WHAT CAN BE RECOVERED UNDER SPARSE ADVERSARIAL CORRUPTION? ASSUMPTION-FREE THEORY FOR LINEAR MEASUREMENTS

Vishal Halder[⋆] *Alexandre Reiffers-Masson*[⋆] *Abdeldjalil Aïssa-El-Bey*[⋆] *Gugan Thoppe*[†]

[⋆] IMT Atlantique, Lab-STICC, UMR CNRS 6285, Brest, France
[†] Indian Institute of Science, Computer Science and Automation, Bengaluru, India

ABSTRACT

Let $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ be an arbitrary, known matrix and $\boldsymbol{e}$ a q-sparse adversarial vector. Given $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star + \boldsymbol{e}$ and q, we seek the smallest set containing $\boldsymbol{x}^\star$—hence the one conveying maximal information about $\boldsymbol{x}^\star$—that is uniformly recoverable from $\boldsymbol{y}$ without knowing $\boldsymbol{e}$. While exact recovery of $\boldsymbol{x}^\star$ via strong (and often impractical) structural assumptions on $\boldsymbol{A}$ or $\boldsymbol{x}^\star$ (e.g., restricted isometry, sparsity) is well studied, recoverability for arbitrary $\boldsymbol{A}$ and $\boldsymbol{x}^\star$ remains open. Our main result shows that the best that one can hope to recover is $\boldsymbol{x}^\star + \ker(\boldsymbol{U})$, where $\boldsymbol{U}$ is the unique projection matrix onto the intersection of rowspaces of all possible submatrices of $\boldsymbol{A}$ obtained by deleting $2q$ rows. Moreover, we prove that every $\boldsymbol{x}$ that minimizes the ℓ_0-norm of $\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}$ lies in $\boldsymbol{x}^\star + \ker(\boldsymbol{U})$, which then gives a constructive approach to recover this set.

***Index Terms*—** sparse adversarial corruption, linear measurement model, robust recovery, robust orthogonal projection

1. INTRODUCTION

Linear measurements are foundational to signal processing and machine learning, yet they are highly susceptible to adversarial corruptions [1–3]. In this paper, we focus on a specific corruption model, called *sparse adversarial corruption,* which models critical real-world scenarios, such as compromised sensors in control systems and smart grids [4–6]. In this model, a limited number of measurements can be arbitrarily and adversarially corrupted, with no bounds on the corruption magnitude when it occurs. Our paper studies the following fundamental question:

Given an arbitrary measurement matrix $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and an unknown signal $\boldsymbol{x}^\star \in \mathbb{R}^n$, what is the set containing the maximum amount of information regarding $\boldsymbol{x}^\star$ that is guaranteed recoverable from $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star + \boldsymbol{e}$ despite a q-sparse adversarial vector $\boldsymbol{e}$?

Prior work has focused on exact recovery of $\boldsymbol{x}^\star$ under sparse corruptions. However, the exact recovery guarantees require strong structural assumptions on $\boldsymbol{A}$, such as restricted isometry [7], nullspace conditions [8,9], or spark criteria [10]. These assumptions are rarely satisfied in practice and are NP-hard to verify [11–13], leaving a critical gap between theoretical guarantees and real-world applicability. We bridge this gap by determining what can be robustly recovered about $\boldsymbol{x}^\star$ without any assumptions on $\boldsymbol{A}$, providing a universal, assumption-free characterization.

To contextualize our contributions, we first review the recoverable information in simpler settings. In the absence of corruption ($\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star$), the most informative set is $\boldsymbol{x}^\star + \ker(\boldsymbol{A})$, since any $\boldsymbol{x}^\star + \boldsymbol{v}$ with $\boldsymbol{v} \in \ker(\boldsymbol{A})$ produces the same measurements. If the corruption support $I \subseteq [m]$ with $|I| \leq q$ were known, then the most informative set would be $\boldsymbol{x}^\star + \ker(\boldsymbol{A}_{[m]\setminus I})$, obtained by discarding the corrupted measurements. In this paper, we establish the analogue of these results when the q-sparse corruption is unknown.

Our Contributions. We begin by introducing the notions of ambiguity sets and robust functions, and establish a necessary and sufficient condition linking them (see Proposition 1), forming a framework for reasoning about adversarial robustness. We then define the solution set associated with a robust function and, in Theorem 1, identify a unique *robust orthogonal projection matrix* $\boldsymbol{U}$. This matrix ensures that (a) $\boldsymbol{U}\boldsymbol{x}^\star$ is the ℓ_2-optimal robust projection of $\boldsymbol{x}^\star$, and (b) the affine solution space $\boldsymbol{x}^\star + \ker(\boldsymbol{U})$ is the inclusion-wise minimal solution set among those generated by all robust functions. We provide an algorithm to compute $\boldsymbol{U}$ (Algorithm 1), and show that pairing it with ℓ_0-decoding gives a constructive recovery procedure. Finally, in Section 4 we validate our theory on a concrete example and study the performance of Algorithm 1 on small scale numerical experiments.

