

A relax-fix-and-exclude algorithm for an MINLP problem with multilinear interpolations

Bruno M. Pacheco^{1*†}, Pedro M. Antunes²,
Eduardo Camponogara², Laio O. Seman², Vinícius R. Rosa³,
Bruno F. Vieira³, Cesar Longhi³

¹CIRRELT and Département d'Informatique et de Recherche
Opérationnelle, Université de Montréal, Montréal, QC, Canada.

²Department of Automation and Systems Engineering, Federal University
of Santa Catarina, Florianópolis, SC, Brazil.

³Petrobras Research Center, Rio de Janeiro, RJ, Brazil.

*Corresponding author(s). E-mail(s):

bruno.machado.pacheco@umontreal.ca;

Contributing authors: antunespedromh@gmail.com;

eduardo.camponogara@ufsc.br; laio.seman@ufsc.br;

viniciusrr@petrobras.com.br; bfv@petrobras.com.br;

cesar.longhi@petrobras.com.br;

[†]The author was affiliated to the Department of Automation and Systems
Engineering of the Federal University of Santa Catarina during the time
this research was being developed.

Abstract

This paper introduces a novel algorithm for Mixed-Integer Nonlinear Programming (MINLP) problems with multilinear interpolations of look-up tables. These problems arise when objectives or constraints contain black-box functions only known at a finite set of evaluations on a predefined grid. We derive a piecewise-linear relaxation for the multilinear interpolants, which require an MINLP formulation. Supported by the fact that our proposed relaxation is as tight as possible, we propose a novel algorithm that iteratively solves the MILP relaxation and refines the solution space through variable fixing and exclusion strategies. This approach ensures convergence to an optimal solution, which we demonstrate, while maintaining computational efficiency. We apply the proposed algorithm

(Relax-Fix-and-Exclude, RFE) to a real-world offshore oil production optimization problem. Compared to the Gurobi solver, RFE was able to solve to optimality 20% more instances within one hour. Furthermore, on medium and large instances that both RFE and Gurobi solved to optimality, our algorithm took, on average, 52% less time than Gurobi to converge, while on all instances that neither solved, RFE always achieved a lower primal-dual gap, on average 91% smaller.

Keywords: Multilinear Interpolation, Relax-Fix-and-Exclude Algorithm, Mixed-Integer Nonlinear Programming (MINLP)

1 Introduction

We consider mixed-integer programming problems where the cost and constraints involve black-box functions—functions known only at a finite number of evaluation points in their domain. Such functions appear in many practical optimization tasks where the underlying models are complex simulations, result from empirical processes, or simply lack a tractable closed-form expression [1–4]. We restrict our attention to the cases where the evaluations of the black-box functions are organized on a rectangular, axis-aligned grid, resulting in the so-called look-up tables [5, 6]. Consequently, to enable optimization over the continuous domain, an interpolation of the look-up table is required¹ to generate a continuous function.

A particularly popular interpolation method in this context is *multilinear interpolation* [7–14]. Unlike piecewise-linear approaches, the resulting interpolant preserves smoothness almost everywhere², does not require arbitrary subdivisions (triangulations) of the domain, and scales efficiently with dimensionality. The downside is that embedding the interpolant in an optimization model introduces multilinear (i.e., non-convex) constraints, making the problem a challenging instance of mixed-integer nonlinear programming.

This paper introduces a novel algorithm for mixed-integer nonlinear programming (MINLP) problems with functions modeled via multilinear interpolation. We explore polyhedral tools for multilinear terms and show that they can be used to derive the convex hull of the nonlinearities arising from the interpolation method. This convex-hull forms the basis of our *Relax-Fix-and-Exclude* algorithm, exploiting it not only to compute tight bounds, but also to guide a depth-first search through variable fixing and generalized no-good cuts.

We demonstrate the performance of the proposed algorithm in an offshore oil production optimization problem. The oil and gas industry is a classical domain application of MINLP over black-box functions, since optimization problems involve complex nonlinear functions representing processes such as reservoir pressure dynamics and multi-phase flow [15]. These functions are often derived from simulations or empirical

¹To cover the entire domain properly, we would need to extrapolate the information on the look-up table. However, it is usual to assume that the optimal is within the boundaries of the table and thus interpolation suffices.

²The interpolant is continuous but not differentiable at the evaluation points, which are finitely many.

data, making them difficult to incorporate directly into optimization problems. Furthermore, the proposed approach is well-suited for these applications, where computational efficiency and accuracy are critical.

The key contribution of this work is the development of a piecewise-linear relaxation for the nonlinear interpolated functions. This relaxation leads to an efficient mixed-integer linear programming (MILP) approximation, which provides both a tractable formulation and a reliable lower bound for the original problem. By incorporating Special Ordered Sets of type 2 (SOS2), the method ensures that adjacent grid points are selected during interpolation, preserving the convexity and efficiency of the optimization process. Furthermore, the algorithm leverages fix-and-exclude strategies to refine the solution space progressively and converge to an optimal solution.

1.1 Related Work

Many real-world optimization problems involve functions with high evaluation costs, which are usually only known at finitely many points of their domain. These problems have been known for a long time [1] and have appeared in a wide range of application areas, such as machine learning [2], chemistry [3], and engineering [4].

A well-studied approach for solving such problems is to perform a piecewise-linear interpolation over the look-up tables of the black-box function [16–18]. In a grid-like partitioning of the domain, the piecewise-linear interpolation requires additional cuts to reduce the partitioning to simplices, such that the model becomes an injective mapping. Although theoretically possible for any number of dimensions, this partitioning has only been proposed for domains of up to three dimension [19]. On top of that, the number of simplices grows exponentially with the domain dimension, requiring exponentially many more integer variables. Finally, the choice of simplices (within the hyperrectangular partition) is not an evident design choice, and it may have a significant impact on the approximation quality, e.g., by overestimating or underestimating the true function.

Multilinear interpolations of look-up tables are a common practice to deal with black-box functions or complex non-linear functions, with applications in physics simulation [7, 8], automotive software [9–11], and even machine learning [12]. For example, Furlan et al. [13] and Martins et al. [14] replace phenomenological models of a chemical process with a look-up table that is interpolated during the simulation, which yielded improved efficiency at minor accuracy loss. In contrast to the piecewise-linear method, the multilinear interpolation is smooth everywhere on the domain except at a finite number of points. Furthermore, it naturally scales to arbitrary domain dimensions and requires no design choice during implementation.

Optimizing with multilinearly interpolated look-up tables, in contrast to piecewise-linear interpolations, introduces multilinear terms, making the problem, at least, a mixed-integer quadratically constrained programming (MIQCP) problem. Optimization with multilinear constraints has been well-studied [20–22] and continues to be an active area of research [23, 24]. In fact, many solvers support multilinear constraints out-of-the-box [25, 26]. However, to the best of our knowledge, our work is the first to study the optimization problem with multilinear interpolations of look-up tables and to propose an exact algorithm for it.

1.2 Contributions

This work brings the following contributions to the technical literature:

- An MINLP approximation for a class of real-world problems involving black-box functions using multilinear interpolation with data from sensitivity analysis, which can be solved efficiently.
- A tight MILP relaxation of multilinear interpolation functions arising from piecewise-linear formulations.
- A Relax-Fix-and-Exclude algorithm with convergence guarantee for solving optimization problems with multilinear interpolation functions, which iteratively solves piecewise-linear based relaxations (*Relax*), optimizes over a reduced domain resulting from variable fixing (*Fix*), and removes this domain from the search space (*Exclude*).
- A computational demonstration of the efficiency of the relax-fix-and-exclude algorithm and comparison with state-of-the-art MINLP solvers (Gurobi, SCIP, and BARON) in an application to oil production in offshore platforms.

1.3 Organization

Section 2 introduces the baseline problem, which consists of a mixed-integer nonlinear program with black-box functions approximated through multilinear interpolation of look-up tables, resulting in an MINLP formulation. Section 3 presents the relax, fix, and exclude steps, which are then combined into the proposed algorithm. Section 5 reports computational results from applying the relax-fix-and-exclude algorithm for oil production optimization in offshore platforms (Section 4), and a comparison with state-of-the-art MINLP solvers Gurobi, SCIP and BARON. Section 6 draws final conclusions and suggests future research directions.

