

**Bi-objective Optimization in Role Mining**

Journal:	<i>ACM Transactions on Privacy and Security</i>
Manuscript ID	Draft
Manuscript Type:	Full-Length Research Paper
Date Submitted by the Author:	n/a
Complete List of Authors:	Crampton, Jason; Royal Holloway University of London, Mathematics Eiben, Eduard; Royal Holloway University of London Gutin, Gregory; Royal Holloway University of London, Computer Science Karapetyan, Daniel; University of Nottingham Majumdar, Diptapriyo ; Indraprastha Institute of Information Technology Delhi
Keywords:	Bi-objective optimization, Role Mining, Generalized Noise Role Mining, Fixed-Parameter Tractability

Bi-objective Optimization in Role Mining

JASON CRAMPTON, Royal Holloway University of London, United Kingdom

EDUARD EIBEN, Royal Holloway University of London, United Kingdom

GREGORY GUTIN, Royal Holloway University of London, United Kingdom

DANIEL KARAPETYAN, University of Nottingham, United Kingdom

DIPTAPRIYO MAJUMDAR, Indraprastha Institute of Information Technology Delhi, India

Role mining is a technique that is used to derive a role-based authorization policy from an existing policy. Given a set of users U , a set of permissions P and a user-permission authorization relation $UPA \subseteq U \times P$, a role mining algorithm seeks to compute a set of roles R , a user-role authorization relation $UA \subseteq U \times R$ and a permission-role authorization relation $PA \subseteq R \times P$, such that the composition of UA and PA is close (in some appropriate sense) to UPA . Role mining is therefore a core problem in the specification of role-based authorization policies. Role mining is known to be hard in general and exact solutions are often impossible to obtain, so there exists an extensive literature on variants of the role mining problem that seek to find approximate solutions and algorithms that use heuristics to find reasonable solutions efficiently.

In this paper, we first introduce the *Generalized Noise Role Mining* problem (GNRM) – a generalization of the *MINNOISE ROLE MINING* problem – which we believe has considerable practical relevance. In particular, GNRM can produce “security-aware” or “availability-aware” solutions. Extending work of Fomin et al., we show that GNRM is fixed parameter tractable, with parameter $r + k$, where r is the number of roles in the solution and k is the number of discrepancies between UPA and the relation defined by the composition of UA and PA . We further introduce a bi-objective optimization variant of GNRM, where we wish to minimize both r and k subject to upper bounds $r \leq \bar{r}$ and $k \leq \bar{k}$, where \bar{r} and \bar{k} are constants. We show that the Pareto front of this bi-objective optimization problem (BO-GNRM) can be computed in fixed-parameter tractable time with parameter $\bar{r} + \bar{k}$. From a practical perspective, a solution to BO-GNRM gives security managers the opportunity to identify a mined policy offering the best trade-off between the number of policy discrepancies and the number of roles.

We then report the results of our experimental work using the integer programming solver Gurobi to solve instances of BO-GNRM. Our key findings are that (a) we obtained strong support that Gurobi’s performance is fixed-parameter tractable, (b) our results suggest that our techniques may be useful for role mining in practice, based on our experiments in the context of three well-known real-world authorization policies. We observed that, in many cases, our solver is capable of obtaining optimal solutions when the values of either k or r are small.

CCS Concepts: • **Security and privacy** → **Access control**; **Formal security models**; • **Theory of computation** → **Parameterized complexity and exact algorithms**.

Additional Key Words and Phrases: Role Mining, Generalized Noise Role Mining, Fixed-Parameter Tractability

Authors’ addresses: Jason Crampton, jason.crampton@rhul.ac.uk, Royal Holloway University of London, Egham Hill, Egham, United Kingdom, TW200EX; Eduard Eiben, Royal Holloway University of London, Egham Hill, Egham, United Kingdom, eduard.eiben@rhul.ac.uk; Gregory Gutin, Royal Holloway University of London, Egham Hill, Egham, United Kingdom, g.gutin@rhul.ac.uk; Daniel Karapetyan, University of Nottingham, Nottingham, United Kingdom, daniel.karapetyan@nottingham.ac.uk; Diptapriyo Majumdar, Indraprastha Institute of Information Technology Delhi, Okhla Industrial Estate Phase-III, New Delhi, Delhi, India, 110020, diptapriyo@iiitd.ac.in.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

ACM Reference Format:

Jason Crampton, Eduard Eiben, Gregory Gutin, Daniel Karapetyan, and Diptapriyo Majumdar. 2024. Bi-objective Optimization in Role Mining. 1, 1 (March 2024), 21 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Role-based access control (RBAC) [19] is a mature, standardized [7] and widely deployed means of enforcing authorization requirements in a multi-user computer system.

Authorization policies, in effect, specify which interactions are authorized between users of and resources provided by a computer system. Such policies are used by the system's access control mechanism to control interactions between users and resources. RBAC policies authorize users for roles and roles for resources (usually referred to as "permissions" in the context of RBAC). RBAC can significantly reduce the administrative burden of specifying and maintaining authorization policies, provided the set of roles is small compared to the number of users and permissions.

The problem of identifying a suitable set of roles for an RBAC system has been studied extensively over the last 25 years. *Role engineering* is a top-down approach that seeks to identify roles by decomposing and analyzing business processes [18]. This approach does not generally scale well and requires substantial human effort [16]. *Role mining*, the bottom-up approach, attempts to discover a set of roles from a given authorization policy that associates users directly with permissions. More formally, the **ROLE MINING PROBLEM** is defined as follows:

ROLE MINING PROBLEM (RMP)

Input: A set of users U , a set of permissions P , a user-permission assignment relation $UPA \subseteq U \times P$, and a natural number r .

Goal: Find a set R of at most r roles, a user-role relation $UA \subseteq U \times R$, a role-permission assignment relation $PA \subseteq R \times P$ such that $(u, p) \in UPA$ if and only if there is $\rho \in R$ such that $(u, \rho) \in UA$ and $(\rho, p) \in PA$.

The value of r , for solutions that are of practical use, will be small compared to the sizes of U and P . However, it may be impossible to find a solution to RMP in which r is sufficiently small. Hence, approximate solutions are often sought, in which r is small and the composition of UA and PA is close, in some suitable sense, to UPA .

A substantial literature now exists on role mining. The problem is known to be hard in general and usually impossible to solve exactly (assuming the number of roles must be small relative to the number of users), so many approximate and heuristic techniques have been developed (see the survey paper of Mitra et al. [16]).

Recent work by Fomin et al. [8] has shown that a particular, well-known variant of the role mining problem is *fixed-parameter tractable* (FPT). Informally, this variant is NP-hard, like many role mining problems, so any exact algorithm to solve the problem is unlikely to be polynomial in the size of the problem's input, unless $P = NP$. However, there exists an algorithm (an *FPT algorithm*) whose running time is exponential in some of the input parameters, but polynomial in the others. Thus, this algorithm may well be effective if the relevant parameters are small in instances of the problem that arise in practice.

Informally, the problem considered by Fomin et al. takes a relation $UPA \subseteq U \times P$ and natural numbers r and k as input. The goal is to find a set of roles R of cardinality less than or equal to r , and relations $UA \subseteq U \times R$ and $PA \subseteq R \times P$ such that $|UPA \Delta (UA \circ PA)| \leq k$ (where $UA \circ PA$ denotes the composition of relations UA and PA and Δ denotes symmetric set difference). In other words, the composition of UA and PA has to be similar (as defined by k) to UPA . The assumption is that k and r will be small parameters. This problem has been studied by the RBAC community and is usually known as the **MINNOISE ROLE MINING PROBLEM (MNRP)** [16].

One potential problem with MNRP is that it doesn't distinguish between (a) an element that is in UPA but not in $UA \circ PA$ (which means some user is no longer authorized for some permission), and (b) an element that is in $UA \circ PA$ and not in UPA (which means that some user is now incorrectly authorized for some permission). We believe that in certain situations it will be important to insist that no additional authorizations are introduced by role mining (what we will refer to as *security-aware* role mining), while in other situations we may require that no authorizations are lost by role mining (*availability-aware* role mining).

In this paper, we introduce the GENNOISE ROLE MINING (GNRM), of which MNRP is a special case. Moreover, SECURITY-AWARE ROLE MINING and AVAILABILITY-AWARE ROLE MINING are also special cases. We extend the results of Fomin et al. by proving that GNRM is also FPT with parameter $k + r$.

Our other theoretical contribution is to introduce a bi-objective optimization version of GNRM, called BO-GNRM, where we wish to minimize both r and k subject to upper bounds \bar{r} and \bar{k} , respectively. We show that the BO-GNRM is FPT with parameter $\bar{r} + \bar{k}$. Solving BO-GNRM would allow an organization to select an appropriate solution, based on the needs to balance the number of roles against the number of deviations from the original authorization matrix. Note that in order to solve BO-GNRM we use a one-objective optimization version of GNRM called OO-GNRM.

We designed a mixed-integer formulation of OO-GNRM and built a solution method based on the Gurobi solver. We then showed that the performance of our solution method is well-aligned with the expectations for an FPT algorithm. Furthermore, we tested our method on real-world instances; in many cases, it proved the optimality of solutions for instances with small k and/or r .

The remainder of this section contains essential background material and defines GNRM and its bi-objective optimization version.

1.1 Parameterized complexity

An instance of a parameterized problem Π is a pair (I, κ) where I is the *main part* and κ is the *parameter*; the latter is usually a non-negative integer. A parameterized problem is *fixed-parameter tractable* (FPT) if there exists a computable function f such that any instance (I, κ) can be solved in time $O(f(\kappa)|I|^c)$, where $|I|$ denotes the size of I and c is an absolute constant. An algorithm to solve the problem with this running time is called an FPT algorithm. The class of all fixed-parameter tractable decision problems is called FPT. The function $f(x)$ may grow exponentially as x increases, but the running time may be acceptable if κ is small for problem instances that are of practical interest. We adopt the usual convention of omitting the polynomial factor in $O(f(\kappa)|I|^c)$ and write $O^*(f(\kappa))$ instead.

1.2 Matrix decomposition and role mining

A *Boolean matrix* is a matrix in which all entries are either 0 or 1. Let \vee and \wedge denote the usual logical operators on the set $\{0, 1\}$. We extend these operators to Boolean matrices in the natural way [14]:

- (1) the sum $\mathbf{A} \vee \mathbf{B}$ of Boolean matrices \mathbf{A} and \mathbf{B} is computed as usual with addition replaced by \vee ;
- (2) the product $\mathbf{A} \wedge \mathbf{B}$ of Boolean matrices \mathbf{A} and \mathbf{B} is computed as usual with multiplication replaced by \wedge and addition by \vee . Thus, if $\mathbf{C} = \mathbf{A} \wedge \mathbf{B}$ then

$$c_{ij} = \bigvee_{p=1}^n a_{ip} \wedge b_{pj},$$

where n is the number of columns in \mathbf{A} and the number of rows in \mathbf{B} .

Henceforth all matrices are Boolean, unless specified otherwise.

