

On the computation of sparse reflexive generalized inverses

Gabriel Oliveira da Ponte

Universidade Federal do Rio de Janeiro
PESC/COPPE, Rio de Janeiro, CEP: 21941-972
pontegabriel@cos.ufrj.br
Advisors: Marcia Fampa and Jon Lee

RESUMO

A pseudo-inversa de Moore-Penrose (M-P) pode ser usada em várias aplicações de álgebra linear; por exemplo, para calcular soluções de mínimos quadrados de sistemas lineares inconsistentes. Entretanto, independentemente de uma dada matriz ser esparsa, sua pseudo-inversa pode ser completamente densa, levando potencialmente a um alto custo computacional e dificuldades numéricas, especialmente ao lidarmos com matrizes de alta dimensão. A pseudo-inversa M-P é caracterizada por quatro propriedades, mas nem todas precisam ser atendidas para algumas aplicações. Nós aplicamos otimização matemática para induzir esparsidade geral e esparsidade estruturada em inversas generalizadas de uma dada matriz que satisfaz somente subconjuntos específicos das propriedades M-P. Utilizamos minimização de norma para induzir esparsidade geral (não-estruturada) e esparsidade (estruturada) em linhas. Nós também investigamos a aplicação de procedimentos de busca local para construir inversas generalizadas com esparsidade estruturada, posto pequeno e com controle da magnitude das entradas.

KEYWORDS. Inversa generalizada. Otimização esparsa. Algoritmo de aproximação

OD - Otimização Discreta. OC- Otimização Contínua. OMO - Otimização Multiobjetivo

ABSTRACT

The well-known Moore-Penrose (M-P) pseudoinverse is used in several linear-algebra applications; for example, to compute least-squares solutions of inconsistent systems of linear equations. Irrespective of whether a given matrix is sparse, its M-P pseudoinverse can be completely dense, potentially leading to high computational burden and numerical difficulties, especially when we are dealing with high-dimensional matrices. The M-P pseudoinverse is uniquely characterized by four properties, but not all of them need to be satisfied for some applications. We apply mathematical optimization to induce general sparsity and structured sparsity on generalized inverses of a given matrix, which satisfy only specific subsets of the M-P properties. We use norm minimization to induce (unstructured) sparsity and (structured) row-sparsity. We also consider the application of local-search procedures that produce generalized inverses with guaranteed structured sparsity, with low rank, and with entries under control.

KEYWORDS. Generalized inverse. Sparse optimization. Approximation algorithm

OD - Discrete Optimization. OC- Continuous Optimization. OMO - Multi objective optimization

1. Introduction

The well-known M-P (Moore-Penrose) pseudoinverse, independently discovered by E.H. Moore and R. Penrose, is used in several linear-algebra applications — for example, to compute least-squares solutions of inconsistent systems of linear equations. If $A = U\Sigma V^\top$ is the real singular-value decomposition of A (see [Golub and Van Loan, 1996], for example), then the M-P pseudoinverse of A can be defined as $A^\dagger := V\Sigma^\dagger U^\top$, where $\Sigma^\dagger$ has the shape of the transpose of the diagonal matrix Σ , and is derived from Σ by taking reciprocals of the non-zero (diagonal) elements of Σ (i.e., the non-zero singular values of A).

The following theorem gives a fundamental characterization of the M-P pseudoinverse.

Theorem 1 ([Penrose, 1955]). *For $A \in \mathbb{R}^{m \times n}$, the M-P pseudoinverse $A^\dagger$ is the unique $H \in \mathbb{R}^{n \times m}$ satisfying:*

$$AHA = A \quad (\text{P1})$$

$$HAH = H \quad (\text{P2})$$

$$(AH)^\top = AH \quad (\text{P3})$$

$$(HA)^\top = HA \quad (\text{P4})$$

The first three M-P properties are particularly important for our purposes.

P1: Following [Rohde, 1964], we say that a *generalized inverse* of A is any H satisfying P1. Note that without P1, even the all-zero matrix (which carries no information about A) satisfies the other M-P properties.

P2: A generalized inverse is *reflexive* if it satisfies P2. Theorem 3.14 in [Rohde, 1964] states that: (i) if H is a generalized inverse of A , then $\text{rank}(H) \geq \text{rank}(A)$, and (ii) a generalized inverse H of A is reflexive if and only if $\text{rank}(H) = \text{rank}(A)$. Therefore, enforcing P2 gives us the lowest possible rank of a generalized inverse of A .