2 Problem Statement

This work is concerned with the following mixed-integer program

$$\begin{aligned}
 C = \min_{\mathbf{x}, \mathbf{y}} \quad & \tilde{f}_0(\mathbf{x}) + \mathbf{a}_0^T \mathbf{y} \\
 \text{s.t.} \quad & \tilde{f}_i(\mathbf{x}) + \mathbf{a}_i^T \mathbf{y} \leq b_i, \quad i = 1, \dots, m \\
 & \mathbf{x} \in X \subseteq \mathbb{R}^n, \mathbf{y} \in \{0, 1\}^p
 \end{aligned} \tag{1}$$

where each

$$\begin{aligned}
 \tilde{f}_i : X &\longrightarrow \mathbb{R} \\
 \mathbf{x} &\longmapsto \tilde{f}_i(\mathbf{x})
 \end{aligned} \tag{2}$$

is a multilinear interpolation of a black-box function $f_i : X \rightarrow \mathbb{R}$, for $i = 0, \dots, m$.

It is assumed that a *look-up table* is available with evaluations of the black-box functions at a finite set of *breakpoints* $\mathcal{B} \subseteq X$. In other words, for each $\hat{\mathbf{x}} \in \mathcal{B}$, we know the value of $f_i(\hat{\mathbf{x}})$ for all $i = 0, \dots, m$.

We assume the breakpoints form a rectangular grid $\mathcal{B} = \mathcal{B}_1 \times \cdots \times \mathcal{B}_n$, such that each $\mathcal{B}_j = \{\hat{x}_j^0, \hat{x}_j^1, \hat{x}_j^2, \dots\}$ is a finite set. Let the breakpoints be indexed in increasing order $\hat{x}_j^0 < \hat{x}_j^1 < \hat{x}_j^2 < \cdots$, for $j = 1, \dots, n$. Further, let \mathcal{I}_j denote the set of *indices* of the breakpoints for the j -th variable, such that we can write $\mathcal{B}_j = \{\hat{x}_j^k\}_{k \in \mathcal{I}_j}$. Then, we will let $\mathcal{I} = \mathcal{I}_1 \times \cdots \times \mathcal{I}_n$ denote the set of *multi-indices* of the breakpoints, such that $\forall \mathbf{k} = (k_1, \dots, k_n) \in \mathcal{I}$, we can write $\hat{\mathbf{x}}^{\mathbf{k}} = (\hat{x}_1^{k_1}, \dots, \hat{x}_n^{k_n}) \in \mathcal{B}$.

2.1 Multilinear Interpolation

To develop an understanding of such an interpolation method, let us consider a single black-box function $f : X \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$. If $n = 1$, we write $x_1 = x$, and we have a traditional linear interpolation over $\mathcal{B} = \{\hat{x}^0, \hat{x}^1, \hat{x}^2, \dots\} \subset \mathbb{R}$. Without loss of generality, take $x \in \text{conv}(\mathcal{B})$ such that $\hat{x}^0 \leq x \leq \hat{x}^1$. Then, we can write the linear interpolation of $f(x)$ as

$$\begin{aligned} \tilde{f}(x) &= f(\hat{x}^0) + \frac{x - \hat{x}^0}{\hat{x}^1 - \hat{x}^0} (f(\hat{x}^1) - f(\hat{x}^0)) \\ &= \underbrace{\left(\frac{\hat{x}^1 - x}{\hat{x}^1 - \hat{x}^0} \right)}_{\xi_0} f(\hat{x}^0) + \underbrace{\left(\frac{x - \hat{x}^0}{\hat{x}^1 - \hat{x}^0} \right)}_{\xi_1} f(\hat{x}^1). \end{aligned} \quad (3)$$

Note that the coefficients ξ_0 and ξ_1 of the linear interpolation are such that $\xi_0 + \xi_1 = 1$, and

$$x = \xi_0 \hat{x}^0 + \xi_1 \hat{x}^1, \quad (4)$$

that is, they match the weights that express x as a convex combination (CC) of two *consecutive* breakpoints. One way to formulate such a convex combination for the one-dimensional case is through the use of an SOS2 variable

$$\begin{aligned} x &= \sum_{k \in \mathcal{I}} \xi_k \hat{x}^k \\ 1 &= \sum_{k \in \mathcal{I}} \xi_k \\ \boldsymbol{\xi} &\in \text{SOS2}(\mathcal{I}), \boldsymbol{\xi} \geq \mathbf{0}, \end{aligned} \quad (5)$$

which ensures that only two consecutive breakpoints will have the respective coefficients assuming non-zero values. We use the notation $\text{SOS2}(\mathcal{I})$ to convey that the SOS2 constraint is applied over the ordering defined by \mathcal{I} . Then, we can write the linear interpolation as

$$\begin{aligned} \lambda_k &= \xi_k, & \forall k \in \mathcal{I}, \\ \tilde{f} &= \sum_{k \in \mathcal{I}} \lambda_k f(\hat{x}^k). \end{aligned} \quad (6)$$

Note that we introduce the redundant λ variables to differentiate ξ_k , the weight of the k -th breakpoint, from λ_k , the weight of the function evaluation at the k -th breakpoint, which happen to coincide in the onedimensional case. in which we introduce the

λ variables to differentiate the weights of the breakpoints from the weights of the evaluations of f .

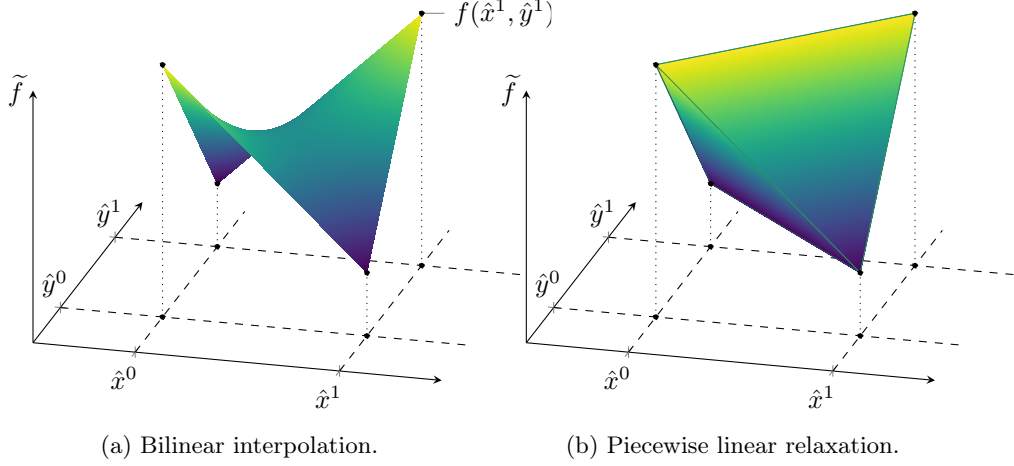


Fig. 1: Example of multilinear interpolation and the respective piecewise-linear relaxation for a black-box function on a two-dimensional domain. The interpolation is illustrated on a single rectangle of the sensitivity analysis grid, namely $[\hat{x}^0, \hat{x}^1] \times [\hat{y}^0, \hat{y}^1]$.

For $n > 1$, the multilinear interpolation over grids can be defined recursively [27]. A bidimensional example is illustrated in Figure 1a. We get into details for the case $n = 3$. To ease the notation, let $x_1 = x$, $x_2 = y$, and $x_3 = z$. Then, for $(x, y, z) \in \text{conv}(\mathcal{B}_x \times \mathcal{B}_y \times \mathcal{B}_z)$, assume, without loss of generality, that $\hat{x}^0, \hat{x}^1 \in \mathcal{B}_x$, $\hat{y}^0, \hat{y}^1 \in \mathcal{B}_y$, and $\hat{z}^0, \hat{z}^1 \in \mathcal{B}_z$ are the vertices of the rectangular prism that contains (x, y, z) , i.e., $\hat{x}^0 \leq x \leq \hat{x}^1$, $\hat{y}^0 \leq y \leq \hat{y}^1$, and $\hat{z}^0 \leq z \leq \hat{z}^1$. By interpolating each dimension recursively, the approximation can be written