Any binary relation $X \subseteq Y \times Z$ may be represented by a matrix X with rows indexed by Y and columns indexed by Z , where $X_{ij} = 1$ iff $(i, j) \in X$. Using matrices we can reformulate the ROLE MINING PROBLEM (RMP) as follows. Given a matrix UPA and an integer r , find a matrix UA with r columns and a matrix PA with r rows such that $UPA = UA \wedge PA$. Thus, role mining may be regarded as a matrix decomposition problem.

1.3 Generalized Noise Role Mining

As previously noted, there is often no solution to RMP if r is small. In such cases, it is helpful to consider an extension of RMP called NOISE ROLE MINING [5, 8, 9, 15, 20], where the input includes a natural number k and our aim is to find a matrix UA with r columns and a matrix PA with r rows such that $d_H(UPA, UA \wedge PA) \leq k$, where $d_H(UPA, UA \wedge PA)$ is the number of entries in which UPA and $UA \wedge PA$ differ (i.e., the *Hamming distance* between them).

NOISE ROLE MINING could be seen as a rather crude approach to the problem of decomposing UPA , as it doesn't distinguish between zeroes in UPA being replaced with ones in $UA \wedge PA$ and ones being replaced with zeroes. In the first case, a user is assigned to a permission that they didn't previously have – a potential security problem. In the second case, a user no longer has a permission that they had been assigned, meaning the user may not be able to perform some of their responsibilities – an availability problem.

Thus, it will be appropriate in many cases to find UA and PA such that either security or availability, as specified by UPA is preserved. Informally, a refinement of NOISE ROLE MINING, then, would be to define AVAILABILITY-PRESERVING ROLE MINING, where we require $UPA \leq UA \wedge PA$, in the sense that every entry in UPA is less than or equal to the corresponding entry in $UA \wedge PA$. In other words, every permission authorized by UPA is also authorized by $UA \wedge PA$. Similarly, we could define SECURITY-PRESERVING ROLE MINING, where we require $UPA \geq UA \wedge PA$.

An even more fine-grained problem – the topic of this paper – is GENERALIZED NOISE ROLE MINING (GNRM), of which NOISE ROLE MINING, AVAILABILITY-PRESERVING ROLE MINING and SECURITY-PRESERVING ROLE MINING are all special cases. In GNRM we specify at the user level whether the decomposition into UA and PA is security-preserving, availability-preserving, neither, or both.

We now introduce some notation to enable us to express GNRM formally. For a positive integer t , let $[t]$ denote $\{1, 2, \dots, t\}$. Let A and B be $m \times n$ Boolean matrices and F be an $m \times n$ label matrix with entries $f_{ij} \in \{\top, \perp\}$. The matrix F is used to define a generalized distance metric between A and B . For any $(i, j) \in [m] \times [n]$ the *F-distance* from entry a_{ij} to entry b_{ij} of matrices A and B is

$$sd_F(a_{ij}, b_{ij}) = \begin{cases} \infty & f_{ij} = \perp \text{ and } a_{ij} \neq b_{ij}, \\ |a_{ij} - b_{ij}| & \text{otherwise} \end{cases}$$

In other words, if we are *not* allowed to change a_{ij} in order to obtain b_{ij} (symbol \perp) and $a_{ij} \neq b_{ij}$ then the *F-distance* from a_{ij} to b_{ij} is ∞ . Otherwise, it is just $|a_{ij} - b_{ij}|$.

We define $sd_F(A, B) = \sum_{i=1}^m \sum_{j=1}^n sd_F(a_{ij}, b_{ij})$. Thus, the *distance* from A to B is finite if and only if for every $(i, j) \in [m] \times [n]$ such that $a_{ij} \neq b_{ij}$ we have $f_{ij} = \top$.

We now define GENERALIZED NOISE ROLE MINING as a parameterized problem.

GENERALIZED NOISE ROLE MINING (GNRM)

Input: An $m \times n$ user-permission assignment matrix UPA , a label matrix \mathbf{F} , and integers $k \geq 0$ and $r \geq 1$.

Parameter: $k + r$

Goal: Is there an $m \times r$ user-role assignment matrix UA , an $r \times n$ role-permission assignment matrix PA such that $\text{sd}_F(\text{UPA}, \text{UA} \wedge \text{PA}) \leq k$? If the answer is yes, then return such matrices UA and PA .

Note that GNRM is parameterized by the sum $k + r$. This is because NOISE ROLE MINING parameterized separately by either k or r is intractable: in particular, for $k = 0$ we have EXACT ROLE MINING which is NP-hard [10]; and for $r = 1$, NOISE ROLE MINING is NP-hard [5, 9].

Note also that GNRM reduces to:

- NOISE ROLE MINING if $\mathbf{f}_{ij} = \top$ for all $(i, j) \in [m] \times [n]$;
- AVAILABILITY-PRESERVING ROLE MINING if $\mathbf{f}_{ij} = \top$ if and only if $\text{UPA}_{ij} = 0$; and
- SECURITY-PRESERVING ROLE MINING if $\mathbf{f}_{ij} = \top$ if and only if $\text{UPA}_{ij} = 1$.

1.4 Bi-objective GNRM

Note that GNRM is a decision problem, but it is clear that in practice GNRM may be viewed as an optimization problem. In such an optimization problem it is natural to minimize two objective functions r and k . It is also natural to impose upper bounds on both r and k as if at least one of them too large then the solution may well be of no practical interest. Thus, let \bar{r} and \bar{k} be upper bounds for r and k , respectively.

This leads to the following BI-OBJECTIVE GNRM problem (BO-GNRM): minimize r and k subject to $r \leq \bar{r}$ and $k \leq \bar{k}$ such that yes-answer matrices UA and PA for GNRM exist. To state BO-GNRM more formally, we will use the terminology below, which is an adaptation of multi-objective optimization terminology, see e.g. [11], for BO-GNRM. We call a pair (r', k') of integers a *feasible solution of BO-GNRM* if given UPA , \mathbf{F} and $r' (\geq 1)$ and $k' (\geq 0)$, the answer for GNRM is yes. A feasible solution (r', k') of BO-GNRM is *Pareto optimal* if there is no feasible solution (r'', k'') such that either $r'' < r'$ and $k'' \leq k'$ or $r'' \leq r'$ and $k'' < k'$. Formally, the goal of BO-GNRM is to find the *Pareto front*, which is the set of all Pareto optimal solutions.

Note that bi-objective optimization has already been used for other access control problems, see e.g. [3, 4].

Paper organization. The rest of the paper is organized in the following way. In Section 2 we describe our FPT algorithms for solving GNRM and BO-GNRM. We describe and discuss our experimental results in Sections 3, 4, and 5. Finally, in Section 6, we conclude the paper with a summary of our contributions and ideas for future work.

A preliminary version of this paper [2] has appeared in proceedings of SACMAT-2022. Sections 2.3, 3, 4 and 5 consist of new material, not published in [2].

2 FPT ALGORITHMS FOR GNRM AND BO-GNRM

Fomin et al. [8] proved that NOISE ROLE MINING parameterized by $k + r$ is FPT. We will extend this result to GNRM by reducing it to the GENERALIZED P-MATRIX APPROXIMATION problem. The reduction is similar to the one used in [8]. However, Fomin et al. [8] solved NOISE ROLE MINING as a decision problem where the aim is only to decide whether the given instance is a yes- or no-instance. In contrast, we solve GENERALIZED NOISE ROLE MINING as a problem where if the given instance is a yes-instance then a solution is also returned.

We first define GENERALIZED P-MATRIX APPROXIMATION and prove that it is FPT. We then explain how this problem is used to establish that GNRM is FPT.

2.1 Generalized P-matrix approximation

Let \mathbf{P} be a $p \times q$ matrix (sometimes called a *pattern* matrix). We say that an $m \times n$ matrix \mathbf{B} is a *P-matrix* if there is a partition $\{I_1, \dots, I_p\}$ of $[m]$ and a partition $\{J_1, \dots, J_q\}$ of $[n]$ such that for every $i \in [p]$, $j \in [q]$, $s \in I_i$, $t \in J_j$, we have $b_{st} = p_{ij}$. Note that, by definition, every set in the partitions of $[m]$ and $[n]$ is non-empty. (Thus, $p \leq m$ and $q \leq n$.) In other words, \mathbf{B} is a P-matrix if \mathbf{P} can be obtained from \mathbf{B} by first permuting rows and columns, then partitioning the resulting matrix into blocks such that in each block \mathbf{L} all entries are of the same value $v(\mathbf{L})$ and finally replacing every block \mathbf{L} by one entry of value $v(\mathbf{L})$.

For an example, let $\mathbf{P} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$. Then \mathbf{Q}_1 and \mathbf{Q}_2 below are both P-matrices: permuting columns 2 and 3 in each matrix, then partitioning (into blocks of equal size for \mathbf{Q}_1 , and between rows 1 and 2 and columns 2 and 3 for \mathbf{Q}_2) and “contracting” gives us \mathbf{P} .

$$\mathbf{Q}_1 = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \quad \mathbf{Q}_2 = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

GENERALIZED P-MATRIX APPROXIMATION

Input: An $m \times n$ matrix \mathbf{A} , a label matrix \mathbf{F} , a $p \times q$ matrix \mathbf{P} , and a nonnegative integer k .

Goal: Is there an $m \times n$ P-matrix \mathbf{B} such that $\text{sd}_{\mathbf{F}}(\mathbf{A}, \mathbf{B}) \leq k$? If the answer is yes, then return such a matrix \mathbf{B} .

Very informally, this problem asks whether there exists a matrix \mathbf{B} that is (almost) the same as \mathbf{A} and contains the rows and columns of \mathbf{P} , and, if so, returns \mathbf{B} . Fomin et al. [8] used the special case of GENERALIZED P-MATRIX APPROXIMATION, where $f_{ij} = \top$ for every $(i, j) \in [m] \times [n]$. It is called the P-MATRIX APPROXIMATION problem. We will use the following two results by Fomin et al. [8].

Observation 2.1. *Let \mathbf{P} be a $p \times q$ matrix. Then, every P-matrix \mathbf{B} has at most p pairwise distinct rows and at most q pairwise distinct columns.*

Proposition 2.2. *Given an $m \times n$ matrix \mathbf{A} and a $p \times q$ matrix \mathbf{P} , there is an algorithm that runs in time $2^{p \log p + q \log q} (nm)^{O(1)}$ and correctly outputs whether \mathbf{A} is a P-matrix.*

If \mathbf{A} has at most $p - 1$ rows or at most $q - 1$ columns, then there is no $m \times n$ matrix \mathbf{B} that is a P-matrix and $\text{sd}_{\mathbf{F}}(\mathbf{A}, \mathbf{B}) \leq k$. In that case, the instance is a no-instance. Let us now assume that \mathbf{A} has at least p rows and at least q columns.

The next lemma was proved in [8] for P-MATRIX APPROXIMATION. Note that replacing \top in $f_{ij} = \top$ by \perp for some entries f_{ij} will only reduce the set of yes-instances of GENERALIZED P-MATRIX APPROXIMATION. Thus, the next lemma follows from its special case in [8].