P3: Following [Xu et al., 2021], we say that H is *ah-symmetric* if it satisfies P3. That is, ah-symmetric means that AH is symmetric. If H is an ah-symmetric generalized inverse (i.e., H satisfies P1 and P3), then $\hat{x} := Hb$ solves the least-squares problem $\min\{\|Ax - b\|_2 : x \in \mathbb{R}^n\}$. So, not all of the M-P properties are required for a generalized inverse to solve this key problem. Especially important in our context, we have that by imposing property P3 on the generalized inverse when solving linear systems, we guarantee least-square solutions. If additionally, P2 is imposed, we have the advantage of working with generalized inverses of minimum rank. Low rank is a desirable property for an ah-symmetric generalized inverse H , in the context of the least-squares application, as it corresponds to a type of “explainability” for the associated linear model $\hat{x} := Hb$.

Even if a given input matrix A is sparse, its M-P pseudoinverse can be dense, leading to a high computational burden, especially in applications with high-dimensional matrices involving many response vectors b . Over the past ten years, sparse-optimization techniques have been widely studied to find sparse matrices H leading to efficiency in calculating $\hat{x} := Hb$. [Dokmanić et al., 2013; Dokmanić and Gribonval, 2017a,b] introduced the idea of seeking to induce sparsity and row-sparsity in H by applying the standard minimization of its (vector) 1-norm and of its 2,1-norm over the set of left inverses (i.e., $HA = I_n$) or right inverses (i.e., $AH = I_m$). [Fuentes et al., 2016, 2020] suggested using relaxations based 1-norm minimization on subsets M-P properties aiming at sparse generalized inverses. [Fampa and Lee, 2018] investigated one such kind of sparse generalized inverse, with particular interest in rank-deficient matrices and suggested a polynomial-time local search with a performance guarantee, aiming at minimizing the 1-norm over block-structured generalized inverses.

The present work started in 2018, as an undergraduate research project and the goal was to follow up the work of [Fampa and Lee, 2018]. Initially, I joined the ongoing project of Prof. Marcia Fampa, Prof. Jon Lee, and Luze Xu, to which I contributed by implementing the proposed algorithms and performing the computational experiments. The outcome of this work is presented in [Xu et al., 2021]. My dissertation reflects the continuity of our research on this topic to date, with new insights developed based on previous results. In the following, I further describe the publications derived from our work.

1. Luze Xu, Marcia Fampa, Jon Lee, Gabriel Ponte. Approximate 1-norm minimization and minimum rank structured sparsity for various generalized inverses via local search. *SIAM Journal on Optimization* 31(3), 1722-1747, 2021. <https://doi.org/10.1137/19M1281514>

We investigate a column block construction method to produce an ah-symmetric reflexive generalized inverse that is structured and has guaranteed sparsity. We provide a theoretically efficient and practical local-search algorithm to column-block construct an approximate 1-norm minimizing ah-symmetric reflexive generalized inverse.

2. Marcia Fampa, Jon Lee, Gabriel Ponte, Luze Xu. Experimental analysis of local searches for sparse reflexive generalized inverses. *Journal of Global Optimization* 81, 1057-1093, 2021. <https://doi.org/10.1007/s10898-021-01087-y>

[Fampa and Lee, 2018] and [Xu et al., 2021] proposed local-search procedures to construct sparse block-structured generalized inverses that satisfy only some of the M-P properties. We have implemented several local-search procedures based on results presented in these two papers and make an experimental analysis of them, considering their application to randomly generated matrices of varied dimensions, ranks, and densities. Further, we carried out a case study on a real-world data set.

3. Marcia Fampa, Jon Lee, Gabriel Ponte. Trading off 1-norm and sparsity against rank for linear models using mathematical optimization. *Open Journal of Mathematical Optimization*, 2(4), 14 p, 2021. <https://doi.org/10.5802/ojmo.6>

We investigate the trade-off between low 1-norm and low rank for generalized inverses that can be used in the computation of least-squares solutions. We propose several algorithmic approaches that start from a 1-norm minimizing generalized inverse that satisfies the two key M-P properties, and gradually decrease its rank, by iteratively imposing the reflexive property. The algorithms iterate until the generalized inverse has the least possible rank. During the iterations, we produce intermediate solutions, trading off low 1-norm (and typically high sparsity) against low rank.