$$\begin{aligned}
 \tilde{f} &= \frac{\hat{z}^1 - z}{\hat{z}^1 - \hat{z}^0} \tilde{f}_{x,y}(\hat{z}^0) + \frac{z - \hat{z}^0}{\hat{z}^1 - \hat{z}^0} \tilde{f}_{x,y}(\hat{z}^1) \\
 \tilde{f}_{x,y}(\hat{z}^{0|1}) &= \frac{\hat{y}^1 - y}{\hat{y}^1 - \hat{y}^0} \tilde{f}_x(\hat{y}^0, \hat{z}^{0|1}) + \frac{y - \hat{y}^0}{\hat{y}^1 - \hat{y}^0} \tilde{f}_x(\hat{y}^1, \hat{z}^{0|1}) \\
 \tilde{f}_x(\hat{y}^{0|1}, \hat{z}^{0|1}) &= \frac{\hat{x}^1 - x}{\hat{x}^1 - \hat{x}^0} \tilde{f}(\hat{x}^0, \hat{y}^{0|1}, \hat{z}^{0|1}) + \frac{x - \hat{x}^0}{\hat{x}^1 - \hat{x}^0} \tilde{f}(\hat{x}^1, \hat{y}^{0|1}, \hat{z}^{0|1}),
 \end{aligned} \tag{7}$$

where $\hat{z}^{0|1}$ denotes either \hat{z}^0 or \hat{z}^1 , and $\hat{y}^{0|1}$ likewise. Expanding the approximation \tilde{f} above in terms of products, we can rewrite the interpolated function \tilde{f} as a summation over products of the basis functions for each variable x , y , and z ,

$$\tilde{f} = \sum_{(k,j,i) \in \{0,1\}^3} \xi_k^x \xi_j^y \xi_i^z f(\hat{x}^k, \hat{y}^j, \hat{z}^i) \tag{8}$$

where

$$\begin{aligned}\xi_k^x &= \begin{cases} \frac{\hat{x}^1 - x}{\hat{x}^1 - \hat{x}^0}, & \text{if } k = 0, \\ \frac{x - \hat{x}^0}{\hat{x}^1 - \hat{x}^0}, & \text{if } k = 1, \end{cases} \\ \xi_j^y &= \begin{cases} \frac{\hat{y}^1 - y}{\hat{y}^1 - \hat{y}^0}, & \text{if } j = 0, \\ \frac{y - \hat{y}^0}{\hat{y}^1 - \hat{y}^0}, & \text{if } j = 1, \end{cases} \\ \xi_i^z &= \begin{cases} \frac{\hat{z}^1 - z}{\hat{z}^1 - \hat{z}^0}, & \text{if } i = 0, \\ \frac{z - \hat{z}^0}{\hat{z}^1 - \hat{z}^0}, & \text{if } i = 1, \end{cases}\end{aligned}\tag{9}$$

and ξ_k^x serves as the basis function for variable x , with k iterating over 0 and 1, and similarly for the other two variables.

Like in the one-dimensional case, each ξ is a coefficient of the CC of breakpoints that results in (x, y, z) , i.e.,

$$\begin{aligned}x &= \xi_0^x \hat{x}^0 + \xi_1^x \hat{x}^1, \\ y &= \xi_0^y \hat{y}^0 + \xi_1^y \hat{y}^1, \\ z &= \xi_0^z \hat{z}^0 + \xi_1^z \hat{z}^1, \\ 1 &= \xi_0^x + \xi_1^x, \\ 1 &= \xi_0^y + \xi_1^y, \\ 1 &= \xi_0^z + \xi_1^z, \\ 0 &\leq \xi_0^x, \xi_1^x, \xi_0^y, \xi_1^y, \xi_0^z, \xi_1^z.\end{aligned}\tag{10}$$

Thus, \tilde{f} is a weighted sum where each term is the product of basis functions for each variable, multiplied by the corresponding function value at that grid point, $f(\hat{x}^k, \hat{y}^j, \hat{z}^i)$. We can rewrite the interpolation (7) as

$$\tilde{f} = \sum_{(k,j,i) \in \{0,1\}^3} \lambda_{k,j,i} f(\hat{x}^k, \hat{y}^j, \hat{z}^i),\tag{11}$$

where $\lambda_{k,j,i} = \xi_k^x \xi_j^y \xi_i^z$, for all $(k, j, i) \in \mathcal{I}$.

More generally, we can express the multilinear interpolation of a function of n variables by using SOS2 vectors $\boldsymbol{\xi}^j$ as coefficients of the CC of the breakpoints \mathcal{B}_j into x_j , for $j = 1, \dots, n$. Precisely, we impose

$$\begin{aligned}x_j &= \sum_{k \in \mathcal{I}_j} \xi_k^j \hat{x}_j^k, \\ 1 &= \sum_{k \in \mathcal{I}_j} \xi_k^j, \\ \boldsymbol{\xi}^j &\in \text{SOS2}(\mathcal{I}_j),\end{aligned}\tag{12}$$

for each $j = 1, \dots, n$. Note that, in the above formulation, an SOS2 constraint is imposed at each vector variable $\boldsymbol{\xi}^j \in \mathbb{R}^{|\mathcal{I}_j|}$ individually. The interpolation weights are

then

$$\lambda_{\mathbf{k}} = \prod_{j=1}^n \xi_{k_j}^j, \quad \forall \mathbf{k} \in \mathcal{I}. \quad (13)$$

Finally, the formulation of \tilde{f} as a multilinear interpolation becomes

$$\tilde{f} = \sum_{\mathbf{k} \in \mathcal{I}} \lambda_{\mathbf{k}} f(\hat{\mathbf{x}}^{\mathbf{k}}). \quad (14)$$

2.2 MINLP Formulation

By considering the formulation of the multilinear interpolant in (14), we can write problem (1) as an MINLP

$$\begin{aligned} C = \min_{\boldsymbol{\xi}, \mathbf{y}} \quad & \tilde{f}_0 + \mathbf{c}^T \mathbf{y} \\ \text{s.t.} \quad & \tilde{f}_i + \mathbf{a}_i^T \mathbf{y} \leq b_i, \quad i = 1, \dots, m, \\ & \tilde{f}_i = \sum_{\mathbf{k} \in \mathcal{I}} \lambda_{\mathbf{k}} f_i(\hat{\mathbf{x}}^{\mathbf{k}}), \quad i = 0, \dots, m, \\ & \lambda_{\mathbf{k}} = \prod_{j=1}^n \xi_{k_j}^j, \quad \forall \mathbf{k} \in \mathcal{I}, \\ & x_j = \sum_{k \in \mathcal{I}_j} \xi_k^j \hat{x}_j^k, \quad j = 1, \dots, n, \\ & 1 = \sum_{k \in \mathcal{I}_j} \xi_k^j, \quad j = 1, \dots, n, \\ & \boldsymbol{\xi}^j \in \text{SOS2}(\mathcal{I}_j), \boldsymbol{\xi}^j \geq \mathbf{0} \quad j = 1, \dots, n, \\ & \mathbf{x} \in X, \mathbf{y} \in \{0, 1\}^p, \end{aligned} \quad (\mathcal{P})$$

where each $f_i(\hat{\mathbf{x}}^{\mathbf{k}})$ is obtained from the look-up table. Variables \mathbf{x} , $\boldsymbol{\lambda}$, and $\tilde{\mathbf{f}}$ are functionally dependent on $\boldsymbol{\xi}$ and are introduced solely for notational convenience. In other words, \mathcal{P} is effectively a problem over the $(\boldsymbol{\xi}, \mathbf{y})$ extended space. Furthermore, we have as many ξ_k^j variables as breakpoints along each j -th dimension, that is, $\boldsymbol{\xi}^j \in \mathbb{R}^{|\mathcal{B}_j|}$, for $j = 1, \dots, n$.

Note that the addition of constraints (12) implies $\mathbf{x} \in \text{conv}(\mathcal{B}) \cap X$. So it is an implicit assumption that the optimal solution to the problem can be expressed as a convex combination of the breakpoints.

3 The Relax-Fix-and-Exclude Algorithm

Problem \mathcal{P} is a nonconvex Mixed-Integer Quadratically-Constrained Quadratic Programming (MIQCQP) problem, which, despite being supported by modern solvers like Gurobi [25], BARON [28, 29], and SCIP [30, 31], is a notoriously hard class of problems. Classical methods for this class of problems are based on sequentially solving

MILP relaxations and NLP subproblems [32]. Our solution approach for \mathcal{P} is based on a tight MILP relaxation, which we derive in Section 3.1. However, beyond just using the relaxation to compute tighter bounds, or refining it with cutting planes in an outer approximation fashion, we opt for an enumeration approach, using the relaxation to guide our exploration of the solution space, similar to a diving heuristic.