Lemma 2.3. *If \mathbf{A} has at least $p + k + 1$ pairwise distinct rows or at least $q + k + 1$ pairwise distinct columns, then output that $(\mathbf{A}, \mathbf{F}, \mathbf{P}, k)$ is a no-instance of GENERALIZED P-MATRIX APPROXIMATION.*

This lemma implies the following reduction/preprocessing rule.

Manuscript submitted to ACM

Reduction Rule 1. Let A be a matrix. If A has at least $p + k + 1$ pairwise distinct rows or at least $q + k + 1$ pairwise distinct columns, then output that (A, F, P, k) is a no-instance of GENERALIZED P-MATRIX APPROXIMATION.

To simplify an instance (A, F, P, k) of GENERALIZED P-MATRIX APPROXIMATION, we can apply the following reduction rule exhaustively. If a row a_i is deleted by the reduction rule, the label matrix F does not change for the other rows. This means that for the reduced instance with matrix A' , the label matrix is F restricted to the rows of A' i.e. $\{a_1, \dots, a_m\} \setminus \{a_i\}$. For simplicity of presentation, the label matrix for A' will still be denoted by F .

We define a second reduction rule that is used to delete superfluous identical rows and columns.

Reduction Rule 2. If A has at least $\max\{p, k\} + 2$ identical rows, then delete one of these identical rows. Similarly, if A has at least $\max\{q, k\} + 2$ identical columns, then delete one of these identical columns.

We say two instances of GENERALIZED P-MATRIX APPROXIMATION are *equivalent* if they are both either yes-instances or no-instances. Fomin et al. proved that any application of Reduction Rule 2 to an instance of P-MATRIX APPROXIMATION returns an equivalent instance of the problem [8, Claim 7]. It is easy to verify that the arguments in their proof of Claim 7 also apply to GENERALIZED P-MATRIX APPROXIMATION.

Applications of the two reductions rules described above either determine that the input instance is a no-instance or produce an equivalent instance with the following properties.

Lemma 2.4. Let (A, F, P, k) be an instance of GENERALIZED P-MATRIX APPROXIMATION. Then, there exists a polynomial-time algorithm that either returns “no-instance” or transforms (A, F, P, k) into an equivalent instance (A', F, P, k) of GENERALIZED P-MATRIX APPROXIMATION. Moreover the following properties are satisfied.

- (1) The matrix A' has at least p rows, at least q columns, at most $(\max\{p, k\} + 1)(p + k)$ rows and at most $(\max\{p, k\} + 1)(p + k)$ columns.
- (2) Given a P-matrix B' such that $\text{sd}_F(A', B') \leq k$, in polynomial time we can compute a P-matrix B such that $\text{sd}_F(A, B) \leq k$.

PROOF. Let (A, P, k) be an input instance of GENERALIZED P-MATRIX APPROXIMATION. As $m \geq p$ and $n \geq q$, if A has at most $p - 1$ rows or has at most $q - 1$ columns, then there is no $m \times n$ P-matrix B such that $\text{sd}_F(A, B) \leq k$. In such a case, we return “no-instance.” Next, we apply Reduction Rule 1 to check the number of pairwise distinct rows as well as the number of pairwise distinct columns in A . If A has $p + k + 1$ pairwise distinct rows or $q + k + 1$ pairwise distinct columns, then we return “no-instance.” After that, we apply Reduction Rule 2 exhaustively and let A' be the obtained matrix. We also obtain a stack S which contains all deleted rows and columns.

We return (A', F, P, k) as the output instance. Clearly, A' has at most $p + k$ pairwise distinct rows and at most $q + k$ pairwise distinct columns. Moreover, A' has at least p rows and at least q columns. Also, A' can have at most $\max\{p, k\} + 1$ pairwise identical rows and at most $\max\{q, k\} + 1$ pairwise identical columns. This means that A' has at most $(\max\{p, k\} + 1)(p + k)$ rows and at most $(\max\{q, k\} + 1)(q + k)$ columns. This completes the proof that property (1) holds.

Suppose that B' is a P-matrix such that $\text{sd}_F(A', B') \leq k$. Note that at any intermediate stage, when a row r (a column c , respectively) was deleted from A , there were at least $(\max\{p, k\} + 1)$ additional rows identical to r (at least $(\max\{q, k\} + 1)$ additional columns identical to c , respectively). Hence, in A' , if a row r or a column c was deleted by Reduction Rule 2, then there are exactly $\max\{p, k\} + 1$ rows identical to r (exactly $\max\{q, k\} + 1$ columns identical to c , respectively). Since $\text{sd}_F(A', B') \leq k$, at most k entries with label \top identical to r (c , respectively) were modified in B' .

Thus, \mathbf{B}' must have at least one row identical to \mathbf{r} (at least one column identical to \mathbf{c} , respectively) if \mathbf{r} (\mathbf{c} , respectively) was deleted by Reduction Rule 2. Therefore, the deleted rows (columns, respectively) are identical to some rows (columns, respectively) which are the same in \mathbf{A}' and \mathbf{B}' . Thus, reinstating the deleted rows and columns using stack S , we obtain matrices \mathbf{A} and \mathbf{B} such that \mathbf{B} is a \mathbf{P} -matrix and $\text{sdf}(\mathbf{A}, \mathbf{B}) \leq k$. \square

Theorem 2.5. GENERALIZED \mathbf{P} -MATRIX APPROXIMATION can be solved in time $2^{p \log p + q \log q} (nm)^{O(1)} ((\max\{p, k\} + 1)(p + k)(\max\{q, k\} + 1)(q + k))^k$.

PROOF. Let $(\mathbf{A}, \mathbf{F}, \mathbf{P}, k)$ be an instance of GENERALIZED \mathbf{P} -MATRIX APPROXIMATION. First, we invoke the polynomial-time algorithm of Lemma 2.4 to either determine that the input instance is a no-instance or generate an instance $(\mathbf{A}', \mathbf{F}, \mathbf{P}, k)$ satisfying properties (1) and (2). Recall that the first property says that \mathbf{A}' has at most $(\max\{p, k\} + 1)(p + k)$ rows, and at most $(\max\{q, k\} + 1)(q + k)$ columns. This means that \mathbf{A}' has at most $(\max\{p, k\} + 1)(p + k)(\max\{q, k\} + 1)(q + k)$ entries. We then consider all possible sets of at most k entries. For every entry of such a set, if the label of an entry is \top , we will modify it. This results in a modified matrix \mathbf{B}' . We then invoke Proposition 2.2 to check whether \mathbf{B}' is a \mathbf{P} -matrix or not. This checking takes $2^{p \log p + q \log q} (nm)^{O(1)}$ -time. If \mathbf{B}' is a \mathbf{P} -matrix, it is a solution to GENERALIZED \mathbf{P} -MATRIX APPROXIMATION for the instance $(\mathbf{A}', \mathbf{F}, \mathbf{B}', k)$ as $\text{sdf}(\mathbf{A}', \mathbf{B}') \leq k$ (since we changed at most k entries in \mathbf{A}'). Then, we make use of property (2) to construct \mathbf{B} satisfying $\text{sdf}(\mathbf{A}, \mathbf{B}) \leq k$ and return \mathbf{B} as a solution of GENERALIZED \mathbf{P} -MATRIX APPROXIMATION for the instance $(\mathbf{A}, \mathbf{F}, \mathbf{B}, k)$. Recall that this step takes polynomial time. Hence, the overall algorithm takes $2^{p \log p + q \log q} (nm)^{O(1)} ((\max\{p, k\} + 1)(p + k)(\max\{q, k\} + 1)(q + k))^k$ time. \square

2.2 GNRM is FPT

We now explain how the algorithm for GENERALIZED \mathbf{P} -MATRIX APPROXIMATION is used to solve GNRM and thus show it is FPT. The basic strategy is to consider all possible pairs of matrices whose product \mathbf{P} could provide the basis for a solution to GNRM. The number of such pairs is bounded above by a function of r . For each such \mathbf{P} , we determine whether the GENERALIZED \mathbf{P} -MATRIX APPROXIMATION instance $(\mathbf{UPA}, \mathbf{F}, \mathbf{P}, k)$ has a solution, in which case we can then compute a solution to the GNRM instance.

Lemma 2.6. Let \mathbf{P} be a $p \times q$ matrix such that $\mathbf{P} = \mathbf{X} \wedge \mathbf{Y}$ for a $p \times r$ matrix \mathbf{X} and an $r \times q$ matrix \mathbf{Y} . Furthermore, consider an $m \times n$ matrix \mathbf{B} which is a \mathbf{P} -matrix. Then, we can in polynomial time obtain an $m \times r$ matrix \mathbf{X}^* , and $r \times n$ matrix \mathbf{Y}^* such that $\mathbf{B} = \mathbf{X}^* \wedge \mathbf{Y}^*$.

PROOF. As \mathbf{B} is a \mathbf{P} -matrix, there are partitions $\{I_1, \dots, I_p\}$ of $[m]$ and $\{J_1, \dots, J_q\}$ of $[n]$ such that for every $i \in [p]$, $j \in [q]$, $s \in I_i$, $t \in J_j$, $\mathbf{b}_{st} = \mathbf{p}_{ij}$.

We initialize $\mathbf{X}^* = \mathbf{X}$ and $\mathbf{Y}^* = \mathbf{Y}$. Consider an entry \mathbf{p}_{ij} . Let \mathbf{x}_i be the i 'th row of \mathbf{X} and \mathbf{y}^j the j 'th column of \mathbf{Y} ; then $\mathbf{x}_i \wedge \mathbf{y}^j = \mathbf{p}_{ij}$. Let $c \in [p]$ and $d \in [q]$ such that $i \in I_c$ and $j \in J_d$. Then, for any $s \in I_c$ and $t \in J_d$, set $\mathbf{b}_{st} = \mathbf{p}_{ij}$. Then, for any $s \in I_c$ and for any $t \in J_d$, we insert \mathbf{x}_i as the s 'th row of \mathbf{X}^* and \mathbf{y}^j as the t 'th column of \mathbf{Y}^* . \square

The *Boolean rank* of a matrix \mathbf{A} , denoted $\text{BRank}(\mathbf{A})$, is the minimum natural number r such that $\mathbf{A} = \mathbf{B} \wedge \mathbf{C}$, where \mathbf{B} and \mathbf{C} are matrices such that the number of columns in \mathbf{B} and the number of rows in \mathbf{C} is r . Thus, a matrix \mathbf{A} has Boolean rank 1 if and only if $\mathbf{A} = \mathbf{x} \wedge \mathbf{y}^T$ for some column-vectors \mathbf{x} and \mathbf{y} . In fact, $\text{BRank}(\mathbf{A}) = r$ if and only if r is the minimum natural number such that $\mathbf{A} = \mathbf{X}^{(1)} \vee \dots \vee \mathbf{X}^{(r)}$, where matrices $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(r)}$ are of Boolean rank 1 [14].

