

Dimensionality Reduction of Affine Variational Inequalities Using Random Projections

Bharat Prabhakar · Ankur A. Kulkarni

Received: date / Accepted: date

Abstract We present a method for dimensionality reduction of an affine variational inequality (AVI) defined over a compact feasible region. Centered around the Johnson Lindenstrauss lemma [16], our method is a randomized algorithm that produces with high probability an approximate solution for the given AVI by solving a lower-dimensional AVI. The algorithm allows the lower dimension to be chosen based on the quality of approximation desired. The algorithm can also be used as a subroutine in an exact algorithm for generating an initial point close to the solution. The lower-dimensional AVI is obtained by appropriately projecting the original AVI on a randomly chosen subspace. The lower-dimensional AVI is solved using standard solvers and from this solution an approximate solution to the original AVI is recovered through an inexpensive process. Our numerical experiments corroborate the theoretical results and validate that the algorithm provides a good approximation at low dimensions and substantial savings in time for an exact solution.

Keywords Dimensionality reduction · Random projection · Affine variational inequality · Johnson Lindenstrauss lemma

1 Introduction

Technological advancements have enabled the collection and storage of a tremendously large amount of data. In parallel, the Internet has been changing the manner in which systems interact. As a result of this, an overwhelming amount of information is being generated and stored. A report by Harvard Magazine claims that, “the total data accumulation of just the past two years – a zettabyte – dwarfs the prior record of human civilization” [29], aptly justifying the name Big Data.

This paper concerns a challenge thrown up by high-dimensional problems in optimization and control that have arisen due to the growing prominence of such ‘big’ or very large data sets [26].

Bharat Prabhakar
Department of Electrical Engineering, Indian Institute of Technology Bombay
E-mail: prabhakar.bharat@iitb.ac.in

Ankur A. Kulkarni
Systems and Control Engineering, Indian Institute of Technology Bombay
E-mail: kulkarni.ankur@iitb.ac.in

Exact algorithms for such problems can be computationally burdensome even if the algorithm has polynomial complexity. Whereas in contexts such as online optimization [33] a conceivable requirement could be not of the exact solution, but rather of a quick approximation, in the spirit of the Latin proverb *bis das si cito das*¹. We are motivated by this specific need where speed is of essence and accuracy can be sacrificed to some extent, if it means large savings in time.

This paper focuses on *affine variational inequalities*. Variational inequalities are a versatile class of problems that generalize convex optimization [10], saddle-point problems, Nash games and generalized Nash games [9, 21, 19], amongst others [10]. We consider affine variational inequalities with compact feasible regions. They are an important subclass, which include, e.g., constrained convex quadratic programming with compact feasible regions. Such quadratic programs are useful problems in their own right and are subproblems for the widely used *sequential quadratic programming* algorithm for nonlinear programs [13, 26]. Affine variational inequalities also capture commonly used models for oligopolistic competition, such as Nash-Cournot games [10].

We present a dimensionality reduction technique for solving high-dimensional affine variational inequalities *approximately*. The method is *probabilistic* in the sense that one can only guarantee that it works with “high probability”. However, in exchange for this, we obtain a substantial saving in time. For a polyhedral set $K \subseteq \mathbb{R}^n$, a matrix $M \in \mathbb{R}^{n \times n}$ and a vector $q \in \mathbb{R}^n$, an affine variational inequality (AVI) $\text{AVI}(K, M, q)$ is the following problem,

$\text{AVI}(K, M, q) \quad \text{Find } x \in K \text{ such that } (y - x)^\top (Mx + q) \geq 0 \quad \forall y \in K.$

Given a high-dimensional (deterministic) AVI, the method derives a *random* AVI from it which is low-dimensional. The lower-dimensional AVI is solved using standard solvers. Using the solution of the lower-dimensional AVI, a candidate solution of the high-dimensional AVI is generated through an inexpensive process. A probabilistic guarantee is obtained on the event that the error satisfies a bound.

The lower-dimensional AVI is obtained by *projecting* the high-dimensional problem on a subspace chosen uniformly at random. We implement this projection by multiplying by a suitably constructed random matrix. By the celebrated *Johnson Lindenstrauss Lemma* (JL lemma) [16] we get that if a set of m points in a high dimensional space, are projected this way to a k -dimensional subspace, then the probability that pairwise distances are at most ϵ -distorted, concentrates. This probability can be made to approach unity by appropriately choosing k . E.g., for this probability to be $1 - \delta$, we get $k = O\left(\frac{\ln(m/\delta)}{\epsilon^2}\right)$. When applied to our setting, approximate distance preservation also allows for approximate preservation of inner products, which translates to an approximate solution of the AVI.

Our main result is that with high probability the (deterministic) preimage (under the projection operation) of the solution of the lower-dimensional AVI approximately solves the given high-dimensional AVI. We recover a random approximation to this deterministic preimage by solving a linear program followed by a norm minimization quadratic program. Thanks to a remarkable result of Candes and Tao [4] we get that the recovered solution approximates the required deterministic preimage with high probability.

This framework also yields the following exact algorithm (with probabilistic guarantees): using the above technique one generates a point that is close to the true solution and this point is supplied to a standard solver as an initial point to obtain an exact solution. It is plausible that this method

¹ Twice you give, if you give quickly.

would improve the run-time for solvers that benefit from the “local” nature of the initial point. We have found this to be the case for the PATH solver [7].

We emphasize that our algorithm does not assume any structure on the AVI such as monotonicity (the only assumption on the AVI is that the set K is compact). It seems plausible that further assumptions on the matrix M or on the set K may improve the theoretical results.

Our numerical computations support the theoretical results. In particular, the exact algorithm (obtained by supplying the approximate solution as initial point) appears promising; in the examples we tried, considerable savings were obtained. For the theoretical results to hold we require the lower-dimensional problem to be of size $O(\ln \eta)$, where η is the number of extreme points of K . If K described by m inequality constraints, then for fixed m , we have $\ln \eta = O(\ln n)$ and in general $\ln \eta$ is at most $O(n)$, where n is dimension of the ambient space of K . Although in the worst case $O(\ln \eta)$ may not be significantly smaller than n , in practice we have found that our algorithm performs well even for small values of the lower dimension.

Conceptually speaking, this work exploits a delicate link between convex analysis and metric embeddings. The JL lemma may be viewed as a metric embedding result [14], with no obvious convex analytic properties. However in a Euclidean space, a metric embedding also implies the ϵ -preservation of inner products, which under convexity allows optimality to be ϵ -preserved.

This paper is organized as follows. Following the introduction, we present some background on the subject and define the desiderata of our algorithm. Section 2 formally introduces concepts pertaining to random projections required for our main results. The algorithm is introduced in Section 3. The proof of correctness is encompassed in Section 4. We discuss some aspects pertaining to the algorithm, including the lower dimension, in Section 5. Section 6 contains numerical results and we conclude in Section 7.

1.1 Background

For a closed convex set $K \subseteq \mathbb{R}^n$ and a continuous function $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$, a variational inequality (VI) $\text{VI}(K, F)$ is the following problem [10],

$\text{VI}(K, F) \quad \text{Find } x \in K \text{ such that } (y - x)^\top F(x) \geq 0 \quad \forall y \in K.$

Solving a VI amounts to ensuring an *angle condition*: the solution is a point $x \in K$ such that $F(x)$ makes an acute angle with all directions ‘ $y - x$ ’ as y ranges over K . The solution of a convex optimization can also be written in this form and is a special case of the VI. The Nash equilibrium of a game is a simultaneous solution of several convex optimization problems and can also be captured by a VI [10]. Recent results [9, 21, 19] have shown that certain equilibria of *generalized* Nash games can also be captured by VIs. Besides these applications, VIs also generalize general equilibrium models, frictional contact problems and problems in finance, even as new applications continue to be considered. We refer the reader to [10] for more on this topic.

Solving $\text{VI}(K, F)$ is also equivalent to finding the zero of the function $\mathbf{F}_K^{\text{nat}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ [10],

$$\mathbf{F}_K^{\text{nat}}(x) = x - \Pi_K(x - F(x)), \tag{1}$$

called the *natural map* of the VI. $\mathbf{F}_K^{\text{nat}}$ is a continuous function and hence $\|\mathbf{F}_K^{\text{nat}}(x^*)\|$ quantifies the “quality” of an approximate solution x^* . Note that $\mathbf{F}_K^{\text{nat}}$ is nonlinear except in rare (possibly uninteresting) cases.

In practice, VIs are solved using a host of techniques, making use of the natural map as well as the *normal map* [10]. We are not explicitly concerned with algorithms for solving VIs, since our method employs an off-the-shelf solver for solving the lower-dimensional AVI. We refer the reader to [11] for more on algorithms and to [8, 7] for implementations.