2. NOTATION

We denote scalars by lowercase letters q, m, n and sets by uppercase letters S, T. We denote vector spaces by calligraphic letters $\mathcal{R}, \mathcal{S}$, and possibly nonlinear functions by script letters $\mathscr{G}$. We write vectors as bold lowercase letters $\boldsymbol{x}, \boldsymbol{y}$ and matrices as bold uppercase letters $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{U}$. For a matrix $\boldsymbol{A}$, we write $\mathrm{rowspan}(\boldsymbol{A})$ for the vector space spanned by its rows, $\ker(\boldsymbol{A})$ for its kernel (null space), and $\mathrm{im}(\boldsymbol{A})$ for its image (column space). We denote the support of a vector $\boldsymbol{v}$ by $\mathrm{supp}(\boldsymbol{v}) = \{i : v_i \neq 0\}$ and its ℓ_0-norm by

$\|\boldsymbol{v}\|_0 = |\operatorname{supp}(\boldsymbol{v})|$. For $T \subseteq [m]$, we denote the submatrix of $\boldsymbol{A}$ with rows indexed by T as $\boldsymbol{A}_T$. We write $\mathcal{R}^\perp$ for the orthogonal complement of a subspace $\mathcal{R}$. We write $\operatorname{span}\{\boldsymbol{v}_i : i \in I\}$ for the span of a set of vectors $\{\boldsymbol{v}_i\}_{i\in I}$. For subspaces $\mathcal{R}_1, \mathcal{R}_2$, their sum is $\mathcal{R}_1 + \mathcal{R}_2 := \{\, \mathbf{v}_1 + \mathbf{v}_2 \mid \mathbf{v}_1 \in \mathcal{R}_1,\ \mathbf{v}_2 \in \mathcal{R}_2 \,\}$.

3. MAIN RESULTS

This section is dedicated to the main theoretical results of our paper. We begin by defining the notion of the ambiguity set, which captures all differences of signal pairs that can be explained by q-sparse corruptions. Formally, we have the following definition:

Definition 1. For $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and integer $q < m/2$, define the *ambiguity set* $S_q^{\boldsymbol{A}} := \{\boldsymbol{v} \in \mathbb{R}^n : \|\boldsymbol{A}\boldsymbol{v}\|_0 \le 2q\}$.

The ambiguity set captures all the $\boldsymbol{v} = \boldsymbol{x}_1 - \boldsymbol{x}_2$ for which there exists $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$ with $\|\boldsymbol{e}_1\|_0, \|\boldsymbol{e}_2\|_0 \le q$ such that $\boldsymbol{A}\boldsymbol{x}_1 + \boldsymbol{e}_1 = \boldsymbol{A}\boldsymbol{x}_2 + \boldsymbol{e}_2$. The $2q$ in Definition 1 arises because any ambiguity comes from a pair of signals, and differences are always computed between two signals, so the maximum support is $2q$. Next, we give a precise definition of functions robust to any q-sparse adversarial corruption to measurements from a matrix $\boldsymbol{A}$.

Definition 2. For $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and integer $q < m/2$, a function $\mathscr{G} : \mathbb{R}^n \to \mathcal{Z}$ (arbitrary codomain) is said to be $(\boldsymbol{A}, q)$-*robust* if for every $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$, and q-sparse $\boldsymbol{e}_1, \boldsymbol{e}_2 \in \mathbb{R}^m$, $\boldsymbol{A}\boldsymbol{x}_1 + \boldsymbol{e}_1 = \boldsymbol{A}\boldsymbol{x}_2 + \boldsymbol{e}_2 \implies \mathscr{G}(\boldsymbol{x}_1) = \mathscr{G}(\boldsymbol{x}_2)$.

Our first result establishes that a function is robust exactly when it produces the same output for signals differing by an element of the ambiguity set.

