}_\nu$. Indeed, the relative boundary of $\mathcal{S}_\nu$ together with all other maximal cones is a finite union of closed polyhedral sets that does not contain $\eta_C(V)$. Since $z_t \in \mathcal{E}_C$ and $z_t \rightarrow \eta_C(V)$, it follows that $z_t \in \text{relint } \mathcal{S}_\nu$ for all sufficiently small $t > 0$. Finally, relative interiors of cones are invariant under positive scaling, and hence

$$F_C(\gamma(tV)) = t^2 z_t \in \text{relint } \mathcal{S}_\nu$$

for all sufficiently small $t > 0$. Since $V \in \mathcal{C}$ was arbitrary, Definition 3.5 applied to the reduced representation shows that $\text{relint } \mathcal{S}_\nu$ is accessible from $\mathcal{M}$ at $U^\star$ along $\mathcal{W}$. $\square$

Having established accessibility of the strata, it remains to verify that the branch quadratics are incompatible on them. As we show next, this incompatibility is in fact automatic: any two distinct strata necessarily carry incompatible branch quadratics. To state the main result, define the *boundary sign vector* $\sigma(\nu) \in \{-1, 0, 1\}^{m_C}$ by $\sigma(\nu)_i = \text{sign}(\nu_i)$ on the saturated coordinates ($|\nu_i| = 1$) and $\sigma(\nu)_i = 0$ on the free coordinates ($|\nu_i| < 1$). Its support is the set of saturated coordinates and its zero set is $\mathcal{I}_<(\nu)$. On the sign cone at ν , it computes the ℓ_1 norm,

$$\|z\|_1 = \langle \sigma(\nu), z \rangle \quad \text{for } z \in \text{relint } \mathcal{S}_\nu. \quad (30)$$

Theorem 5.11 (Branching criterion via cone selection) *Fix a critical point U^* of f_{LR} satisfying $U^*U^{*\top} = X^*$. Under Assumption 5.5, suppose there exist two distinct vertices ν_1, ν_2 of Λ_C , a linear subspace $\mathcal{W} \subseteq \ker \mathcal{J}_C$, and two nonempty relatively open cones $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{W}$ such that*

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j), \quad \forall V \in \mathcal{C}_j \quad j = 1, 2.$$

Then f_{LR} admits no C^2 identifiable manifold at U^ for $0 \in \partial f(U^*)$; that is, U^* is inactive.*

Proof By Theorem 3.9, it suffices to show that the activity hypothesis implies accessibility of two strata and incompatibility of their branch quadratics. We therefore work under the activity hypothesis throughout the remainder of the proof.

Since $P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j)$ for $V \in \mathcal{C}_j$, Lemma 5.10 shows that the strata $\mathcal{Z}_j := \text{relint } \mathcal{S}_{\nu_j}$, $j = 1, 2$, are accessible. It remains to prove the incompatibility of the branch quadratics on $\mathcal{Z}_1$ and $\mathcal{Z}_2$. Fix a piece τ_j active on $\mathcal{Z}_j$, i.e. $\tau_j \in \mathcal{I}_{\mathcal{Z}_j}$, and let $\zeta_{\mathcal{Z}_j, \mathcal{W}} := \zeta_{\tau_j}|_{\mathcal{W}}$ denote the corresponding branch quadratic, as in Lemma 3.7, with ζ_{τ_j} defined in (8). Since $f_{\text{LR}} = g \circ F$ with g a finite maximum of linear functions, the Hessian term for g_{τ_j} in (8) vanishes, so

$$\zeta_{\tau_j}(V) = \langle \tau_j, \eta_{\mathcal{M}}(V) \rangle = \langle \tau_{j,C}, \eta_C(V) \rangle + \underbrace{\langle -\text{sign}(\varepsilon_N), \eta_N(V) \rangle}_{:= q_N(V)},$$

where $\tau_{j,C} := (\tau_j)_{i \in \Omega_C}$ and $\eta_C := (\eta_{\mathcal{M}})_{i \in \Omega_C}$ are the restrictions of τ_j and $\eta_{\mathcal{M}}$ to the clean indices, and $\eta_N := (\eta_{\mathcal{M}})_{i \in \Omega_N}$ is the restriction to the noisy indices. Note that q_N is independent of the choice of active piece τ_j . By (27) and Proposition 5.8,

$$\eta_C(V) = \mathcal{Q}_C(V) + \mathcal{J}_C[Z] = \mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V)),$$

The assumption $P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j)$ for $V \in \mathcal{C}_j$ combined with Lemma 5.9 implies $\mathcal{S}_{\Lambda_C}(\mathcal{Q}_C(V)) \in \text{relint } \mathcal{S}_{\nu_j}$ for $V \in \mathcal{C}_j$, and hence $\eta_C(V) \in \text{relint } \mathcal{S}_{\nu_j}$ for $V \in \mathcal{C}_j$. In particular $\eta_C(V)$ vanishes on the free coordinates of ν_j , while $\tau_{j,C}$ agrees with $\sigma(\nu_j)$ on the saturated ones. Consequently,

$$\langle \tau_{j,C}, \eta_C(V) \rangle = \langle \sigma(\nu_j), \eta_C(V) \rangle =: \psi_{\nu_j}(V), \quad \forall V \in \mathcal{C}_j, \quad j = 1, 2,$$

and therefore

$$\zeta_{\tau_j}(V) = q_N(V) + \psi_{\nu_j}(V), \quad \forall V \in \mathcal{C}_j, \quad j = 1, 2.$$

Since $\zeta_{\mathcal{Z}_j, \mathcal{W}}$ and ζ_{τ_j} are both quadratic forms on $\mathcal{W}$ agreeing on the nonempty relatively open cone $\mathcal{C}_j$, Lemma 3.6 shows they agree on all of $\mathcal{W}$:

$$\zeta_{\mathcal{Z}_j, \mathcal{W}}(V) = q_N(V) + \psi_{\nu_j}(V), \quad \forall V \in \mathcal{W}, \quad j = 1, 2.$$

Hence $\zeta_{\mathcal{Z}_1, \mathcal{W}} - \zeta_{\mathcal{Z}_2, \mathcal{W}} = \psi_{\nu_1} - \psi_{\nu_2}$ on $\mathcal{W}$, and it suffices to show $\psi_{\nu_1} \not\equiv \psi_{\nu_2}$ on $\mathcal{W}$. For $V \in \mathcal{C}_1$, since $\eta_C(V) \in \text{relint } \mathcal{S}_{\nu_1}$, its support equals the saturated coordinates of ν_1 with matching sign, so

$$\psi_{\nu_1}(V) = \langle \sigma(\nu_1), \eta_C(V) \rangle = \|\eta_C(V)\|_1, \quad \forall V \in \mathcal{C}_1.$$

Since $\sigma(\nu_2) \in \{-1, 0, +1\}^{m_C}$, by Hölder's inequality,

$$\psi_{\nu_2}(V) = \langle \sigma(\nu_2), \eta_C(V) \rangle \leq \|\eta_C(V)\|_1, \quad \forall V \in \mathcal{C}_1,$$

with equality only if $\sigma(\nu_2)_i = \text{sign}(\eta_C(V)_i)$ for every $i \in \text{supp}(\eta_C(V))$, that is, only if every coordinate saturated by ν_1 is also saturated by ν_2 with the same sign, forcing $\eta_C(V) \in \mathcal{S}_{\nu_2}$. But $\mathcal{S}_{\nu_1}$ and $\mathcal{S}_{\nu_2}$ are distinct maximal cones of the polyhedral fan due to Proposition 5.6, so they meet only along a common proper face, and $\text{relint } \mathcal{S}_{\nu_1}$ is disjoint from $\mathcal{S}_{\nu_2}$. Since $\eta_C(V) \in \text{relint } \mathcal{S}_{\nu_1}$ for $V \in \mathcal{C}_1$, we conclude $\eta_C(V) \notin \mathcal{S}_{\nu_2}$, and thus $\psi_{\nu_2}(V) < \psi_{\nu_1}(V)$ strictly for all $V \in \mathcal{C}_1$. As $\mathcal{C}_1 \subseteq \mathcal{W}$ is nonempty, this shows $\psi_{\nu_1} \not\equiv \psi_{\nu_2}$ on $\mathcal{W}$. Thus the activity hypothesis implies the two incompatible accessible branches required by Theorem 3.9, completing the proof. $\square$

Theorem 5.11 reduces the branching criterion to a cone-selection test in $\mathcal{N}_C$. The next subsection verifies that this test succeeds with high probability under appropriate random models.

5.2 Probabilistic Analysis

In this subsection, we show that the cone-selection test—and hence the branching criterion—succeeds with high probability for robust low-rank recovery under the random models considered here. Before turning to the formal analysis, we develop geometric intuition for why this should occur. This intuition applies directly to nonnegative sparse regression and symmetric matrix sensing. Asymmetric matrix sensing requires an analogous but slightly sharper argument, which we present in Appendix B.10.

Suppose for now that the true solution U^* is a critical point of f_{LR} and that Assumption 5.5 (the interior and nondegeneracy conditions) holds. Suppose further, purely for the sake of intuition, that there exist $q \geq 2$ distinct directions $V_1, \dots, V_q \in \ker \mathcal{J}_C$ whose projected second-order residuals $P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$ are spherically symmetric on $\mathcal{N}_C$, i.e., their normalized directions are uniformly distributed on the unit sphere. Under this idealized randomness assumption, the cone-selection test succeeds unless all of these projected residuals lie in the relative interior of a single normal cone $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ for some vertex ν of Λ_C .⁵

To estimate the probability of this event, first recall from Proposition 5.6 that the clean multiplier polytope Λ_C is full-dimensional in $\mathcal{N}_C$. Hence its vertex normal cones $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$, $\nu \in \text{Vert}(\Lambda_C)$ are pointed [55, Theorem 2.3.2]. A pointed cone is always contained in some half-space. Consequently, for any two indices $1 \leq i < j \leq q$, the probability that both $P_{\mathcal{N}_C} \mathcal{Q}_C(V_i)$ and $P_{\mathcal{N}_C} \mathcal{Q}_C(V_j)$ fall into the *same* normal cone $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ is at most $1/2$. Consequently, the probability that all q projected residuals land in the same cone is at most $2^{-(q-1)}$, and the cone selection test succeeds whenever this does not happen, giving

$$\mathbb{P}(\text{cone-selection test succeeds}) \geq 1 - 2^{-(q-1)}.$$

This intuition is formalized in the following proposition.

Proposition 5.12 (Spherically symmetric projected residuals imply successful cone selection) *Fix a critical point U^* of f_{LR} satisfying $U^* U^{*\top} = X^*$. Suppose that Assumption 5.5 holds and that $\dim \mathcal{N}_C \geq 1$. Let $q \geq 2$, and suppose there exist distinct nonzero directions $V_1, \dots, V_q \in \ker \mathcal{J}_C$ such that, conditional on $\mathcal{J}_C$ and Λ_C , the projected residuals $P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$ are independent and have continuous, spherically symmetric distributions on $\mathcal{N}_C$. Then, conditional on $\mathcal{J}_C$ and Λ_C , with probability at least $1 - 2^{-(q-1)}$, the cone-selection test in Theorem 5.11 succeeds. Specifically, there exist two distinct vertices ν_1, ν_2 of Λ_C , a linear subspace $\mathcal{W} \subseteq \ker \mathcal{J}_C$, and two nonempty relatively open cones $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{W}$ such that*

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j) \quad \text{for every } V \in \mathcal{C}_j, \quad j = 1, 2.$$

⁵By Theorem 5.11, the cone-selection test in fact requires $P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ not just for a single V , but for all V in some relatively open cone $\mathcal{C}$. We suppress this distinction here for the sake of intuition.

The proof is presented in Appendix B.5. For this result to materialize, three conditions must hold: (1) the true solution U^* is a critical point of f_{LR} ; (2) Assumption 5.5 (nonempty relative interior and nondegeneracy of Λ_C) holds; and (3) the projected residuals $P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$ are independent and spherically symmetric.