Our work is essentially about efficiently solving VIs. There is a large body of work on solving large-dimensional optimization problems and variational inequalities with certain sparsity structure. These include, for instance, Benders' decomposition [3] and the host of applications it spawned to stochastic optimization problems [31, 20] and complementarity problems (see, e.g., [28] and references therein). In a somewhat similar direction lie the series of works on splitting methods (see, e.g., [24] and related works), and more recent decomposition methods, e.g., [23]. These lines of research exploit the structure of the problem to decompose the larger problem into smaller subproblems. The algorithms so developed are exact algorithms with deterministic guarantees. In contrast, our algorithm does not assume any sparsity, but it is an approximate algorithm which works under probabilistic guarantees. Further, as mentioned in the introduction, it may be used as a subroutine in an exact algorithm.

To the best of our knowledge our work is the first application of random projections for dimensionality reduction in VIs. However, it has been preceded by many random projection-based algorithms in the computer science community (see, e.g., the monograph by Vempala [32]). In the operations research and control community, there are two applications of random projections we are aware of. First [1], where a low-rank approximation is used to approximately find the zero of a linear equation, and second [18] where a similar approximation is used within a Newton method-based stochastic approximation. Our work differs from [1] in two fundamental ways. First, we seek a low-dimensional approximation (rather than a low-rank one), and second, solving a VI reduces to solving the nonlinear equation (1) whereas [1] critically relies on linearity.

1.2 Problem definition

We now formally define the problem we aim to solve. The objective of this paper is to solve $\text{AVI}(K, M, q)$, i.e., where $K \subseteq \mathbb{R}^n$ is a compact polyhedron (a polytope), $M \in \mathbb{R}^{n \times n}$ and $q \in \mathbb{R}^n$. The case we are interested in is where n is large. We seek an approximation algorithm that satisfies the following requirements

1. The most expensive step in the algorithm must involve solving a lower dimensional problem, i.e., the algorithm must operate in a lower-dimensional space.
2. The lower dimensional problem should also be an affine variational inequality.
3. The algorithm may be approximate, i.e., the candidate solution generated by the algorithm need not solve the problem exactly, but it should be a good approximation.
4. The guarantee for the algorithm need not be deterministic, i.e., probabilistic guarantees on the solution would suffice.

With this specified, we now proceed with the main contents of the paper, beginning with an overview of random projections in the following section.

2 Random Projections

In this section we review some results from the theory of random projections. Random projection is a particular case of an embedding of one metric space into another, and is as such is a part of

a deeper mathematical study [14]. We limit our survey here to operational aspects and to results relevant to our algorithm.

2.1 How to randomly project

Random projection involves the projection of vectors lying in a higher dimensional space to a *randomly* chosen lower-dimensional subspace. Note that this projection need not be Euclidean, i.e., the subspace need not be aligned with the basis vectors from the original space. A vector is projected by multiplying the vector by a suitable random matrix; the choice of this matrix specifies the type of randomness introduced in the projection.

Several different methods of constructing the random projection matrix have been studied in the literature. In this paper, we project the vectors on a *uniformly* random k -dimensional subspace. Such a subspace can be realized by choosing a uniformly random orthonormal matrix [32]. Below we show how such a matrix can be constructed.

2.1.1 Constructing a uniformly random orthonormal matrix

An $n \times k$ -dimensional real valued orthonormal matrix R is *uniformly random* if R is uniformly distributed over the manifold, called Stiefel manifold [30], of real $n \times k$ matrices such that $R^\top R = I$. We construct our $n \times k$ -dimensional random projection matrix R as follows.

- R 1.** Construct a matrix R_1 with each entry chosen independently from the distribution $N(0, \frac{1}{k})$.
- R 2.** Orthonormalize the columns of R_1 using Gram-Schmidt process (QR-decomposition) and form the required matrix R using these resultant vectors as columns.

We first observe that the matrix R_1 above is full rank almost surely. This is formalized in the following lemma.

Lemma 1 *The $n \times k$ -dimensional random matrix R_1 (where $n > k$), obtained in the step **R 1** in the construction described above, has a rank equal to k with probability 1.*

Proof Consider the matrix Y formed by normalizing each column of R_1 to a unit vector (each column is not a 0-vector almost surely). It suffices to show that Y has rank k . Let the columns of Y be denoted by $Y^k = (Y_1, \dots, Y_k)$. Recall that each vector $Y_i, 1 \leq i \leq k$ is a uniformly distributed point on the unit sphere S^{n-1} [17]. We prove the claim by induction. The first vector Y_1 is linearly independent with probability 1, as it is not a zero vector almost surely. Assume that $1 < r \leq k$ and that the first $r-1$ vectors are linearly independent with probability 1. Then with probability 1 these $r-1$ vectors span a subspace V of dimension $r-1$, which intersects S^{n-1} in an $(r-2)$ -dimensional “subsphere” S_V^{r-2} . This subsphere forms a measure zero set under the uniform probability measure on S^{n-1} . Therefore the probability that Y_r lies in this subsphere is zero. It follows that with probability 1 the vectors Y_1, \dots, Y_{r-1}, Y_r are linearly independent. Hence, by induction, $\text{rank}(Y) = k$.

From Lemma 1 it follows that the matrix R produced in Step **R 2** is also of full column rank. To show that R is indeed uniformly distributed on the Stiefel manifold, we invoke the “real” counterpart of Lemma 2 from [30].

Lemma 2 ([30]) *Let H be an $k \times n$ real valued standard Gaussian matrix with $n \geq k$. Denote its QR-decomposition by $H = \hat{Q}\hat{R}$. The upper triangular matrix \hat{R} is independent of \hat{Q} , which is uniformly distributed over the manifold of $k \times n$ matrices such that $\hat{Q}\hat{Q}^\top = I$.*

This completes the construction of a matrix R that is uniformly distributed on the Steifel manifold and has full column rank.

2.1.2 Projecting vectors and matrices

Consider a column vector $x \in \mathbb{R}^n$ (throughout this paper a vector is automatically to be assumed as a column vector), then the projection of x is given by,

$$y = \sqrt{\frac{n}{k}} R^\top x, \quad (2)$$

The constant $\sqrt{\frac{n}{k}}$ is multiplied to ensure that the expected length of y remains equal to that of x (this property is needed by the Johnson Lindenstrauss lemma as we shall see ahead). Note that the value of constant being multiplied may vary depending upon the construction of the random projection matrix.

For an arbitrary random vector y and an arbitrary random matrix R , a vector x such that $y = \sqrt{\frac{n}{k}} R^\top x$ would, in general, be sample-path dependent (i.e., random). However if y is indeed a projection of a *deterministic* vector x , we call such an x its *deterministic preimage*.

Definition 1 If a vector y is a projection of some deterministic vector x as in (2), its deterministic preimage is defined to be any deterministic vector x_o , such that $\sqrt{\frac{n}{k}} R^\top x_o = y$.

The columns of a matrix $Z \in \mathbb{R}^{n \times d}$ can be thought of as a collection of d n -dimensional vectors. Thus, the projection of Z is given by,

$$Y = \sqrt{\frac{n}{k}} R^\top Z,$$

where $Y \in \mathbb{R}^{k \times d}$.

2.2 The Johnson Lindenstrauss lemma

The Johnson Lindenstrauss lemma [16] is a landmark result that shows that for any finite set of points there exists a mapping such that the distance between any pair of points is approximately equal to the distance between their images under the mapping. A key improvement [12] obtained later showed that if any finite set of points are projected *randomly* by multiplication with a random orthonormal matrix, their pairwise distances are approximately preserved in the above sense. While the original JL lemma only provides the existence of a distance preserving mapping, at the expense of a probabilistic guarantee, the result from [12] provides a *construction* of this mapping.

For the purpose of our paper, we do not need this exact result but only its precursor below which concerns approximately preserving the norm of a single vector; this latter result can be found in [6]. Since the lemma below differs from the exact result in [6], we provide a sketch of the proof.

Lemma 3 *Let $R \in \mathbb{R}^{n \times k}$ be a random matrix as constructed in Section 2.1.1, and let $f(u) = \sqrt{\frac{n}{k}} R^\top u$ for $u \in \mathbb{R}^n$. Then for any $0 < \epsilon < 1$,*

$$P((1 - \epsilon)\|u\|^2 \leq \|f(u)\|^2 \leq (1 + \epsilon)\|u\|^2) \geq 1 - 2e^{-(\epsilon^2/2 - \epsilon^3/3)k/2}$$

Proof Let $u \in \mathbb{R}^n$ and without loss of generality assume $\|u\| = 1$. The vector $\sqrt{\frac{n}{k}} R^\top u$ is the projection of u on a k -dimensional subspace chosen uniformly at random. As observed in [6] the distribution of $R^\top u$ is the same as the distribution of the projection of a unit vector chosen uniformly at random to a fixed subspace. Thus, let X_1, \dots, X_n be n independent Gaussian $N(0, 1)$ random variables and define $Y = \frac{1}{\|X\|}(X_1, \dots, X_n)$. As argued in Lemma 1, Y is a uniformly random vector lying on the unit sphere S^{n-1} . Fix the projection subspace to be one spanned by the first k coordinates and let $Z \in \mathbb{R}^k$ be the projection of Y on the first k coordinates. Define $L := \|Z\|^2$. Clearly, $\mathbb{E}[L] = k/n$. From [6] we get,

$$P(L \leq (1 - \epsilon)\mathbb{E}[L]) \leq e^{-\frac{k\epsilon^2}{4}} \quad \text{and} \quad P(L \geq (1 + \epsilon)\mathbb{E}[L]) \leq e^{-\frac{k(\epsilon^2/2 - \epsilon^3/3)}{2}}.$$

Now since $R^\top u$ has the same distribution as Z , $\|f(u)\|^2$ has the same distribution as $\frac{L}{\mathbb{E}[L]}$, from which the result follows.