4. Gabriel Ponte, Marcia Fampa, Jon Lee, Luze Xu. On computing sparse generalized inverses. *Operations Research Letters* 52, 2024. <https://doi.org/10.1016/j.orl.2023.107058>

We show that a 2,1-norm minimizing generalized inverse satisfies two additional M-P properties, including one needed for computing least-squares solutions. We present formulations related to finding row-sparse generalized inverses that can be solved very efficiently, which we verify numerically.

I gave talks about this work in:

1. 10th Week of Academic Integration at UFRJ, 2019;

2. LII Brazilian Symposium of Operations Research (SOBRAPO), 2020;
3. XLII Giulio Massarani Scientific, Technological, Artistic and Cultural Initiation Journey (JICTAC 2020 - Special Edition), 2021;
4. 11th Week of Academic Integration at UFRJ, 2021;

and I received the following awards:

1. Honorable Mention at 10th Week of Academic Integration at UFRJ, 2019;
2. SOBRAPO Undergraduate Research Project Award, 2020;
3. Honorable Mention at 11th Week of Academic Integration at UFRJ, 2021.

Notation. We use vector-norm notation on matrices: we write $\|H\|_1$ to mean $\|\text{vec}(H)\|_1$. We denote the identity matrix of order n by I_n . Matrix dot product is indicated by $\langle X, Y \rangle = \text{trace}(X^\top Y)$. For matrices, $\det(\cdot)$ denotes determinant, $\text{rank}(\cdot)$ denotes rank. We use $A[S, T]$ is the submatrix of A having row indices S and column indices T , with $S = \cdot$ (resp. $T = \cdot$) indicating all rows (resp. columns); further, $A_{i\cdot} := A[\{i\}, \cdot]$, $A_{\cdot j} := A[\cdot, \{j\}]$, and $A_{ij} := A[\{i\}, \{j\}]$. For $A \in \mathbb{R}^{m \times n}$, $\|A\|_{p,q} := \|(\|A_{1\cdot}\|_p, \dots, \|A_{m\cdot}\|_p)\|_q$ ($p, q \geq 0$).

In what follows, we describe the major outcomes from my dissertation.

2. Approximation algorithm with sparsity guarantees

In this section, let A be an arbitrary $m \times n$ real matrix with $\text{rank}(A) = r$. We seek to obtain a solution to $\min\{\|H\|_1 : P1 + P2 + P3\}$ (that is, a 1-norm minimizing ah-symmetric reflexive generalized inverse). As we have mentioned, ah-symmetric generalized inverses play a key role in solving least-squares problems. We develop an approximation approach for this problem with guaranteed structured sparsity, with low rank, and with entries under control.

Proposition 1. (see, for example, [Fuentes et al., 2016, Prop. 4.3]) If H satisfies $P1$ and $P3$, then $AH = AA^\dagger$.

Note from Proposition 1 that if H is an ah-symmetric generalized inverse, then $AH = AA^\dagger$, where $A^\dagger$ is the M-P pseudoinverse. Therefore, $P2$ ($HAH = H$) becomes a linear constraint $HAA^\dagger = H$, which implies that $\min\{\|H\|_1 : P1 + P2 + P3\}$ can be cast as an LP. However, the extreme solutions of this (linear program) LP only have a guaranteed bound of $mr + (m-r)(n-r)$ for the number of nonzeros, while the extreme solutions of $\min\{\|H\|_1 : P1 + P3\}$ have at most mr nonzeros.

Proposition 2. Suppose that $A \in \mathbb{R}^{m \times n}$ has rank r .

1. Extreme solutions of the LP for $\min\{\|H\|_1 : P1 + P3\}$ have at most mr nonzeros. Furthermore, the bound is sharp for all $m \geq n \geq r \geq 1$.
2. Extreme solutions of the LP for $\min\{\|H\|_1 : P1 + P2 + P3\}$ have at most $mr + (m-r)(n-r)$ nonzeros.

We seek to do better than what Proposition 2, part 2 provides. We want fewer nonzeros, and we want block structure. To get these properties, we give a new *column block construction*, producing an ah-symmetric reflexive generalized inverse that has at most mr nonzeros.