Proposition 1. *Let $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, integer $q < m/2$, and $S_q^{\boldsymbol{A}} := \{\boldsymbol{v} \in \mathbb{R}^n : \|\boldsymbol{A}\boldsymbol{v}\|_0 \le 2q\}$. A function $\mathscr{G} : \mathbb{R}^n \to \mathcal{Z}$ is $(\boldsymbol{A}, q)$-robust iff $\mathscr{G}(\boldsymbol{x} + \boldsymbol{v}) = \mathscr{G}(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathbb{R}^n$ and $\boldsymbol{v} \in S_q^{\boldsymbol{A}}$.*

Proof. ($\Rightarrow$) For $\boldsymbol{v} \in S_q^{\boldsymbol{A}}$, let $T = \operatorname{supp}(\boldsymbol{A}\boldsymbol{v})$ with $|T| \le 2q$. Partition T to get $T_1 \cup T_2$ with $|T_i| \le q$. Next, define $\boldsymbol{e}'$ and $\boldsymbol{e}$ by $(\boldsymbol{e}')_j = (\boldsymbol{A}\boldsymbol{v})_j$ if $j \in T_1$, 0 otherwise, and $(\boldsymbol{e})_j = -(\boldsymbol{A}\boldsymbol{v})_j$ if $j \in T_2$, 0 otherwise. Then, for any $\boldsymbol{x} \in \mathbb{R}^n$ and $\boldsymbol{v} \in S_q^{\boldsymbol{A}}$, we have $\boldsymbol{A}(\boldsymbol{x} + \boldsymbol{v}) + \boldsymbol{e} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}'$ with $\|\boldsymbol{e}\|_0, \|\boldsymbol{e}'\|_0 \le q$, so robustness implies $\mathscr{G}(\boldsymbol{x} + \boldsymbol{v}) = \mathscr{G}(\boldsymbol{x})$.

($\Leftarrow$) Let $\boldsymbol{x}$ and $\boldsymbol{x}'$ be such that $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} = \boldsymbol{A}\boldsymbol{x}' + \boldsymbol{e}'$ with $\|\boldsymbol{e}\|_0, \|\boldsymbol{e}'\|_0 \le q$. Then, $\boldsymbol{A}(\boldsymbol{x}' - \boldsymbol{x}) = \boldsymbol{e} - \boldsymbol{e}'$ has support $\le 2q$, which implies $\boldsymbol{x}' - \boldsymbol{x} \in S_q^{\boldsymbol{A}}$. Therefore, by hypothesis, $\mathscr{G}(\boldsymbol{x}) = \mathscr{G}(\boldsymbol{x} + \boldsymbol{x}' - \boldsymbol{x}) = \mathscr{G}(\boldsymbol{x}')$, as desired. □

We now provide a way to quantify the amount of information we can obtain from any $(\boldsymbol{A}, q)$-robust function $\mathscr{G}$.

Definition 3. Let $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and $q < m/2$ be an integer. Then, for any $(\boldsymbol{A}, q)$-robust function $\mathscr{G}$ and any $\boldsymbol{x}^\star \in \mathbb{R}^n$, the *$\mathscr{G}$-solution set containing $\boldsymbol{x}^\star$* is $X(\mathscr{G}, \boldsymbol{x}^\star) := \{\boldsymbol{x} \in \mathbb{R}^n : \mathscr{G}(\boldsymbol{x}) = \mathscr{G}(\boldsymbol{x}^\star)\}$. Clearly, $X(\mathscr{G}, \boldsymbol{x}^\star) = \mathscr{G}^{-1}(\mathscr{G}(\boldsymbol{x}^\star))$, the preimage of $\mathscr{G}(\boldsymbol{x}^\star)$ under $\mathscr{G}$.

We now state and prove our main results: Theorems 1 and 2. Theorem 1 shows that, for any $\boldsymbol{A}$ and q, there exists a linear function $\mathscr{G}^\star(\boldsymbol{x}) = \boldsymbol{U}\boldsymbol{x}$ such that (1) $\boldsymbol{U}$ is an orthogonal projection matrix, (2) $\mathscr{G}^\star$ is $(\boldsymbol{A}, q)$-robust, and (3) $\mathscr{G}^*$ yields an inclusion-wise minimal solution set: $X(\mathscr{G}^*, \boldsymbol{x}^\star) \subseteq X(\mathscr{G}, \boldsymbol{x}^\star)$ for every $(\boldsymbol{A}, q)$-robust $\mathscr{G}$ and every $\boldsymbol{x}^\star \in \mathbb{R}^n$. Seeking an inclusion-wise minimal solution set is equivalent to finding the set that provides the maximum information about $\boldsymbol{x}^\star$.