Theorem 2.7. GENERALIZED NOISE ROLE MINING admits an $O^*(2^{O(r2^r + rk)})$ -time algorithm.

PROOF. Let \mathbf{B} be an $m \times n$ -matrix and let r be the Boolean rank of \mathbf{B} . Thus, there are r matrices $\mathbf{B}^{(1)}, \dots, \mathbf{B}^{(r)}$, each of Boolean rank 1, such that $\mathbf{B} = \mathbf{B}^{(1)} \vee \dots \vee \mathbf{B}^{(r)}$, where for each $i \in [r]$, $\mathbf{B}^{(i)} = \mathbf{x}^i \wedge (\mathbf{y}^i)^T$ for some column-vectors \mathbf{x}^i and \mathbf{y}^i . It can be shown by induction on r that \mathbf{B} has at most 2^r distinct rows and at most 2^r distinct columns.¹ Therefore, \mathbf{B} is of Boolean rank at most r if and only if there is a $p \times q$ matrix \mathbf{P} of Boolean rank at most r for $p = \min\{2^r, m\}$ and $q = \min\{2^r, n\}$ such as \mathbf{B} is a \mathbf{P} -matrix.

Moreover, an $m \times n$ -matrix \mathbf{B} is of rank r if r is the minimum natural number such that $\mathbf{B} = \mathbf{C} \wedge \mathbf{D}$, where \mathbf{C} is an $m \times r$ -matrix and \mathbf{D} is an $r \times n$ -matrix. Hence, GENERALIZED NOISE ROLE MINING can be reformulated as follows: Decide whether there is a $p \times q$ -pattern matrix \mathbf{P} of Boolean rank r and an $m \times n$ \mathbf{P} -matrix \mathbf{B} such that $\text{sd}_F(\text{UPA}, \mathbf{B}) \leq k$ and if \mathbf{B} does exist then find matrices \mathbf{UA} and \mathbf{PA} of sizes $m \times r$ and $r \times n$, respectively, such that $\mathbf{B} = \mathbf{UA} \wedge \mathbf{PA}$.

Thus, to solve GENERALIZED NOISE ROLE MINING with input $(\text{UPA}, \mathbf{F}, k)$, we can use the following algorithm:

1. Generate all pairs (\mathbf{X}, \mathbf{Y}) of matrices of sizes $p \times r$ and $r \times q$, respectively, and for each such pair compute $\mathbf{P} = \mathbf{X} \wedge \mathbf{Y}$;
2. For each \mathbf{P} , solve GENERALIZED \mathbf{P} -MATRIX APPROXIMATION for the instance $(\text{UPA}, \mathbf{F}, \mathbf{P}, k)$. If $(\text{UPA}, \mathbf{F}, \mathbf{P}, k)$ is a yes-instance, then using the algorithm of Lemma 2.6 return matrices \mathbf{UA} and \mathbf{PA} of sizes $m \times r$ and $r \times n$ such that $\mathbf{B} = \mathbf{UA} \wedge \mathbf{PA}$, where \mathbf{B} is the solution of the instance $(\text{UPA}, \mathbf{F}, \mathbf{P}, k)$;
3. If all instances above are no-instances of GENERALIZED \mathbf{P} -MATRIX APPROXIMATION, return “no-instance.”

It remains to evaluate the running time of the above algorithm. Since $p \leq 2^r$ and $q \leq 2^r$, there are at most $2^{O(r2^r)}$ pairs (\mathbf{X}, \mathbf{Y}) , and we can compute all matrices \mathbf{P} in time $2^{O(r2^r)}$. Thus, the running time of the algorithm is dominated by that of Step 2. The running time of Step 2 is upper bounded by the number of matrices \mathbf{P} (it is equal to $2^{O(r2^r)}$) times the maximum running time of solving GENERALIZED \mathbf{P} -MATRIX APPROXIMATION on an instance $(\text{UPA}, \mathbf{F}, \mathbf{P}, k)$ and computing \mathbf{UA} and \mathbf{PA} , if the instance is a yes-instance. By Lemma 2.6, Theorem 2.5 and the bounds $p \leq 2^r, q \leq 2^r$, the maximum running time is upper bounded by $O^*(2^{O(r2^r+rk)})$. It remains to observe that $2^{O(r2^r)} \cdot O^*(2^{O(r2^r+rk)}) = O^*(2^{O(r2^r+rk)})$. \square

2.3 Solving BO-GNRM

To design an algorithm for computing the Pareto front of BO-GNRM, let us consider a related problem, the ONE-OBJECTIVE GNRM problem (OO-GNRM): compute the minimum value $k_{\min}(r)$ of k for every $r \in [\bar{r}]$ such that $k_{\min}(r) \leq \bar{k}$. In other words, given r , $k_{\min}(r)$ is the smallest number of discrepancies for a solution containing r roles. OO-GNRM can be easily solved by running the $O^*(2^{O(\bar{r}2^{\bar{r}}+\bar{r}\bar{k})})$ -time algorithm of Theorem 2.7 for $k \in \{0, 1, \dots, \bar{k}\}$.² (Note that if $k_{\min}(r) > \bar{k}$ then there is no solution to BO-GNRM for that value of r .) This allows us to compute the Pareto front \hat{P} of BO-GNRM as follows:

$$\hat{P} = \{(r, k_{\min}(r)) : r \in [\bar{r}], 0 \leq k_{\min}(r) \leq \bar{k}, k_{\min}(r) < k_{\min}(r-1) \text{ if } r \geq 2\} \quad (1)$$

Figure 1 illustrates the notions introduced above for $\bar{k} = 6$ and $\bar{r} = 11$. In particular, there is no solution such that $r = 1$ and k is less than the maximum value allowed (6 in this case), and no solution for $r = 2$. In contrast, we can find solutions for $r \in \{3, 4, 5\}$ with $k_{\min}(r) = 5$. Hence $(3, 5)$ belongs to the Pareto front. Similarly $(6, 3)$, $(8, 2)$ and $(9, 0)$ belong to the Pareto front. An organization can decide which point on the Pareto front is preferable: for example,

¹For $r = 1$, since $\mathbf{B}^{(1)} = \mathbf{x}^1 \wedge (\mathbf{y}^1)^T$, $\mathbf{B}^{(1)}$ has at most two distinct rows, \mathbf{y}^1 and the all-zero row, and has at most two distinct columns, \mathbf{x}^1 and the all-zero column. Now let $r \geq 2$. By induction hypothesis, $\mathbf{B} = \mathbf{B}^{(\leq r-1)} \vee \mathbf{B}^{(r)}$, where $\mathbf{B}^{(\leq r-1)}$ has at most 2^{r-1} rows and columns and $\mathbf{B}^{(r)}$ at most two rows and columns. Since every row (column, respectively) of \mathbf{B} is the disjunction of the corresponding rows (columns, respectively) of $\mathbf{B}^{(\leq r-1)}$ and $\mathbf{B}^{(r)}$, the number of distinct rows (columns, respectively) in \mathbf{B} is at most $2^{r-1} \cdot 2 = 2^r$.

²We could introduce a different one-objective optimization version of GNRM, where we minimize r . Our choice of minimizing k is explained in Section 3.1.

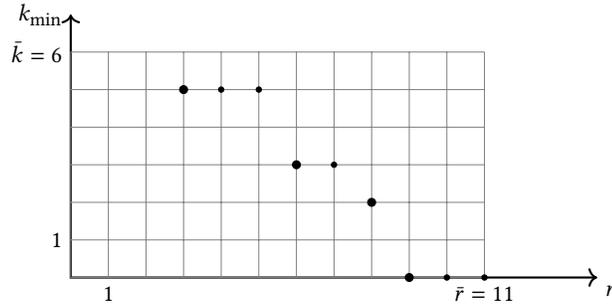


Fig. 1. Solutions of OO-GNRM and BO-GNRM, where all the circles are solutions of OO-GNRM and all the large circles form \hat{P} . Note that points $(1, k_{\min}(1))$ and $(2, k_{\min}(2))$ are not depicted since $k_{\min}(2) > \bar{k}$.

having no discrepancies between the original policy and the mined policy using nine roles versus only six roles but accepting three discrepancies.

By the arguments above, we have the following:

Theorem 2.8. *There is an $O^*(2^{O(\bar{r}2^{\bar{r}} + \bar{r}\bar{k})})$ -time algorithm for constructing the Pareto front of BO-GNRM. Thus, BO-GNRM is FPT with parameter $\bar{r} + \bar{k}$.*

3 GNRM SOLVER

To solve the BO-GNRM problem, we use a general-purpose solver. We could have implemented a bespoke FPT algorithm to solve BO-GNRM, based on our results in the preceding sections. However, we believe using a general-purpose solver is likely to be more useful in practice. First, the formulation of the problem as an integer program is quite intuitive, and therefore easier to understand and maintain than a bespoke algorithm that relies on some relatively complex theory. Second, general-purpose solvers may perform well on instances of a hard problem that is known to be FPT [13]. For example, an intelligent general-purpose solver might be able to automatically identify and apply reductions during the pre-solve process leading to an FPT-like behaviour.

In this section, we describe our approach to solving BO-GNRM, the integer programming formulation of OO-GNRM, and in Section 4 we confirm that the empirical behaviour of this solver is consistent with that expected of an FPT algorithm. We then apply our solver in Section 5 to real-world instances to solve BO-GNRM.

3.1 The choice of OO-GNRM

Our approach to solving the BO-GNRM is to decompose it into multiple OO-GNRM instances. We have two options: (i) find $k_{\min}(r)$ for each r , or (ii) find $r_{\min}(k)$ for each k . (Here, $r_{\min}(k)$ is the smallest value of r that makes the instance of GNRM satisfiable for a given k .)

Note that the solution size (the number of decision variables in the solution representation) depends on r but does not depend on k . This makes it easier technically and computationally to fix r and minimise k . Also note that the value of $k_{\min}(r) \in O(mn)$ whereas $r_{\min}(k) \in O(m)$. Indeed, we observed in our experiments that $k_{\min}(r)$ can reach large values for small r whereas $r_{\min}(k)$ is relatively small even for $k = 0$. Considering the above observations, we chose approach (i), i.e. our solver takes r as a parameter and searches for $k_{\min}(r)$. Then, to obtain the Pareto front for the BO-GNRM, we use formula (1).

3.2 Formulation of GNRM

We present a mixed integer programming formulation of the OO-GNRM in order to solve the problem using a general-purpose solver. Our formulation consists of two matrices of Boolean variables $ua_{i,\ell}$, $i \in [m]$, $\ell \in [r]$, and $pa_{\ell,j}$, $\ell \in [r]$, $j \in [n]$. We also use a matrix of auxiliary Boolean variables $d_{i,j}$, $i \in [m]$, $j \in [n]$, representing the discrepancies between $UA \wedge PA$ and UPA . We use the values 0 and 1 to represent the values \perp and \top , respectively, in the F matrix. Figure 2 describes the formulation in detail.