It is important to note that unlike in the deterministic result of Johnson and Lindenstrauss, the projection *does not* depend on the vector being projected (the projection only has to be uniformly random). Further, since random projection here is a linear operator, the above result can easily be converted into a *distance preservation* result by applying it to a vector corresponding to the difference between two different vectors and then preserving its norm. By applying this projection to a set of vectors and then taking a union bound leads us to the complete version of JL lemma (see e.g., [32]).

Finally note that, the probability in Lemma 3 is for the event that the norm of the projected vector lies within a factor of $(1 \pm \epsilon)$ of the norm of original vector. Whereas in *expectation*, the value of both these norms is the same, i.e., $\mathbb{E}[\|f(u)\|^2] = \|u\|^2$. This underlines the nature of the JL lemma as a concentration phenomenon.

2.3 Preserving inner products via JL lemma

We now present the following lemma claiming that the inner products are approximately preserved under random projection. It follows from the norm preserving JL lemma.

Lemma 4 *Let $R \in \mathbb{R}^{n \times k}$ be a random matrix as constructed in Section 2.1.1. Define a mapping $f(u) = \sqrt{\frac{n}{k}} R^\top u$ for $u \in \mathbb{R}^n$. Then for any two vectors $u, v \in \mathbb{R}^n$, and $0 < \epsilon < 1$,*

$$P(|u^\top v - f(u)^\top f(v)| \leq \epsilon \|u\| \|v\|) \geq 1 - 4e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}$$

Proof Consider two vectors $u' + v'$ and $u' - v'$, such that $u' = \frac{u}{\|u\|}$ and $v' = \frac{v}{\|v\|}$. We try to preserve the norms of both these vectors *simultaneously*. Let A and B denote the events

$$\begin{aligned} A &:= \{(1 - \epsilon)\|u' + v'\|^2 \leq \|f(u' + v')\|^2 \leq (1 + \epsilon)\|u' + v'\|^2\} \\ B &:= \{(1 - \epsilon)\|u' - v'\|^2 \leq \|f(u' - v')\|^2 \leq (1 + \epsilon)\|u' - v'\|^2\} \end{aligned}$$

We need to find a lower bound on $P(A \cap B)$, which is equivalent to finding an upper bound on $P(A^c \cup B^c)$. We have, $P(A \cap B) = 1 - P(A^c \cup B^c) \geq 1 - P(A^c) - P(B^c)$, where the last inequality follows from the union bound on $P(A^c \cup B^c)$. From Lemma 3, we have $P(A) \geq 1 - 2e^{-(\epsilon^2 - \epsilon^3)k/4}$ and $P(B) \geq 1 - 2e^{-(\epsilon^2 - \epsilon^3)k/4}$. Consequently,

$$P(A \cap B) \geq 1 - 4e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}. \quad (3)$$

Now since, $4f(u')^\top f(v') = \|f(u' + v')\|^2 - \|f(u' - v')\|^2$ we have from the definition of A, B , that under the event $A \cap B$,

$$\begin{aligned} 4f(u')^\top f(v') &\geq (1 - \epsilon)\|u' + v'\|^2 - (1 + \epsilon)\|u' - v'\|^2 \\ &= 4(u')^\top v' - 2\epsilon(\|u'\|^2 + \|v'\|^2). \end{aligned}$$

Since u' and v' are unit vectors, we get $f(u')^\top f(v') \geq (u')^\top v' - \epsilon$. Similarly, under the event $A \cap B$,

$$4f(u')^\top f(v') \leq 4(u')^\top v' + 2\epsilon(\|u'\|^2 + \|v'\|^2),$$

whereby, $f(u')^\top f(v') \leq (u')^\top v' + \epsilon$. Since f is linear,

$$\{|u^\top v - f(u)^\top f(v)| \leq \epsilon\|u\|\|v\|\} \supseteq A \cap B.$$

Now using (3), we get the result.

The above lemma talks about preserving only one inner product. The following lemma generalizes the result and ensures the preservation of a finitely many inner products simultaneously.

Lemma 5 *For each $i = 1, \dots, m$, let u_i, v_i be vectors \mathbb{R}^n . Also, let $R \in \mathbb{R}^{n \times k}$ be a random matrix as constructed in Section 2.1.1, and define a mapping $f(x) = \sqrt{\frac{n}{k}}R^\top x$ for $x \in \mathbb{R}^n$. Then for $0 < \epsilon < 1$,*

$$P(|u_i^\top v_i - f(u_i)^\top f(v_i)| \leq \epsilon\|u_i\|\|v_i\| \quad \forall i = 1, \dots, m) \geq 1 - 4me^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}$$

Proof Let A_i denote the event $\{|u_i^\top v_i - f(u_i)^\top f(v_i)| \leq \epsilon\|u_i\|\|v_i\|\}$ for $i = 1, \dots, m$. Thus, from Lemma 2.4 we get for all $i = 1, \dots, m$,

$$P(A_i^c) \leq 4e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2},$$

Similar to the analysis in the proof of Lemma 4, by the union bound,

$$\begin{aligned} P(A_1 \cap A_2 \cdots \cap A_m) &= 1 - P(A_1^c \cup A_2^c \cdots \cup A_m^c), \\ &\geq 1 - [P(A_1^c) + P(A_2^c) + \cdots + P(A_m^c)], \\ &\geq 1 - 4me^{-(\epsilon^2/2 - \epsilon^3/3)k/2}, \end{aligned}$$

as required.

This completes the preliminaries pertaining to random projections. In the following section we present our algorithm.

3 Algorithm

Algorithm A below is our dimensionality reduction algorithm for solving a typically high-dimensional affine variational inequality with a compact feasible region. It constructs a lower-dimensional AVI and using the solution obtained to this lower dimensional problem, generates an approximate solution to our original problem. Our algorithm probabilistically guarantees that the solution vector it generates, solves the original problem approximately.

Algorithm A

1. **Input:** AVI(K, M, q), i.e., $M \in \mathbb{R}^{n \times n}, q \in \mathbb{R}^n, K = \{x | Ax \leq b\}$ where $A \in \mathbb{R}^{m \times n}, b \in \mathbb{R}^m$.
2. Fix the error parameter $\epsilon \in (0, 1)$ and the success confidence parameter $\delta \in (0, 1]$. Pick the lower dimension value k such that $k \geq \frac{2 \ln(\frac{4\eta}{\delta})}{(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})}$, where $\eta = |\text{ext}(K)|$ and $\text{ext}(K)$ denotes the set of all extreme points of K .
3. Construct an $n \times k$ dimensional random matrix R as described in Section 2.1.1.
4. Construct the corresponding lower-dimensional AVI($\tilde{K}, \tilde{M}, \tilde{q}$). Define $\tilde{q} = \sqrt{\frac{n}{k}} R^\top q$ and $\tilde{M} = R^\top M R$, and \tilde{K} to be the projection of the polytope K , i.e., $\tilde{K} = \{\tilde{x} \in \mathbb{R}^k | \tilde{x} = \sqrt{\frac{n}{k}} R^\top x, x \in K\}$.
5. Solve AVI($\tilde{K}, \tilde{M}, \tilde{q}$) to obtain a lower-dimensional solution \tilde{x} .
6. Solve the following ℓ_1 norm minimization problem (can be converted into a linear program) to obtain a vector x^* :

$$x^* \in \arg \min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{subject to} \quad \sqrt{\frac{n}{k}} R^\top x = \tilde{x}$$

7. Project x^* on K to obtain the final random approximate solution $x^\#$, i.e.,

$$x^\# = \Pi_K(x^*) \triangleq \arg \min_{x \in K} \|x^* - x\|^2$$

8. **Output:** $x^\#$.

Our main result is as follows.