Theorem 1. *Consider a matrix $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and an integer $q < m/2$. Also, let $\mathcal{R} = \bigcap_{T\subseteq[m],\, |T|=m-2q} \operatorname{rowspan}(\boldsymbol{A}_T)$ and $\boldsymbol{U}$ the orthogonal projector onto $\mathcal{R}$. That is, suppose $\boldsymbol{U}$ is the unique symmetric idempotent matrix with $\operatorname{range}(\boldsymbol{U}) = \mathcal{R}$ and $\ker(\boldsymbol{U}) = \mathcal{R}^\perp$. Then, the $\boldsymbol{x} \to \boldsymbol{U}\boldsymbol{x}$ map is $(\boldsymbol{A}, q)$-robust. Moreover, for any $(\boldsymbol{A}, q)$-robust $\mathscr{G} : \mathbb{R}^n \to \mathcal{Z}$ and $\boldsymbol{x}^\star \in \mathbb{R}^n$,*

$$\{\boldsymbol{x} \in \mathbb{R}^n : \mathscr{G}(\boldsymbol{x}) = \mathscr{G}(\boldsymbol{x}^\star)\} \supseteq \boldsymbol{x}^\star + \ker(\boldsymbol{U}), \tag{1}$$

with equality for $\mathscr{G}(\boldsymbol{x}) = \boldsymbol{U}\boldsymbol{x}$.

Proof. By the fundamental theorem of linear algebra [14], $\operatorname{rowspan}(\boldsymbol{A}_T) = (\ker(\boldsymbol{A}_T))^\perp$, and since $(\cap \mathcal{V}_i)^\perp = \sum \mathcal{V}_i^\perp$ (see [15, Ex. 5, §12]), we obtain $\mathcal{R}^\perp = \sum_{\substack{T\subseteq[m],\\ |T|=m-2q}} \ker(\boldsymbol{A}_T)$, where the sum of vector spaces is as defined in Section 2. For $\boldsymbol{v} \in S_q^{\boldsymbol{A}}$ with $Q = \operatorname{supp}(\boldsymbol{A}\boldsymbol{v})$ and $|Q| \le 2q$, we have $\boldsymbol{v} \in \ker(\boldsymbol{A}_{Q^c})$, so

$$S_q^{\boldsymbol{A}} = \bigcup_{\substack{Q\subseteq[m],\\ |Q|\le 2q}} \ker(\boldsymbol{A}_{Q^c}), \quad \operatorname{span}(S_q^{\boldsymbol{A}}) = \sum_{\substack{Q\subseteq[m],\\ |Q|\le 2q}} \ker(\boldsymbol{A}_{Q^c}) \tag{2}$$

(since the span of a union of subspaces is their sum, see [15, Thm. 3, §11]). Writing $T = Q^c$ implies $|T| \ge m - 2q$, and using $\ker(\boldsymbol{A}_{T_2}) \subseteq \ker(\boldsymbol{A}_{T_1})$ for $T_1 \subseteq T_2$, it follows that

$$\operatorname{span}(S_q^{\boldsymbol{A}}) = \sum_{\substack{T\subseteq[m],\\ |T|=m-2q}} \ker(\boldsymbol{A}_T) = \mathcal{R}^\perp. \tag{3}$$

Thus,

$$\ker(\boldsymbol{U}) = \operatorname{span}(S_q^{\boldsymbol{A}}), \tag{4}$$

which implies that $\boldsymbol{U}\boldsymbol{v} = 0$ whenever $\boldsymbol{v} \in S_q^{\boldsymbol{A}}$. Proposition 1 now shows that the map $\mathscr{G}^\star(\boldsymbol{x}) = \boldsymbol{U}\boldsymbol{x}$ is $(\boldsymbol{A}, q)$-robust. To prove the final statement, let $\boldsymbol{x}^\star \in \mathbb{R}^n$ and $\mathscr{G}$ be $(\boldsymbol{A}, q)$-robust. Any $\boldsymbol{w} \in \operatorname{span}(S_q^{\boldsymbol{A}})$ can be written as $\boldsymbol{w} = \sum_{j=1}^r c_j \boldsymbol{v}_j$ with $\boldsymbol{v}_j \in S_q^{\boldsymbol{A}}$ and $c_j \in \mathbb{R}$. Since $c_j\boldsymbol{v}_j \in S_q^{\boldsymbol{A}}$, we apply Proposition 1 iteratively: first $\mathscr{G}(\boldsymbol{x}^\star + c_1\boldsymbol{v}_1) = \mathscr{G}(\boldsymbol{x}^\star)$, then $\mathscr{G}(\boldsymbol{x}^\star + c_1\boldsymbol{v}_1 + c_2\boldsymbol{v}_2) = \mathscr{G}(\boldsymbol{x}^\star + c_1\boldsymbol{v}_1) = \mathscr{G}(\boldsymbol{x}^\star)$, and so on until $\mathscr{G}(\boldsymbol{x}^\star + \sum_{j=1}^r c_j\boldsymbol{v}_j) = \mathscr{G}(\boldsymbol{x}^\star)$. Hence, $\boldsymbol{x}^\star + \operatorname{span}(S_q^{\boldsymbol{A}}) \subseteq \{\boldsymbol{x} \in \mathbb{R}^n : \mathscr{G}(\boldsymbol{x}) = \mathscr{G}(\boldsymbol{x}^\star)\}$. □