A naïve formulation of the problem would include a matrix of Boolean variables $_$ to represent the product $UA \wedge PA$ and link these variables to the $ua_{i,\ell}$ and $pa_{\ell,j}$ variables. Then it would be easy to formulate the constraints and link the decision variables to variables $d_{i,j}$. However, knowing the values of $upa_{i,j}$ and $f_{i,j}$ for a specific pair (i, j) , we can formulate the constraints more compactly. Thus, the formulation described in Figure 2 defines the constraints separately for each combination of values of $upa_{i,j}$ and $f_{i,j}$.³

Minimise $\sum_{i \in [m]} \sum_{j \in [n]} d_{i,j}$,	(2)	The objective is to minimise the discrepancies.
For all $i \in [m]$ and $j \in [n]$ such that $f_{i,j} = 0$ and $upa_{i,j} = 0$:		Since the discrepancy is not allowed, either $ua_{i,\ell}$ or $pa_{\ell,j}$ has to be zero for every ℓ .
$ua_{i,\ell} + pa_{\ell,j} \leq 1 \quad \forall \ell \in [r]$,	(3)	
For all $i \in [m]$ and $j \in [n]$ such that $f_{i,j} = 0$ and $upa_{i,j} = 1$:		We need $ua_{i,\ell} = pa_{\ell,j} = 1$ for some ℓ . We enforce that at least one of $x_{i,j,1} \dots x_{i,j,r}$ is 1 and also if $x_{i,j,\ell} = 1$ then $ua_{i,\ell} = pa_{\ell,j} = 1$.
$x_{i,j,\ell} \leq ua_{i,\ell} \quad \forall \ell \in [r]$,	(4)	
$x_{i,j,\ell} \leq pa_{\ell,j} \quad \forall \ell \in [r]$,	(5)	
$\sum_{\ell \in [r]} x_{i,j,\ell} \geq 1 \quad \forall \ell \in [r]$,	(6)	
For all $i \in [m]$ and $j \in [n]$ such that $f_{i,j} = 1$ and $upa_{i,j} = 0$:		If both $ua_{i,\ell}$ and $pa_{\ell,j}$ are ones for some ℓ then this is a discrepancy.
$ua_{i,\ell} + pa_{\ell,j} \leq 1 + d_{i,j} \quad \forall \ell \in [r]$,	(7)	
For all $i \in [m]$ and $j \in [n]$ such that $f_{i,j} = 1$ and $upa_{i,j} = 1$:		If either $ua_{i,\ell} = 0$ or $pa_{\ell,j} = 0$ for every ℓ , this is a discrepancy. Auxiliary variable $x_{i,j,\ell}$ is forced to 0 if either $ua_{i,\ell} = 0$ or $pa_{\ell,j} = 0$. If $x_{i,j,\ell} = 0$ for every ℓ then we force $d_{i,j} = 1$.
$ua_{i,\ell} \geq x_{i,j,\ell} \quad \forall \ell \in [r]$,	(8)	
$pa_{\ell,j} \geq x_{i,j,\ell} \quad \forall \ell \in [r]$,	(9)	
$\sum_{\ell \in [r]} x_{i,j,\ell} \geq 1 - d_{i,j} \quad \forall \ell \in [r]$,	(10)	
$x_{i,j,\ell} \in \{0, 1\}$	$\forall i \in [m], \forall j \in [n], \forall \ell \in [r]$,	(11) Auxiliary variables, see the cases where $upa_{i,j} = 1$.
$d_{i,j} \in \{0, 1\}$	$\forall i \in [m], \forall j \in [n]$,	(12) Indicates whether there is a discrepancy in the corresponding element.
$ua_{i,j} \in \{0, 1\}$	$\forall i \in [m], \forall j \in [r]$,	(13) Defines the UA matrix.
$pa_{i,j} \in \{0, 1\}$	$\forall i \in [r], \forall j \in [n]$,	(14) Defines the PA matrix.

Fig. 2. CSP formulation of GNRM.

³The compact formulation performed significantly better than the naïve formulation in our experiments, so we only report results for the compact configuration.

3.3 The choice of the solver

For our conference paper, we used the CP-SAT solver from Google OR-Tools to solve the GNRM problem (the decision version of BO-GNRM). For this paper, we considered three options:

- (1) Solving GNRM using CP-SAT to identify the minimum value of k that makes the instance satisfiable (https://developers.google.com/optimization/cp/cp_solver);
- (2) Solving OO-GNRM using the linear optimisation solver from Google OR-Tools (<https://developers.google.com/optimization/lp>); and
- (3) Solving OO-GNRM using the Gurobi mixed integer programming solver (<https://www.gurobi.com/solutions/gurobi-optimizer>).

Following experimentation, we concluded that Gurobi is considerably faster than the other approaches. For example, for an instance of size $m \times n = 25 \times 25$ and $r = 5$ (with $k_{\min}(5) = 22$), the Gurobi solver was 86 times faster than the CP-SAT-solver-based approach and 50 times faster than the linear optimisation solver from OR-Tools. Thus, all the reported experiments in this paper are conducted with Gurobi. We also attempted a few modifications of the formulation in Figure 2. Specifically, we tried several approaches to symmetry breaking as well as adding simple custom cuts. None of these changes improved the performance though; we assume that the internal mechanisms of Gurobi are intelligent enough to identify all the simple properties of our formulation and exploit them effectively. We also found the default parameter values of Gurobi to be effective. Apart from the time limit, the only Gurobi parameter that we adjusted in some experiments was *MIPFocus*; we set it to 1 to intensify the search for feasible solutions when the solver was used as a heuristic.

4 DOES THE GUROBI-BASED SOLVER HAVE FPT-LIKE RUNTIME?

In this section, we test the hypothesis that the Gurobi-based solver is capable of exploiting the FPT structure of GNRM. Since our solver addresses the OO-GNRM, we focus on testing if its running time is FPT-like with respect to $k_{\min} = k_{\min}(r)$ when r is fixed.

We say that a solver has FPT-like running time if its empirical running time scales polynomially with the size of the problem instance. As we talk about empirical running time, we focus on ‘typical’ instances rather than the worst case. However, the concept of a ‘typical’ instance is vague and brings difficulties to the experimental set-up. In the rest of this section, we introduce a new methodology to identify the scaling of the empirical running time of an algorithm and use it to show that the Gurobi-based solver has FPT-like running time.

Note that using real-world instances in such a study is generally impractical for the following reasons:

- Such a study requires a large number of instances whereas real-world benchmark sets are usually very limited;
- To draw conclusions about the scaling behaviour of a solver, we need instances with a wide range of parameters which might not be present in real-world benchmark sets; and
- We need to ensure that the instances have consistent difficulty.

The first two points can easily be addressed by using synthetic instances produced by a pseudo-random instance generator. The third point, however, remains a challenge even if we use a pseudo-random instance generator. Indeed, changes in some parameters of the generator such as the instance size may affect the hardness of the instance. In decision problems, this behaviour can often be explained by shifting between the under- and over-subscribed instance regions. Then, to keep the hardness of the instances consistent, it is necessary to adjust some other generator parameters. This approach was used in several other studies, e.g. [1, 12, 13] focusing on the phase-transition region (the region

between the under- and over-subscribed instances). However, we could not use this approach in this study as phase transition is undefined for optimisation problems.

Instead, our approach in this study relies on the median running times across wide ranges of diverse instances of similar size and value of the parameter. Thus, we measure how the median running time scales with the instance size and parameter value.

4.1 Pseudo-random instance generator

We adopted a pseudo-random instance generator used in earlier experimental work by Vaidya et al. [21]. The generator takes an integer $r_0 > 0$ as a parameter and produces an instance of $\text{UPA} \in \mathbb{R}^{m \times n}$ by creating random $\text{UA} \in \mathbb{R}^{m \times r_0}$ and $\text{PA} \in \mathbb{R}^{r_0 \times n}$ matrices and multiplying them together: $\text{UPA} = \text{UA} \wedge \text{PA}$. This means that we know an upper bound on the number of roles we need to mine. We used the following settings in our generator:

- the number of roles per user was randomly chosen for each user from the interval $[0, r_u]$, where r_u is an instance generator parameter; and
- the number of permissions per role was randomly chosen for each role from $[0, \lfloor 0.25n \rfloor]$.

Thus, the parameters of our generator are as follows:

- m is the number of users in the instance;
- n is the number of permissions in the instance;
- r_0 is the overall number of roles in the UA matrix used to produce the UPA matrix.
- r_u is the maximum number of roles per user; $r_u \leq r_0$; and
- r is the number of roles that need to be mined; this value does not affect the UPA matrix but is used by the solver.

4.2 Data collection process

All our computations were performed on a machine based on two Xeon E5-2630 v2 CPUs (2.60 GHz), with 32 GB of RAM. We used Gurobi 10.0. It was restricted to one thread for the experiments in Section 4, with up to 12 experiments running in parallel, whereas the number of threads was unrestricted for the experiments in Section 5 but only one experiment was conducted at a time.

In order to prove that the Gurobi-based solver has FPT-like running time, we would need to generate instances of various sizes but with fixed parameters r and k_{\min} , and plot the running time against the instance size. It was easy to fix the value of r ; we set it to 5 for all the experiments in this section. However, k_{\min} is actually not a parameter of the instance generator; it is the objective value. Thus, we developed the following methodology.

First, we generated a large set of instances for a wide range of instance generator parameters. By solving each instance, we obtained the value of k_{\min} for that instance, which then allowed us to select instances based on their values of k_{\min} . For example, we could see how the running time scaled with the size of the instance for instances with $k_{\min} = 12$.

This approach, however, has a fundamental issue. Since we used a wide range of instance parameters, it was inevitable that some of the instances were prohibitively hard. We could put a time limit on the solver but that would skew the results; the harder instances would not be represented in our set. Our workaround was to add a constraint to the solver restricting k_{\min} to values up to 30; this made any instances with k_{\min} above 30 infeasible and so we disregarded them. Note that this did not skew the results as the set of instances for each $k_{\min} \leq 30$ was unrestricted.

As we expected that the hardness of an instance mainly depends on r and k_{\min} but not as much on the other instance generator parameters, restricting the values of r and k_{\min} was sufficient to avoid overly hard instances thus making data collection feasible.

As we mentioned above, the value of r was fixed to 5. The values of the other instance generator parameters were randomly sampled from the following ranges: $n, m \in \{10, 11, \dots, 70\}$, $r_0 \in \{5, 6, \dots, 15\}$, and $r_u \in \{1, 2, \dots, 5\}$. The selection of the parameters was based on typical values used in the literature and to ensure that the *authorisation density* (the proportion of non-zero entries in the UPA matrix) was mainly in the range 5–35%. (All but one of the real-world datasets commonly used in role mining research have authorisation densities less than 20% [17, Table 1].)

We set $f_{i,j} = upa_{i,j}$ to prioritise the security considerations; other settings of $f_{i,j}$ are studied in Section 5.2.