Theorem 1 *Let $\epsilon \in (0, 1), \delta \in (0, 1]$. Let AVI(K, M, q) be an AVI where K is a compact polyhedron and consider Algorithm A. Suppose \tilde{x} is the lower-dimensional solution obtained in Step 5 and let $x_o \in K$ be a deterministic preimage of \tilde{x} . Then the following claims hold.*

1. *If the lower dimension k satisfies $k \geq \frac{2 \ln(\frac{4\eta}{\delta})}{(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})}$, then with probability strictly greater than $(1 - \delta)$, x_o solves AVI(K, M, q) approximately, i.e.,*

$$(y - x_o)^\top (q + Mx_o) \geq \hat{\epsilon}, \quad \forall y \in K,$$

where $\hat{\epsilon} = -\frac{(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})n}{2 \ln(\frac{4\eta}{\delta})} \|M\| B - \epsilon \cdot D \|q\| - \epsilon \cdot D \|M\| B$, where $\|M\|$ denotes the ℓ_2 induced norm of M , $D := \max_{x_1, x_2 \in K} \|x_1 - x_2\|$ is the diameter of K and $B := \max_{x \in K} \|x\|$.

2. Let $x^\#$ be the output generated by the algorithm A in Step 8, then with probability at least $1 - O(n^{-1/\alpha})$:

$$\|x_o - x^\#\|_2 \leq CB' \cdot (k/\ln(n))^{-1/2},$$

where $\alpha > 0$ is a sufficiently small number (less than an absolute constant) and C is a constant depending only on α and $B' = \max_{x \in K} \|x\|_1$.

There are two parts to the above result. The first part shows that the deterministic preimage x_o of the lower-dimensional solution \tilde{x} *approximately* solves the given problem $\text{AVI}(K, M, q)$ in the sense that the inner product $(y - x_o)^\top (q + Mx_o)$ as y ranges over K is at least $\hat{\epsilon}$. Here $\hat{\epsilon}$ is a negative number that can be made arbitrarily close to zero by choosing ϵ to be small enough; for an exact solution to the AVI we would require $\hat{\epsilon} \geq 0$. Notice though that this result in itself is not computationally useful since x_o cannot be computed from a single sample of \tilde{x} . Each run of the algorithm produces a particular sample path: any process of “inverting” the relation $\sqrt{\frac{n}{k}} R^\top x_o = \tilde{x}$ would in general yield an x_o that is sample-path dependent (and hence R -dependent). Steps 6-8 of the algorithm produce a random approximation $x^\#$ to (the deterministic approximate solution) x_o . The second part of the above theorem establishes that with high probability, these steps produce a good approximation.

In the following section, we prove Theorem 1 and establish the correctness of the algorithm.

4 Correctness of the Algorithm

There are two parts that need to be established to show the correctness of the algorithm. The first part showing the point x_o is an approximate solution will be proved using the JL lemma. The second part, showing that the recovered solution $x^\#$ approximates x_o will be showed using a result of Candes and Tao [4]. Before we proceed with this analysis, we note that the lower-dimensional problem is indeed an AVI and that it admits a solution.

The lower dimensional problem is a variational inequality, $\text{VI}(\tilde{K}, \tilde{F})$ where the mapping $\tilde{F}(\tilde{x}) \equiv \tilde{q} + \tilde{M}\tilde{x}$ is affine. \tilde{K} is the projection of the polytope K . To show that this VI is an AVI, it suffices to show that \tilde{K} is also a polytope. For any set S we denote the set of its extreme points by $\text{ext}(S)$.

Lemma 6 *Let $K \subseteq \mathbb{R}^n$, $M \in \mathbb{R}^{n \times n}$, $q \in \mathbb{R}^n$ be a polytope and suppose $\text{AVI}(K, M, q)$ is provided as an input to Algorithm A . Then the following are true,*

1. *The set \tilde{K} generated in Step 4 of Algorithm A is also a polytope.*
2. *$\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$ admits a solution.*

Proof To show part 1, it suffices to show that there is a finite set of points in \tilde{K} such that every point in \tilde{K} can be expressed as a convex combination of these points. To this end let $\text{ext}(K)$ denote the set of extreme points of K . Since K is a polytope, $\eta := |\text{ext}(K)| < \infty$. Then from Step 4 of Algorithm A , we have, $\forall \tilde{x} \in \tilde{K}$, $\exists x \in K$ such that $\tilde{x} = R^\top x$. Since K itself is a polytope, $\exists \alpha_1, \dots, \alpha_\eta \geq 0$, $\sum_i \alpha_i = 1$ such that $x = \sum_i \alpha_i x_i$, where x_1, \dots, x_η are the extreme points of K . This implies, $\tilde{x} = \sum_i \alpha_i \sqrt{\frac{n}{k}} R^\top x_i$. Thus, any vector $\tilde{x} \in \tilde{K}$ can be expressed as a convex combination of the points in the set $S = \{\sqrt{\frac{n}{k}} R^\top x \mid x \in \text{ext}(K)\}$. Since S is finite, \tilde{K} is a polytope. Consequently, \tilde{K} is also compact. Standard results [10] now show that $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$ admits a solution.

4.1 Problem reformulation

For a point to be a solution of a VI, by definition, (*uncountably*) *infinitely many* inequalities must simultaneously hold. However for an AVI on a polytope, a reduction to a finite set of inequalities is possible. This important reformulation is formalized in the following lemma.

Lemma 7 *Let $K \subseteq \mathbb{R}^n$ be a polytope and $q \in \mathbb{R}^n$, $M \in \mathbb{R}^{n \times n}$. A vector $x^* \in K$ is a solution of $\text{AVI}(K, M, q)$ if and only if:*

$$(x_e - x^*)^\top (q + Mx^*) \geq 0, \quad \forall x_e \in \text{ext}(K), \quad (4)$$

where $\text{ext}(K)$ denotes the set of all the extreme points of K .

Proof “ \implies ” Let $x^* \in K$ solve $\text{AVI}(K, q, M)$. Thus,

$$(y - x^*)^\top (q + Mx^*) \geq 0, \quad \forall y \in K.$$

Put $y = x_e$ where $x_e \in \text{ext}(K)$. Thus,

$$(x_e - x^*)^\top (q + Mx^*) \geq 0, \quad \forall x_e \in \text{ext}(K).$$

Hence, x^* satisfies (4).

“ \impliedby ” Let $\text{ext}(K) = x_1, \dots, x_\eta$ where $\eta = |\text{ext}(K)|$. Consider an arbitrary vector $y \in K$. Since K is a polytope, there exists $\{\alpha_i\}_{i=1}^\eta$, such that $\sum_{i=1}^\eta \alpha_i = 1$, $\alpha_i \geq 0, \forall i \in 1, \dots, \eta$, and $y = \sum_{i=1}^\eta \alpha_i x_i$. Since x^* solves (4), we get for all i ,

$$\alpha_i (x_i - x^*)^\top (q + Mx^*) \geq 0.$$

Summing all the inequalities over $i = 1, \dots, \eta$ and using that $\sum_{i=1}^\eta \alpha_i = 1$, we get that

$$(y - x^*)^\top (q + Mx^*) \geq 0.$$

Since $y \in K$ was an arbitrary vector, the above inequality is true for every $y \in K$. Hence, x^* also solves $\text{AVI}(K, M, q)$.

4.2 Error analysis

The lower dimensional problem $(\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q}))$ as constructed in the algorithm is the following.

$\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$	Find $\tilde{x} \in \tilde{K}$ such that $(y - \tilde{x})^\top (\tilde{M}\tilde{x} + \tilde{q}) \geq 0 \quad \forall y \in \tilde{K}$.
---	---

We have the following claim about the deterministic preimage of the solution of the problem $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$.

Lemma 8 *Let $\epsilon \in (0, 1)$ and suppose $\text{AVI}(K, M, q)$ is provided as input to Algorithm A. Let \tilde{x} be the solution of $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$ generated by Step 5 of Algorithm A and let x_o be its deterministic preimage. Then with probability greater than $1 - 4\eta e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}$, x_o satisfies*

$$(x_e - x_o)^\top (q + Mx_o) \geq \frac{n}{k} (x_e - x_o)^\top R R^\top (M - M') x_o - \epsilon \cdot \|x_e - x_o\| \|q + Mx_o\|$$

for every $x_e \in \text{ext}(K)$, where $M' = M R R^\top$ and $\eta = |\text{ext}(K)|$.