The intuition behind our approach is that for smooth functions and a sufficiently dense set of breakpoints, the solution to the MILP relaxation will lie on the same hyperrectangle of the domain as the optimal solution to the original MINLP. Thus, we use a solution of the MILP relaxation to determine which NLP subproblem to solve by *fixing* the discrete decisions of MINLP problem, restricting it to a hyperrectangle of the domain (Section 3.2). Then, we *exclude* said hyperrectangle from the solution space of the MILP relaxation, such that re-optimizing it will result in either a tighter bound or a certificate of optimality. Furthermore, the new solution to the MILP relaxation will lie in a different hyperrectangle, focusing on a new NLP subproblem.

3.1 MILP Relaxation

The only nonlinear terms in \mathcal{P} are the non-convex constraints (13) over the ξ -space, which come from the multilinear interpolation of the black-box functions. The tightest possible linear relaxation for each constraint³ in (13) is precisely its convex hull. This is because the set defined by each (13) is the graph of a multilinear constraint, and the convex envelope of a multilinear function is polyhedral [20, 23]. We derive the desired MILP relaxation of \mathcal{P} by exploiting the traditional SOS2 formulation for piecewise-linear functions [21, 33], and we show that a simple modification yields precisely the desired convex envelopes. In other words, we show that the resulting MILP relaxation is as tight as possible, as fixing the discrete variables always results in the convex hull of the MINLP problem (under the same fixing values).

Recall that our breakpoints \mathcal{B} are distributed in a rectangular grid, such that the domain variables \mathbf{x} always lie in a hyperrectangle of consecutive breakpoints. In a piecewise-linear *interpolation* over \mathcal{B} , the SOS2 constraints over ξ are used to express \mathbf{x} as a CC of the hyperrectangle's vertices, exactly as in our MINLP formulation (\mathcal{P}) of the multilinear interpolants \tilde{f} . However, in the piecewise-linear case, \mathbf{x} is further constrained to be a CC of *at most* n vertices of any given hyperrectangle of the domain, effectively further partitioning the domain into simplexes [19]. Figure 2 illustrates the piecewise-linear interpolation for the case of a function over two variables. Note that the domain is naturally partitioned in rectangles by the breakpoints, then it is further subdivided into triangles (2D simplexes) whose vertices are the n breakpoints used in the CC.

We derive a piecewise-linear *relaxation* by removing the constraint that the domain variable must be a CC of at most n consecutive breakpoints, i.e., allowing it to be expressed as a convex combination of all $2n$ consecutive breakpoints. This allows the interpolant variable \tilde{f} to assume any value between all possible piecewise-linear interpolations, thus, the convex envelope. Fig. 1b illustrates the piecewise-linear relaxation

³By “linear relaxation” of a constraint, we mean a polyhedron that contains the set defined by the constraint.

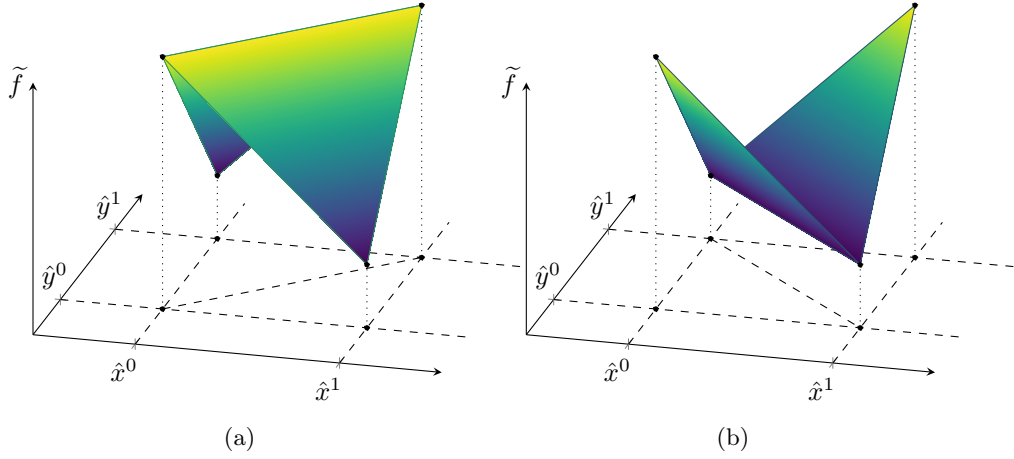


Fig. 2: Example of piecewise-linear interpolation with the same look-up table but different triangulations of the domain.

of the function illustrated in Fig. 2. Note how the facets of the convex envelope in Fig. 1b match exactly the graphs of the interpolants in Fig. 2a and 2b.

Since our multilinear interpolation is already formulated through SOS2 constraints, we need only to replace constraints (13) by

$$\xi_{\hat{k}_j}^j = \sum_{\substack{\mathbf{k} \in \mathcal{I}: \\ k_j = \hat{k}_j}} \lambda_{\mathbf{k}}, \quad \forall \hat{k}_j \in \mathcal{I}_j, j = 1, \dots, n. \quad (15)$$

Then, our MILP relaxation of \mathcal{P} is

$$\begin{aligned}
\overline{C} &= \min_{\boldsymbol{\xi}, \boldsymbol{\lambda}, \mathbf{y}} \quad \widetilde{f}_0 + \mathbf{c}^T \mathbf{y} \\
\text{s.t.} \quad & \widetilde{f}_i + \mathbf{a}_i^T \mathbf{y} \leq b_i, & i = 1, \dots, m \\
& \widetilde{f}_i = \sum_{\mathbf{k} \in \mathcal{I}} \lambda_{\mathbf{k}} f_i(\hat{\mathbf{x}}^{\mathbf{k}}), & i = 0, \dots, m, \\
& \xi_{k_j}^j = \sum_{\substack{\mathbf{k} \in \mathcal{I}: \\ k_j = \hat{k}_j}} \lambda_{\mathbf{k}}, & \forall \hat{k}_j \in \mathcal{I}_j, j = 1, \dots, n, \\
& x_j = \sum_{k \in \mathcal{I}_j} \xi_k^j \hat{x}_j^k, & j = 1, \dots, n, \\
& 1 = \sum_{k \in \mathcal{I}_j} \xi_k^j, & j = 1, \dots, n, \\
& \boldsymbol{\xi}^j \in \text{SOS2}(\mathcal{I}_j), & j = 1, \dots, n, \\
& \boldsymbol{\lambda} \in [0, 1]^{|\mathcal{I}|}, \\
& \mathbf{x} \in X, \mathbf{y} \in \{0, 1\}^p.
\end{aligned} \tag{$\overline{\mathcal{P}}$}$$

Differently from problem \mathcal{P} , $\boldsymbol{\lambda}$ is not functionally dependent on $\boldsymbol{\xi}$ in $\overline{\mathcal{P}}$, that is, $\overline{\mathcal{P}}$ is effectively a problem over the $(\boldsymbol{\xi}, \boldsymbol{\lambda}, \mathbf{y})$ extended space. Furthermore, the number of λ variables grows exponentially with n , as there must be one λ variable for each element in $\mathcal{B} = \mathcal{B}_1 \times \dots \times \mathcal{B}_n$. However, at most 2^n of them are non-null in any feasible solution, once the SOS2 constraints over the ξ variables force every λ variable that is not a vertex of the selected hyperrectangle to be 0.

Problem $\overline{\mathcal{P}}$ contains the tightest possible relaxation of the multilinear constraints of \mathcal{P} , given that each $\lambda_{\mathbf{k}}$ variable must be a convex combination of the product of the $\xi_{k_1}^1, \dots, \xi_{k_n}^n$ variables at their extremes [20, 23]. For completeness, we demonstrate through Proposition 1 that the solution space of $\overline{\mathcal{P}}$ is indeed a relaxation of the solution space of \mathcal{P} , which is a necessary condition for the correctness of Algorithm 1.