Theorem 2. For $A \in \mathbb{R}^{m \times n}$, let $r := \text{rank}(A)$. For any T , an ordered subset of r elements from $\{1, \dots, n\}$, let $\hat{A} := A[:, T]$ be the $m \times r$ submatrix of A formed by columns T . If $\text{rank}(\hat{A}) = r$, let $\hat{H} := \hat{A}^\dagger = (\hat{A}^\top \hat{A})^{-1} \hat{A}^\top$. The $n \times m$ matrix H with all rows equal to zero, except rows T , which are given by $\hat{H}$, is an ah-symmetric reflexive generalized inverse of A .

For $r := \text{rank}(A)$, we will efficiently find an ah-symmetric reflexive generalized inverse following our column block construction that is within a factor $r(1 + \epsilon)$ of the 1-norm of the ah-symmetric reflexive generalized inverse having minimum 1-norm.

Definition 1. Let A be an arbitrary $m \times n$, rank- r matrix, and let S be an ordered subset of r elements from $\{1, \dots, m\}$ such that these r rows of A are linearly independent. For T an ordered subset of r elements from $\{1, \dots, n\}$, and fixed $\epsilon \geq 0$, if $|\det(A[S, T])|$ cannot be increased by a factor of more than $1 + \epsilon$ by swapping an element of T with one from its complement, then we say that $A[S, T]$ is a $(1 + \epsilon)$ -local maximizer for the absolute determinant on the set of $r \times r$ nonsingular submatrices of $A[S, :]$.

Lemma 1. Let T be an ordered subset of r elements from $\{1, \dots, n\}$ and $\hat{A} := A[:, T]$ be the $m \times r$ submatrix of an $m \times n$ matrix A formed by columns T , and $\text{rank}(\hat{A}) = r$. There exists an $m \times n$ matrix W and a skew-symmetric $m \times m$ matrix U such that

$$\hat{A}^\top W A^\top + \hat{A}^\top U = E,$$

where $E := \text{sign}(\hat{A}^\dagger)$. Furthermore, $\langle A, W \rangle = \|\hat{A}^\dagger\|_1$.

Theorem 3. Let A be an arbitrary $m \times n$, rank- r matrix, and let S be an ordered subset of r elements from $\{1, \dots, m\}$ such that these r rows of A are linearly independent. Choose $\epsilon \geq 0$, and let $\tilde{A} := A[S, T]$ be a $(1 + \epsilon)$ -local maximizer for the absolute determinant on the set of $r \times r$ nonsingular submatrices of $A[S, :]$. Then the $n \times m$ matrix H constructed by Theorem 2 over $\hat{A} := A[:, T]$, is an ah-symmetric reflexive generalized inverse of A satisfying $\|H\|_1 \leq r(1 + \epsilon)\|H_{\text{opt}}^r\|_1$, where H_{opt}^r is a 1-norm minimizing ah-symmetric reflexive generalized inverse of A .

3. Inducing sparsity with optimization problems

It is well known that structured sparsity is more useful than sparsity, because of computational efficiency and for explainability. In the context of the least-squares application of an ah-symmetric generalized inverse H , it is desirable to have few rows that have non-zeros, as then the associated linear model is more explainable. Because no generalized inverse of A can have rank less than $\text{rank}(A)$, every generalized inverse of A has at least $\text{rank}(A)$ rows with non-zeros. So a row-sparse generalized inverse of A is one having or nearly having $\text{rank}(A)$ rows with non-zeros.

Let $A \in \mathbb{R}^{m \times n}$ with rank r and $A =: U\Sigma V^\top$ be the singular-value decomposition of A , where $U \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices ($U^\top U = I_m$, $V^\top V = I_n$), and $\Sigma \in \mathbb{R}^{m \times n}$ with

$$\Sigma =: \begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{matrix} r \times r & r \times (n-r) \\ (m-r) \times r & (m-r) \times (n-r) \end{matrix}$$

with D being a diagonal matrix with rank r . Let $H \in \mathbb{R}^{n \times m}$ and $\Gamma := V^\top H U$, where

$$\Gamma =: \begin{bmatrix} X & Y \\ Z & W \end{bmatrix}, \quad \begin{matrix} r \times r & r \times (m-r) \\ (n-r) \times r & (n-r) \times (m-r) \end{matrix}$$

then $H = I_n H I_m = (V V^\top) H (U^\top U) = V \Gamma U^\top$.