Condition (1) has already been established in the literature for all three models considered here. We now show that condition (2) can be verified through a *whitening procedure*. Define the linearized design map at U^* by

$$\mathcal{T}(U^*) : \mathbb{R}^{n \times k} \rightarrow \mathbb{S}^n, \quad \mathcal{T}(U^*)[V] := \mathcal{L} \left(U^* V^\top + V U^{*\top} \right).$$

When the base point U^* is clear from context, we suppress it and simply write $\mathcal{T}$. The definitions of the clean and noisy Jacobians then give $\mathcal{J}_C = \mathcal{A}_C \circ \mathcal{T}$ and $\mathcal{J}_N = \mathcal{A}_N \circ \mathcal{T}$, and hence, $\mathcal{J}_C^* = \mathcal{T}^* \circ \mathcal{A}_C^*$ and $\mathcal{J}_N^* = \mathcal{T}^* \circ \mathcal{A}_N^*$. In particular, both adjoint Jacobians take values in $\text{Im } \mathcal{T}^*$. We write

$$\dim \text{Im } \mathcal{T}^* = \dim \text{Im } \mathcal{T} =: d_\star.$$

Assumption 5.13 (Whitening operator) There exists a linear isomorphism $\mathcal{H} : \text{Im } \mathcal{T}^* \rightarrow \mathbb{R}^{d_\star}$, called the *whitening operator*, such that the whitened adjoint Jacobians $\tilde{\mathcal{J}}_C^* = \mathcal{H} \circ \mathcal{J}_C^*$ and $\tilde{\mathcal{J}}_N^* = \mathcal{H} \circ \mathcal{J}_N^*$ satisfy

$$\tilde{\mathcal{J}}_C^*[e_i] \sim N(0, I_{d_\star}) \text{ i.i.d. for } i \in \Omega_C, \quad \tilde{\mathcal{J}}_N^*[e_j] \sim N(0, I_{d_\star}) \text{ i.i.d. for } j \in \Omega_N,$$

where e_i denotes the i -th standard basis vector of $\mathbb{R}^{m_C}$ or $\mathbb{R}^{m_N}$, as appropriate.

Our next lemma shows that this whitening operator exists for Gaussian measurement operators $\mathcal{A}$, which underlie symmetric and asymmetric matrix sensing as well as nonnegative sparse regression.

Lemma 5.14 (Existence of a whitening operator) Fix a ground-truth factor U^* and let $\mathcal{T}$ be its associated linearized design map. Suppose that $A_1, \dots, A_m$ are independent of the outlier vector ε . Suppose further that there exist a linear isometry

$$\iota_\star : \text{Im } \mathcal{T} \rightarrow \mathbb{R}^{d_\star}, \quad d_\star := \dim \text{Im } \mathcal{T},$$

and a positive-definite matrix $\Sigma \succ 0$ such that

$$\iota_\star(P_{\text{Im } \mathcal{T}} A_1), \dots, \iota_\star(P_{\text{Im } \mathcal{T}} A_m) \stackrel{\text{i.i.d.}}{\sim} N(0, \Sigma).$$

Then Assumption 5.13 holds. In particular, a valid whitening operator is

$$\mathcal{H} := \Sigma^{-1/2} \circ \iota_\star \circ (\mathcal{T}^*|_{\text{Im } \mathcal{T}})^{-1} : \text{Im } \mathcal{T}^* \rightarrow \mathbb{R}^{d_\star}.$$

Moreover, the hypothesis above holds for the Gaussian measurement ensembles underlying nonnegative sparse regression, symmetric matrix sensing, and asymmetric matrix sensing.

Its proof is presented in Appendix B.6. To see the utility of the whitening procedure, note that since $\mathcal{H}$ is an isomorphism from $\text{Im } \mathcal{T}^*$ onto $\mathbb{R}^{d_\star}$ and both $\text{Im } \mathcal{J}_C^*$ and $\text{Im } \mathcal{J}_N^*$ belong to $\text{Im } \mathcal{T}^*$, applying $\mathcal{H}$ preserves the solution set defining Λ_C :

$$\Lambda_C = \{v \in [-1, 1]^{m_C} : \mathcal{J}_C^*[v] = \mathcal{J}_N^*[\text{sign}(\varepsilon_N)]\}$$

$$\begin{aligned}
&= \{v \in [-1, 1]^{m_C} : \mathcal{H} \circ \mathcal{J}_C^*[v] = \mathcal{H} \circ \mathcal{J}_N^*[\text{sign}(\varepsilon_N)]\} \\
&= \left\{v \in [-1, 1]^{m_C} : \tilde{\mathcal{J}}_C^*[v] = \tilde{\mathcal{J}}_N^*[\text{sign}(\varepsilon_N)]\right\}.
\end{aligned}$$

The latter characterization makes it considerably easier to establish the nonempty relative interior and nondegeneracy of Λ_C , owing to the isotropic Gaussian structure of $\tilde{\mathcal{J}}_C^*$ and $\tilde{\mathcal{J}}_N^*$, as we show next.

Proposition 5.15 (Interior and nondegeneracy under whitening) *Suppose Assumption 5.13 holds, $m \geq c_1 d_\star$ for some constant $c_1 > 2$, and $m_C > m/2$. Then, with probability at least $1 - \exp(-c_2 d_\star)$ for some constant $c_2 > 0$, Assumption 5.5 holds:*

- (I) Λ_C meets the interior of the cube, i.e., $\Lambda_C \cap (-1, 1)^{m_C} \neq \emptyset$;
- (N) for every vertex ν of Λ_C , the restriction $\mathcal{J}_{C, \mathcal{I}_{<}(\nu)}$ of $\mathcal{J}_C$ to the free coordinates of ν has full rank, i.e., $\text{rank } \mathcal{J}_{C, \mathcal{I}_{<}(\nu)} = \text{rank } \mathcal{J}_C$.

The proof is given in Appendix B.7. With Proposition 5.15 in hand, the proofs of Theorems 5.1, 5.2, and 5.3 separate into common and model-specific steps. For a fixed ground-truth factor $U^\star$, criticality follows from Ma and Fattahi [52, Corollaries 1 and 2] and the corresponding diagonal specialization for nonnegative sparse regression. Lemma 5.14 and Proposition 5.15 then verify Assumption 5.5 with high probability.

For nonnegative sparse regression and symmetric matrix sensing, Proposition 5.12 reduces the remaining work to constructing distinct nonzero directions $V_1, \dots, V_q \in \ker \mathcal{J}_C$ such that, conditional on $\mathcal{J}_C$ and Λ_C , the projected residuals $P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$ are independent and spherically symmetric on $\mathcal{N}_C$. The asymmetric model uses the same common steps, but its final cone-selection argument relies on a pair of complementary directions. Table 1 summarizes these model-specific constructions. Each displayed direction lies in $\ker \mathcal{J}_C$ because it acts either on an off-support coordinate or along a null direction of the ground-truth factor. Once inactivity is established for a fixed factor, Lemma 2.10 extends it to all (balanced) ground-truth factors. The full model-specific proofs are given in Appendices B.8, B.9, and B.10.

Model	$d_\star$	Kernel directions used for cone selection
Nonnegative sparse regression	s	Let $S = \text{supp}(u^\star)$ and enumerate $S^c = \{i_a\}_{a=1}^{n-s}$. Set $V_a = e_{i_a}$.
Symmetric matrix sensing	$nr - \frac{r(r-1)}{2}$	Choose a unit vector $w \in \ker U^\star$ and an orthonormal basis $\{u_j\}_{j=1}^{n-r}$ of $\text{col}(U^\star)^\perp$. Set $V_j = u_j w^\top$.
Asymmetric matrix sensing	$r(n_1 + n_2 - r)$	Choose unit vectors $w \in \ker U_1^\star = \ker U_2^\star$, $u \in \text{col}(U_1^\star)^\perp$, and $v \in \text{col}(U_2^\star)^\perp$. Set $V_1 = (uw^\top, 0)$ and $V_2 = (0, vw^\top)$.

Table 1: Model-specific dimensions and directions used in the cone-selection argument. In each row, the displayed directions belong to $\ker \mathcal{J}_C$.

6 Conclusion

We studied when critical points of finite-max composite problems fail to admit a C^2 identifiable manifold. After characterizing the locally minimal identifiable set, we derived a branching criterion based on a second-order incompatibility: if two strata accessible along a common family of tangent directions induce different branch quadratics, then no identifiable manifold can exist. This criterion detects inactivity in settings where first-order conditions are inconclusive.

Our applications show that inactivity is not merely pathological. It arises almost surely at common interpolators of random finite-scenario minimax problems and with high probability in overparameterized robust low-rank recovery, encompassing both global minimizers and strict saddles. These results suggest that finite identification and the resulting smooth local dynamics cannot be taken for granted in modern nonsmooth problems. Developing algorithmic guarantees that do not rely on identifiable manifolds, and identifying weaker geometric structures that remain available at inactive points, are natural directions for future work.

Acknowledgments

This research is supported, in part, by the NSF CAREER grant CCF-2337776 and ONR grant N00014-26-1-2074.

References

- [1] Lee, J.D., Simchowitz, M., Jordan, M.I., Recht, B.: Gradient descent only converges to minimizers. In: Conference on Learning Theory, pp. 1246–1257 (2016). PMLR
- [2] Lee, J.D., Panageas, I., Piliouras, G., Simchowitz, M., Jordan, M.I., Recht, B.: First-order methods almost always avoid strict saddle points: Jd lee et al. *Mathematical programming* **176**(1), 311–337 (2019)
- [3] Lewis, A.S.: Active sets, nonsmoothness, and sensitivity. *SIAM Journal on Optimization* **13**(3), 702–725 (2002)
- [4] Hare, W.L., Lewis, A.S.: Identifying active constraints via partial smoothness and prox-regularity. *Journal of Convex Analysis* **11**(2), 251–266 (2004)
- [5] Hare, W.L., Lewis, A.S.: Identifying active manifolds. *Algorithmic Operations Research* **2**(2), 75–82 (2007)
- [6] Wright, S.J.: Identifiable surfaces in constrained optimization. *SIAM Journal on Control and Optimization* **31**(4), 1063–1079 (1993)
- [7] Burke, J.V., Moré, J.J.: On the identification of active constraints. *SIAM Journal on Numerical Analysis* **25**(5), 1197–1211 (1988)

- [8] Burke, J.: On the identification of active constraints ii: The nonconvex case. *SIAM Journal on Numerical Analysis* **27**(4), 1081–1102 (1990)
- [9] Bonnans, J.F., Shapiro, A.: *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research and Financial Engineering. Springer, New York (2000)
- [10] Rockafellar, R.T., Wets, M., Wets, R.J.B.: *Variational Analysis*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin (2009)
- [11] Dontchev, A.L.: *Lectures on Variational Analysis*. Applied Mathematical Sciences. Springer, Cham (2022). <https://books.google.com/books?id=GuZcEAAAQBAJ>
- [12] Davis, D., Drusvyatskiy, D.: Proximal methods avoid active strict saddles of weakly convex functions. *Foundations of Computational Mathematics* **22**(2), 561–606 (2022)
- [13] Bianchi, P., Hachem, W., Schechtman, S.: Stochastic subgradient descent escapes active strict saddles on weakly convex functions. *Mathematics of Operations Research* **49**(3), 1761–1790 (2024)
- [14] Drusvyatskiy, D., Lewis, A.S.: Optimality, identifiability, and sensitivity. *Mathematical Programming* **147**(1-2), 467–498 (2014)
- [15] Lewis, A.S., Wright, S.J.: A proximal method for composite minimization. *Mathematical Programming* **158**(1), 501–546 (2016)
- [16] Liang, J., Fadili, J., Peyré, G.: Activity identification and local linear convergence of forward–backward-type methods. *SIAM Journal on Optimization* **27**(1), 408–437 (2017)
- [17] Lee, S., Wright, S.J., Bottou, L.: Manifold identification in dual averaging for regularized stochastic online learning. *Journal of Machine Learning Research* **13**(6) (2012)
- [18] Drusvyatskiy, D., Ioffe, A.D., Lewis, A.S.: Generic minimizing behavior in semi-algebraic optimization. *SIAM Journal on Optimization* **26**(1), 513–534 (2016)
- [19] Charisopoulos, V., Chen, Y., Davis, D., Díaz, M., Ding, L., Drusvyatskiy, D.: Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *Foundations of Computational Mathematics* **21**(6), 1505–1593 (2021)
- [20] Wright, J., Ma, Y.: Dense error correction via ℓ^1 -minimization. *IEEE Transactions on Information Theory* **56**(7), 3540–3560 (2010)
- [21] Wiesenmann, W., Kuhn, D., Sim, M.: Distributionally robust convex optimization.