Proposition 1. *Problem $\overline{\mathcal{P}}$ is a relaxation of \mathcal{P} .*

Proof. Because the only change from \mathcal{P} to $\overline{\mathcal{P}}$ is the removal of constraint (13) and the addition of constraint (15), it suffices to show that any solution of \mathcal{P} satisfies (15). Thus, let $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}, \tilde{\boldsymbol{\xi}}, \tilde{\boldsymbol{\lambda}})$ be a feasible solution of \mathcal{P} . Consider first the case $j = 1$ and take any $\hat{k}_1 \in \mathcal{I}_1$. Because $\tilde{\boldsymbol{\lambda}}_{\mathbf{k}}$ respects (13) for any $\mathbf{k} \in \mathcal{I}$, we have

$$\begin{aligned}
\sum_{\substack{\mathbf{k} \in \mathcal{I}: \\ k_1 = \hat{k}_1}} \tilde{\lambda}_{\mathbf{k}} &= \sum_{\substack{\mathbf{k} \in \mathcal{I}: \\ k_1 = \hat{k}_1}} \prod_{j=1}^n \tilde{\xi}_{k_j}^j \\
&= \tilde{\xi}_{k_1}^1 \sum_{k_2 \in \mathcal{I}_2} \tilde{\xi}_{k_2}^2 \sum_{k_3 \in \mathcal{I}_3} \tilde{\xi}_{k_3}^3 \dots \sum_{k_n \in \mathcal{I}_n} \tilde{\xi}_{k_n}^n \\
&= \tilde{\xi}_{k_1}^1,
\end{aligned} \tag{16}$$

where the last equality comes from $\tilde{\xi}$ satisfying

$$\sum_{k_j \in \mathcal{I}_j} \tilde{\xi}_{k_j}^j = 1, j = 1, \dots, n \quad (17)$$

as it is feasible for \mathcal{P} . Therefore, the first constraint of (15) is satisfied, given that the choice for \hat{k}_1 was arbitrary. *Mutatis mutandis*, the same can be shown for $j = 2, \dots, n$, implying that all constraints in (15) are satisfied by a feasible solution of \mathcal{P} . \square

One interesting characteristic of the proposed piecewise relaxation is that it can also be used to relax multilinear constraints. For example, suppose that function f_m in problem (1) is not a black-box function, but rather a known function of the form

$$f_m(\mathbf{x}) = \sum_{(i,j)} a_{ij} x_i x_j. \quad (18)$$

Then, its multilinear interpolation is exact and the relaxation of f_m in $\overline{\mathcal{P}}$ will match precisely the piecewise McCormick envelopes [34] at the partitioning induced by the breakpoints of the look-up table, since that also characterizes the convex hull in such cases [35].

3.2 Fixing

In our algorithm, we find candidate solutions by solving NLP subproblems that stem from fixing all discrete decisions of \mathcal{P} . This is a usual step in mixed-integer programming which consists of fixing the assignment of the integer variables. In \mathcal{P} , the SOS2 constraints make this process more involved, since the discrete decisions are not explicit.

We begin by observing that the SOS2 constraints $\xi^j \in \text{SOS2}(\mathcal{I}_j)$ can be equivalently formulated as

$$\xi^j \leq \zeta^j \text{ and } \zeta^j \in \text{SOS2}(\mathcal{I}_j), \quad (19)$$

for $\zeta^j \in \{0, 1\}^{|\mathcal{B}_j|}$ and $j = 1, \dots, n$. Each ζ^j can be understood as the discrete decisions associated with the SOS2 constraint over ξ^j , that is, the selection of the pair of consecutive breakpoints along the j -th axis.

Now, take $\overline{\mathbf{y}} \in \{0, 1\}^p$ and $\overline{\zeta}^j \in \text{SOS2}(\mathcal{I}_j) \cap \{0, 1\}^{|\mathcal{B}_j|}$, for every $j = 1, \dots, n$. Then, it is easy to see that fixing $\mathbf{y} = \overline{\mathbf{y}}$ and $\xi^j \leq \zeta^j$, $j = 1, \dots, n$, reduces \mathcal{P} to an NLP. More formally, let P denote the feasible region of \mathcal{P} , and define

$$P(\overline{\zeta}, \overline{\mathbf{y}}) \triangleq \left\{ (\xi, \mathbf{y}) \in P : \mathbf{y} = \overline{\mathbf{y}}, \xi^j \leq \overline{\zeta}^j, j = 1, \dots, n \right\} \quad (20)$$

If each $\overline{\zeta}^j \in \text{SOS2}(\mathcal{I}_j)$, then at most two dimensions of ξ^j are *not* fixed to 0 in $P(\overline{\zeta}, \overline{\mathbf{y}})$, i.e., the dimensions over which ξ^j can assume positive value. Further, by projecting $P(\overline{\zeta}, \overline{\mathbf{y}})$ back onto the x -space through $x_j = \sum_{k \in \mathcal{I}_j} \xi_k^j \hat{x}_j^k$, we see that $x_j \in [\hat{x}_j^{\bar{k}}, \hat{x}_j^{\bar{k}+1}]$, where \bar{k} is the index of the first positive value of $\overline{\zeta}^j$. In other words, our fixing strategy

is equivalent to restricting each x_j to lie between two consecutive breakpoints in \mathcal{B}_j , thus, restricting \mathbf{x} to lie in a single hyperrectangle of the partitioning induced by the breakpoints in \mathcal{B} of the functions' domain.

Now let $(\tilde{\xi}, \tilde{\mathbf{y}}) \in P$ be a feasible solution for \mathcal{P} . We can define an NLP subproblem for which $(\tilde{\xi}, \tilde{\mathbf{y}})$ is feasible by selecting a binary vector $\tilde{\zeta}^j \in \text{SOS2}(\mathcal{I}_j)$ such that $\tilde{\xi}^j \leq \tilde{\zeta}^j$, $j = 1, \dots, n$. However, $\tilde{\zeta} = \lceil \tilde{\xi} \rceil$ satisfies this condition. Thus,

$$(\tilde{\xi}, \tilde{\mathbf{y}}) \in P(\lceil \tilde{\xi} \rceil, \tilde{\mathbf{y}}), \forall (\tilde{\xi}, \tilde{\mathbf{y}}) \in P. \quad (21)$$

Further, let \bar{P} be the feasible region of our MILP relaxation $\bar{\mathcal{P}}$, and take a feasible solution $(\bar{\xi}, \bar{\lambda}, \bar{\mathbf{y}}) \in \bar{P}$. Then, a consequence of (21) is that

$$(\bar{\xi}, \bar{\mathbf{y}}) \in P \iff (\bar{\xi}, \bar{\mathbf{y}}) \in P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}}). \quad (22)$$

Recall that our goal is to fix all discrete decisions of our MINLP subproblem. The rounding procedure in (21) provides a way to fix the discrete decisions in the SOS2 constraints without formulating them as binary variables. In other words, we do *not* need additional binary variables (e.g., ζ) to formulate the NLP subproblems. The fixing operation is illustrated in Figure 3b, with P , the feasible regions of \mathcal{P} , and $P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$, the feasible region of one of the NLP subproblems.

3.3 Excluding

Suppose we solve our MILP relaxation $\bar{\mathcal{P}}$ and find an optimal solution $(\bar{\xi}, \bar{\lambda}, \bar{\mathbf{y}}) \in \bar{P}$ with cost \bar{C} . We then use fixing and solve the associated NLP subproblem, finding an optimum $(\tilde{\xi}, \tilde{\mathbf{y}}) \in P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}}) \subseteq P$, with cost C . If $\bar{C} = C$, then we have a certificate of optimality for our candidate solution $(\tilde{\xi}, \tilde{\mathbf{y}})$. On the other hand, if $\bar{C} < C$, better solutions for \mathcal{P} might only exist in $P \setminus P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$, since our candidate is optimal for the NLP subproblem. Therefore, a tighter dual bound than \bar{C} can be computed by *excluding* $P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ from the feasible region of our MILP relaxation.

Let $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ be defined with respect to \bar{P} akin to how $P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ is defined with respect to P in (20), that is, through the imposition of the same fixing constraints. Then, we have

$$P \setminus P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}}) \subseteq \bar{P} \setminus \bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}}). \quad (23)$$

So instead of excluding $P(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ from our MILP relaxation $\bar{\mathcal{P}}$, we can exclude $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$.

A common way of excluding a given discrete decision from the solution space is through a no-good cut [36, 37]. Extending such a combinatorial cut to the SOS2 variables results in the strict inequality

$$\sum_{j=1}^n \sum_{k \in \mathcal{I}_j} \left(1 - \lceil \bar{\xi}_k^j \rceil\right) \xi_k^j + \sum_{l=1}^p \left\{ \bar{y}_l (1 - y_l) + (1 - \bar{y}_l) y_l \right\} > 0, \quad (24)$$

which indeed implies precisely the desired result. However, imposing a strict inequality may lead to unreachable optima and numerical instabilities.