At most one instance per combination of parameters was generated. In total, we produced 110 697 instances of which 56 220 had $k_{\min} \leq 30$ and, thus, were included in this study. For each instance, we recorded the running time of the solver, thus our dataset included the instance parameters, the value of k_{\min} and the solver running time. In Section 4.3, we describe how this dataset was used to build a model of the solver running time. This in turn was used to support our hypothesis that the Gurobi-based solver has FPT-like running time.

4.3 Running time model

Our hypothesis is that the median running time of the Gurobi-based solver can be approximated as a product of two functions: a function of k_{\min} only and a function of the instance size only (as we fixed r in this experiment, it is not a part of the model). If the second function is a polynomial then we can claim that the Gurobi-based solver has FPT-like running time, meaning that it is suitable for reasonably large instances as long as the value of k_{\min} is small.

We came up with two candidates for the definition of the instance size:

- The size mn of the UPA matrix; and
- The number of non-zero elements in the UPA matrix, which we denote by σ .

There are arguments for each of these candidate definitions; indeed, the size of the MIP formulation in Figure 2 depends, in various ways, on both mn and σ . Thus, we tested both definitions when fitting running time models. In our experience, the second definition gives a much better fit, hence we use σ as the definition of the instance size in the rest of the paper.

Let $f(k_{\min}, \sigma)$ be our model of the median running time of the Gurobi-based solver. While it is possible to fit $f(k_{\min}, \sigma)$ to all the data points, the result would be affected by the imbalances in our set of instances; for example, the lower values of k_{\min} are represented better in our dataset, and this would be reflected in the model. Thus, we designed the following process to balance the dataset. For each $k_{\min} \in \{0, 1, \dots, 30\}$ and $j \in \{1, 2, \dots\}$, we calculated the median running time $t_{k_{\min}, j}$ over all the instances of size σ such that $100(j-1) \leq \sigma < 100j$ and with the given value of k . We then fit $f(k_{\min}, \sigma)$ to $(k_{\min}, 100j - 50, t_{k_{\min}, j})$ for all the combinations of k_{\min} and j where the number of instances within the corresponding range is at least 10.

The running times across our experiments vary from milliseconds to minutes; a residual of a few seconds is acceptable for larger and harder instances but is a poor approximation for small and easy instances. In other words, when fitting $f(k_{\min}, \sigma)$, we are concerned with the relative error, not the absolute error. Thus, we use the logarithmic scale for the running times, i.e. we fit $\log f(k_{\min}, \sigma)$ to $\log t_{k_{\min}, j}$.

According to our initial assumption, the model was a product of two functions: $f(k_{\min}, \sigma) = f_k(k_{\min}) \cdot f_\sigma(\sigma)$. By visualising the experimental data, we established that $f_k(k_{\min})$ was exponential in k_{\min} as can be clearly seen in Figure 5,

however the slope of the curves was different for different values of σ . In other words, f_k turned out to be a function of both k_{\min} and σ . By estimating the slopes, we established that a good model for f_k was $f_k(k_{\min}, \sigma) = c_1 \cdot 2^{k_{\min}/(c_2\sigma+c_3)}$.

A reasonable fit for $f_\sigma(\sigma)$ could be a parabola or an exponential function with the base close to 1. We tested both hypotheses and found that a quadratic function fits significantly better than exponential.

Thus, our final model is as follows:

$$f(k_{\min}, \sigma) = c_1 \cdot 2^{k_{\min}/(c_2\sigma+c_3)} \cdot (\sigma^2 + c_4\sigma + c_5), \quad (15)$$

where c_1, c_2, \dots, c_5 are coefficients. This model gives a close fit for the experimental data except for small values of k_{\min} . Thus, we removed $k_{\min} < 2$ when fitting the data and obtained the following values of the coefficients: $c_1 \approx 1.11 \cdot 10^{-4}$, $c_2 \approx 3.29 \cdot 10^{-2}$, $c_3 \approx 4.14$, $c_4 \approx -77.3$ and $c_5 \approx 7.23 \cdot 10^3$.

Figure 3 shows the fit of $f(k_{\min}, \sigma)$ in three dimensions, whereas Figures 4 and 5 show slices through the space of k_{\min} and σ . Except for small k_{\min} , the model accurately predicts the aggregated running times. It even predicts that, for small σ and large k_{\min} , the running time slightly decreases as σ increases, see Figure 4.

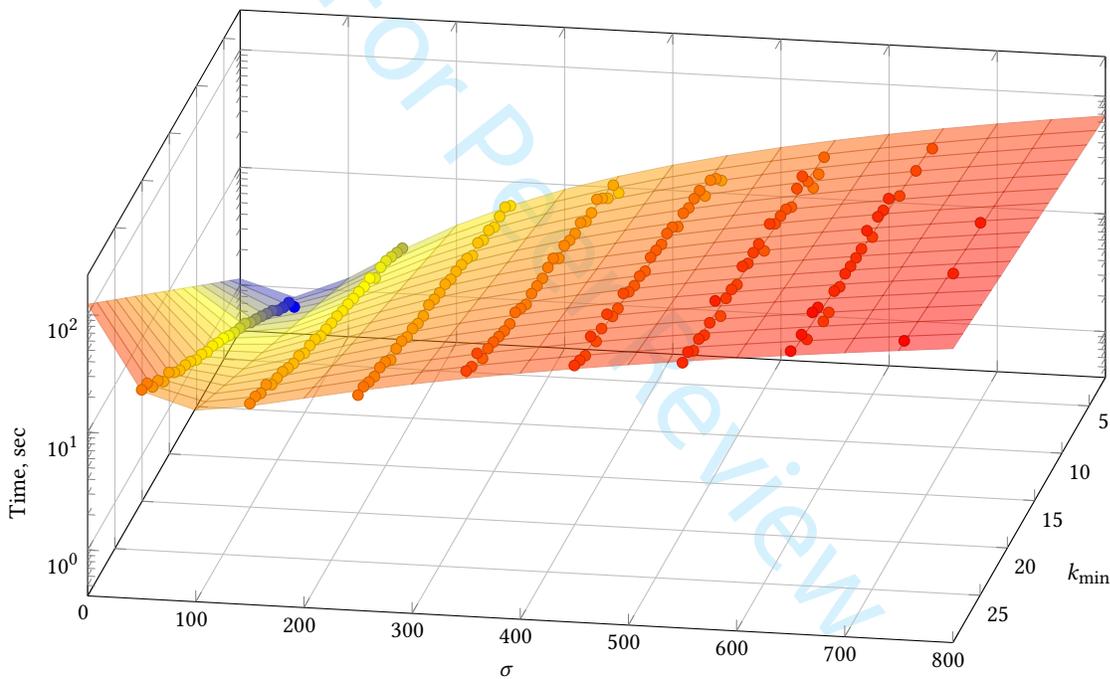


Fig. 3. This graph demonstrates how our model $f(k_{\min}, \sigma)$ (surface) fits the data aggregated into $t_{k_{\min}, j}$ (scatter plot). The colour represents the time (the value along the vertical axis).

So far, we focused on the median running times. To see how far the actual running times fall from the predicted ones, we produced one more visualisation, see Figure 6. It is designed to demonstrate that our $f_\sigma(\sigma)$ is a good fit for the data across all values of $k_{\min} \geq 2$. To compensate for the differences caused by the variation of k_{\min} , we divided the running times by $f_k(k_{\min}, \sigma)$. One can see that all the instances fall into a relatively narrow interval around our model $f_\sigma(\sigma)$, except for small σ . Similarly to the case of small k_{\min} , we assume that small σ may significantly change the behaviour of Gurobi.

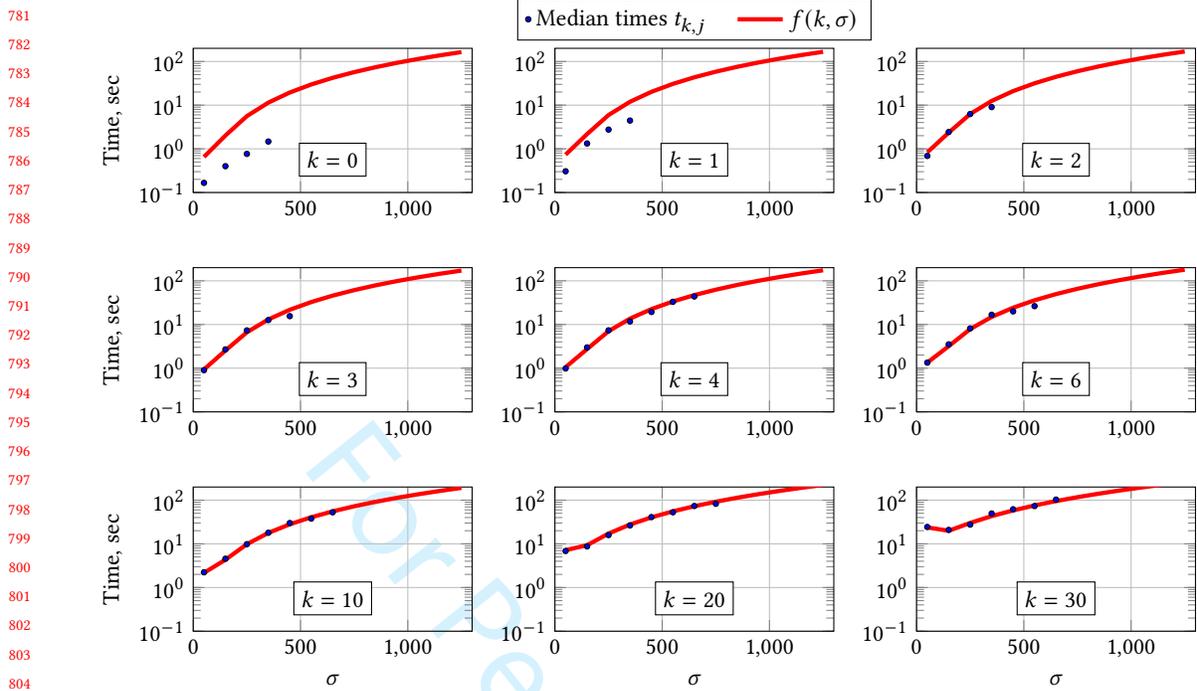


Fig. 4. The aggregated data $t_{k,j}$ and the best fit model $f(k, \sigma)$ sliced along the k axis.

Nevertheless, it is clear from this visualisation that a good model for $f_\sigma(\sigma)$ has to be a polynomial; an exponential dependence would be a straight line in this semi-logarithmic plot. Our model accurately approximates the data for all σ and k_{\min} except for several smallest values.

4.4 Conclusions about the FPT-like behaviour

We conclude that (15) is a good fit to the experimental data for $k_{\min} \geq 2$. The instances for $k_{\min} < 2$ are being solved notably faster than our model predicts. We hypothesise that this behaviour is linked to the use of heuristics within the Gurobi solver, however further investigation is needed to confirm this.