Operations research **62**(6), 1358–1376 (2014)

- [22] Wright, S., Nocedal, J., *et al.*: Numerical optimization. Springer Science **35**(67-68), 7 (1999)
- [23] Lemaréchal, C., Oustry, F., Sagastizábal, C.: The u-lagrangian of a convex function. Transactions of the American mathematical Society **352**(2), 711–729 (2000)
- [24] Drusvyatskiy, D., Lewis, A.S.: Optimality, identifiability, and sensitivity. arXiv preprint arXiv:1207.6628 (2012)
- [25] Calamai, P.H., Moré, J.J.: Projected gradient methods for linearly constrained problems. Mathematical programming **39**(1), 93–116 (1987)
- [26] Wright, S.J.: Modifying sqp for degenerate problems. SIAM Journal on Optimization **13**(2), 470–497 (2002)
- [27] Wright, S.J.: Constraint identification and algorithm stabilization for degenerate nonlinear programs. Mathematical Programming **95**(1), 137–160 (2003)
- [28] Díaz, M., Lehmann, P.I., Lu, H., Yang, J.: Active set identification and rapid convergence for degenerate primal-dual problems. arXiv preprint arXiv:2602.10436 (2026)
- [29] Díaz, M., Jiang, L., Labassi, A.G.: Preconditioned subgradient method for composite optimization: overparameterization and fast convergence. arXiv preprint arXiv:2509.11486 (2025)
- [30] Chi, Y., Lu, Y.M., Chen, Y.: Nonconvex optimization meets low-rank matrix factorization: An overview. IEEE Transactions on Signal Processing **67**(20), 5239–5269 (2019)
- [31] Ding, L., Jiang, L., Chen, Y., Qu, Q., Zhu, Z.: Rank overspecified robust matrix recovery: Subgradient method and exact recovery. Advances in Neural Information Processing Systems **34**, 26767–26778 (2021)
- [32] Tong, T., Ma, C., Chi, Y.: Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. Journal of Machine Learning Research **22**(150), 1–63 (2021)
- [33] Ma, J., Fattahi, S.: Global convergence of sub-gradient method for robust matrix recovery: Small initialization, noisy measurements, and over-parameterization. Journal of Machine Learning Research **24**(96), 1–84 (2023)
- [34] Ma, J., Liang, G., Fattahi, S.: Understanding the implicit regularization of gradient descent in over-parameterized models. arXiv preprint arXiv:2505.17304 (2025)

- [35] Ma, C., Wang, K., Chi, Y., Chen, Y.: Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval and matrix completion. In: International Conference on Machine Learning, pp. 3345–3354 (2018). PMLR
- [36] Zhang, G., Fattahi, S., Zhang, R.Y.: Preconditioned gradient descent for over-parameterized nonconvex burer–monteiro factorization with global optimality certification. *Journal of Machine Learning Research* **24**(163), 1–55 (2023)
- [37] Qian, W., Zhang, Y., Chen, Y.: Structures of spurious local minima in k-means. *IEEE Transactions on Information Theory* **68**(1), 395–422 (2021)
- [38] Boumal, N.: *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, Cambridge (2023)
- [39] Borwein, J.M., Lewis, A.S.: Partially finite convex programming, part i: Quasi relative interiors and duality theory. *Mathematical Programming* **57**(1), 15–48 (1992)
- [40] Lee, J.: *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer, New York (2012). <https://books.google.com/books?id=xygVcKGP5NwC>
- [41] Dries, L.: *Tame Topology and O-minimal Structures*. 150 184, vol. v. 248. Cambridge University Press, Cambridge (1998)
- [42] Coste, M.: *An Introduction to O-minimal Geometry*. Dottorato di ricerca in matematica / Università di Pisa, Dipartimento di Matematica. Istituti Editoriali e Poligrafici Internazionali, Pisa (2000)
- [43] Loi, T.L.: Verdier stratifications and [wf]-stratification in o-minimal structures. arXiv preprint math/9704232 (1997)
- [44] Lewis, A.S., Tian, T.: Identifiability, the kl property in metric spaces, and sub-gradient curves. *Foundations of Computational Mathematics* **25**(3), 905–942 (2025)
- [45] Care, A., Garatti, S., Campi, M.C.: Scenario min-max optimization and the risk of empirical costs. *SIAM Journal on Optimization* **25**(4), 2061–2080 (2015)
- [46] Sagawa, S., Koh, P.W., Hashimoto, T.B., Liang, P.: Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. arXiv preprint arXiv:1911.08731 (2019)
- [47] Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. arXiv preprint arXiv:1611.03530 (2016)
- [48] Bartlett, P.L., Long, P.M., Lugosi, G., Tsigler, A.: Benign overfitting in linear

- regression. *Proceedings of the National Academy of Sciences* **117**(48), 30063–30070 (2020)
- [49] Wright, J., Ma, Y.: *High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications*. Cambridge University Press, Cambridge (2022)
- [50] Burer, S., Monteiro, R.D.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical programming* **95**(2), 329–357 (2003)
- [51] Ma, J., Fattahi, S.: Blessing of depth in linear regression: Deeper models have flatter landscape around the true solution. *Advances in Neural Information Processing Systems* **35**, 34334–34346 (2022)
- [52] Ma, J., Fattahi, S.: Can learning be explained by local optimality in robust low-rank matrix recovery? *Mathematics of Operations Research* **51**(3), 2209–2239 (2025)
- [53] Aravkin, A., Burke, J.V., He, D.: On the global minimizers of real robust phase retrieval with sparse noise. *IEEE Transactions on Information Theory* **67**(3), 1886–1896 (2020)
- [54] Rockafellar, R.T.: *Convex Analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, Princeton (1997)
- [55] Cox, D.A., Little, J.B., Schenck, H.K.: *Toric Varieties*. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI (2024)
- [56] Vershynin, R.: *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2018)

A Derivation of the Identifiable Set in Example 2.12

Our goal is to show that the minimal identifiable set

$$\mathcal{M}_{\bar{v}} := \bigcup_{\lambda \in \Lambda_{\bar{v}}} \{z \in \mathbb{R}^m : \text{supp } \lambda \subseteq I(z)\}, \quad \Lambda_{\bar{v}} := \left\{ \lambda \in \Delta_{\mathcal{I}} : \bar{v} = \sum_{\sigma} \lambda_{\sigma} \sigma \right\}$$

coincides with

$$\mathcal{S}_{\bar{v}} := \{z \in \mathbb{R}^m : z_i = 0 \text{ if } |\bar{v}_i| < 1, z_i \geq 0 \text{ if } \bar{v}_i = 1, z_i \leq 0 \text{ if } \bar{v}_i = -1\}.$$

Fix a coordinate i . Since $\lambda \in \Delta_{\mathcal{I}}$ is a probability distribution over $\mathcal{I} = \{\pm 1\}^m$, the constraint $\bar{v}_i = \sum_{\sigma} \lambda_{\sigma} \sigma_i$ expresses $\bar{v}_i$ as a weighted average of $\sigma_i \in \{-1, +1\}$.

Setting $\alpha_i^+ := \sum_{\sigma \in \mathcal{I}_i^+} \lambda_\sigma$, $\alpha_i^- := \sum_{\sigma \in \mathcal{I}_i^-} \lambda_\sigma$, and $\mathcal{I}_i^\pm := \{\sigma \in \mathcal{I} : \sigma_i = \pm 1\}$, we have $\alpha_i^+ + \alpha_i^- = 1$ and $\bar{v}_i = \alpha_i^+ - \alpha_i^-$, so

$$|\bar{v}_i| = 1 \iff \text{supp } \lambda \subseteq \mathcal{I}_i^{\text{sign } \bar{v}_i}, \quad |\bar{v}_i| < 1 \iff \text{supp } \lambda \cap \mathcal{I}_i^+ \neq \emptyset \text{ and } \text{supp } \lambda \cap \mathcal{I}_i^- \neq \emptyset. \quad (31)$$

Proof of $\mathcal{M}_{\bar{v}} \subseteq \mathcal{S}_{\bar{v}}$. Suppose $z \in \mathcal{M}_{\bar{v}}$ and pick $\lambda \in \Lambda_{\bar{v}}$ with $\text{supp } \lambda \subseteq I(z)$. For coordinate i , if $z_i > 0$ then every active sign pattern has $\sigma_i = +1$, so $\text{supp } \lambda \subseteq \mathcal{I}_i^+$, which by (31) forces $\bar{v}_i = 1$. Likewise, $z_i < 0$ forces $\bar{v}_i = -1$. Finally, if $|\bar{v}_i| < 1$ then by (31) the support of λ must intersect both $\mathcal{I}_i^+$ and $\mathcal{I}_i^-$, which requires both signs to be active at z_i , hence $z_i = 0$. Thus $z \in \mathcal{S}_{\bar{v}}$.

Proof of $\mathcal{S}_{\bar{v}} \subseteq \mathcal{M}_{\bar{v}}$. Suppose $z \in \mathcal{S}_{\bar{v}}$. Since $\bar{v} \in [-1, 1]^m = \text{conv}\{\pm 1\}^m$, there exists $\lambda \in \Lambda_{\bar{v}}$; we claim $\text{supp } \lambda \subseteq I(z)$ for any such λ . Indeed, fix a sign pattern $\sigma \in \text{supp } \lambda$ and a coordinate i . If $z_i \neq 0$, then $\bar{v}_i = \text{sign}(z_i)$ by definition of $\mathcal{S}_{\bar{v}}$, so (31) gives $\sigma_i = \text{sign}(z_i)$, meaning σ is active at z_i . If $z_i = 0$, both signs are active at z_i regardless of σ_i . Hence $\sigma \in I(z)$, so $\text{supp } \lambda \subseteq I(z)$ and $z \in \mathcal{M}_{\bar{v}}$.

B Omitted Proofs

B.1 Proof of Lemma 2.10

Define the lifted transformation

$$\tilde{\mathcal{T}} : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{X} \times \mathbb{R}, \quad \tilde{\mathcal{T}}(x, \alpha) := (\mathcal{T}x, \alpha).$$

Since $\mathcal{T}$ is orthogonal and $h \circ \mathcal{T} = h$, the map $\tilde{\mathcal{T}}$ is orthogonal and leaves the epigraph of h invariant: $\tilde{\mathcal{T}}(\text{epi } h) = \text{epi } h$. The limiting normal cone is equivariant under orthogonal transformations. Hence

$$N_{\text{epi } h}(\mathcal{T}x, h(\mathcal{T}x)) = \tilde{\mathcal{T}} N_{\text{epi } h}(x, h(x)).$$

Since $\tilde{\mathcal{T}}(v, -1) = (\mathcal{T}v, -1)$, the definition of the limiting subdifferential gives

$$\begin{aligned} v \in \partial h(x) &\iff (v, -1) \in N_{\text{epi } h}(x, h(x)) \\ &\iff (\mathcal{T}v, -1) \in N_{\text{epi } h}(\mathcal{T}x, h(\mathcal{T}x)) \\ &\iff \mathcal{T}v \in \partial h(\mathcal{T}x). \end{aligned}$$

Therefore, $\partial h(\mathcal{T}x) = \mathcal{T}(\partial h(x))$. In particular, since $\mathcal{T}0 = 0$,

$$0 \in \partial h(x) \iff 0 \in \partial h(\mathcal{T}x),$$

which proves the invariance of criticality.

We next prove the invariance of identifiable sets. Suppose that $\mathcal{M}$ is identifiable for h at x for v , and consider sequences $y_\nu \rightarrow \mathcal{T}x$ and $w_\nu \rightarrow \mathcal{T}v$ with $w_\nu \in \partial h(y_\nu)$. Set $x_\nu := \mathcal{T}^{-1}y_\nu$ and $v_\nu := \mathcal{T}^{-1}w_\nu$. By continuity of $\mathcal{T}^{-1}$, we have $x_\nu \rightarrow x$ and $v_\nu \rightarrow v$. The subdifferential equivariance established above also gives $v_\nu \in \partial h(x_\nu)$.

Identifiability of $\mathcal{M}$ therefore implies $x_\nu \in \mathcal{M}$ for all sufficiently large ν , and hence $y_\nu = \mathcal{T}x_\nu \in \mathcal{TM}$ for all sufficiently large ν . Thus, $\mathcal{TM}$ is identifiable for h at $\mathcal{T}x$ for $\mathcal{T}v$. Applying the same argument to $\mathcal{T}^{-1}$ proves the converse.