This theorem shows that the robustness of a matrix $\boldsymbol{A}$ against unknown q-sparse adversaries is characterized exactly by the deletion of $2q$ rows of $\boldsymbol{A}$. In particular, the theorem identifies the subspace

$$\mathcal{R} = \bigcap_{\substack{T\subseteq[m],\\ |T|=m-2q}} \operatorname{rowspan}(\boldsymbol{A}_T) \tag{5}$$

as the object governing robustness. We will refer to this intersection as the *robust subspace*, and to its orthogonal projector $\boldsymbol{U}$—constructed in Theorem 1—as the *robust orthogonal projection matrix*. In Section 4, we present an algorithm for computing $\boldsymbol{U}$.

Example 1. Let

$$\boldsymbol{A} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix}^\top, \qquad q = 1.$$

Both $\boldsymbol{x}_1 = [1,2]^\top$ and $\boldsymbol{x}_2 = [1,1]^\top$ can explain the same corrupted measurement, each with different 1-sparse corruptions $\boldsymbol{e}_1 = [0,0,0,-1,0]^\top$, $\boldsymbol{e}_2 = [0,0,0,0,1]^\top$, respectively. In this case, only the first coordinate of $\boldsymbol{x}^\star$ is determined, suggesting that the robust subspace is $\operatorname{rowspan}([1,0])$. Indeed, by Theorem 1, the robust orthogonal projection matrix $\boldsymbol{U} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ projects onto $\operatorname{rowspan}([1,0])$, and, $\boldsymbol{U}\boldsymbol{x}^\star$ is exactly the robust part of the signal—its first coordinate.

Remark. Our results in Theorem 1 are universal in the sense that they fully characterize what can be deterministically recovered without any assumptions on $\boldsymbol{A}$. They also generalize prior results on full signal recovery under structured measurement matrices. Indeed, Theorem 1 recovers the full-recovery condition stated in Proposition 2 in [16] as a special case when $\boldsymbol{A}$ satisfies conditions equivalent to the robust orthogonal projection being identity, i.e., $\boldsymbol{U} = \mathbb{I}_n$.

Given Theorem 1, which establishes that the inclusion-wise minimal solution set that is robust to q adversarial corruptions is $\boldsymbol{x}^\star + \ker(\boldsymbol{U})$, the central question now becomes how to extract this set. Clearly, this set is equivalent to $\{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{U}\boldsymbol{x} = \boldsymbol{U}\boldsymbol{x}^\star\}$. Our next result shows how we can extract this set.

Theorem 2. *Let $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star + \boldsymbol{e}$ with $\|\boldsymbol{e}\|_0 \le q$, and let $\boldsymbol{U}$ be the matrix as defined in Theorem 1. Then every $\hat{\boldsymbol{x}} \in \arg\min_{\boldsymbol{x}\in\mathbb{R}^n} \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_0$ satisfies $\boldsymbol{U}\hat{\boldsymbol{x}} = \boldsymbol{U}\boldsymbol{x}^\star$.*

Proof. Let $\hat{\boldsymbol{e}} = \boldsymbol{y} - \boldsymbol{A}\hat{\boldsymbol{x}}$. Since $\hat{\boldsymbol{x}}$ minimizes $\|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_0$, we have $\|\hat{\boldsymbol{e}}\|_0 \le \|\boldsymbol{e}\|_0 \le q$. Thus $\operatorname{supp}(\hat{\boldsymbol{e}} - \boldsymbol{e}) \subseteq \operatorname{supp}(\hat{\boldsymbol{e}}) \cup \operatorname{supp}(\boldsymbol{e})$, giving $\|\hat{\boldsymbol{e}} - \boldsymbol{e}\|_0 \le 2q$. But $\hat{\boldsymbol{e}} - \boldsymbol{e} = (\boldsymbol{y} - \boldsymbol{A}\hat{\boldsymbol{x}}) - (\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}^\star) = \boldsymbol{A}(\boldsymbol{x}^\star - \hat{\boldsymbol{x}})$, so $\boldsymbol{x}^\star - \hat{\boldsymbol{x}} \in S_q^{\boldsymbol{A}}$. By (4), $S_q^{\boldsymbol{A}} \subseteq \ker(\boldsymbol{U})$, hence $\boldsymbol{U}\hat{\boldsymbol{x}} = \boldsymbol{U}\boldsymbol{x}^\star$. □

Theorem 2 demonstrates that ℓ_0-decoding achieves the theoretical guarantee in Theorem 1: any minimizer $\hat{\boldsymbol{x}}$ satisfies $\boldsymbol{U}\hat{\boldsymbol{x}} = \boldsymbol{U}\boldsymbol{x}^\star$, exactly recovering the orthogonal projection of $\boldsymbol{x}^\star$ in the robust subspace $\operatorname{im}(\boldsymbol{U})$. Thus, the robust orthogonal projection matrix $\boldsymbol{U}$ and ℓ_0-decoding together yield a constructive procedure to recover all recoverable information about $\boldsymbol{x}^\star$: For a given $\boldsymbol{A}, q$ and $\boldsymbol{y}$, first compute $\hat{\boldsymbol{x}}$ using the ℓ_0-decoder, then compute $\boldsymbol{U}$ using Algorithm 1, and return $\hat{\boldsymbol{x}} + \ker(\boldsymbol{U})$.