An algorithm is formally called FPT if its worst-case time complexity is $O(f_k(k_{\min}) \cdot f_\sigma(\sigma))$, where $f_\sigma(\sigma)$ is a polynomial. In our model of the *median* running times, $f_\sigma(\sigma)$ is indeed polynomial but f_k depends not only on k_{\min} but also on σ . However, the dependence on σ is inverted; the larger the σ , the smaller the $f_k(k_{\min}, \sigma)$. Thus, we can substitute $\sigma = 1$ to (15) to obtain an upper bound for the median running times of the Gurobi solver of the standard FPT form: $f(k_{\min}, \sigma) = O(2^{k_{\min}/(c_2+c_3)} \cdot (\sigma^2 + c_4\sigma + c_5))$.

An interesting research question is to understand why the exponent in $f_k(k_{\min}, \sigma)$ is inversely proportional to σ . If we ignore c_3 , we see that the exponent in $f_k(k_{\min}, \sigma)$ is the proportion of values in the UPA matrix that are adjusted by the role mining: k_{\min}/σ . Our theory does not explain this phenomenon but it is an interesting observation that deserves future research; for example, there could be parameterisation of the problem based on this ratio.

Overall, these results confirm that general-purpose solvers can be effective on FPT problems. Moreover, they suggest that our solver might be appropriate for relatively large instances, provided the value of the parameter k_{\min} is small

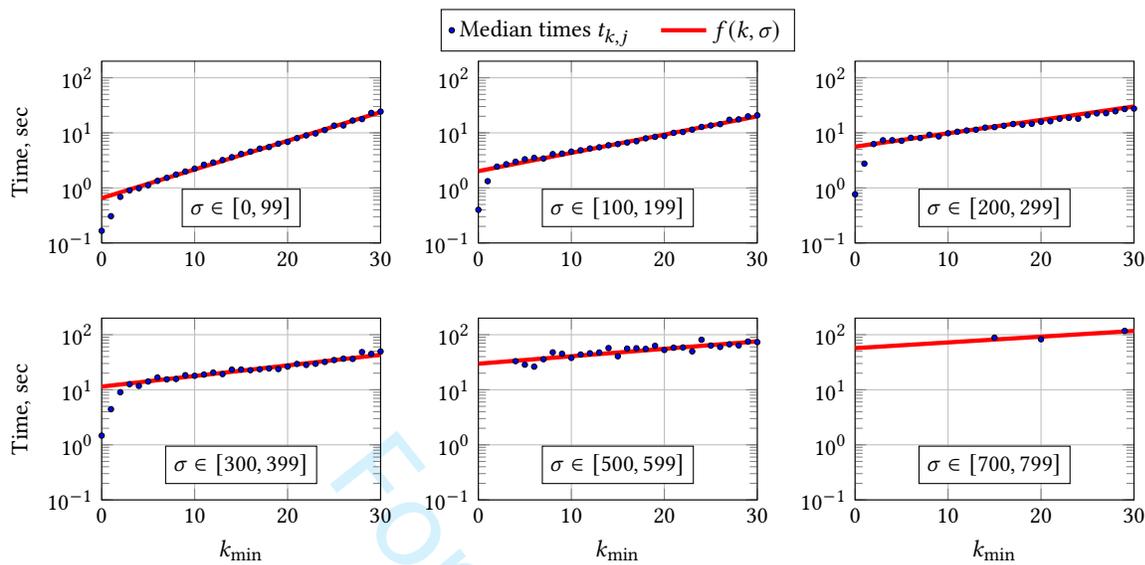


Fig. 5. The aggregated data $t_{k_{\min},j}$ and the best fit model $f(k_{\min}, \sigma)$ sliced along the σ axis.

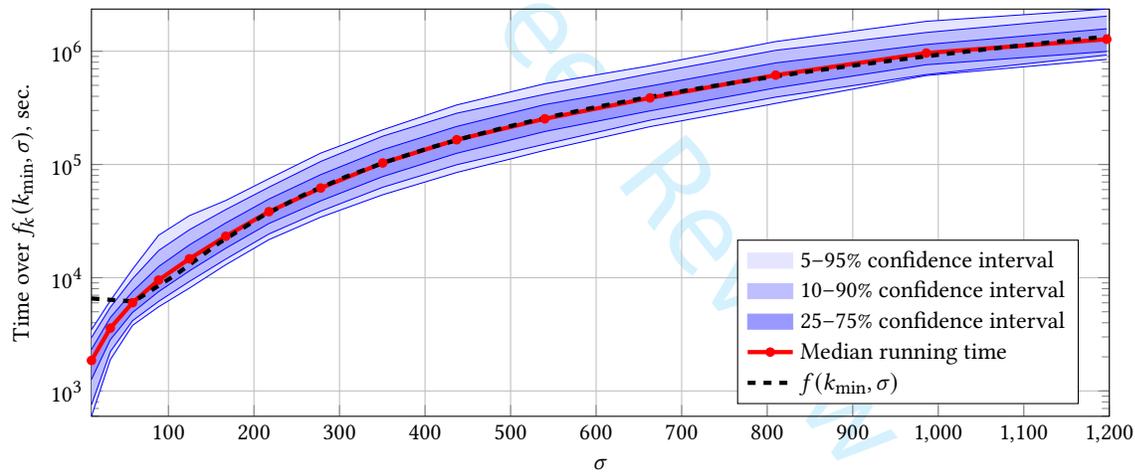


Fig. 6. This graph demonstrates the fit of $f_{\sigma}(\sigma)$ to the data for all $2 \leq k_{\min} \leq 30$. The times are divided by $f_k(k_{\min}, \sigma)$ to compensate for the varying k_{\min} . Blue areas show the distribution of the data; blue lines are percentiles of the running times.

enough. Further research may build a model of the solver's running time as a function of three parameters: σ , k_{\min} and r .

5 COMPUTATIONAL EXPERIMENTS WITH REAL-WORLD INSTANCES

To test our solver and study the trade-off between k_{\min} and r in BO-GNRM, we use a set of real-world benchmark instances [6]. This set includes nine instances of various sizes. The largest instance *americas large* has $m = 3485$,

$n = 10\,127$, $\sigma = 185\,294$ and around 400 roles according to the best heuristical solutions in the literature. Our approach is impractical for such large instances as the size of the formulation would be prohibitive. However, some other instances in this benchmark set are more manageable. The details of the instances we use in this section are given in Table 1.

Instance	m	n	σ	density	r_{\min}
Domino	79	231	730	4.0%	20
Firewall 2	325	590	36 428	19.0%	10
Healthcare	46	46	1 486	70.2%	14

Table 1. Real-world instances used in this section. The last column gives the minimum number of roles r_{\min} [6]; this is for RMP (i.e., $k_{\min} = 0$).

In classic role mining, the UPA matrix is sufficient to define the instance. In GNRM, we also need to define the value of r (the number of roles to be extracted) and the matrix F .

In our experiments, we use three types of the F matrix:

Security: $f_{i,j} = upa_{i,j}$; the new roles may remove a permission from a user but can never add a new permission;

Availability: $f_{i,j} = 1 - upa_{i,j}$; the new roles may add a permission to a user but can never remove an existing permission;

Noise: $f_{i,j} = 1$ for every i and j ; the new roles may both remove and add permissions.

In Section 5.1, we focus on the Security instances. We compare all three types of the F matrix in Section 5.2.

Our Gurobi-based solver can be used as an exact solver, but it can also serve as a heuristic if we specify the time budget. In this section, we use it as a heuristic, however in some cases it proves the optimality of the solutions within the provided time budget.

5.1 Performance of the solver

Figure 7 shows the trade-off between k_{\min} and r for several real-world instances. We experimented with three time budgets: 100 sec, 1 000 sec and 10 000 sec. All the solutions that were proven optimal by Gurobi within the time budget were marked with circles. Thus, for example, being given 10000 seconds, the solver proved optimality of solutions for $k_{\min} = 1, 2, 3, 4, 19$ for the *Domino* instance. Note that the vertical axes in these plots are logarithmic while k_{\min} can take any integer value starting from $k_{\min} = 0$. Thus, we added value $k_{\min} = 0$ to the logarithmic scale.

Generally, increasing the value of r decreases the value of k ; indeed, the *optimal* value of $k_{\min}(r)$ is a monotonically decreasing function of r . However, there are exceptions to this rule in Figure 7, particularly for smaller time budgets. In many cases, increasing the value of r makes the problem harder as the formulation size increases, and the solver might not be able to find a near-optimal solution within the time budget. In extreme cases, the solver only finds the trivial solution $ua_{i,\ell} = ua_{\ell,j} = 0$ for every i, j and ℓ .

Note that, formally speaking, Figure 7 does not show solutions to BO-GNRM. In practice, one would want to remove the dominated solutions to obtain the Pareto front and set the upper bounds \bar{k} and \bar{r} . These steps, however, would hide some of the aspects of the solver's behaviour, hence we decided to include all the solutions here and below.

Naturally, lower time budgets produce worse solutions in general. However, the difference between the 1 000 sec time budget and the 10 000 sec time budget is often insignificant. This hints at the proximity of the obtained solutions to the optimal solutions.

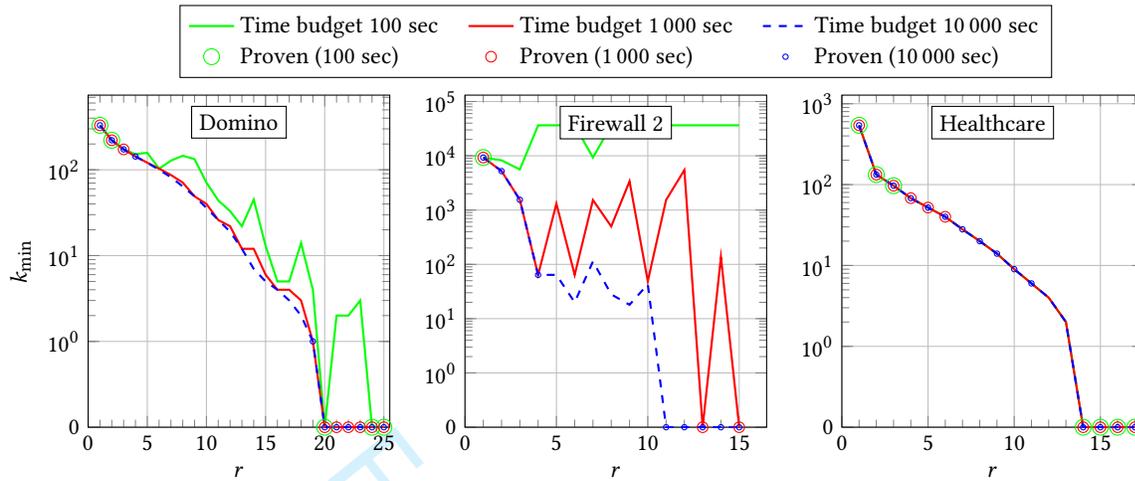


Fig. 7. The trade-off between k_{\min} and r for several real-world instances (Security type). Each line corresponds to the values obtained by our Gurobi-based solver within the given time budget (100, 1000 or 10000 seconds). When the time budget was sufficient to prove optimality of the solution, a corresponding mark was added. Note that we added value $k = 0$ to the vertical axis.