It remains to verify that the manifold and restricted-smoothness properties are also preserved. Suppose that $\mathcal{M}$ is a C^2 manifold around x . Locally, it can be written as $\mathcal{M} = \Phi^{-1}(0)$ for some C^2 map Φ whose derivative is surjective. Around $\mathcal{T}x$, we then have

$$\mathcal{TM} = (\Phi \circ \mathcal{T}^{-1})^{-1}(0).$$

Because $\mathcal{T}^{-1}$ is a linear isomorphism, the derivative of $\Phi \circ \mathcal{T}^{-1}$ remains surjective. Hence $\mathcal{TM}$ is a C^2 embedded manifold. Moreover, $\mathcal{T}|_{\mathcal{M}} : \mathcal{M} \rightarrow \mathcal{TM}$ is a C^∞ diffeomorphism, and the invariance of h yields

$$h|_{\mathcal{TM}} = (h|_{\mathcal{M}}) \circ (\mathcal{T}|_{\mathcal{M}})^{-1}.$$

Thus, $h|_{\mathcal{M}}$ is C^2 if and only if $h|_{\mathcal{TM}}$ is C^2 . It follows that $\mathcal{M}$ is a C^2 identifiable manifold at x for v if and only if $\mathcal{TM}$ is a C^2 identifiable manifold at $\mathcal{T}x$ for $\mathcal{T}v$. The invariance of inactivity follows immediately. $\square$

B.2 Proof of Proposition 5.4

Let $\widehat{\Lambda} := \{v \in \partial\|F(U^*)\|_1 : DF(U^*)^*[v] = 0\}$ be the full multiplier set at U^* and let $\widehat{\mathcal{S}}_v \subseteq \mathbb{R}^m$ be the sign cone (3). By Proposition 3.1 applied to $g = \|\cdot\|_1$, the set

$$\widehat{\mathcal{M}} := \bigcup_{v \in \widehat{\Lambda}} F^{-1}(\widehat{\mathcal{S}}_v)$$

is locally minimal identifiable for f_{LR} at U^* for $0 \in \partial f_{\text{LR}}(U^*)$. For corrupted indices $i \in \Omega_N$, $F_i(U^*) = -\varepsilon_i \neq 0$. Hence, after shrinking to a neighborhood $\mathcal{O}$ of U^* , the corrupted signs remain fixed:

$$\text{sign}(F_N(U)) = -\text{sign}(\varepsilon_N) \quad \text{for all } U \in \mathcal{O}.$$

Therefore the corrupted block of every multiplier $v \in \widehat{\Lambda}$ is fixed at $\text{sign}(F_N(U)) = -\text{sign}(\varepsilon_N)$. Writing $v = (\bar{v}, -\text{sign}(\varepsilon_N))$ with $\bar{v} \in \mathbb{R}^{m_C}$, the criticality equation becomes

$$\mathcal{J}_C^*[\bar{v}] - \mathcal{J}_N^*[\text{sign}(\varepsilon_N)] = 0,$$

so the clean block $\bar{v}$ lies in Λ_C as defined in (21). Since the corrupted signs are already pinned to $-\text{sign}(\varepsilon_N)$ on $\mathcal{O}$, membership $F(U) \in \widehat{\mathcal{S}}_v$ depends only on the clean indices; as a result, $F(U) \in \widehat{\mathcal{S}}_v$ if and only if $F_C(U) \in \mathcal{S}_{\bar{v}}$. Consequently,

$$\widehat{\mathcal{M}} \cap \mathcal{O} = F_C^{-1} \left(\bigcup_{v \in \Lambda_C} \mathcal{S}_v \right) \cap \mathcal{O} = F_C^{-1}(\mathcal{E}_C) \cap \mathcal{O},$$

which is the claimed locally minimal identifiable set. $\square$

B.3 Proof of Lemma 5.7

Proof We prove each inclusion separately.

Establishing $T_{U^}\mathcal{M} \subseteq \ker \mathcal{J}_C$.* Let $V \in T_{U^*}\mathcal{M}$ and, for an open interval I containing 0, let $\gamma_V : I \rightarrow \mathcal{M}$ be a C^1 curve with $\gamma_V(0) = U^*$ and $\dot{\gamma}_V(0) = V$. Since $\mathcal{M} = F_C^{-1}(\mathcal{E}_C)$, we have $F_C(\gamma_V(t)) \in \mathcal{E}_C$ for all $t \in I$. Since $\mathcal{E}_C$ is a closed cone, dividing by $t > 0$ preserves membership, so $\frac{1}{t}F_C(\gamma_V(t)) \in \mathcal{E}_C$. Letting $t \downarrow 0$ and using $F_C(U^*) = 0$ gives

$$\mathcal{J}_C[V] = \lim_{t \downarrow 0} \frac{F_C(\gamma_V(t)) - F_C(\gamma_V(0))}{t} = \lim_{t \downarrow 0} \frac{F_C(\gamma_V(t)) - F_C(U^*)}{t} = \lim_{t \downarrow 0} \frac{F_C(\gamma_V(t))}{t} \in \mathcal{E}_C.$$

On the other hand, $\mathcal{J}_C[V] \in \text{Im } \mathcal{J}_C = \mathcal{R}_C$ by definition. By Proposition 5.8 applied with $a = 0$, the intersection $\mathcal{E}_C \cap \mathcal{R}_C = \{0\}$. Hence $\mathcal{J}_C[V] = 0$, so $V \in \ker \mathcal{J}_C$.

Establishing $\ker \mathcal{J}_C \subseteq T_{U^}\mathcal{M}$.* Let $V \in \ker \mathcal{J}_C$. We must find a C^1 curve $\gamma_V : I \rightarrow \mathcal{M}$ with $\gamma_V(0) = U^*$ and $\dot{\gamma}_V(0) = V$, which by definition of $\mathcal{M}$ requires $F_C(\gamma_V(t)) \in \mathcal{E}_C$ for all $t \in I$. Consider the straight-line curve $t \mapsto U^* + tV$. Since $F_C(U^*) = 0$ and $\mathcal{J}_C[V] = 0$, we have

$$a_V(t) := F_C(U^* + tV) = t^2 \mathcal{Q}_C(V),$$

where $\mathcal{Q}_C(V)$ is defined as (26). This need not lie in $\mathcal{E}_C$, so we correct it. By Proposition 5.8, for any $a \in \mathbb{R}^{m_C}$ there is a unique point $\mathcal{S}_{\Lambda_C}(a) \in \mathcal{E}_C$ with $\mathcal{S}_{\Lambda_C}(a) - a \in \mathcal{R}_C$. Define the residual correction

$$\Gamma(a) := \mathcal{S}_{\Lambda_C}(a) - a \in \mathcal{R}_C.$$

We seek a correction $\Delta(t) = O(t^2)$ such that $\gamma_V(t) := U^* + tV + \Delta(t)$ satisfies $F_C(\gamma_V(t)) \in \mathcal{E}_C$. Specifically, we write

$$F_C(\gamma_V(t)) = a_V(t) + \mathcal{J}_C[\Delta(t)] + \rho_{V,\Delta}(t),$$

where $\rho_{V,\Delta}(t) := F_C(U^* + tV + \Delta(t)) - a_V(t) - \mathcal{J}_C[\Delta(t)]$ is the remainder. Moreover, due to the C^2 -smoothness of F , we obtain $\|\rho_{V,\Delta}(t)\| = O(|t|\|\Delta(t)\| + \|\Delta(t)\|^2)$

For $F_C(\gamma_V(t))$ to lie in $\mathcal{E}_C$, we need

$$\mathcal{J}_C[\Delta(t)] = \Gamma(a_V(t) + \rho_{V,\Delta}(t)).$$

Since Γ takes values in $\mathcal{R}_C = \text{Im } \mathcal{J}_C$, and $\mathcal{J}_C$ restricted to $\mathcal{L} := (\ker \mathcal{J}_C)^\perp$ is a bijection onto $\mathcal{R}_C$, every right-hand side can be realized by a unique $\Delta(t) \in \mathcal{L}$. This gives the fixed-point equation

$$\Delta(t) = \underbrace{(\mathcal{J}_C|_{\mathcal{L}})^{-1} \left(\Gamma(a_V(t) + \rho_{V,\Delta}(t)) \right)}_{=: \mathcal{T}_t(\Delta(t))}.$$

We show $\mathcal{T}_t$ has a fixed point in the ball $\mathbb{B}_t := \{\Delta \in \mathcal{L} : \|\Delta\| \leq Mt^2\}$ for sufficiently large $M > 0$. By Proposition 5.8 and Lemma 5.9, $\mathcal{S}_{\Lambda_C}$ is piecewise linear, hence Lipschitz with $\Gamma(0) = 0$. For $\Delta(t) \in \mathbb{B}_t$, the remainder satisfies $\rho_{V,\Delta}(t) = O(t^3)$, so $a_V(t) + \rho_{V,\Delta}(t) = O(t^2)$, and Lipschitzness of Γ gives $\Gamma(a_V(t) + \rho_{V,\Delta}(t)) = O(t^2)$. Choosing M large enough and taking $|t|$ small enough, we have $\mathcal{T}_t(\mathbb{B}_t) \subseteq \mathbb{B}_t$, that is, $\mathcal{T}_t$ maps $\mathbb{B}_t$ to itself. To see that $\mathcal{T}_t$ is a contraction on $\mathbb{B}_t$, for $\Delta_1(t), \Delta_2(t) \in \mathbb{B}_t$, the definition of $\mathcal{T}_t$, the Lipschitzness of Γ , and C^2 -smoothness of F_C give

$$\begin{aligned} \|\mathcal{T}_t(\Delta_1(t)) - \mathcal{T}_t(\Delta_2(t))\| &= O\left(\|\rho_{V,\Delta_1}(t) - \rho_{V,\Delta_2}(t)\|\right) \\ &= O\left((|t| + \|\Delta_1(t)\| + \|\Delta_2(t)\|)\|\Delta_1(t) - \Delta_2(t)\|\right) \\ &= O(|t| + Mt^2)\|\Delta_1(t) - \Delta_2(t)\|. \end{aligned}$$

which is strictly less than $\|\Delta_1(t) - \Delta_2(t)\|$ for small $|t|$. By the Banach fixed-point theorem, there exists $\Delta(t) \in \mathbb{B}_t$ with $\Delta(t) = \mathcal{T}_t(\Delta(t))$ and $\|\Delta(t)\| \leq Mt^2 = O(t^2)$, so $F_C(U^* + tV + \Delta(t)) \in \mathcal{E}_C$ by construction.

We have constructed points $U^* + tV + \Delta(t) \in \mathcal{M}$ with $\Delta(t) = O(t^2)$, but we have not shown $\Delta(t)$ is smooth. Instead, we use the local defining map of $\mathcal{M}$. Let Φ be a local C^2 defining map for $\mathcal{M}$ near U^* , so $\mathcal{M} = \Phi^{-1}(0)$ locally. Since $U^* + tV + \Delta(t) \in \mathcal{M}$, we have $\Phi(U^* + tV + \Delta(t)) = 0$. Taylor expanding and using $\|\Delta(t)\| = O(t^2)$ gives

$$0 = t D\Phi(U^*)[V] + D\Phi(U^*)[\Delta(t)] + O(t^2) = t D\Phi(U^*)[V] + O(t^2).$$

Dividing by t and letting $t \downarrow 0$ yields $D\Phi(U^*)[V] = 0$, which means $V \in T_{U^*}\mathcal{M}$. Since $V \in \ker \mathcal{J}_C$ was arbitrary, $\ker \mathcal{J}_C \subseteq T_{U^*}\mathcal{M}$. Combining both steps, $T_{U^*}\mathcal{M} = \ker \mathcal{J}_C$. $\square$

B.4 Proof of Proposition 5.8

We first prove that the intersection $(a + \mathcal{R}_C) \cap \mathcal{E}_C$ is non-empty. Due to Proposition 5.6, Λ_C is a full-dimensional polytope in, and contained in, the affine subspace $\bar{v} + \mathcal{N}_C$. Therefore, the union of the normal cones of Λ_C at its vertices covers $\mathcal{N}_C$, that is, $\mathcal{N}_C \subseteq \bigcup_{\nu \in \text{Vert}(\Lambda_C)} N_{\Lambda_C}(\nu)$ [55, Proposition 2.3.6]. Hence $P_{\mathcal{N}_C}a$ lies in $\mathcal{N}_C \cap N_{\Lambda_C}(\nu)$ for at least one vertex ν . By Proposition 5.6, $N_{\Lambda_C}(\nu) = \mathcal{R}_C + \mathcal{S}_\nu$, so there are $r \in \mathcal{R}_C$ and $s \in \mathcal{S}_\nu$ with $P_{\mathcal{N}_C}a = r + s$. Since $a - P_{\mathcal{N}_C}a \in \mathcal{R}_C$, we have $s - a \in \mathcal{R}_C$, and therefore $s \in (a + \mathcal{R}_C) \cap \mathcal{E}_C$, establishing its nonemptiness.