We propose a weaker cut based on the extended formulation of the SOS2 constraints using additional binary variables. Consider the extended formulation of the SOS2 constraints in (19). Then, the no-good cut

$$\sum_{j=1}^n \sum_{k \in \mathcal{I}_j} \left(1 - \lceil \bar{\xi}_k^j \rceil\right) \zeta_k^j + \sum_{l=1}^p \left\{ \bar{y}_l (1 - y_l) + (1 - \bar{y}_l) y_l \right\} \geq 1 \quad (25)$$

excludes *at least* the interior of $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ from our MILP relaxation. Note that, due to the SOS2 constraints over each ζ^j , we restrict the cut to the indices $k \in \mathcal{I}_j$ for which $\bar{\xi}_k^j = 0$. Nevertheless, differently from the fixing step, the addition of cut (25) to our MILP relaxation *requires* the extended formulation of the SOS2 constraints, that is, the addition of the ζ binary variables to $\bar{\mathcal{P}}$.

To see the effect of adding cut (25) to the MILP relaxation $\bar{\mathcal{P}}$, note that its feasible region can be expressed as

$$\bar{P} = \bigcup_{\substack{\mathbf{y} \in \{0,1\}^p, \\ \zeta^j \in \text{SOS2}(\mathcal{I}_j), j=1, \dots, n}} \bar{P}(\zeta, \mathbf{y}), \quad (26)$$

which is in line with the classic understanding of the feasible region of MILP problems as the union of finitely many polyhedra. Thus, the no-good cut excludes $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ from the union. Because there might be other polyhedra that intersect $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ at its border, the cut is only guaranteed to exclude the interior from the feasible region. However, iteratively solving the relaxation and adding such cuts is guaranteed to eventually exclude the entire solution space.

Finally, we note that most formulations for the SOS2 constraints found in the literature [16, 17] can be used with only minor modifications to cut (25).

3.4 The Algorithm

The proposed relax-fix-and-exclude approach is formalized in Algorithm 1. Our algorithm progressively explores the solution space of the MINLP problem \mathcal{P} guided by the solutions of its MILP relaxation $\bar{\mathcal{P}}$. An optimal solution to $\bar{\mathcal{P}}$ gives a dual bound, which we use as a stopping criterion. If the criterion has not been met, we construct the NLP subproblem by *fixing* the binary and SOS2 variables. An optimal solution to the NLP subproblem yields a candidate solution and, possibly, a tighter primal bound. Then, the hyperrectangle containing the feasible region of the NLP subproblem is *excluded* from the MILP relaxation. An optimal solution to the updated MILP relaxation yields a tighter dual bound and a new hyperrectangle to explore. Figure 3 illustrates the main steps of the algorithm. A proof of the algorithm's convergence is established through Theorem 1.

The algorithm's guarantees rely on the availability of a global solver for the NLP subproblem. The fact that the solution found at Line 4 is a global optimum of the subproblem is what allows us to exclude $\bar{P}(\lceil \bar{\xi} \rceil, \bar{\mathbf{y}})$ from our MILP relaxation and, thus, guide our exploration of the search space and generate tighter dual bounds.

We note, however, that this is not such a strong assumption, as although the NLP subproblem may be nonconvex, it is smooth and bounded. In fact, many contemporary solvers (e.g., Gurobi [25], BARON [28, 29], and SCIP [30, 31]) already provide global optimality guarantees for these types of problems, exploiting spatial branch-and-bound and domain reduction techniques [38].

Theorem 1. *Algorithm 1 terminates in a finite number of steps and returns either (i) ∞ , if the problem is infeasible, (ii) $-\infty$ if the problem is unbounded, or (iii) an optimal solution to \mathcal{P} .*

Proof. First, we note that the feasible region of the NLP subproblems define an exact cover for the feasible region of \mathcal{P} . More precisely, by the definition of $P(\boldsymbol{\zeta}, \mathbf{y})$ in (20), we have, as already presented in Section 3.3,

$$P = \bigcup_{\substack{\mathbf{y} \in \{0,1\}^p, \\ \boldsymbol{\zeta}^j \in \text{SOS2}(\mathcal{I}_j), j=1, \dots, n}} P(\boldsymbol{\zeta}, \mathbf{y}), \quad (27)$$

which is a direct consequence of the partitioning induced in the x -space from the breakpoints $\mathcal{B}_1, \dots, \mathcal{B}_n$. To see this, recall that fixing $\boldsymbol{\zeta}$ is equivalent to restraining \mathbf{x} to lie in a particular hyperrectangle of \mathbb{R}^n defined by consecutive breakpoints (see Section 3.2). Similarly, the family of sets $\bar{P}(\boldsymbol{\zeta}, \mathbf{y})$ also defines an exact cover for \bar{P} .

The finite termination of the algorithm is a direct consequence of the exclude step at Line 9, which guarantees that the same subproblem will never be visited twice. Putting it into terms, we have that, for any $(\bar{\boldsymbol{\xi}}, \bar{\boldsymbol{\lambda}}, \bar{\mathbf{y}}) \in \bar{P}$,

$$(\boldsymbol{\xi}, \boldsymbol{\lambda}, \mathbf{y}) \in \bar{P} \setminus \bar{P}(\lceil \bar{\boldsymbol{\xi}} \rceil, \bar{\mathbf{y}}) \implies \lceil \boldsymbol{\xi} \rceil \neq \lceil \bar{\boldsymbol{\xi}} \rceil. \quad (28)$$

Thus, once the multilinear interpolation implies that our feasible region is bounded, we will eventually run out of subproblems (or hyperrectangles) to explore, resulting in $\bar{P} = \emptyset$ and inducing $\bar{C} = \infty$ at Line 10, which meets the termination condition.

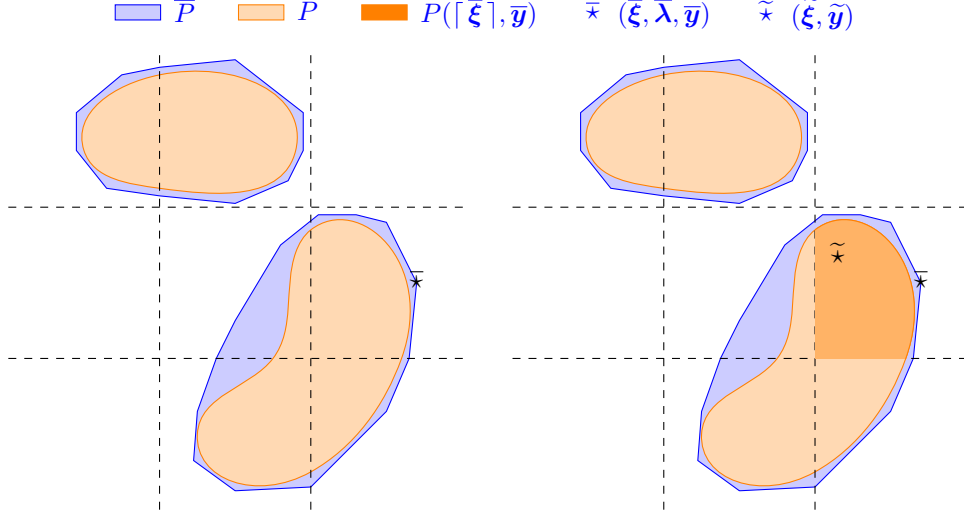
Consider now the case in which the algorithm terminates and no candidate solution is found, i.e., the algorithm returns $C = \infty$. This implies that the termination condition of the while loop was met through $\bar{C} = \infty$. But for any $(\bar{\boldsymbol{\xi}}, \bar{\boldsymbol{\lambda}}, \bar{\mathbf{y}}) \in \bar{P}$, we have $P(\bar{\boldsymbol{\xi}}, \bar{\mathbf{y}}) \subseteq \text{Proj}_{\boldsymbol{\xi}\mathbf{y}}(\bar{P}(\bar{\boldsymbol{\xi}}, \bar{\mathbf{y}}))$ and

$$P \setminus P(\bar{\boldsymbol{\xi}}, \bar{\mathbf{y}}) \subseteq \text{Proj}_{\boldsymbol{\xi}\mathbf{y}}(\bar{P} \setminus \bar{P}(\bar{\boldsymbol{\xi}}, \bar{\mathbf{y}})), \quad (29)$$

where $\text{Proj}_{\boldsymbol{\xi}\mathbf{y}}(\cdot)$ denotes the projection in the $(\boldsymbol{\xi}, \mathbf{y})$ -space. So, once no feasible solution was found during the iterations of the while loop, we have a guarantee that $P = \emptyset$.