Considering the notation above, we present in the following, results related to the M-P pseudoinverse properties that will be used to prove our main results in the next section. We wish to emphasize that the results that follow refer to the singular-value decomposition of A and associated structural notation for Σ and Γ above.

Lemma 2. (Structure of generalized inverses). *P1 is equivalent to $X = D^{-1}$.*

Lemma 3. (Structure of reflexive generalized inverses). *If P1 is satisfied, then P2 is equivalent to $ZDY = W$.*

Lemma 4. (Structure of ah-symmetric generalized inverses). *If P1 is satisfied, then P3 is equivalent to $Y = 0$.*

We consider the optimization problem

$$\begin{aligned} \min \quad & f_{\text{row}}(\|H_{1\cdot}\|_2, \|H_{2\cdot}\|_2, \dots, \|H_{n\cdot}\|_2) \\ \text{s.t.} \quad & AHA = A, \end{aligned} \quad (P_1^{f_{\text{row}}})$$

where $f_{\text{row}} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is increasing in each argument.

Theorem 4. *Suppose that H is an optimal solution to $P_1^{f_{\text{row}}}$, then $Y = 0, W = 0$.*

Corollary 1. *Suppose that H is an optimal solution to $P_1^{f_{\text{row}}}$. Then, P2 and P3 are satisfied.*

Next, aiming at sparse and block-structured ah-symmetric reflexive generalized inverses, we consider a special case of $P_1^{f_{\text{row}}}$, where f_{row} is the 1-norm and, therefore, we minimize the 2,1-norm of H .

$$\begin{aligned} z(P_1^{2,1}) &:= \min \|H\|_{2,1} \\ \text{s.t.} \quad & AHA = A. \end{aligned} \quad (P_1^{2,1})$$

Minimizing the 1-norm of a matrix has been widely applied in the literature as a surrogate to minimizing the 0-norm, or the sparsity. So we also consider the two following alternative formulations to obtain H .

$$\begin{aligned} z((P_1^{2,1})^1) &:= \min \|H\|_1 \\ \text{s.t.} \quad & AHA = A, \\ & \|H\|_{2,1} \leq z(P_1^{2,1}). \end{aligned} \quad ((P_1^{2,1})^1)$$

$$\begin{aligned} z(P_{123}^1) &:= \min \|H\|_1 \\ \text{s.t.} \quad & AHA = A, \\ & HAH = A, \\ & AH = (AH)^\top. \end{aligned} \quad (P_{123}^1)$$

We have that the solution $H \in \mathbb{R}^{n \times m}$ of the three formulations $P_1^{2,1}$, $(P_1^{2,1})^1$ and P_{123}^1 , can be written as $H = VTU^\top$, where

$$\Gamma := \begin{bmatrix} D^{-1} & 0 \\ r \times r & r \times (m-r) \\ Z & 0 \\ (n-r) \times r & (n-r) \times (m-r) \end{bmatrix}.$$

Therefore, formulations $P_1^{2,1}$, $(P_1^{2,1})^1$ and P_{123}^1 can be reformulated *much more tractably* as, respectively,

$$z(\mathcal{P}_1^{2,1}) := \min \sum_{i=1}^n \left\| \mathbf{e}_i^\top V \begin{bmatrix} D^{-1} \\ Z \end{bmatrix} \right\|_2, \quad (\mathcal{P}_1^{2,1})$$

$$\begin{aligned} z((\mathcal{P}_1^{2,1})^1) &:= \min \left\| V \begin{bmatrix} D^{-1} & 0 \\ Z & 0 \end{bmatrix} U^\top \right\|_1 \\ \text{s.t. } &\sum_{i=1}^n \left\| \mathbf{e}_i^\top V \begin{bmatrix} D^{-1} \\ Z \end{bmatrix} \right\|_2 \leq z(\mathcal{P}_1^{2,1}), \end{aligned} \quad ((\mathcal{P}_1^{2,1})^1)$$

$$z(\mathcal{P}_{123}^1) := \min \left\| V \begin{bmatrix} D^{-1} & 0 \\ Z & 0 \end{bmatrix} U^\top \right\|_1. \quad (\mathcal{P}_{123}^1)$$

4. Concluding remarks

Generalized inverses have a wide variety of uses in matrix algebra and its applications. Sparsity of a generalized inverse is highly preferred for efficiency in its use; structured sparsity and low rank (=reflexivity) are both preferred for explainability. (Approximate) 1-norm minimization is useful for keeping entries under control and for inducing sparsity. Ah-symmetric generalized inverses have the key use in solving least-squares problems. Reflexive generalized inverses have low rank (same as the input matrix), and this is usually preferred in applications.