Next, we show that $(a + \mathcal{R}_C) \cap \mathcal{E}_C$ is a unique point.

- *Uniqueness within a vertex cone.* For a vertex ν , we write $\mathcal{I}_\nu := \mathcal{I}_<(\nu)$ for its free coordinates and $P_{\mathcal{I}_\nu}$ as the projection onto those coordinates. Suppose $z, z' \in (a + \mathcal{R}_C) \cap \mathcal{S}_\nu$, then $z - z' \in \mathcal{R}_C$. Moreover, since $\mathcal{J}_{C, \mathcal{I}_\nu} = P_{\mathcal{I}_\nu} \circ \mathcal{J}_C$ has rank equal to $\dim \mathcal{R}_C$, the restricted projection $P_{\mathcal{I}_\nu}|_{\mathcal{R}_C} : \mathcal{R}_C \rightarrow \mathbb{R}^{\mathcal{I}_\nu}$ is injective. Since $z - z' \in \mathcal{R}_C$ and $P_{\mathcal{I}_\nu}(z - z') = 0$, we have $z = z'$.
- *Uniqueness across vertex cones.* Take

$$z_1 \in (a + \mathcal{R}_C) \cap \mathcal{S}_{\nu_1}, \quad z_2 \in (a + \mathcal{R}_C) \cap \mathcal{S}_{\nu_2}.$$

Then $z_1 - z_2 \in \mathcal{R}_C$, while $\nu_1 - \nu_2 \in \mathcal{N}_C$, so

$$0 = \langle z_1 - z_2, \nu_1 - \nu_2 \rangle = (\langle z_1, \nu_1 \rangle - \langle z_1, \nu_2 \rangle) + (\langle z_2, \nu_2 \rangle - \langle z_2, \nu_1 \rangle).$$

Note that $z_1 \in \mathcal{S}_{\nu_1}$; hence, due to the definition of $\mathcal{S}_{\nu_1}$ in (3), $\langle z_1, \nu_1 \rangle - \langle z_1, \nu_2 \rangle \geq 0$. Similarly, $\langle z_2, \nu_2 \rangle - \langle z_2, \nu_1 \rangle \geq 0$. Combined with the above equality, $\langle z_1, \nu_1 \rangle = \langle z_1, \nu_2 \rangle$ and $\langle z_2, \nu_2 \rangle = \langle z_2, \nu_1 \rangle$. Therefore, $z_1, z_2 \in \mathcal{S}_{\nu_1} \cap \mathcal{S}_{\nu_2}$, and in particular $z_1, z_2 \in (a + \mathcal{R}_C) \cap \mathcal{S}_{\nu_1}$. From the uniqueness within a vertex cone, we have $z_1 = z_2$. Finally $a - a' \in \mathcal{R}_C$ gives $a + \mathcal{R}_C = a' + \mathcal{R}_C$, so $\mathcal{S}_{\Lambda_C}$ depends only on $P_{\mathcal{N}_C}a$.

It remains to show that $\mathcal{S}_{\Lambda_C}$ is a two-sided inverse of $P_{\mathcal{N}_C}|_{\mathcal{E}_C} : \mathcal{E}_C \rightarrow \mathcal{N}_C$, and that it is positively homogeneous. We first verify both compositions $\mathcal{S}_{\Lambda_C} \circ P_{\mathcal{N}_C}|_{\mathcal{E}_C}$ and $P_{\mathcal{N}_C}|_{\mathcal{E}_C} \circ \mathcal{S}_{\Lambda_C}$ are identity. Let $z \in \mathcal{E}_C$. Since $z \in (z + \mathcal{R}_C) \cap \mathcal{E}_C$, uniqueness gives $\mathcal{S}_{\Lambda_C}(z) = z$. Hence $\mathcal{S}_{\Lambda_C}(P_{\mathcal{N}_C}z) = \mathcal{S}_{\Lambda_C}(z) = z$, since, as established earlier, $\mathcal{S}_{\Lambda_C}(z)$ depends only on $P_{\mathcal{N}_C}z$. Therefore, $\mathcal{S}_{\Lambda_C} \circ P_{\mathcal{N}_C}|_{\mathcal{E}_C}$ is identity. Next, let $w \in \mathcal{N}_C$. By

definition, $\mathcal{S}_{\Lambda_C}(w) \in (w + \mathcal{R}_C) \cap \mathcal{E}_C$, so $\mathcal{S}_{\Lambda_C}(w) - w \in \mathcal{R}_C$. Projecting onto $\mathcal{N}_C$ and using $\mathcal{R}_C \perp \mathcal{N}_C$ gives $P_{\mathcal{N}_C} \mathcal{S}_{\Lambda_C}(w) = w$. Therefore, $P_{\mathcal{N}_C}|_{\mathcal{E}_C} \circ \mathcal{S}_{\Lambda_C}$ is identity.

To establish positive homogeneity, note that both $\mathcal{E}_C$ and $\mathcal{R}_C$ are cones, hence invariant under scaling by $\lambda > 0$. Thus $\lambda \mathcal{S}_{\Lambda_C}(a) \in (\lambda a + \mathcal{R}_C) \cap \mathcal{E}_C$, and uniqueness gives $\mathcal{S}_{\Lambda_C}(\lambda a) = \lambda \mathcal{S}_{\Lambda_C}(a)$. At $\lambda = 0$, we have $\mathcal{S}_{\Lambda_C}(0) = 0$ since $0 \in \mathcal{R}_C \cap \mathcal{E}_C$. This completes the proof. $\square$

B.5 Proof of Proposition 5.12

Let $\mathcal{O}_q$ be the event that there are distinct indices $\ell_1, \ell_2 \in [q]$ and distinct vertices $\nu_1, \nu_2 \in \text{Vert}(\Lambda_C)$ such that

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V_{\ell_j}) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j), \quad j = 1, 2.$$

We first show that the following bound holds almost surely

$$\mathbb{P}(\mathcal{O}_q \mid \mathcal{J}_C, \Lambda_C) \geq 1 - 2^{-(q-1)}. \quad (32)$$

By Proposition 5.6, Λ_C is full-dimensional in its affine space. Hence its vertex normal cones $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$, $\nu \in \text{Vert}(\Lambda_C)$, are full-dimensional and pointed, cover $\mathcal{N}_C$, and have pairwise disjoint relative interiors [55, Theorem 2.3.2 and Proposition 2.3.6]. Their boundaries therefore have spherical measure zero. For each $\nu \in \text{Vert}(\Lambda_C)$, set

$$\alpha_\nu := \mathbb{P} \left\{ P_{\mathcal{N}_C} \mathcal{Q}_C(V_1) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu) \mid \mathcal{J}_C, \Lambda_C \right\}.$$

Spherical symmetry and continuity identify α_ν with the normalized spherical measure of $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$, so

$$\sum_{\nu \in \text{Vert}(\Lambda_C)} \alpha_\nu = 1, \quad 0 < \alpha_\nu \leq \frac{1}{2}.$$

Indeed, $0 < \alpha_\nu$ follows from the full-dimensionality of each normal cone, while $\alpha_\nu \leq \frac{1}{2}$ follows from pointedness: the spherical section of $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ is disjoint from its antipodal image, and the two sets have the same spherical measure.

Up to the boundary-null event, $\mathcal{O}_q^c$ is the event that all q residuals lie within the same $\text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$ for some $\nu \in \text{Vert}(\Lambda_C)$. Conditional independence therefore gives

$$\mathbb{P}(\mathcal{O}_q^c \mid \mathcal{J}_C, \Lambda_C) = \sum_{\nu \in \text{Vert}(\Lambda_C)} \alpha_\nu^q \leq \left(\max_{\nu \in \text{Vert}(\Lambda_C)} \alpha_\nu \right)^{q-1} \sum_{\nu \in \text{Vert}(\Lambda_C)} \alpha_\nu \leq 2^{-(q-1)}.$$

Taking the complementary event proves (32).

It remains to verify that this event implies the success of the cone-selection test of Theorem 5.11. Choose the distinct indices $\ell_1, \ell_2 \in [q]$ and distinct vertices ν_1, ν_2 supplied by the event $\mathcal{O}_q$, so that

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V_{\ell_j}) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j), \quad j = 1, 2,$$

and set $\mathcal{W} := \text{span}\{V_{\ell_1}, V_{\ell_2}\}$. For $j = 1, 2$, define

$$\mathcal{C}_j := \left\{ V \in \mathcal{W} \setminus \{0\} : P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j) \right\}.$$

Each $\mathcal{C}_j$ is nonempty because it contains V_{ℓ_j} . It is relatively open in $\mathcal{W} \setminus \{0\}$ because $V \mapsto P_{\mathcal{N}_C} \mathcal{Q}_C(V)$ is continuous and $\text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j)$ is open in $\mathcal{N}_C$. It is also a cone because this map is homogeneous of degree two. Thus $\mathcal{W}, \mathcal{C}_1, \mathcal{C}_2, \nu_1, \nu_2$ satisfy the conditions of Theorem 5.11, thereby completing the proof. $\square$

B.6 Proof of Lemma 5.14

Since $\ker \mathcal{T}^* = (\text{Im } \mathcal{T})^\perp$, we have

$$\ker(\mathcal{T}^*|_{\text{Im } \mathcal{T}}) = \text{Im } \mathcal{T} \cap (\text{Im } \mathcal{T})^\perp = \{0\}.$$

Thus $\mathcal{T}^*|_{\text{Im } \mathcal{T}}$ is injective. Moreover,

$$\dim \text{Im } \mathcal{T} = \text{rank } \mathcal{T} = \text{rank } \mathcal{T}^* = \dim \text{Im } \mathcal{T}^*,$$

so $\mathcal{T}^*|_{\text{Im } \mathcal{T}} : \text{Im } \mathcal{T} \rightarrow \text{Im } \mathcal{T}^*$ is a linear isomorphism. Since $\iota_\star$ and $\Sigma^{-1/2}$ are also linear isomorphisms, so is $\mathcal{H}$. For every $A \in \mathbb{S}^n$, the orthogonal decomposition $A = P_{\text{Im } \mathcal{T}} A + P_{(\text{Im } \mathcal{T})^\perp} A$ and the identity $(\text{Im } \mathcal{T})^\perp = \ker \mathcal{T}^*$ give $\mathcal{T}^*[A] = \mathcal{T}^*[P_{\text{Im } \mathcal{T}} A]$. Consequently,

$$(\mathcal{T}^*|_{\text{Im } \mathcal{T}})^{-1} \mathcal{T}^*[A] = P_{\text{Im } \mathcal{T}} A.$$

In particular, for every measurement index i ,

$$\mathcal{H}(\mathcal{T}^*[A_i]) = \Sigma^{-1/2} \iota_\star (\mathcal{T}^*|_{\text{Im } \mathcal{T}})^{-1} \mathcal{T}^*[A_i] = \Sigma^{-1/2} \iota_\star (P_{\text{Im } \mathcal{T}} A_i).$$

The hypothesis therefore yields

$$\mathcal{H}(\mathcal{T}^*[A_1]), \dots, \mathcal{H}(\mathcal{T}^*[A_m]) \stackrel{\text{i.i.d.}}{\sim} N(0, I_{d_\star}).$$

Recall that $\mathcal{J}_C^*[e_i] = \mathcal{T}^*[A_i]$ for $i \in \Omega_C$ and $\mathcal{J}_N^*[e_j] = \mathcal{T}^*[A_j]$ for $j \in \Omega_N$. Conditional on ε , the index sets Ω_C and Ω_N are fixed. Since ε is independent of the measurement matrices, conditioning on ε does not alter their joint distribution. Hence

$$\{\mathcal{H} \circ \mathcal{J}_C^*[e_i] : i \in \Omega_C\} \quad \text{and} \quad \{\mathcal{H} \circ \mathcal{J}_N^*[e_j] : j \in \Omega_N\}$$

are independent standard Gaussian vectors in $\mathbb{R}^{d_\star}$. This proves Assumption 5.13.