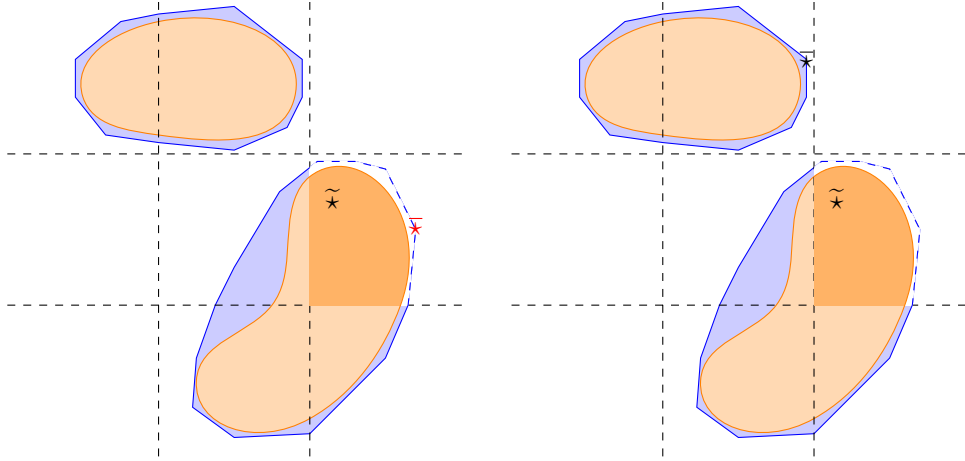
Suppose now the algorithm terminates with a candidate solution that is not optimal. Thus, there must exist a solution $(\boldsymbol{\xi}', \mathbf{y}') \in P$ with an associated cost $C' < C$. However, once we assume the NLP solver used in the fix step (at Line 4) has global optimality guarantees, $(\boldsymbol{\xi}', \mathbf{y}')$ must not be feasible for any of the NLP subproblems visited. In other words, it must be true that $(\boldsymbol{\xi}', \mathbf{y}') \in \text{Proj}_{\boldsymbol{\xi}\mathbf{y}}(\bar{P}(\bar{\boldsymbol{\xi}}, \bar{\mathbf{y}}))$, where \bar{P} is the feasible region of the MILP at the termination, i.e., after all exclusions. But then we reach a

contradiction, as the termination condition implies $C' \geq \overline{C} \geq C$. Thus, the candidate solution returned by the algorithm must be optimal. \square



(a) Before the first iteration, a solution \bar{x} for the MILP relaxation \bar{P} is computed (Line 2)

(b) In the loop, at Line 4, the NLP subproblem over $P([\tilde{\xi}], \tilde{y})$ is constructed by *fixing* the integer variables of P with the values of \bar{P} . The subproblem is solved to optimality (\tilde{x}), which we note results in a feasible solution for P .



(c) The exclude step (Line 9) adds a cut to \bar{P} that removes the interior of the hyperrectangle containing $P([\tilde{\xi}^*, \tilde{y}^*])$. Note that \bar{P} is now an MILP relaxation over the “unexplored” region of P .

(d) Finally, a new optimal is computed for \bar{P} which is a lower bound of the unexplored region, and, thus, induces a stopping criteria.

Fig. 3: Illustration of fixing and excluding in Algorithm 1. The orange and blue regions represent, respectively, the feasible regions of P and \bar{P} , and the dashed lines represent the partitioning induced by the look-up table. Points $(\bar{x}, \bar{y}, \bar{\xi}, \bar{\lambda})$ and $(\tilde{x}, \tilde{y}, \tilde{\xi}, \tilde{\lambda})$ represent optimal solutions for \bar{P} and the subproblem $P([\tilde{\xi}], \tilde{y})$.

Algorithm 1 The Relax-Fix-and-Exclude algorithm. We assume the availability of global solvers both for MILP and NLP that return $-\infty$ as cost if the problem is unbounded and ∞ if the problem is infeasible.

Require: Breakpoints $\mathcal{B}_j = \{\hat{x}_j^k\}_{k \in \mathcal{I}_j}$ for $j = 1, \dots, n$, and look-up tables $\{(\hat{\mathbf{x}}, f_i(\hat{\mathbf{x}})) : \hat{\mathbf{x}} \in \mathcal{B}_1 \times \dots \times \mathcal{B}_n\}$ for $i = 0, \dots, n$.

- 1: $C \leftarrow \infty$ (infeasible)
- 2: Solve the MILP relaxation: find an optimal solution $(\bar{\xi}, \bar{\lambda}, \bar{y}) \in \bar{P}$ with cost \bar{C}
- 3: **while** $\bar{C} < C$ **do**
- 4: (Fix) Solve the NLP subproblem: find an optimal solution $(\tilde{\xi}, \tilde{y}) \in P(\lceil \bar{\xi} \rceil, \bar{y})$ with cost \tilde{C}
- 5: **if** $\tilde{C} < C$ **then**
- 6: $C \leftarrow \tilde{C}$
- 7: $(\xi^*, y^*) \leftarrow (\tilde{\xi}, \tilde{y})$
- 8: **end if**
- 9: (Exclude) Add cut (25) to \bar{P} using $\lceil \bar{\xi} \rceil$ and \bar{y}
- 10: Solve the restricted MILP relaxation: update the optimal solution $(\bar{\xi}, \bar{\lambda}, \bar{y}) \in \bar{P}$ and the cost \bar{C}
- 11: **end while**
- 12: **if** $|C| < \infty$ **then return** (ξ^*, y^*)
- 13: **else return** C
- 14: **end if**

In Algorithm 1, the relax-fix-and-exclude approach for the MINLP problem has three potential outcomes:

1. *Infeasibility* (∞): The algorithm outputs ∞ if either the MILP relaxation is infeasible⁴, or all NLP subproblems are infeasible. Because the union of the feasible regions of the NLP subproblems are precisely the feasible region of \mathcal{P} , the algorithm returns ∞ if, and only if, the original MINLP is infeasible.
2. *Optimal Solution*: An optimal solution is reached when the lower bound \bar{C} from the MILP relaxation converges to (or surpasses) the cost C of the best-known feasible solution (the incumbent). This convergence implies that no further improvement is possible, confirming the incumbent as the global optimum of the problem.
3. *Unbounded Solution* ($-\infty$): If, at any point, the NLP solution within a fixed region reveals an unbounded solution (where $C = -\infty$), the algorithm returns $-\infty$, signifying that the original MINLP problem is unbounded. This outcome occurs when certain regions yield solutions that are unbounded in the objective due to either the problem structure or insufficient constraints.

⁴Assuming that $\bar{C} < C$ evaluates to false when both variables take ∞ .

4 Oil Production Optimization

Oil Production Optimization (OPO) [39] is essential for maximizing economic gains in the petroleum industry while adhering to operational and regulatory constraints. Short-term production optimization focuses on control decisions with a planning horizon from hours to days, manipulating the production system to achieve an optimal steady-state operation [40]. Such optimization problems involve the complexities of fluid flow physics and the subtleties of various artificial-lifting technologies (such as gas-lift systems and electrical submersible pumps) while addressing the discrete decisions of having multiple sources [41].

4.1 Problem Statement

The problem’s objective rests on maximizing the oil production of a subsea system consisting of sources (oil wells), connections (flowlines, pipelines, and manifolds), and a single sink (offshore platform). Sources and sink establish pressure references: reservoir pressure, in the case of wells, and separator pressure at the platform on the topside. Choke valves and lift-gas injection⁵ increase and decrease (resp.) the pressure at specific network points. The operators aim to maximize oil production by determining which wells will be open and which will be closed, the optimal opening of the valves, and the rate of lift-gas injected into the oil wells, while honoring flow assurance constraints and maintaining network balance [41]. We refer the reader to Müller et al.[40] for a recent work on short-term production optimization of complex production systems, considering flow assurance constraints, flow routing decisions, and multiple artificial lifting techniques, among other features.

Figure 4 depicts an offshore platform that produces from satellite wells (connected directly to the platform) and wells connected to a subsea manifold, from which the mixed flow is transferred by risers to the platform. At the topside, the production is processed by separators that split the flow into streams of i) water treated before being reinjected or discarded, ii) oil transferred to onshore terminals by tankers, and iii) gas which is exported by subsea pipelines, used for artificial lifting or consumed by turbo generators.