Proof Let \tilde{x} be a solution to $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$ and x_o be its deterministic preimage. This implies,

$$(\tilde{y} - \tilde{x})^\top (\tilde{q} + \tilde{M}\tilde{x}) \geq 0, \forall \tilde{y} \in \tilde{K}.$$

Let $M' = MRR^\top$, then it follows from definition of K that,

$$\sqrt{\frac{n}{k}}(R^\top(y - x_o))^\top \sqrt{\frac{n}{k}}R^\top(q + M'x_o) \geq 0, \quad \forall y \in K.$$

Consequently, for each $x_e \in \text{ext}(K)$,

$$\sqrt{\frac{n}{k}}(R^\top(x_e - x_o))^\top \sqrt{\frac{n}{k}}R^\top(q + M'x_o) \geq 0.$$

On adding $(\sqrt{\frac{n}{k}}(R^\top(x_e - x_o))^\top \sqrt{\frac{n}{k}}R^\top(M - M')x_o)$ on both sides we get that for any $x_e \in \text{ext}(K)$,

$$\begin{aligned} \sqrt{\frac{n}{k}}(R^\top(x_e - x_o))^\top \sqrt{\frac{n}{k}}R^\top(q + Mx_o) &\geq \sqrt{\frac{n}{k}}(R^\top(x_e - x_o))^\top \sqrt{\frac{n}{k}}R^\top(M - M')x_o \\ &= \frac{n}{k}(x_e - x_o)^\top RR^\top(M - M')x_o. \end{aligned} \quad (5)$$

Observe that the left hand side of (5) is the an inner product of the random projection of vectors $(x_e - x_o)$ and $q + Mx_o$. Hence by Lemma 5, with probability greater than $1 - 4\eta e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}$, x_o satisfies,

$$\begin{aligned} (x_e - x_o)^\top (q + Mx_o) \\ \geq \sqrt{\frac{n}{k}}(R^\top(x_e - x_o))^\top \sqrt{\frac{n}{k}}R^\top(q + Mx_o) - \epsilon \|x_e - x_o\| \|q + Mx_o\|, \quad \forall x_e \in \text{ext}(K). \end{aligned}$$

Now using (6), we get the result.

For $\text{AVI}(K, M, q)$ and R, x_o, ϵ as in Lemma 8 and a point $x_e \in \text{ext}(K)$, define the quantities,

$$\mu_1(x_e) := \frac{n}{k}(x_e - x_o)^\top RR^\top(M - M')x_o, \quad (7)$$

$$\mu_2(x_e) := -\epsilon \cdot \|x_e - x_o\| \cdot \|q + Mx_o\|. \quad (8)$$

Below we derive bounds on μ_1, μ_2 .

To this end, recall that any matrix $A \in \mathbb{R}^{m \times n}$ admits a unique *pseudoinverse*, which is a matrix $A^+ \in \mathbb{R}^{n \times m}$ that satisfies a set of conditions [22]. Furthermore, if A has full column rank, then $A^+ = (A^\top A)^{-1}A^\top$. By Lemma 1 and 2 the random matrix R constructed in Steps **R 1** - **R 2** has full column rank, whereby $(R^\top R)^{-1}R^\top = R^+$ is the pseudoinverse of R . Furthermore, by the singular value decomposition [15], since R has rank k , there exist orthogonal matrices $U \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{k \times k}$ (i.e., $U^\top U = I, V^\top V = I$), such that

$$R = U\Sigma V^\top,$$

where $\Sigma = \begin{pmatrix} S & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{n \times k}$, $S = \text{diag}(\sigma_1, \dots, \sigma_k) \in \mathbb{R}^{k \times k}$, and $\sigma_1 \geq \dots \geq \sigma_k > 0$. This is because, out of the singular values of R (diagonal entries of Σ), exactly k (the rank of R) values $\sigma_1, \dots, \sigma_k$ must be nonzero [15]. Furthermore, the pseudoinverse $R^+ = R^\top$ of R is given as [22],

$$R^+ = V\Sigma^+U^\top,$$

where $\Sigma^+ = \begin{pmatrix} S^{-1} & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{k \times n}$.

Finally, recall that the ℓ_2 -induced norm on a matrix is *unitarily invariant* [15, p. 346, 357], i.e., if $P \in \mathbb{R}^{m \times m}$ and $Q \in \mathbb{R}^{m \times m}$ are two orthogonal matrices, i.e., $P^\top P = I$ and $Q^\top Q = I$ then for any $A \in \mathbb{R}^{m \times m}$, $\|PAQ\| = \|A\|$.

Lemma 9 *Let $\text{AVI}(K, M, q)$ and R, x_o, ϵ be as in Lemma 8, let $x_e \in \text{ext}(K)$ and let μ_1, μ_2 be defined as in (7) and (8). Then the following bounds hold:*

1. $\mu_1(x_e) \geq -\frac{n}{k}D \cdot \|M\| \cdot B$
2. $\mu_2(x_e) \geq -\epsilon D \|q\| - \epsilon D \|M\| B$

where $D := \max_{x_1, x_2 \in K} \|x_1 - x_2\|$ and $B := \max_{x \in K} \|x\|$.

Proof By Cauchy-Schwartz inequality,

$$\begin{aligned} \mu_1(x_e) &\geq -\frac{n}{k} \|x_e - x_o\| \cdot \|RR^\top (M - M') x_o\| \\ &\geq -\frac{n}{k} \cdot D \cdot \|M\| \cdot \|RR^\top\| \cdot \|(I - RR^\top)\| \cdot B \end{aligned}$$

where $D = \max_{x_1, x_2 \in K} \|x_1 - x_2\|$ and $B = \max_{x \in K} \|x\|$. Let the singular value decomposition of R be $R = U\Sigma V^\top$, where U, Σ, V are as above. Consequently,

$$\|RR^\top\| = \|U\Sigma V^\top V\Sigma^+ U^\top\| \stackrel{(a)}{=} \|U^\top U \Sigma \Sigma^+ U^\top U\| = \|\Sigma \Sigma^+\|,$$

where (a) follows from noting that $V^\top V = I$ and the unitary invariance of the ℓ_2 -induced norm. Similarly,

$$\|(I - RR^\top)\| = \|I - U\Sigma V^\top V\Sigma^+ U^\top\| = \|U^\top U - U^\top U \Sigma \Sigma^+ U^\top U\| = \|I - \Sigma \Sigma^+\|.$$

Since R has rank k (cf. Lemma 1), $\Sigma \Sigma^+$ is an $n \times n$ matrix with only k of its diagonal entries as 1 and the remaining all entries being 0. Therefore $\|\Sigma \Sigma^+\| = 1$. Likewise, $\|I - \Sigma \Sigma^+\| = 1$. This gives the required bound on μ_1 .

For the other term μ_2 we have by triangle inequality and by definitions of D, B ,

$$\mu_2(x_e) \geq -\epsilon \cdot D \cdot (\|q\| + \|Mx_o\|) \geq -\epsilon \cdot D \cdot \|q\| - \epsilon \cdot D \cdot \|M\| \cdot B,$$

which completes the proof.

Recall that D, B in the above lemma are finite since K is compact (cf. Lemma 6). In the following section, we complete the proof of Theorem 1.

4.3 Proof of Theorem 1

To prove Theorem 1, we require the following result due to Candes and Tao [4] on optimal recovery from random measurements.

Theorem 2 Suppose that $f \in \mathbb{R}^n$ obeys $\|f\|_1 \leq C_1$, and let $\alpha > 0$ be a sufficiently small number (less than an absolute constant). Assume that we are given k random measurements $y_i = \langle f, \Psi_i \rangle$, where $i \in \Omega$, $|\Omega| = k$ and $\{\Psi_i\}$ is a set of k -uniformly random orthonormal vectors. Then with probability 1, we have a unique minimizer $f^\#$ to the following problem:

$$\min_{x \in \mathbb{R}^n} \|x\|_1 \quad \text{subject to} \quad y_i = \langle x, \Psi_i \rangle, \quad \forall i = 1, \dots, k.$$

Furthermore, with probability at least $1 - O(n^{-1/\alpha})$, we have the approximation

$$\|f - f^\#\|_2 \leq C \cdot C_1 \cdot (k/\ln(n))^{-1/2}.$$

Here, C is a fixed constant depending on α but not on anything else. The implicit constant in $O(n^{-1/\alpha})$ is allowed to depend on α .

We now complete the proof of our main result, Theorem 1.

Proof (Proof of Theorem 1) Let \tilde{x} be generated by Step 5 by solving the lower-dimensional AVI($\tilde{K}, \tilde{M}, \tilde{q}$) and let $x_o \in K$ be a deterministic preimage of \tilde{x} . From Lemma 8 and Lemma 9 that with probability greater than $p := (1 - 4\eta e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2})$, x_o satisfies

$$(x_e - x_o)^\top (q + Mx_o) \geq -\frac{n}{k} \|M\|B - \epsilon \cdot D\|q\| - \epsilon \cdot D\|M\|B, \quad \forall x_e \in \text{ext}(K). \quad (9)$$

Consider any $y \in K$. There exist $\alpha_1, \dots, \alpha_\eta \geq 0$ such that $\sum_i \alpha_i = 1$ such that $y = \sum_i \alpha_i x_i$ where x_1, \dots, x_η are the extreme points of K . Multiplying the inequality in (9) corresponding to each $x_i \in \text{ext}(K)$ by α_i , and adding over all i , we get that under the event that (9) is true, the event

$$(y - x_o)^\top (q + Mx_o) \geq -\frac{n}{k} \|M\|B - \epsilon \cdot D\|q\| - \epsilon \cdot D\|M\|B, \quad \forall y \in K, \quad (10)$$

is true. Consequently, x_o satisfies (10) with probability at least p . Since δ is the confidence parameter, we set $p > (1 - \delta)$. This necessitates that $k > \frac{2 \ln(\frac{4\eta}{\delta})}{(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})}$. Consequently, with probability strictly greater than $(1 - \delta)$, x_o satisfies,

$$(y - x_o)^\top (q + Mx_o) \geq -\frac{\left(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3}\right) n}{2 \ln(\frac{4\eta}{\delta})} \|M\|B - \epsilon \cdot D\|q\| - \epsilon \cdot D\|M\|B, \quad \forall y \in K.$$

This proves the first statement of the theorem.