It remains to verify the final assertion. In each of the three models, the measurement matrices are isotropic Gaussian random elements, up to a deterministic scaling, on a linear subspace $\mathcal{G}$ containing $\text{Im } \mathcal{T}$:

- for nonnegative sparse regression, $\mathcal{G}$ is the subspace of diagonal matrices and $A_i = \text{diag}(a_i)$ is standard Gaussian on $\mathcal{G}$;

- for symmetric matrix sensing, $\mathcal{G} = \mathbb{S}^n$ and $A_i \sim \text{GOE}(n)$ is standard Gaussian on $\mathbb{S}^n$ under the Frobenius inner product;
- for asymmetric matrix sensing, $\mathcal{G}$ is the subspace of symmetric matrices with zero diagonal blocks, and

$$A_i = \frac{1}{2} \begin{bmatrix} 0 & \Gamma_i \\ \Gamma_i^\top & 0 \end{bmatrix}$$

is Gaussian on $\mathcal{G}$ with covariance $\frac{1}{2}I_{\mathcal{G}}$.

Since $\text{Im } \mathcal{T} \subseteq \mathcal{G}$, orthogonal projection onto $\text{Im } \mathcal{T}$ shows that

$$\iota_\star(P_{\text{Im } \mathcal{T}} A_i) \stackrel{\text{i.i.d.}}{\sim} N(0, cI_{d_\star}),$$

where $c = 1$ for nonnegative sparse regression and symmetric matrix sensing, and $c = \frac{1}{2}$ for asymmetric matrix sensing. Thus the hypothesis holds with $\Sigma = cI_{d_\star}$ in all three cases. $\square$

B.7 Proof of Proposition 5.15

Proof of Condition (I). Write $h := \tilde{\mathcal{J}}_N^*[\text{sign}(\varepsilon_N)]$ and $\mathcal{Z}_C := \{\tilde{\mathcal{J}}_C^*[v] : v \in [-1, 1]^{m_C}\}$. Since $\Lambda_C = \{v \in [-1, 1]^{m_C} : \tilde{\mathcal{J}}_C^*[v] = h\}$, showing $\Lambda_C \cap (-1, 1)^{m_C} \neq \emptyset$ reduces to showing h lies in the interior of $\mathcal{Z}_C$. It suffices to exhibit $\tau > 0$ such that the ball $\mathcal{B}(\tau) := \{z : \|z\|_2 < \tau\}$ satisfies $\mathcal{B}(\tau) \subseteq \mathcal{Z}_C$ and $h \in \mathcal{B}(\tau)$. Define

$$\tau := \inf_{\|u\|_2=1} \left\| \tilde{\mathcal{J}}_C[u] \right\|_1$$

Suppose, for contradiction, that some $z \in \mathcal{B}(\tau)$ satisfies $z \notin \mathcal{Z}_C$. Since $\mathcal{Z}_C$ is closed and convex, the separating hyperplane theorem gives a unit vector u with

$$\langle u, z \rangle > \sup_{w \in \mathcal{Z}_C} \langle u, w \rangle.$$

But

$$\langle u, z \rangle \leq \|z\|_2 < \tau \leq \left\| \tilde{\mathcal{J}}_C[u] \right\|_1 = \sup_{\|y\|_\infty \leq 1} \langle \tilde{\mathcal{J}}_C[u], y \rangle = \sup_{w \in \mathcal{Z}_C} \langle u, w \rangle,$$

where the middle equality writes the ℓ_1 norm as a supremum over the dual ball and the last equality uses $\mathcal{Z}_C = \tilde{\mathcal{J}}_C^*([-1, 1]^{m_C})$. This contradicts the strict inequality above, so $\mathcal{B}(\tau) \subseteq \mathcal{Z}_C$.

Next, we show $\|h\| < \tau$. Note for $m_N = 0$, we have $h = 0$, and since $\tau > 0$ by the above construction, this inequality holds trivially. Let $\mu := \mathbb{E}|Z| = \sqrt{2/\pi}$ for $Z \sim N(0, 1)$. Since $\tilde{\mathcal{J}}_C$ is an $m_C \times d_\star$ standard Gaussian matrix, Vershynin [56, Theorem 9.6.3] gives

$$\mathbb{E} \tau \geq \mu m_C - c\sqrt{m_C d_\star}$$

for a universal constant $c > 0$. Under the assumption $m_C \geq c_2 d_\star$ with $c_2 := 4c^2/\mu^2$, this simplifies to $\mathbb{E} \tau \geq (\mu/2)m_C$. The map $X \mapsto \inf_{\|u\|_2=1} \|Xu\|_1$ is $\sqrt{m_C}$ -Lipschitz

with respect to the Frobenius norm, so by concentration for Lipschitz functions of Gaussian matrices [56, Theorem 5.2.2],

$$\mathbb{P}\{\tau \geq \mathbb{E} \tau - t\} \geq 1 - \exp\left(-\frac{t^2}{2m_C}\right).$$

Setting $t = (\mu/4)m_C$ and using $\mathbb{E} \tau \geq (\mu/2)m_C$ gives

$$\mathbb{P}\left\{\tau > \frac{\mu}{4}m_C\right\} \geq 1 - \exp\left(-\frac{m_C}{32}\right).$$

Since $\tilde{\mathcal{J}}_N^*$ is an $d_\star \times m_N$ standard Gaussian matrix and $h = \tilde{\mathcal{J}}_N^*[\text{sign}(\varepsilon_N)]$, we have $h \sim N(0, m_N I_{d_\star})$. As $x \mapsto \|x\|_2$ is 1-Lipschitz, the same concentration argument (applied to the Gaussian vector $h/\sqrt{m_N} \sim N(0, I_{d_\star})$) yields

$$\mathbb{P}\left\{\|h\|_2 < 2\sqrt{m_N d_\star}\right\} \geq 1 - \exp\left(-\frac{d_\star}{2}\right).$$

If $m_C \geq c_3 d_\star$ with $c_3 := 64/\mu^2$, then $(\mu/4)m_C \geq 2\sqrt{m_N d_\star}$ (using $m_N \leq m_C$, since $m_C > m/2$). A union bound gives

$$\mathbb{P}(\|h\|_2 < \tau) \geq 1 - \exp\left(-\frac{m_C}{32}\right) - \exp\left(-\frac{d_\star}{2}\right),$$

provided $m_C \geq \max\{c_2, c_3\}d_\star$. On this event, $h \in \mathcal{B}(\tau) \subseteq \mathcal{Z}_C$, so $\Lambda_C \cap (-1, 1)^{m_C} \neq \emptyset$, completing the proof of Condition (I).

Proof of Condition (N). Since $m_C > m/2$ and $m \geq c_1 d_\star$ with $c_1 > 2$, we have $m_C \geq d_\star$. Hence the Gaussian matrix $\tilde{\mathcal{J}}_C^*$ has rank $d_\star$ almost surely. Because the whitening operator $\mathcal{H}$ is an isomorphism and $\tilde{\mathcal{J}}_C^* = \mathcal{H} \circ \mathcal{J}_C^*$,

$$\text{rank } \mathcal{J}_C = \text{rank } \mathcal{J}_C^* = \text{rank } \tilde{\mathcal{J}}_C^* = d_\star \quad \text{almost surely.}$$

We first establish that, almost surely,

$$|\mathcal{I}_<(v)| \geq d_\star \quad \text{for every } v \in \Lambda_C. \quad (33)$$

Fix $\mathcal{D} \subseteq [m_C]$ with $|\mathcal{D}| < d_\star$. Let $\tilde{\mathcal{J}}_{C, \mathcal{D}}$ denote the restriction of $\tilde{\mathcal{J}}_C$ to the coordinates in $\mathcal{D}$, and write $\mathcal{D}^c := [m_C] \setminus \mathcal{D}$. Define

$$\mathcal{B}_{\mathcal{D}, \tau} := \left\{ \text{there exists } v \in \Lambda_C \text{ such that } \mathcal{I}_<(v) = \mathcal{D} \text{ and } v_{\mathcal{D}^c} = \tau \right\}.$$

The event on which (33) fails is exactly

$$\left\{ \text{there exists } v \in \Lambda_C \text{ with } |\mathcal{I}_<(v)| < d_\star \right\} = \bigcup_{\substack{\mathcal{D} \subseteq [m_C] \\ |\mathcal{D}| < d_\star}} \bigcup_{\tau \in \{\pm 1\}^{\mathcal{D}^c}} \mathcal{B}_{\mathcal{D}, \tau}. \quad (34)$$

This is a finite union, so it suffices to prove $\mathbb{P}(\mathcal{B}_{\mathcal{D},\tau}) = 0$ for every fixed choice of $\mathcal{D}$ and τ . Fix such a pair $(\mathcal{D}, \tau)$. Recall $h = \tilde{\mathcal{J}}_N^*[\text{sign}(\varepsilon_N)]$. On the event $\mathcal{B}_{\mathcal{D},\tau}$, the free coordinates $v_{\mathcal{D}}$ of the corresponding point $v \in \Lambda_C$ satisfy

$$\tilde{\mathcal{J}}_{C,\mathcal{D}}^*[v_{\mathcal{D}}] = h - \tilde{\mathcal{J}}_{C,\mathcal{D}^c}^*[\tau]. \quad (35)$$

Let $e_1, \dots, e_{m_C}$ be the standard basis of $\mathbb{R}^{m_C}$. Since $|\mathcal{D}| < d_{\star} \leq m_C$, the complement $\mathcal{D}^c$ is nonempty; fix $j \in \mathcal{D}^c$. Equation (35) can hold only if

$$\tilde{\mathcal{J}}_C^*[e_j] \in \tau_j \left(h - \sum_{i \in \mathcal{D}^c \setminus \{j\}} \tau_i \tilde{\mathcal{J}}_C^*[e_i] \right) + \text{Im } \tilde{\mathcal{J}}_{C,\mathcal{D}}^* =: \mathcal{L}_{\mathcal{D},\tau,j}.$$

Let $\mathcal{F}_j := \sigma(h, \tilde{\mathcal{J}}_C^*[e_i] : i \in [m_C] \setminus \{j\})$ be the σ -algebra generated by the random variables h and all columns of $\tilde{\mathcal{J}}_C^*$ except the j th. Since $j \notin \mathcal{D}$, conditional on $\mathcal{F}_j$, the set $\mathcal{L}_{\mathcal{D},\tau,j}$ is a fixed affine subspace of $\mathbb{R}^{d_{\star}}$ of dimension at most $|\mathcal{D}| < d_{\star}$. Due to Assumption 5.13, $\tilde{\mathcal{J}}_C^*[e_j] \sim N(0, I_{d_{\star}})$ and is independent of $\mathcal{F}_j$. Therefore,

$$\mathbb{P}(\mathcal{B}_{\mathcal{D},\tau} \mid \mathcal{F}_j) \leq \mathbb{P}\{\tilde{\mathcal{J}}_C^*[e_j] \in \mathcal{L}_{\mathcal{D},\tau,j} \mid \mathcal{F}_j\} = 0 \quad \text{almost surely.}$$

Taking expectations gives $\mathbb{P}(\mathcal{B}_{\mathcal{D},\tau}) = 0$. Since this holds for every term in the finite union (34), the union has probability zero, proving (33).

Now fix a vertex ν of Λ_C and set $\mathcal{D} := \mathcal{I}_{<}(\nu)$. We claim that $\tilde{\mathcal{J}}_{C,\mathcal{D}}^*$ is injective. Otherwise, there exists a nonzero vector $u_{\mathcal{D}} \in \ker \tilde{\mathcal{J}}_{C,\mathcal{D}}^*$. Extend it to $u \in \mathbb{R}^{m_C}$ by setting $u_{\mathcal{D}^c} = 0$. Since $|\nu_i| < 1$ for every $i \in \mathcal{D}$, there exists $\epsilon > 0$ such that

$$\nu + \epsilon u, \nu - \epsilon u \in [-1, 1]^{m_C}.$$

Moreover, $\tilde{\mathcal{J}}_C^*[u] = \tilde{\mathcal{J}}_{C,\mathcal{D}}^*[u_{\mathcal{D}}] = 0$. It follows that $\tilde{\mathcal{J}}_C^*[\nu \pm \epsilon u] = \tilde{\mathcal{J}}_C^*[\nu] = h$. Thus, $\nu + \epsilon u$ and $\nu - \epsilon u$ are two distinct points of Λ_C whose midpoint is ν , contradicting the assumption that ν is a vertex. Therefore, $\tilde{\mathcal{J}}_{C,\mathcal{D}}^*$ is injective, which implies $|\mathcal{D}| \leq d_{\star}$. Combining this with (33) gives $|\mathcal{D}| = d_{\star}$. Consequently,

$$\text{rank } \tilde{\mathcal{J}}_{C,\mathcal{D}}^* = \text{rank } \tilde{\mathcal{J}}_{C,\mathcal{D}} = d_{\star}.$$

Since $\mathcal{H}$ is an isomorphism on the image of $\mathcal{J}_{C,\mathcal{D}}^*$, we conclude that

$$\text{rank } \mathcal{J}_{C,\mathcal{D}} = \text{rank } \mathcal{J}_{C,\mathcal{D}}^* = \text{rank } \tilde{\mathcal{J}}_{C,\mathcal{D}}^* = d_{\star} = \text{rank } \mathcal{J}_C.$$

Since this argument applies to every vertex ν , Condition (N) holds almost surely.