We presented a local-search algorithm that efficiently produces reflexive ah-symmetric generalized inverses. Our algorithm produces generalized inverses with guaranteed structured sparsity, with low rank (same as the input matrix), and with entries under control (by approximate 1-norm minimization). To our knowledge, no other method has been proposed in the literature with all these nice properties.

We showed that the local-search procedures presented in [Fampa and Lee, 2018; Xu et al., 2021] can be successfully implemented to construct sparse, block-structured reflexive generalized inverses with different properties. We find that the performance (1-norm achieved) is much better than tight worst-case guarantees. Overall, we find that the search procedures are very robust in terms of many of the algorithmic choices that need to be made. For scaling purposes, we found that it is necessary to be mindful of the numerics and of economizing when seeking local improvements, and calculating initial solutions efficiently proves to be a surprisingly difficult practical issue.

We have proposed algorithmic approaches to numerically analyze the trade-off between 1-norm minimization and low rank of generalized inverses, using the fact that a generalized inverse has minimum rank if and only if it satisfies the reflexive property. The algorithms start with a minimum 1-norm ah-symmetric generalized inverse and iteratively impose the reflexive property, gradually increasing 1-norm and decreasing rank, until the minimum rank is obtained. The intermediate solutions from these algorithms can represent the best trade-off between 1-norm and sparsity against rank in numerical applications.

We showed that a 2,1-norm minimizing generalized inverse satisfies two additional M-P pseudoinverse properties, including the one needed for computing least-squares solutions. We present mathematical optimization formulations related to finding row-sparse generalized inverses that can be solved very efficiently, and compare their solutions numerically to generalized inverses constructed by other methodologies, also aiming at sparsity and row-sparse structure.

As future work, we wish to extend the work in [Fampa and Lee, 2018; Xu et al., 2021] and present an approximation algorithm for the 2,1-norm. We also wish to exploit the reformulations presented in [Ponte et al., 2024] to apply Alternating Direction Method of Multipliers and Linearized Bregman algorithms for fast computation of 1- and 2,1-norm minimizing solutions.

References

- Dokmanić, I. and Gribonval, R. Beyond Moore-Penrose Part I: generalized inverses that minimize matrix norms. <http://arxiv.org/abs/1706.08349>, 2017a.
- Dokmanić, I. and Gribonval, R. Beyond Moore-Penrose Part II: the sparse pseudoinverse. <https://hal.inria.fr/hal-01547283/file/pseudo-part2.pdf>, 2017b.
- Dokmanić, I., Kolundžija, M., and Vetterli, M. (2013). Beyond Moore-Penrose: sparse pseudoinverse. In *ICASSP 2013*, pp. 6526–6530. IEEE.
- Fampa, M. and Lee, J. (2018). On sparse reflexive generalized inverses. *Operations Research Letters*, 46(6):605–610.
- Fuentes, V., Fampa, M., and Lee, J. (2020). Diving for sparse partially-reflexive generalized inverses. In Le Thi, H.A., et al., editor, *WCGO 2019*, p. 89–98.
- Fuentes, V. K., Fampa, M., and Lee, J. (2016). Sparse pseudoinverses via LP and SDP relaxations of Moore-Penrose. In *CLAIO 2016*, p. 343–350.
- Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations (3rd Ed.)*. Johns Hopkins University Press, Baltimore, MD, USA.
- Penrose, R. (1955). A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51:406–413.
- Ponte, G., Fampa, M., Lee, J., and Xu, L. (2024). On computing sparse generalized inverses. *Operations Research Letters*, 52:107058.
- Rohde, C. A. (1964). *Contributions to the theory, computation and application of generalized inverses*. PhD thesis, North Carolina State University, Raleigh. https://www4.stat.ncsu.edu/~boos/library/mimeo.archive/ISMS_1964_392.pdf.
- Xu, L., Fampa, M., Lee, J., and Ponte, G. (2021). Approximate 1-norm minimization and minimum-rank structured sparsity for various generalized inverses via local search. *SIAM Journal on Optimization*, 31(3):1722–1747.