To prove the second statement, let x^* be the vector generated by Step 6 of Algorithm A and let $\alpha > 0$ be a small number as required by Theorem 2. By Theorem 2, we have with probability at least $1 - O(n^{-1/\alpha})$,

$$\|x_0 - x^*\| \leq C\|x_0\|_1 \leq CB' \left(\frac{k}{\ln(n)} \right)^{-1/2}, \quad (11)$$

where $B' := \max_{x \in K} \|x\|_1$. Let $x^\#$ be the output of Algorithm A. Since $x_o \in K$ and since the ℓ_2 -projection on the closed convex set K in Step 7 is non-expansive [10], we have

$$\|x_o - x^\#\| = \|\Pi_K(x_o) - \Pi_K(x^*)\| \leq \|x_o - x^*\|. \quad (12)$$

Combining (11) and (12), we get the second statement. The proof is complete.

With this we conclude the theoretical portion of the paper. In the following section we point out some remarks about the algorithm, following which we present numerical results.

5 Some remarks about the algorithm

5.1 Solving the lower dimensional AVI

Although the lower-dimensional problem $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$ is indeed an AVI, there is a practical difficulty in processing it. AVI solvers typically require the polyhedron defining the constraints to be given in its half-space representation and we do not have direct access to the half-space representation of \tilde{K} . Following are two approaches to this issue.

First approach: We augment extra variables and convert the lower-dimensional AVI to a sparse larger-dimensional AVI. Define a new augmented variable $v = (\tilde{x}, x) \in \mathbb{R}^{n+k}$. Instead of $\text{AVI}(\tilde{K}, \tilde{M}, \tilde{q})$, the solver is supplied problem $\text{AVI}(\hat{K}, \hat{M}, \hat{q})$, where

$$\hat{q} = \begin{bmatrix} \tilde{q} \\ 0_{n \times 1} \end{bmatrix}, \quad \hat{M} = \begin{bmatrix} \tilde{M} & 0_{k \times n} \\ 0_{n \times k} & 0_{n \times n} \end{bmatrix}, \quad \text{and} \quad \hat{K} = \{v \in \mathbb{R}^{n+k} \mid Cv = 0, Dv \leq b\},$$

where $C = [I \mid -R^\top]$, $D = [0 \mid A]$. Although the problem now becomes a larger dimensional one, there is an enormous amount of sparsity in the formulation which can potentially be exploited by AVI solvers to solve this problem rather quickly. We have found this to be the case with the PATH solver and have used it in our numerical results.

Second approach: Alternatively, one may try and deduce the half-space representation of \tilde{K} . To do so, first, complete the random projection matrix R by filling in the remaining entries (denote these entries by $\Delta^{n \times (n-k)}$) by picking the remaining set of $(n-k)$ uniformly random orthonormal vectors in exactly the same manner as described in Section 2.1.1. Let this complete $n \times n$ matrix be denoted by \hat{R} . If $K = \{x \in \mathbb{R}^n \mid Ax \leq b\}$, let $\hat{K} = \{y \in \mathbb{R}^n \mid \hat{A}y \leq b\}$, where $\hat{A} = \sqrt{\frac{k}{n}} A((\hat{R}^\top)^{-1})$. Note that $\hat{R} = [R \Delta^{n \times (n-k)}]$, which implies that the first k components of a vector $y \in \hat{K}$ comprise the vector $\tilde{x} = \sqrt{\frac{n}{k}} R^\top x$ which is an element of \tilde{K} .

Thus the required polytope \tilde{K} is the Euclidean projection of \hat{K} on the first k components. To obtain the half-space representation for this polytope one may make use of a method such as Fourier-Motzkin elimination [5]. Unfortunately, Fourier-Motzkin elimination is known to have poor complexity and could potentially nullify any advantages of dimensionality reduction.

Converting the half-space representation of a polyhedron to its vertex-representation and vice-versa is a fundamental combinatorial problem. Indeed, a side-story of our algorithm is the bringing to fore of the combinatorial nature of polyhedra and indeed of the AVI, which has otherwise been suppressed in “continuous” optimization efforts.

5.2 On the lower-dimension

The lower dimension k is required to be of the order of $\ln(\eta/\delta)/\epsilon^2$. By introducing slack variables if necessary, we may assume without loss of generality that the polytope K is represented as $\{x' \mid A'x' = b', x'_I \geq 0\}$ where $x'_I = (x'_i)_{i \in I}$ and that the given $\text{AVI}(K, M, q)$ is specified in the space of x' . Suppose n the dimension of x' . If A' is full row rank, then $\eta = |\text{ext}(K)| \leq \binom{n}{m}$ where m is the number of rows of A' (this follows from the argument used for bounding the number of basis feasible solutions in a linear program; see example [2, Ch. 3]). Consequently, for fixed m , we have $\ln(\eta) = O(\ln n)$ and the lower dimension $k = O(\ln n/\delta)/\epsilon^2$. In general, by Stirling’s approximation,

n	k	m	Natural Map Residual	Angle	Difference Norm	Major Low	Minor Low	Major High	Minor High
100	5	10	8.68	142.20	5.34	4	9	7	2468
100	10	10	9.22	141.75	2.50	4	9	7	2468
100	30	10	7.84	141.38	1.80	4	11	7	2468
100	50	10	7.68	141.05	5.90	5	15	7	2468
100	70	10	6.43	142.14	7.27	5	15	7	2468
100	90	10	4.87	141.71	11.71	7	239	7	2468
150	5	15	12.12	141.55	6.09	4	12	3	1736
150	10	15	11.26	140.69	1.30	4	12	3	1736
150	30	15	11.05	141.94	2.96	5	16	3	1736
150	50	15	10.51	141.65	1.58	5	15	3	1736
150	80	15	9.03	141.70	2.83	6	21	3	1736
150	110	15	7.42	142.21	1.42	5	160	3	1736
150	135	15	5.27	140.33	6.36	7	113	3	1736
200	5	16	12.68	141.68	1.46	5	15	2	2201
200	10	16	12.39	141.59	1.47	4	13	2	2201
200	30	16	12.33	142.22	2.13	5	19	2	2201
200	50	16	12.32	141.61	1.90	5	22	2	2201
200	80	16	11.65	141.37	2.18	6	29	2	2201
200	110	16	9.92	141.44	1.29	5	21	2	2201
200	140	16	8.74	140.94	3.74	6	53	2	2201
200	180	16	5.51	141.28	3.13	8	128	2	2201
250	5	25	15.03	141.87	2.27	4	14	3	3031
250	10	25	14.75	141.22	4.22	4	17	3	3031
250	30	25	14.98	141.75	2.43	4	19	3	3031
250	50	25	14.16	141.28	3.82	5	22	3	3031
250	70	25	13.84	141.65	2.33	5	28	3	3031
250	100	25	13.20	141.03	6.31	6	45	3	3031
250	130	25	12.13	140.69	2.42	5	28	3	3031
250	160	25	10.88	141.25	2.33	5	27	3	3031
250	190	25	9.49	141.31	13.29	6	230	3	3031
250	225	25	2.58	141.49	496.15	8	140	3	3031

Table 1: Performance of Algorithm A on randomly generated test problems where each entry of M, q, A, b was chosen from $N(0, 1)$.

$\ln \binom{n}{m} \simeq nH\left(\frac{m}{n}\right)$ where $H(t) \equiv -t \ln t - (1-t) \ln(1-t)$ is the entropy function [25, p. 2], and hence $\ln \binom{n}{m}$ is at most n (this is an approximate statement, the quality of which depends on the quality of Stirling's approximation).

The term η appears because to solve the AVI, we have to satisfy η -many inequalities, and η -many inner products have to be *simultaneously* preserved in the JL lemma. Notice, however, that each fixed $y \in K$ can be written as a convex combination of $n+1$ points from $\text{ext}(K)$ (by Caratheodory's theorem [27]). Thus for each $y \in K$, we only need to simultaneously preserve $n+1$ inner products. Thus for any $y \in K$ the inequality

$$(y - x_o)^\top (Mx_o + q) \geq 0,$$

holds with probability $1 - 4(n+1)e^{-(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3})k/2}$, whereby for this event to hold with probability $> 1 - \delta$, we need k to be only of the order of $\ln \left(\frac{4(n+1)}{\delta} \right) / \epsilon^2$. It is for showing that “ $(y - x_o)^\top (Mx_o + q) \geq 0$ ” holds *for all* $y \in K$ that one requires $k \sim \ln(\eta/\delta)/\epsilon^2$. One may interpret this issue also to be a manifestation of the combinatorial nature of the polyhedron.