Finally, combining the high-probability event for Condition (I) with the almost-sure event for Condition (N), and adjusting the universal constants, gives

$$\mathbb{P}\{\text{Conditions (I) and (N) hold}\} \geq 1 - \exp(-c_2 d_{\star}).$$

This completes the proof. $\square$

B.8 Proof of Theorem 5.1

Fix a ground-truth factor u^* , and set $S := \text{supp}(u^*) = \text{supp}(x^*)$ and $q := n - s \geq 2$.

Criticality and affine-cube geometry. Using the argument of Ma and Fattahi [52, Corollary 1], specialized to diagonal Gaussian measurement matrices, u^* is a critical point of f_{sp} with probability at least $1 - e^{-cs}$ whenever $m \geq c_1 s$, where $c, c_1 > 0$ depend only on p . For nonnegative sparse regression, the linearized design map is $\mathcal{T} : \mathbb{R}^n \rightarrow \mathbb{S}^n$ with $\mathcal{T}[V] = 2 \text{diag}(u^* \odot V)$. Its image is the space of diagonal matrices supported on S , and therefore $d_* = \dim \text{Im } \mathcal{T} = s$. Lemma 5.14 and Proposition 5.15 then imply that Assumption 5.5 holds with probability at least $1 - e^{-c's}$, after increasing c_1 if necessary. A union bound, followed by an adjustment of constants, shows that with probability at least $1 - e^{-c_2 s}$, the factor u^* is critical and Assumption 5.5 holds simultaneously.

Isotropic kernel directions. Enumerate $S^c = [n] \setminus S$ as $\{i_1, \dots, i_q\}$ and, for each $a \in [q]$, set $V_a := e_{i_a}$. Since $u_{i_a}^* = 0$,

$$\mathcal{T}[V_a] = 2 \text{diag}(u^* \odot V_a) = 0,$$

and hence $V_a \in \ker \mathcal{J}_C$. Moreover,

$$\mathcal{Q}_C(V_a) = \mathcal{A}_C(V_a V_a^\top) = (a_{\ell, i_a})_{\ell \in \Omega_C}.$$

Thus $\mathcal{Q}_C(V_1), \dots, \mathcal{Q}_C(V_q)$ are independent standard Gaussian vectors in $\mathbb{R}^{m_C}$. More importantly, this remains true after conditioning on $\mathcal{J}_C$ and Λ_C . To see this, note that $\mathcal{J}_C$ depends only on the on-support coordinates $\{a_{\ell, S} : \ell \in \Omega_C\}$, while Λ_C is determined by these coordinates, their noisy counterparts $\{a_{\ell, S} : \ell \in \Omega_N\}$, and ε . In contrast, $\mathcal{Q}_C(V_a)$ depends only on the off-support coordinates $\{a_{\ell, i_a} : \ell \in \Omega_C\}$. Independence of the Gaussian coordinates and independence of ε from the measurements therefore give the claimed conditional independence.

Since $\mathcal{N}_C$ is determined by $\mathcal{J}_C$, it follows that, conditional on $\mathcal{J}_C$ and Λ_C ,

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$$

are independent standard Gaussian vectors on $\mathcal{N}_C$; in particular, their distributions are continuous and spherically symmetric. Moreover, $m_C > m/2 > s$ after increasing c_1 if necessary, and $\text{rank } \mathcal{J}_C = s$ almost surely, so $\dim \mathcal{N}_C = m_C - s \geq 1$. All hypotheses of Proposition 5.12 are therefore satisfied, implying that, conditional on $\mathcal{J}_C$ and Λ_C , the cone-selection test succeeds with probability at least $1 - 2^{-(q-1)}$. Theorem 5.11 then implies that u^* is inactive.

Probability bound and uniformity. Combining the preceding high-probability events and recalling that $q = n - s$ gives

$$\mathbb{P}\{u^* \text{ is an inactive critical point}\} \geq 1 - e^{-c_2 s} - 2^{-(n-s-1)}.$$

Finally, any other ground-truth factor $\hat{u}^*$ satisfying $\hat{u}^* \odot \hat{u}^* = x^*$ has the form $\hat{u}^* = Du^*$ for a diagonal sign matrix D . The map $u \mapsto Du$ is orthogonal and leaves f_{sp} invariant. Lemma 2.10 therefore shows that, on the same event, every ground-truth factor is an inactive critical point. No union bound over the ground-truth factors is needed. This proves the theorem. $\square$

B.9 Proof of Theorem 5.2

The proof follows the same structure as that of Theorem 5.1. Fix a ground-truth factor U^* satisfying $U^*U^{*\top} = X^*$, and set $q := n - r \geq 2$.

Criticality and affine-cube geometry. By Ma and Fattahi [52, Corollary 1], U^* is a critical point of f_{sym} with probability at least $1 - e^{-c_1 nr}$ whenever $m \geq c_1 nr$, where $c, c_1 > 0$ depend only on p . In the symmetric model, $\mathcal{L}$ is the identity and the linearized design map is $\mathcal{T}[V] = U^*V^\top + VU^{*\top}$. To compute the dimension of its image, choose an orthogonal matrix $Q \in \mathbb{R}^{k \times k}$ such that $U^*Q = (\bar{U}, 0)$ for some full-rank $\bar{U} \in \mathbb{R}^{n \times r}$. Then

$$\text{Im } \mathcal{T} = \{U^*QQ^\top V^\top + VQQ^\top U^{*\top} : V \in \mathbb{R}^{n \times k}\} = \{\bar{U}H^\top + H\bar{U}^\top : H \in \mathbb{R}^{n \times r}\}.$$

The kernel of $H \mapsto \bar{U}H^\top + H\bar{U}^\top$ is $\{\bar{U}K : K^\top = -K\}$. The rank-nullity theorem gives

$$d_* = \dim \text{Im } \mathcal{T} = nr - \frac{r(r-1)}{2}.$$

Lemma 5.14 and Proposition 5.15 now imply that Assumption 5.5 holds with probability at least $1 - e^{-c'd_*}$, after increasing c_1 if necessary. Since $d_* \geq nr/2$, a union bound and an adjustment of constants show that, with probability at least $1 - e^{-c_2 nr}$, the factor U^* is critical and Assumption 5.5 holds simultaneously.

Isotropic kernel directions. Choose a unit vector $w \in \ker U^*$, which exists because $k \geq r + 1$, and let $u_1, \dots, u_q$ be an orthonormal basis of $\text{col}(U^*)^\perp$. For each $j \in [q]$, set $V_j := u_j w^\top$. Since $U^*w = 0$,

$$\mathcal{T}[V_j] = U^*V_j^\top + V_jU^{*\top} = 0,$$

and hence $V_j \in \ker \mathcal{J}_C$. Moreover, because $\|w\| = 1$,

$$\mathcal{Q}_C(V_j) = \mathcal{A}_C(V_j V_j^\top) = \mathcal{A}_C(u_j u_j^\top) = (\langle \Gamma_\ell, u_j u_j^\top \rangle)_{\ell \in \Omega_C}.$$

The matrices $u_j u_j^\top$, $j \in [q]$, form an orthonormal family in $(\text{Im } \mathcal{T})^\perp$. Gaussian isotropy therefore shows that $\mathcal{Q}_C(V_1), \dots, \mathcal{Q}_C(V_q)$ are independent standard Gaussian vectors in $\mathbb{R}^{m_C}$. This conclusion remains valid after conditioning on $\mathcal{J}_C$ and Λ_C : the pair $(\mathcal{J}_C, \Lambda_C)$ is determined by the clean and corrupted measurements restricted to $\text{Im } \mathcal{T}$, together with ε , whereas the vectors $\mathcal{Q}_C(V_j)$ depend on mutually orthogonal measurement components in $(\text{Im } \mathcal{T})^\perp$. Independence of orthogonal Gaussian components and independence of ε from the measurements give the claim.

Since $\mathcal{N}_C$ is determined by $\mathcal{J}_C$, it follows that, conditional on $\mathcal{J}_C$ and Λ_C ,

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V_1), \dots, P_{\mathcal{N}_C} \mathcal{Q}_C(V_q)$$

are independent standard Gaussian vectors on $\mathcal{N}_C$; in particular, their distributions are continuous and spherically symmetric. Furthermore, $m_C \geq (1-p)m$. After increasing c_1 if necessary, Gaussian injectivity on the fixed space $\text{Im } \mathcal{T}$ gives $\text{rank } \mathcal{J}_C = d_\star$ almost surely and

$$\dim \mathcal{N}_C = m_C - d_\star \geq 1.$$

All hypotheses of Proposition 5.12 are therefore satisfied. Conditional on $\mathcal{J}_C$ and Λ_C , with probability at least $1 - 2^{-(q-1)}$, the cone-selection test in Theorem 5.11 succeeds. That theorem then implies that $U^\star$ is inactive.

Probability bound and uniformity. Combining the preceding high-probability events and recalling that $q = n - r$ gives

$$\mathbb{P}\{U^\star \text{ is an inactive critical point}\} \geq 1 - e^{-c_2 n r} - 2^{-(n-r-1)}.$$

Finally, every other ground-truth factor $\hat{U}^\star$ satisfying $\hat{U}^\star(\hat{U}^\star)^\top = X^\star$ has the form $\hat{U}^\star = U^\star Q$ for some orthogonal $Q \in \mathbb{R}^{k \times k}$. The map $U \mapsto UQ$ is orthogonal and leaves f_{sym} invariant. Lemma 2.10 therefore shows that, on the same event, every ground-truth factor is an inactive critical point. No union bound over the ground-truth factors is needed. This proves the theorem. $\square$

B.10 Proof of Theorem 5.3

Rather than constructing $q \geq 2$ distinct directions $V_1, \dots, V_q \in \ker \mathcal{J}_C$ with spherically symmetric projected second-order residuals and invoking the resulting success probability $1 - 2^{-(q-1)}$, we use a sharper argument. We construct two directions $V_1, V_2 \in \ker \mathcal{J}_C$ whose projected second-order residuals take the form g and $-g$, where $g \neq 0$ almost surely. Since the boundaries of the vertex normal cones have measure zero and each such cone is pointed, g and $-g$ lie in the relative interiors of two distinct vertex normal cones almost surely. Consequently, the cone-selection test succeeds almost surely, improving the probability $1 - 2^{-(q-1)}$ to one.