4. ALGORITHM

In this section, we introduce a computation scheme (Algorithm 1) that takes as input the matrix A and an upper bound q on the number of adversarial measurements, and returns the robust orthogonal projection matrix, defined in Theorem 1. We also study its complexity.

Algorithm 1 Computing the robust orthogonal projector

Require: $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, integer $q < m/2$

Output: Orthogonal projector $\boldsymbol{U}$ onto the robust subspace $\bigcap_{T\subseteq[m],\,|T|=m-2q} \operatorname{rowspan}(\boldsymbol{A}_T)$

1: $\mathscr{T} \leftarrow \{T \subseteq [m] : |T| = m - 2q\}$, $\boldsymbol{C} \leftarrow 0_{n\times n}$
2: **for** $T \in \mathscr{T}$ **do**
3: $\quad\boldsymbol{A}_T \leftarrow$ submatrix of $\boldsymbol{A}$ with rows in T
4: $\quad\boldsymbol{B}_T \leftarrow$ orthonormal basis of $\ker(\boldsymbol{A}_T)$ via $\operatorname{svd}(\boldsymbol{A}_T)$
5: $\quad\boldsymbol{C} \leftarrow \boldsymbol{C} + \boldsymbol{B}_T\boldsymbol{B}_T^\top$
6: **end for**
7: $[\boldsymbol{Q}, \boldsymbol{\Lambda}] \leftarrow \operatorname{eig}(\boldsymbol{C})$
8: $\boldsymbol{Z} \leftarrow$ columns of $\boldsymbol{Q}$ with eigenvalue $\lambda = 0$
9: **return** $\boldsymbol{U} \leftarrow \boldsymbol{Z}\boldsymbol{Z}^\top$

Claim 1. *Algorithm 1 outputs the orthogonal projector onto the robust subspace $\bigcap_{T\subseteq[m],\,|T|=m-2q} \operatorname{rowspan}(\boldsymbol{A}_T)$.*

Proof. Let $\mathcal{R} = \bigcap_{T\subseteq[m],\,|T|=m-2q} \operatorname{rowspan}(\boldsymbol{A}_T)$. For each T, let $\boldsymbol{P}_T := \boldsymbol{B}_T\boldsymbol{B}_T^\top$ be the orthogonal projector onto $\ker(\boldsymbol{A}_T)$, so $\operatorname{im}(\boldsymbol{P}_T) = \ker(\boldsymbol{A}_T)$. Then $\operatorname{im}(\boldsymbol{C}) = \operatorname{im}(\sum_T \boldsymbol{P}_T) = \sum_T \operatorname{im}(\boldsymbol{P}_T) = \sum_T \ker(\boldsymbol{A}_T)$, since in finite dimensions the image of a sum equals the sum of images. Hence $\ker(\boldsymbol{C}) = (\operatorname{im}(\boldsymbol{C}))^\perp = (\sum_T \ker(\boldsymbol{A}_T))^\perp$. By the orthogonal-complement identity $(\sum_i V_i)^\perp = \cap_i V_i^\perp$ [15, Ex. 5, §12], and by the fundamental theorem of linear algebra $\operatorname{rowspan}(\boldsymbol{A}_T) = (\ker(\boldsymbol{A}_T))^\perp$ [14] , we have $\ker(\boldsymbol{C}) = \cap_T(\ker(\boldsymbol{A}_T))^\perp = \cap_T \operatorname{rowspan}(\boldsymbol{A}_T) = \mathcal{R}$. Therefore the zero-eigenvectors of $\boldsymbol{C}$ span $\mathcal{R}$, and $\boldsymbol{U} = \boldsymbol{Z}\boldsymbol{Z}^\top$ is the orthogonal projector onto $\mathcal{R}$. □

Algorithm 1 is inherently combinatorial, iterating over all subsets $T \subseteq [m]$ of size $m - 2q$, of which there are $\binom{m}{2q}$, growing exponentially with m in the worst case. Computing the SVD for each $\boldsymbol{A}_T \in \mathbb{R}^{(m-2q)\times n}$ costs $O((m-2q)n^2)$, and forming $\boldsymbol{C}$ and its eigen-decomposition adds $O(n^3)$. Hence, the overall complexity is dominated by $\binom{m}{2q}$, making the algorithm exponential in m. This combinatorial structure is tied to the NP-hardness of recovering signals under arbitrary q-sparse corruption, analogous to the subset-sum problem (see §I.B in [7]). While not polynomial-time, the algorithm is conceptually valuable as an exact deterministic procedure to compute the robust orthogonal projector, achieving the set inclusion-wise minimality guarantee established in Theorem 1.