Another way one may interpret this matter is via the V-representation of polyhedra [34]. In the V-representation the shape of K is defined by its η extreme points and η is thus the indicator of the “complexity” of this shape. On the other hand in the H-representation, the shape of the polyhedron is determined by the number of half-spaces necessary to describe it.

n	k	m	Natural Map Residual	Angle	Difference Norm	Major Low	Minor Low	Major High	Minor High
100	5	10	2.54	131.67	1.08	4	9	2	1132
100	10	10	2.55	123.58	1.04	4	9	2	1132
100	30	10	2.44	122.26	1.97	5	11	2	1132
100	50	10	2.01	125.92	1.01	5	12	2	1132
100	70	10	1.82	119.70	2.76	8	17	2	1132
100	90	10	1.15	107.13	125.55	7	911	2	1132
150	5	15	3.06	124.51	1.33	4	11	7	3404
150	10	15	3.17	123.56	2.02	4	12	7	3404
150	30	15	3.13	118.31	2.82	5	14	7	3404
150	50	15	3.04	120.46	3.37	6	18	7	3404
150	80	15	2.58	120.34	2.70	6	21	7	3404
150	110	15	2.02	108.59	4.24	7	16	7	3404
150	135	15	1.36	109.33	4.93	5	47	7	3404
200	5	16	3.59	119.97	1.03	4	9	3	2332
200	10	16	3.59	123.81	1.09	5	13	3	2332
200	30	16	3.75	115.21	12.22	6	214	3	2332
200	50	16	3.63	120.52	1.24	5	15	3	2332
200	80	16	3.35	118.66	1.41	6	16	3	2332
200	110	16	3.05	114.44	1.27	6	18	3	2332
200	140	16	2.35	113.24	165.33	8	33	3	2332
200	180	16	1.62	106.38	1.84	8	1287	3	2332
250	5	25	4.14	119.08	1.44	4	14	3	2929
250	10	25	4.03	119.81	1.32	4	16	3	2929
250	30	25	4.15	117.70	2.12	5	13	3	2929
250	50	25	4.08	116.17	1.70	6	20	3	2929
250	70	25	3.90	118.32	1.50	6	29	3	2929
250	100	25	3.61	115.63	9.64	7	34	3	2929
250	130	25	3.39	114.22	2.15	8	37	3	2929
250	160	25	2.97	111.25	4.20	8	364	3	2929
250	190	25	2.57	109.61	2.92	8	79	3	2929
250	225	25	1.79	104.41	1.49	7	1381	3	2929

Table 2: Performance of Algorithm A on randomly generated test problems where each entry of M, q, A, b was chosen from $U[0, 1]$.

5.3 Construction of \widetilde{M}

Notice that while \widetilde{q} and \widetilde{K} are projections (in the sense of Section 2.1) of q and K , respectively, \widetilde{M} is *not* a projection of M . Rather \widetilde{M} may be viewed as the projection of a matrix that is akin to a least squares approximation to M . We explain this below.

For the lower-dimensional problem to be an AVI, we require that in addition to \widetilde{M} being a projection of some matrix, \widetilde{M} must also be compatible with right-multiplication by $\widetilde{x} = \sqrt{\frac{n}{k}} R^\top x$. Thus \widetilde{M} must be of the form $\widetilde{M} = R^\top M'$, where $M' = X R^\top$ and X is the matrix to be determined. Now since the eventual error depends on $\|M - M'\|$, in our construction we further let M' to be of the form $M Z R^\top$ where now Z is to be determined. To minimize the error, we require Z to be such that $Z R^\top$ is as close to the identity as possible, i.e., Z must solve the following “least-squares” optimization problem:

$$\min_{Z \in \mathbb{R}^{n \times k}} \|Z R^\top - I\|^2.$$

The solution of the above problem is $Z = R(R^\top R)^{-1} = R$, since R is orthonormal. This gives the construction of \widetilde{M} as $\widetilde{M} = R^\top M R$. In addition to being suitable for minimizing the error, we note that \widetilde{M} enjoys the property that \widetilde{M} is positive semidefinite if M is positive definite.

n	k	m	Natural Map Residual	Angle	Difference Norm	Major Low	Minor Low	Major High	Minor High
150	30	15	11.53	142.05	13.61	8	39	3	1736
150	30	15	9.98	142.24	1.11	3	8	3	1736
150	30	15	9.22	140.96	1.03	4	8	3	1736
150	30	15	11.67	143.69	2.12	6	19	3	1736
150	30	15	9.84	140.42	1.04	3	9	3	1736
150	30	15	10.30	141.12	5.62	7	24	3	1736
150	30	15	9.91	143.82	1.13	4	14	3	1736
150	30	15	10.14	143.04	1.13	4	16	3	1736
150	30	15	10.53	141.18	1.23	3	9	3	1736
150	30	15	11.25	140.83	1.61	4	11	3	1736

Table 3: Results for Algorithm A with different samples of the random projection (for a test problem generated by taking each entry of M, q, A, b from $N(0, 1)$).

6 Numerical Results

We now present numerical results to show how our algorithm performs in practice.

To keep the trials generic, test problems were generated randomly. The input to the algorithm was given as $\text{AVI}(K, M, q)$ where we took $K = \{x \in \mathbb{R}^n | Ax \leq b, L \leq x \leq U\}$. The bounds L, U on x were introduced to ensure that K is compact. To generate a random input problem, all the entries of q, M, A and b were generated randomly with entries drawn from distributions $N(0, 1)$ and $U[0, 1]$; $-L, U$ were set to be large. We applied our algorithm with various choices of the lower dimension k , for each value of the higher dimension n . All AVIs were solved using the PATH solver [7].

The sections below show the performance of the algorithm for two applications. Section 6.1 discusses results for obtaining an approximate solution to the given AVI (i.e., Algorithm A). In Section 6.2, we apply Algorithm A to solve the given AVI exactly by using Algorithm A to generate an initial point for the solver.

6.1 Performance for an approximate solution

Table 1 contains results pertaining to Algorithm A. To benchmark the performance of the algorithm, the given AVI was approximated by Algorithm A and also solved exactly. Since the projections were random, the results reported for each test problem are the average of the results over 10 random choices of R . Table 3 shows the behavior across different random trials for a representative test problem.

The columns of Tables 1, 2 and 3 contain the following entries.

- n – dimension of the $\text{AVI}(K, M, q)$ given as input to Algorithm A; m – number of rows of the matrix A in the definition of K ; k – dimension to which we project the given AVI.
- Natural Map Residual = $\frac{\|F_K^{\text{nat}}(x^\#)\|}{\|x^\#\| + 1}$, where $x^\#$ is the output of Algorithm A. This quantity should be 0 for $x^\#$ to be a solution. We normalize by $(\|x^\#\| + 1)$ in order to allow for comparisons across difference values of n, k .
- Angle – the largest angle between $q + Mx^\#$ and a vector $y - x^\#$ as y ranges over K , where $x^\#$ is as above. Define $\beta = \min_{y \in K} (y - x^\#)^\top (q + Mx^\#)$ and let y^* be the corresponding minimizer, then

$$\text{Angle} = \arccos \left(\frac{\beta}{\|y^* - x^\#\| \cdot \|q + Mx^\#\|} \right).$$

n	k	m	Major High	Minor High	Major Additional	Minor Additional	Major Total	Minor Total
100	5	10	40	9885	16	3699	19	3708
100	10	10	40	9885	2	690	6	700
100	25	10	40	9885	9	2478	14	2494
100	50	10	40	9885	10	2784	15	2802
100	75	10	40	9885	16	4088	21	4107
125	5	12	12	4354	4	1821	9	1834
125	10	12	12	4354	7	2873	12	2886
125	25	12	12	4354	11	4223	16	4236
125	50	12	12	4354	7	2693	12	2712
125	75	12	12	4354	8	3392	14	3414
125	100	12	12	4354	8	3185	14	3205
150	5	15	19	7356	5	2574	10	2587
150	10	15	19	7356	8	3940	12	3952
150	30	15	19	7356	5	2521	9	2535
150	60	15	19	7356	4	2252	9	2272
150	75	15	19	7356	7	3468	13	3495
150	90	15	19	7356	4	2004	12	2037
150	120	15	19	7356	4	2011	10	2101

Table 4: Performance of the exact algorithm (test problems generated by choosing each entry of M, q, A, b independently from $N(0, 1)$)

- Difference Norm – norm of the difference between $x^\#$ and the exact solution computed by solving the AVI directly (denoted \bar{x}), normalized by the norm of the exact solution, i.e.,

$$\text{Difference Norm} = \frac{\|x^\# - \bar{x}\|}{\|\bar{x}\| + 1}.$$

- Major/Minor Low - number of major/minor iterations reported by PATH [7] to obtain a solution using our algorithm
- Major/Minor High - number of major/minor iterations reported by PATH [7] to solve the higher dimensional problem directly

For every tuple (n, m, k) , 10 independent simulations were carried out, realizing a different random matrix R each time. All the parameters in Table 1 and Table 2 are the average values over the set of 10 trials. The results are as given in Table 1 (for test problems generated from $N(0, 1)$) and Table 2 (for test problems generated from $U[0, 1]$).