5.2 Comparison of instance types

So far, we experimented with the Security instances only. In this section, we will study how changing the instance type (Security, Availability, and Noise) affects the properties of the instances. We will use the same real-world UPA matrices but we will solve instances with various F matrices.

Figure 8 shows the trade-off between r and k for each instance type, for several real-world instances. As in Section 5.1, the vertical axes are logarithmic with added $k = 0$. We use the 10 000 seconds time budget for this experiment to be as close to the optimal solutions as we can.

Noise instances are most flexible, hence, as expected, their k does not exceed that of the Security and Availability instances. The gap between Noise instances and Security/Availability instances depends on the UPA matrix. For *Domino*, the gap between Security and Noise instances is minimal (the ratio between their values of k_{\min} does not exceed 1.025) whereas the gap between Availability and Noise instances is often significant (the ratio goes up to 53 for small r). Intuitively, this makes sense considering that the *Domino* UPA matrix is very sparse; to ensure that every value ‘1’ is preserved, one may need to sacrifice many ‘0’s. For the *Healthcare* instances, we observe that solutions to both Security and Availability instances are relatively close to the solutions to the Noise instances; the ratio never exceeds 2.3. This is consistent with the observation that the *Healthcare* UPA matrix is relatively dense. It is difficult to make any conclusions for the *Firewall 2* UPA matrix as some of the solutions are likely to be relatively far from optimal; this is evident from the jumps in the Security and Availability curves.

Finally, judging by the ability of the solver to prove optimality within the 10 000 seconds time budget, the instances with small r or k_{\min} are easier than the instances with both r and k_{\min} being relatively large. In other words, we found that the solver performs well on instances with either of the parameters being small whereas our theoretical results suggest that parameterisation by either of them is not efficient. Our interpretation of these results is that the problem is hard even for small k_{\min} or r but it is particularly hard (from the practical point of view) when both k_{\min} and r are large.

20

Crampton et al.

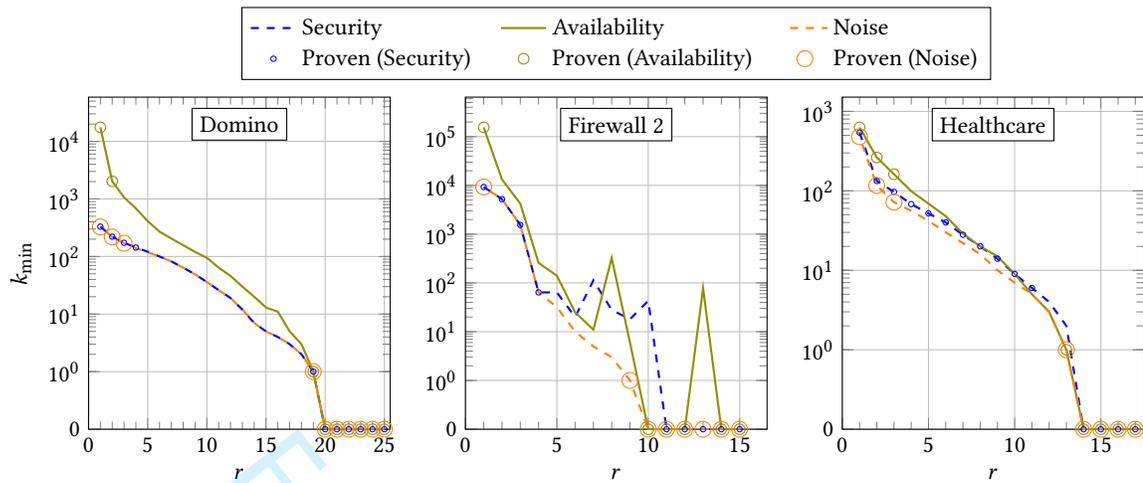


Fig. 8. The effect of the instance type (Security, Availability or Noise) on the trade-off between k_{\min} and r . The time budget is 10 000 sec. Note that we added value $k_{\min} = 0$ to the vertical axis.

6 CONCLUDING REMARKS

This paper introduces the GENERALIZED NOISE ROLE MINING problem (GNRM). We believe this is a useful contribution to the literature on role mining, not least because it allows us to define security- and availability-aware role mining problems. We extended GNRM to BO-GNRM, a bi-objective optimization variant of GNRM that allows us to find an optimal balance between the number of roles r and the number k of deviations from the original authorization matrix.

We have shown that GNRM and BO-GNRM are fixed-parameter tractable, which means that they can be solved in a reasonable amount of time, provided problem instances only require solutions in which $r + k$ ($\bar{r} + \bar{k}$, respectively) are relatively small. Algorithms for role mining do not need to be particularly fast, but they cannot be exponential in the size of the input, given the size of typical instances. Knowing that algorithms exist that do solve role mining problems, subject to certain constraints on the solution, provides grounds for cautious optimism about the feasibility of solving real-world role mining problems.

Our experimental work provides further cause for optimism. In particular, we show that our solution method based on a general-purpose solver efficiently exploits the FPT structure of the problem allowing us to solve several real-world instances, sometimes even proving solution optimality. Also, the results on security-aware role mining suggest that it is possible to find solutions relatively quickly if we require those solutions to preserve the security of the original configuration, something that is generally desirable.

The work in this paper provides plenty of scope for future work. In particular, we would like to explore whether our work on matrix decomposition and FPT results can be extended to other variants of role mining (see [16, Section 2]). More generally, this paper along with [3, 4] demonstrates usefulness of multi-objective optimization in access control. Thus, it may be useful to apply multi-objective optimization to other problems in the area.

REFERENCES

- [1] CRAMPTON, J., EIBEN, E., GUTIN, G., KARAPETYAN, D., AND MAJUMDAR, D. Valued authorization policy existence problem. In *the 26th ACM Symposium on Access Control Models and Technologies* (2021).

Manuscript submitted to ACM

- [2] CRAMPTON, J., EIBEN, E., GUTIN, G. Z., KARAPETYAN, D., AND MAJUMDAR, D. Generalized noise role mining. In *SACMAT '22: The 27th ACM Symposium on Access Control Models and Technologies, New York, NY, USA, June 8 - 10, 2022* (2022), S. Dietrich, O. Chowdhury, and D. Takabi, Eds., ACM, pp. 91–102.
- [3] CRAMPTON, J., GUTIN, G. Z., KARAPETYAN, D., AND WATRIGANT, R. The bi-objective workflow satisfiability problem and workflow resiliency. *J. Comput. Secur.* 25, 1 (2017), 83–115.
- [4] CURREY, J., MCKINSTRY, R., DADGAR, A., AND GRITTER, M. Informed privilege-complexity trade-offs in RBAC configuration. In *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies, SACMAT 2020, Barcelona, Spain, June 10-12, 2020* (2020), J. Lobo, S. D. Stoller, and P. Liu, Eds., ACM, pp. 119–130.
- [5] DAN, C., HANSEN, K. A., JIANG, H., WANG, L., AND ZHOU, Y. Low rank approximation of binary matrices: Column subset selection and generalizations. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018* (2018), I. Potapov, P. G. Spirakis, and J. Worrell, Eds., vol. 117 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 41:1–41:16.
- [6] ENE, A., HORNE, W., MILOSAVLJEVIC, N., RAO, P., SCHREIBER, R., AND TARJAN, R. E. Fast exact and heuristic methods for role minimization problems. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies* (New York, NY, USA, 2008), SACMAT '08, Association for Computing Machinery, p. 1–10.
- [7] FERRAILOLO, D. F., SANDHU, R. S., GAVRILA, S. I., KUHN, D. R., AND CHANDRAMOULI, R. Proposed NIST standard for role-based access control. *ACM Trans. Inf. Syst. Secur.* 4, 3 (2001), 224–274.
- [8] FOMIN, F. V., GOLOVACH, P. A., AND PANOLAN, F. Parameterized low-rank binary matrix approximation. *Data Min. Knowl. Discov.* 34, 2 (2020), 478–532.
- [9] GILLIS, N., AND VAVASIS, S. A. On the complexity of robust PCA and ℓ_1 -norm low-rank matrix approximation. *Mathematics of Operations Research* 43, 4 (2018), 1072–1084.
- [10] GREGORY, D. A., PULLMAN, N. J., JONES, K. F., AND LUNDGREN, J. R. Biqule coverings of regular bigraphs and minimum semiring ranks of regular matrices. *J. Comb. Theory, Ser. B* 51, 1 (1991), 73–89.
- [11] KAISA, M. *Nonlinear Multiobjective Optimization*, vol. 12 of *International Series in Operations Research & Management Science*. Kluwer Academic Publishers, Boston, USA, 1999.
- [12] KARAPETYAN, D., AND GUTIN, G. Solving the workflow satisfiability problem using general purpose solvers. *IEEE Transactions on Dependable and Secure Computing* (2022).
- [13] KARAPETYAN, D., PARKES, A. J., GUTIN, G., AND GAGARIN, A. Pattern-based approach to the workflow satisfiability problem with user-independent constraints. *Journal of Artificial Intelligence Research* 66 (2019), 85–122.
- [14] KIM, K. H. *Boolean Matrix Theory and Applications*, vol. 70 of *Monographs and Textbooks in Pure and Applied Mathematics*. Marcel Dekker, New York, 1982.
- [15] LU, H., VAIDYA, J., AND ATLURI, V. Optimal boolean matrix decomposition: Application to role engineering. In *ICDE* (2008), IEEE Computer Society, pp. 297–306.
- [16] MITRA, B., SURAL, S., VAIDYA, J., AND ATLURI, V. A survey of role mining. *ACM Comput. Surv.* 48, 4 (2016), 50:1–50:37.
- [17] MOLLOY, I. M., LI, N., LI, T., MAO, Z., WANG, Q., AND LOBO, J. Evaluating role mining algorithms. In *SACMAT* (2009), ACM, pp. 95–104.
- [18] NEUMANN, G., AND STREMBECK, M. A scenario-driven role engineering process for functional RBAC roles. In *SACMAT* (2002), ACM, pp. 33–42.
- [19] SANDHU, R. S., COYNE, E. J., FEINSTEIN, H. L., AND YUMAN, C. E. Role-based access control models. *Computer* 29, 2 (1996), 38–47.
- [20] UZUN, E., ATLURI, V., LU, H., AND VAIDYA, J. An optimization model for the extended role mining problem. In *DBSec* (2011), vol. 6818 of *Lecture Notes in Computer Science*, Springer, pp. 76–89.
- [21] VAIDYA, J., ATLURI, V., WARNER, J., AND GUO, Q. Role engineering via prioritized subset enumeration. *IEEE Transactions on Dependable and Secure Computing* 7, 3 (2010), 300–314.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

1
2
3
4 **Cover Letter**
5

6 Dear editors,
7

8 I'd like to submit my paper entitled "Bi-objective Optimization in Role Mining" for possible
9 publication in TOPS.
10

11 Yours sincerely,
12

13 Gregory Gutin
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Peer Review