Proposition B.1 (Two-sector branching implies successful cone selection) *Fix a critical point $U^\star$ of f_{LR} satisfying $U^\star U^{\star\top} = X^\star$. Suppose that Assumption 5.5 holds and that $\dim \mathcal{N}_C \geq 1$. Suppose further that there exist linearly independent directions $V_1, V_2 \in \ker \mathcal{J}_C$ and a random vector $g \in \mathcal{N}_C$ such that, conditional on $\mathcal{J}_C$ and Λ_C , the law of g is absolutely continuous with respect to Lebesgue measure on $\mathcal{N}_C$ and, almost surely,*

$$P_{\mathcal{N}_C} \mathcal{Q}_C(aV_1 + bV_2) = abg \quad \text{for all } a, b \in \mathbb{R}. \quad (36)$$

Then, conditional on $\mathcal{J}_C$ and Λ_C , the cone-selection test in Theorem 5.11 succeeds almost surely. Specifically, there exist two distinct vertices ν_1, ν_2 of Λ_C , a linear subspace $\mathcal{W} \subseteq \ker \mathcal{J}_C$, and two nonempty relatively open cones $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{W}$ such that

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_j) \quad \text{for all } V \in \mathcal{C}_j, \quad j = 1, 2.$$

Proof By Proposition 5.6, Λ_C is full-dimensional in its affine space. Hence its vertex normal cones $N_{\Lambda_C}^{\mathcal{N}_C}(\nu)$, $\nu \in \text{Vert}(\Lambda_C)$, are full-dimensional and pointed, cover $\mathcal{N}_C$, and have pairwise disjoint relative interiors [55, Theorem 2.3.2 and Proposition 2.3.6]. The union of their boundaries therefore has Lebesgue measure zero in $\mathcal{N}_C$. By the absolute continuity of the conditional law of g , almost surely both g and $-g$ are nonzero and avoid these boundaries. Hence there are unique vertices $\nu_1, \nu_2 \in \text{Vert}(\Lambda_C)$ such that $g \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_1)$ and $-g \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_2)$. The vertices are distinct; otherwise the pointed cone $N_{\Lambda_C}^{\mathcal{N}_C}(\nu_1)$ would contain both g and $-g$, and hence the nonzero line $\text{span } g$, which is impossible. Set $\mathcal{W} := \text{span}\{V_1, V_2\}$ and define

$$\mathcal{C}_1 := \{aV_1 + bV_2 : ab > 0\}, \quad \mathcal{C}_2 := \{aV_1 + bV_2 : ab < 0\}.$$

Because V_1 and V_2 are linearly independent, $\mathcal{C}_1$ and $\mathcal{C}_2$ are nonempty relatively open cones in $\mathcal{W}$. For every $V = aV_1 + bV_2 \in \mathcal{C}_1$, equation (36) yields

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) = abg \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_1),$$

because $ab > 0$. Similarly, for every $V = aV_1 + bV_2 \in \mathcal{C}_2$, we have $ab < 0$ and therefore

$$P_{\mathcal{N}_C} \mathcal{Q}_C(V) = |ab|(-g) \in \text{relint } N_{\Lambda_C}^{\mathcal{N}_C}(\nu_2).$$

Thus $\mathcal{W}$, $\mathcal{C}_1$, $\mathcal{C}_2$, ν_1 , and ν_2 satisfy the cone-selection conditions in Theorem 5.11. This completes the proof. $\square$

Using Proposition B.1 in lieu of Proposition 5.12, we prove Theorem 5.3.

Proof of Theorem 5.3 The proof follows the same structure as those of Theorems 5.1 and 5.2. Fix a balanced ground-truth factor pair (U_1^*, U_2^*) satisfying $U_1^* U_2^{*\top} = X^*$ and $U_1^{*\top} U_1^* = U_2^{*\top} U_2^*$, and introduce the stacked factor

$$U^* := \begin{bmatrix} U_1^* \\ U_2^* \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times k}.$$

The objective $f_{\text{asym}}(U_1, U_2)$ is precisely the robust low-rank objective (15) evaluated at the stacked factor $U = [U_1^\top U_2^\top]^\top$.

Criticality and affine-cube geometry. By Ma and Fattahi [52, Corollary 2], the stacked factor U^* , equivalently the pair (U_1^*, U_2^*) , is a critical point with probability at least $1 - e^{-c(n_1+n_2)r}$ whenever $m \geq c_1(n_1+n_2)r$, where $c, c_1 > 0$ depend only on p . For a perturbation

$$H = \begin{bmatrix} H_1 \\ H_2 \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times k},$$

the linearized design map defined in Section 5 is

$$\mathcal{T}[H] = \mathcal{L}(U^* H^\top + H U^{*\top}) = \begin{bmatrix} 0 & U_1^* H_2^\top + H_1 U_2^{*\top} \\ U_2^* H_1^\top + H_2 U_1^{*\top} & 0 \end{bmatrix}.$$

Write $\mathcal{T}_{12}[H] := U_1^* H_2^\top + H_1 U_2^{*\top}$ for its upper-right block. Since $\mathcal{T}[H]$ is uniquely determined by $\mathcal{T}_{12}[H]$, we have $\dim \text{Im } \mathcal{T} = \dim \text{Im } \mathcal{T}_{12}$. Because the factor pair is balanced and X^* has rank r , the common Gram matrix $U_1^{*\top} U_1^* = U_2^{*\top} U_2^*$ has rank r . Choose an orthogonal matrix $Q \in \mathbb{R}^{k \times k}$ such that $U_1^* Q = (\bar{U}_1, 0)$ and $U_2^* Q = (\bar{U}_2, 0)$, where $\bar{U}_1 \in \mathbb{R}^{n_1 \times r}$ and $\bar{U}_2 \in \mathbb{R}^{n_2 \times r}$ have full column rank. Since right multiplication by Q is a linear isomorphism on the perturbation space,

$$\text{Im } \mathcal{T}_{12} = \left\{ \bar{U}_1 H_2^\top + H_1 \bar{U}_2^\top : H_1 \in \mathbb{R}^{n_1 \times r}, H_2 \in \mathbb{R}^{n_2 \times r} \right\}.$$

The kernel of the map $(H_1, H_2) \mapsto \bar{U}_1 H_2^\top + H_1 \bar{U}_2^\top$ is $\{(\bar{U}_1 K, -\bar{U}_2 K^\top) : K \in \mathbb{R}^{r \times r}\}$. The rank-nullity theorem therefore gives

$$d_\star = \dim \text{Im } \mathcal{T} = \dim \text{Im } \mathcal{T}_{12} = r(n_1 + n_2) - r^2 = r(n_1 + n_2 - r).$$

Lemma 5.14 and Proposition 5.15 now imply that Assumption 5.5 holds with probability at least $1 - e^{-c'd_\star}$, after increasing c_1 if necessary. Since $r \leq \min\{n_1, n_2\}$, we have $d_\star \geq \frac{1}{2}(n_1 + n_2)r$. A union bound and an adjustment of constants thus show that, with probability at least $1 - e^{-c_2(n_1+n_2)r}$, the stacked factor $U^\star$ is critical and Assumption 5.5 holds simultaneously.

Two-sector kernel directions. The balancing identity gives $\ker U^\star = \ker U_1^\star = \ker U_2^\star$. Choose unit vectors $w \in \ker U^\star$, $u \in \text{col}(U_1^\star)^\perp$, and $v \in \text{col}(U_2^\star)^\perp$, which exist because $k \geq r + 1$, $n_1 \geq r + 1$, and $n_2 \geq r + 1$. Define the stacked directions

$$V_1 := \begin{bmatrix} uw^\top \\ 0_{n_2 \times k} \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times k}, \quad V_2 := \begin{bmatrix} 0_{n_1 \times k} \\ vw^\top \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times k}.$$

The directions V_1 and V_2 are linearly independent. Moreover,

$$\mathcal{T}_{12}[V_1] = u(U_2^\star w)^\top = 0, \quad \mathcal{T}_{12}[V_2] = (U_1^\star v)^\top = 0.$$

Hence $V_1, V_2 \in \ker \mathcal{T} \subseteq \ker \mathcal{J}_C$. For $a, b \in \mathbb{R}$, the two blocks of $aV_1 + bV_2$ are $auw^\top$ and $bvw^\top$, and therefore

$$\mathcal{L}\left((aV_1 + bV_2)(aV_1 + bV_2)^\top\right) = ab \begin{bmatrix} 0 & uv^\top \\ vu^\top & 0 \end{bmatrix}.$$

It follows that the clean second-order residual satisfies

$$\mathcal{Q}_C(aV_1 + bV_2) = ab z_C, \quad z_C := (\langle \Gamma_\ell, uv^\top \rangle)_{\ell \in \Omega_C}.$$

Furthermore, $uv^\top \in (\text{Im } \mathcal{T}_{12})^\perp$, because $U_1^{\star\top} u = 0$ and $U_2^{\star\top} v = 0$. Consequently, Gaussian isotropy shows that z_C is a standard Gaussian vector in $\mathbb{R}^{m_C}$ independent of the measurement components that determine $\mathcal{J}_C$ and $\mathcal{J}_N$. Since ε is independent of the measurements, z_C is independent of $(\mathcal{J}_C, \mathcal{J}_N, \varepsilon)$ and therefore of $(\mathcal{J}_C, \Lambda_C)$. Since $m_C \geq (1-p)m$, after increasing c_1 if necessary we have $m_C > d_\star$. Gaussian injectivity on the fixed space $\text{Im } \mathcal{T}$ then gives $\text{rank } \mathcal{J}_C = d_\star$ almost surely, and hence

$$\dim \mathcal{N}_C = m_C - d_\star \geq 1.$$

Because $\mathcal{N}_C$ is determined by $\mathcal{J}_C$, it follows that, conditional on $\mathcal{J}_C$ and Λ_C , $g := P_{\mathcal{N}_C} z_C$ is a standard Gaussian vector on $\mathcal{N}_C$ and therefore has a density there. Projecting the second-order identity onto $\mathcal{N}_C$ yields

$$P_{\mathcal{N}_C} \mathcal{Q}_C(aV_1 + bV_2) = ab g \quad \text{for all } a, b \in \mathbb{R}.$$

Thus all hypotheses of Proposition B.1 are satisfied. That proposition shows that the cone-selection conditions in Theorem 5.11 hold almost surely, and the theorem implies that $U^\star$, equivalently $(U_1^\star, U_2^\star)$, is inactive.

Probability bound and uniformity. Combining the preceding events gives

$$\mathbb{P}\{U^\star \text{ is an inactive critical point}\} \geq 1 - e^{-c_2(n_1+n_2)r}.$$

Finally, let

$$\hat{U}^\star := \begin{bmatrix} \hat{U}_1^\star \\ \hat{U}_2^\star \end{bmatrix}$$

be any other balanced ground-truth factor. Any two balanced factorizations of $X^\star$ differ by a common right-orthogonal transformation, so $\hat{U}^\star = U^\star \hat{Q}$ for some orthogonal matrix $\hat{Q} \in \mathbb{R}^{k \times k}$. The map $U \mapsto U\hat{Q}$ is orthogonal and leaves the stacked objective, and hence f_{asym} , invariant. Lemma 2.10 therefore shows that, on the same event, every balanced ground-truth factor is an inactive critical point. No union bound over the balanced factorizations is needed. This proves the theorem. $\square$

C Details of the Low-Rank Matrix-Sensing Experiment

Figure 2 uses the symmetric matrix-sensing formulation (17) with exact measurements. We set $n = k = 100$, $r = 2$, and $m = 200$. In each of 50 independent trials, we generate an orthogonal matrix Q from a standard Gaussian matrix, let Q_r contain its first two columns, and define

$$X^* = Q_r \text{diag}(10, 1) Q_r^\top.$$

Thus, any ground-truth factor U^* with $U^* U^{*\top} = X^*$ yields $\|U^*\|_F^2 / \|U^*\|_{\text{op}}^2 = 1.1$. The sensing matrices Γ_i have i.i.d. standard normal entries, and $y_i = \langle \Gamma_i, X^* \rangle$. We compare the ℓ_1 -loss

$$\phi_0(U) = \frac{1}{m} \sum_{i=1}^m |\langle \Gamma_i, U U^\top - X^* \rangle|$$

with $\phi_\lambda(U) = \phi_0(U) + \lambda \langle G, U \rangle$, where $\lambda = 0.20$ and $G = \tilde{G} / \|\tilde{G}\|_F$ for a standard normal matrix $\tilde{G}$. The paired trajectories start from the same initialization $U_0 = U^* + E$, where the entries of E are independently drawn from $N(0, (5 \cdot 10^{-5})^2)$. For $\beta \in \{0, \lambda\}$, we perform 5000 iterations

$$\begin{aligned} U_{t+1}^{(\beta)} &= U_t^{(\beta)} - \frac{0.01}{\sqrt{t+1}} g_t^{(\beta)}, \\ g_t^{(\beta)} &= \frac{1}{m} \sum_{i=1}^m \text{sign} \left(\left\langle \Gamma_i, U_t^{(\beta)} U_t^{(\beta)\top} - X^* \right\rangle_F \right) (\Gamma_i + \Gamma_i^\top) U_t^{(\beta)} + \beta G, \end{aligned}$$

Within each trial, the two trajectories share the ground truth, measurements, and initialization and differ only by the linear perturbation. We record the relative recovery error and effective rank every ten iterations and report all trajectories together with their coordinate-wise medians in Figure 2. All simulations use NumPy's `default_rng`, with seed $2 + 1009j$ for trial $j = 0, \dots, 49$.