6.1.1 Key observations

The first observation to be made from Tables 1 and 2 is that the algorithm produces only an approximation, which is evident from the fact that the natural map residual is non-zero and “Angle” is greater than 90° . However, for each n , the normalized natural map residual decreases monotonically with the value of lower dimension k . Clearly, as we increase the value of the lower dimension, the natural map residual decreases towards zero. Furthermore, in Table 2, one sees a decrease in the “Angle” with increasing k , for each n . Interestingly, this decrease is not seen in Table 1. We do not know of a way of explaining this.

The most important observation is that for every case, the number of minor iterations (minor low) consumed by our algorithm is significantly lower than the case where the high dimensional

problem is attempted to solve directly (minor high). Though results corresponding to major iterations are inconclusive on a whole, this validates our approach of finding a quick approximation. Notice that “Difference norm” is *not* monotonic. This may be due to possible non-uniqueness of the solution of the AVI.

Recall that Tables 1 and 2 are average values of multiple trials of R . For a fixed problem and a fixed n , k and m , a representative set of results from different choices of R are reported in Table 3. One can see that the variation across different trials is relatively small and the average values reported Tables 1 and 2 are representative.

6.2 Performance for an exact solution

For these results, the computation was carried out in two steps. First, an approximate solution was obtained using our algorithm, like in the previous section. Next, this vector was supplied as an initial point into the AVI solver and an exact solution to the original AVI was computed. The original problem was independently solved using the solver directly with a random initial point, and the performance was compared. The results for this case have been tabulated in terms of the following parameters:

- n, k, m and Major/Minor High as in Section 6.1
- Major/Minor Total - total number of major/minor iterations reported by PATH to obtain a final solution with our approximate solution as an initial point
- Major/Minor Additional - number of extra major/minor iterations reported by PATH to solve the higher dimensional problem *after* the initial point has been supplied from our algorithm

6.2.1 Key observations

Notice that the solution obtained in this case exactly solves our original high dimensional AVI. Minor Total is lower in all the test cases than the corresponding number for when the problem is solved directly. The additional iterations, both major and minor, consumed on the higher dimensional problem, are significantly lower than the corresponding values for when the problem is solved directly. This implies that when our recovered solution is supplied as an initial point to the solver, it computes an exact solution faster than the direct case, where an initial point is generated randomly.

7 Conclusions

Motivated by emerging problems in ‘Big Data’, this paper has presented a new method for dimensionality reduction of AVIs with compact feasible regions. The method yields with high probability an approximate solution to the given AVI by solving an AVI of a lower dimension; the latter is formed by appropriately projecting the given AVI on a lower-dimensional space. Using the approximate solution as an initial point, the method can also be used to ‘hot start’ a solver for the given problem and thereby find an exact solution. We presented numerical results to demonstrate that the method is indeed effective in practice.

Acknowledgments

The authors would like to thank Dr. Dinesh Garg of IBM Research, India for his inputs on this topic.

References

1. Barman, K., Borkar, V.S.: A note on linear function approximation using random projections. *Systems & Control Letters* doi:10.1016/j.sysconle.2008.03.003 (2008)
2. Bazaraa, M.S., Jarvis, J.J., Sherali, H.D.: *Linear programming and network flows*. John Wiley & Sons (2011)
3. Benders, J.F.: Partitioning procedures for solving mixed-variables programming problems. *Numer. Math.* **4**, 238–252 (1962)
4. Candes, E.J., Tao, T.: Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on* **52**(12), 5406–5425 (2006)
5. Dantzig, G.B., Curtis Eaves, B.: Fourier-Motzkin elimination and its dual. *Journal of Combinatorial Theory, Series A* **14**(3), 288–297 (1973)
6. Dasgupta, S., Gupta, A.: An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms* **22**(1), 60–65 (2003)
7. Dirkse, S., Ferris, M.C., Munson, T.: The PATH solver. URL <http://pages.cs.wisc.edu/ferris/path.html>
8. Dirkse, S.P., Ferris, M.C.: The PATH solver: A non-monotone stabilization scheme for mixed complementarity problems. CS-TR-1993-1179 (Department of Computer Science, University of Wisconsin at Madison, 1993). URL citeseer.ist.psu.edu/dirkse95path.html
9. Facchinei, F., Fischer, A., Piccialli, V.: On generalized Nash games and variational inequalities. *Operations Research Letters* **35**(2), 159–164 (2007)
10. Facchinei, F., Pang, J.S.: *Finite-Dimensional Variational Inequalities and Complementarity Problems I*, first edn. Springer, New York (2003)
11. Facchinei, F., Pang, J.S.: *Finite-Dimensional Variational Inequalities and Complementarity Problems II*, first edn. Springer, New York (2003)
12. Frankl, P., Maehara, H.: The Johnson-Lindenstrauss lemma and the sphericity of some graphs. *Journal of Combinatorial Theory, Series B* **44**(3), 355–362 (1988)
13. Gill, P.E., Murray, W., Saunders, M.A.: SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM J. Optim.* **12**(4), 979–1006 (electronic) (2002)
14. Goodman, J.E., O'Rourke, J.: *Handbook of discrete and computational geometry*. CRC press (2004)
15. Horn, R.A., Johnson, C.R.: *Matrix analysis*. Cambridge University Press, Cambridge (1990)
16. Johnson, W.B., Lindenstrauss, J.: Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics* **26**(189-206), 1 (1984)
17. Knuth, D.E.: *The art of computer programming*, 3rd edn., vol. 2. *Seminumerical Algorithms* (1998)
18. Kulkarni, A.A., Borkar, V.S.: Finite dimensional approximation and Newton-based algorithm for stochastic approximation in Hilbert space. *Automatica* **45**(12), 2815–2822 (2009). DOI <http://dx.doi.org/10.1016/j.automatica.2009.09.031>
19. Kulkarni, A.A., Shanbhag, U.V.: On the variational equilibrium as a refinement of the generalized Nash equilibrium. *Automatica* **48**(1), 45–55 (2012)
20. Kulkarni, A.A., Shanbhag, U.V.: Recourse-based stochastic nonlinear programming: properties and Benders-SQP algorithms. *Computational Optimization and Applications* (appeared in *Online first* in 2010) **51**(1) (2012)
21. Kulkarni, A.A., Shanbhag, U.V.: Revisiting generalized Nash games and variational inequalities. *Journal of Optimization Theory and Applications* **154**(1), 1–12 (2012)
22. Laub, A.J.: *Matrix analysis for scientists and engineers*. SIAM (2005)
23. Luna, J.P., Sagastizábal, C., Solodov, M.: A class of Dantzig–Wolfe type decomposition methods for variational inequality problems. *Mathematical Programming* **143**(1-2), 177–209 (2014)
24. Luo, Z.Q., Tseng, P.: Error bound and convergence analysis of matrix splitting algorithms for the affine variational inequality problem. *SIAM Journal on Optimization* **2**(1), 43–54 (1992)
25. MacKay, D.J.: *Information theory, inference, and learning algorithms*, vol. 7. Citeseer (2003)
26. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag, New York (1999)
27. Rockafellar, R.T.: *Convex Analysis*. Princeton University Press, Princeton, NJ (1997). Reprint of the 1970 original, Princeton Paperbacks.

-
28. Shanbhag, U.V.: Decomposition and sampling methods for stochastic equilibrium problems. Ph.D. thesis, Department of Management Science and Engineering (Operations Research), Stanford University (2006)
 29. Shaw, J.: Why “Big Data” is a big deal. *Harvard Magazine* **Mar-April**, 30–35, 74–75 (2014)
 30. Tulino, A.M., Verdú, S.: Random matrix theory and wireless communications, vol. 1. Now Publishers Inc (2004)
 31. Van Slyke, R.M., Wets, R.: L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.* **17**, 638–663 (1969)
 32. Vempala, S.: The random projection method., *DIMACS series*, vol. 65. AMS (2004)
 33. Wang, Y., Boyd, S.: Fast model predictive control using online optimization. *Control Systems Technology, IEEE Transactions on* **18**(2), 267–278 (2010)
 34. Ziegler, G.M.: Lectures on polytopes, vol. 152. Springer (1995)