5. NUMERICAL ILLUSTRATIONS

We illustrate our theoretical results with a concrete example and experimental simulations. Consider a network of five links, where each measurement corresponds to a path through the network and records the sum of the link values along that path (e.g., total delay or packet loss). The goal is to robustly recover key link properties (e.g., average delay) despite a small number of adversarially corrupted measurements. This setup exemplifies a broad class of systems in which only a limited number of linear measurements are available. Recovery in such systems is vulnerable to even a few corrupted measurements, necessitating methods that are robust to sparse adversarial corruptions. A natural instance of this class is *network tomography* [17], where link properties are inferred from path-level measurements.

Example 2. Consider recovering $\boldsymbol{x}^\star$ from $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star + \boldsymbol{e}$ in a network tomography setup with

$$\boldsymbol{A} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix},$$

where columns correspond to the five network links and rows correspond to measured paths. Here, links 1–3 form an upstream group and links 4–5 form a downstream group, reflecting constraints, e.g., from privacy-oriented networks. We assume that there is only one adversary ($q = 1$). Prior works demand a tall $\boldsymbol{A}$ (e.g., Proposition 2 in [16]), unavailable here, rendering full recovery of $\boldsymbol{x}^\star$ impossible. Our results quantify the precise recovery tradeoff when these conditions fail. Theorem 1 shows $\boldsymbol{U}\boldsymbol{x}^\star$ is the best robust projection recoverable. To analytically calculate $\boldsymbol{U}$, first, we notice that $\mathrm{rowspan}(\boldsymbol{A}) = \mathrm{span}(\boldsymbol{u}, \boldsymbol{v})$, where $\boldsymbol{u} = [1, 1, 1, 0, 0]^\top$, $\boldsymbol{v} = [0, 0, 0, 1, 1]^\top$, and deleting any $2q = 2$ rows preserves this span. With $\boldsymbol{B} = [\boldsymbol{u}, \boldsymbol{v}]$, we compute

$$\boldsymbol{U} = \boldsymbol{B}(\boldsymbol{B}^\top\boldsymbol{B})^{-1}\boldsymbol{B}^\top = \tfrac{1}{3}\boldsymbol{u}\boldsymbol{u}^\top + \tfrac{1}{2}\boldsymbol{v}\boldsymbol{v}^\top,$$

which explicitly equals

$$\boldsymbol{U} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

Thus, with one adversarial measurement, Theorem 1 gives the robustly recoverable quantities as the block averages $\frac{1}{3}\sum_{i=1}^{3} x_i^\star$ and $\frac{1}{2}\sum_{i=4}^{5} x_i^\star$.

m	n	q	Time (ms)
8	8	1	0.39 ± 0.056
8	8	3	0.203 ± 0.025
8	16	1	0.476 ± 0.098
8	16	3	0.232 ± 0.033
8	32	1	0.564 ± 0.054
8	32	3	0.357 ± 0.058
16	8	1	2.061 ± 0.097
16	8	3	127.777 ± 2.549
16	8	7	0.783 ± 0.063
16	16	1	3.649 ± 0.163
16	16	3	189.256 ± 3.391
16	16	7	1 ± 0.281
16	32	1	5.131 ± 0.289
16	32	3	234.563 ± 5.302
16	32	7	1.437 ± 0.181

Table 1. Runtime in milliseconds (mean $\pm$ std) over 10 runs.

We validate the output of a MATLAB implementation of Algorithm 1 on Example 2, and confirm that the algorithm matches our hand-calculated result. We then do a runtime analysis[1] of this MATLAB implementation on random normal matrices, where we sweep across $m \in \{8, 16\}$, $n \in \{8, 16, 32\}$, and $q \in \{1, 3, 7\}$. The results are reported in Table 1. We note that the runtime peaks at $q \approx m/4$, which the exponential $\binom{m}{2q}$ theoretical complexity also suggests. Our simulation results also confirm practical viability of our algorithm in small-scale systems.

6. CONCLUSION

We characterize the maximal recoverable information about an unknown signal $\boldsymbol{x}^\star$ from adversarially q-sparse corrupted measurements $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^\star + \boldsymbol{e}$. This information corresponds to the inclusion-wise minimal solution set of signals, precisely $\boldsymbol{x}^\star + \ker(\boldsymbol{U})$, where $\boldsymbol{U}$ is the orthogonal projector onto the intersection of rowspaces of all submatrices of $\boldsymbol{A}$ formed by deleting any $2q$ rows. Combining $\boldsymbol{U}$ with the ℓ_0-decoder enables a constructive recovery procedure for this set. We provide an algorithm to compute $\boldsymbol{U}$, leaving randomized sampling-based variants (e.g., [18]), and extensions of our results to random signals (e.g., [19]) for future work. Our results advance beyond prior full-recovery methods reliant on structural assumptions, delivering a universal, exact, and assumption-free characterization of adversarial robustness while laying the groundwork for scalable ℓ_1-based approximations.

[1] All simulations were run on a machine with 11th-Gen Intel Core i7-1185G7 (3.00 GHz), 32.0 GB RAM (4267 MT/s), Intel Iris Xe integrated graphics, and approximately 954 GB storage. Reported runtimes are wall-clock times.

7. REFERENCES

[1] Peter J Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[2] Peter J Rousseeuw, "Least median of squares regression," *Journal of the American statistical association*, vol. 79, no. 388, pp. 871–880, 1984.

[3] Yeshwanth Cherapanamjeri, Sandeep Silwal, David P Woodruff, Fred Zhang, Qiuyi Zhang, and Samson Zhou, "Robust algorithms on adaptive inputs from bounded adversaries," in *International Conference on Learning Representations*. International Conference on Learning Representations, 2023.

[4] Saurabh Amin, Alvaro A Cárdenas, and S Shankar Sastry, "Safe and secure networked control systems under denial-of-service attacks," in *International workshop on hybrid systems: computation and control*. Springer, 2009, pp. 31–45.

[5] Yao Liu, Peng Ning, and Michael K Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 1, pp. 1–33, 2011.

[6] Yuanxin Li, Yue Sun, and Yuejie Chi, "Low-rank positive semidefinite matrix recovery from corrupted rank-one measurements," *IEEE Transactions on Signal Processing*, vol. 65, no. 2, pp. 397–408, 2016.

[7] Emmanuel J Candes and Terence Tao, "Decoding by linear programming," *IEEE transactions on information theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[8] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore, "Compressed sensing and best k-term approximation," *Journal of the American mathematical society*, vol. 22, no. 1, pp. 211–231, 2009.

[9] Hamza Fawzi, Paulo Tabuada, and Suhas Diggavi, "Secure estimation and control for cyber-physical systems under adversarial attacks," *IEEE Transactions on Automatic control*, vol. 59, no. 6, pp. 1454–1467, 2014.

[10] David L Donoho and Michael Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization," *Proceedings of the National Academy of Sciences*, vol. 100, no. 5, pp. 2197–2202, 2003.

[11] Andreas M Tillmann and Marc E Pfetsch, "The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing," *IEEE Transactions on Information Theory*, vol. 60, no. 2, pp. 1248–1259, 2013.

[12] Afonso S Bandeira, Edgar Dobriban, Dustin G Mixon, and William F Sawin, "Certifying the restricted isometry property is hard," *IEEE transactions on information theory*, vol. 59, no. 6, pp. 3448–3450, 2013.

[13] Jonathan Weed, "Approximately certifying the restricted isometry property is hard," *IEEE Transactions on Information Theory*, vol. 64, no. 8, pp. 5488–5497, 2017.

[14] Gilbert Strang, "The fundamental theorem of linear algebra," *The American Mathematical Monthly*, vol. 100, no. 9, pp. 848–855, 1993.

[15] Paul R. Halmos, *Finite-Dimensional Vector Spaces*, Undergraduate Texts in Mathematics. Springer New York, NY, 1958.

[16] Hamza Fawzi, Paulo Tabuada, and Suhas Diggavi, "Secure state-estimation for dynamical systems under active adversaries," in *2011 49th annual allerton conference on communication, control, and computing (allerton)*. IEEE, 2011, pp. 337–344.

[17] Liang Ma, Ting He, Kin K Leung, Don Towsley, and Ananthram Swami, "Efficient identification of additive link metrics via network tomography," in *2013 IEEE 33rd International Conference on Distributed Computing Systems*. IEEE, 2013, pp. 581–590.

[18] Moritz Hardt and Ankur Moitra, "Algorithms and hardness for robust subspace recovery," in *Conference on Learning Theory*. PMLR, 2013, pp. 354–375.

[19] Swetha Ganesh, Alexandre Reiffers-Masson, and Gugan Thoppe, "Online learning with adversaries: A differential-inclusion analysis," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 1288–1293.