4.2 Vertical Lift Performance (VLP) curve

The Vertical Lift Performance (VLP) curve is a fundamental tool used to model the pressure drop across various components in an oil production system. It describes the relationship between the flow rate of fluids and the pressure required to transport them through these components. The VLP curve encapsulates the complex interactions of multiphase fluid flow, accounting for factors like fluid properties, gas-oil ratio (GOR), water cut (WCT), component geometry, and operational conditions. For this reason, it is often impractical to model VLP curves analytically in optimization problems; instead, high-fidelity simulations are employed to approximate the complex nonlinear relationships inherent in multiphase fluid flow within the network. In other words,

⁵In this work, lift-gas injection is the only artificial lifting method modeled.

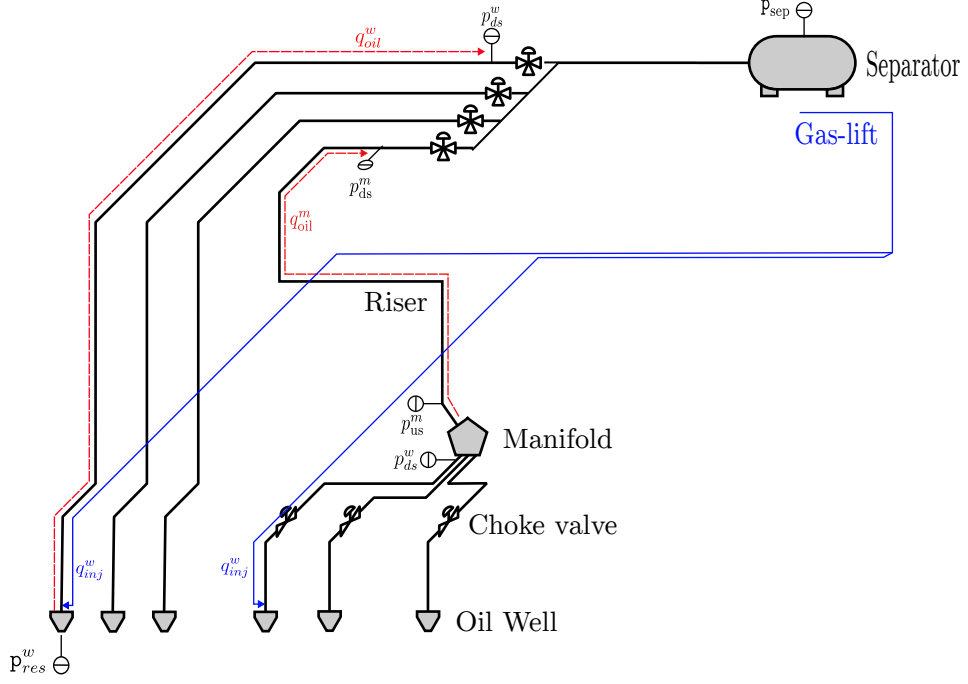


Fig. 4: Platform diagram based on Müller et al.[40]. The platform coordinates the production from satellite wells, which flow directly to the platform, and wells connected to a subsea manifold, which collects the production and then sends it to the platform through riser pipelines. At the platform, the mixed production is processed by separators that split the production flow into water, oil, and gas streams.

practical considerations deem the VLP curve a black-box function accessible only through expensive, single-point evaluations.

In this work, we employ the VLP curve to compute the *upstream* pressure of a component as a function of its downstream pressure, the liquid flow rate, the mixture components (shares of water, oil, and gas), and the rate of lift-gas injected. For a component c of the system, we formulate the VLP curve as a function:

$$f_{us}^c : \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+ \times [0, 1] \longrightarrow \mathbb{R}_+ \quad (30)$$

$$q_{liq}^c, p_{ds}^c, iglr^c, gor^c, wct^c \longmapsto p_{us}^c,$$

in which q_{liq}^c is the liquid flow rate at that component, p_{ds}^c is its downstream pressure, $iglr^c$ is the injected-gas to liquid ratio, and gor^c and wct^c are, respectively, the gas-oil ratio and the water-cut ratio of the flowing mixture. Note that, for every component, p_{ds}^c and $iglr^c$ are directly controlled. The former is done by opening and closing choke valves, and the latter by lift-gas injection.

VLP curves are computed for each oil well in the system, taking as reference the reservoir pressure (p_{us}^c) and the pressure at the choke valve just upstream of the nearest

connection (p_{ds}^c). Therefore, for satellite wells, which are connected directly to the production platform, p_{ds}^c is the pressure at the separator inlet to which the well is connected, whereas for manifold wells, which are connected to subsea manifolds, p_{ds}^c is the pressure at its manifold valve. VLP is also computed for the manifolds to account for the pressure drop along the riser pipeline, such that the upstream pressure is the pressure after the wells' valves, and the downstream pressure is the pressure at the valve connected to the separator.

4.3 Network Modeling

Let \mathcal{W} be the set of all oil wells and \mathcal{M} , the set of all subsea manifolds. Further specializing, let \mathcal{W}_{sat} be the set of all satellite wells (i.e., those that are connected directly to the production platform) and $\mathcal{W}_{\text{man}} = \bigcup_{m \in \mathcal{M}} \mathcal{W}_{\text{man}}^m$ be the set of manifold wells, where $\mathcal{W}_{\text{man}}^m$ are the wells connected to manifold m . Then, the oil production optimization problem is formulated as

$$\begin{aligned} \max \quad & \sum_{w \in \mathcal{W}} q_{\text{oil}}^w \\ \text{s.t.} \quad & \left. \begin{aligned} \text{wct}^c &= q_{\text{water}}^c / q_{\text{liq}}^c \\ q_{\text{oil}}^c &= q_{\text{liq}}^c - q_{\text{water}}^c \\ \text{gor}^c &= q_{\text{gas}}^c / q_{\text{oil}}^c \\ \text{iglr}^c &= q_{\text{inj}}^c / q_{\text{liq}}^c \end{aligned} \right\} \quad \forall c \in \mathcal{W} \cup \mathcal{M} \end{aligned} \quad (31a)$$

$$\left. \begin{aligned} q_{\text{oil}}^m &= \sum_{w \in \mathcal{W}_{\text{man}}^m} q_{\text{oil}}^w \\ q_{\text{water}}^m &= \sum_{w \in \mathcal{W}_{\text{man}}^m} q_{\text{water}}^w \\ q_{\text{gas}}^m &= \sum_{w \in \mathcal{W}_{\text{man}}^m} q_{\text{gas}}^w \\ q_{\text{inj}}^m &= \sum_{w \in \mathcal{W}_{\text{man}}^m} q_{\text{inj}}^w \end{aligned} \right\} \quad \forall m \in \mathcal{M} \quad (31b)$$

$$p_{\text{us}}^c = f_{\text{us}}^c(q_{\text{liq}}^c, p_{\text{ds}}^c, \text{iglr}^c, \text{gor}^c, \text{wct}^c) \quad \forall c \in \mathcal{W} \cup \mathcal{M} \quad (31c)$$

$$q_{\text{liq}}^w = \text{PI}^w \cdot (\bar{p}_{\text{res}}^w \cdot y^w - p_{\text{us}}^w) \quad \forall w \in \mathcal{W} \quad (31d)$$

$$y^m \leq \sum_{w \in \mathcal{W}_{\text{man}}^m} y^w \quad \forall m \in \mathcal{M} \quad (31e)$$

$$y^w \leq y^m \quad \forall w \in \mathcal{W}_{\text{man}}^m, \forall m \in \mathcal{M} \quad (31f)$$

$$\sum_{w \in \mathcal{W}} q_{\text{liq}}^w \leq \bar{q}_{\text{liq}}^{\max} \quad (31\text{g})$$

$$\sum_{w \in \mathcal{W}} q_{\text{inj}}^w \leq \bar{q}_{\text{inj}}^{\max} \quad (31\text{h})$$

$$\Delta p_{\text{choke}}^w = p_{\text{ds}}^w - \bar{p}_{\text{sep}} \quad \forall w \in \mathcal{W}_{\text{sat}} \quad (31\text{i})$$

$$\Delta p_{\text{choke}}^w = p_{\text{ds}}^w - p_{\text{us}}^m \quad \forall w \in \mathcal{W}_{\text{man}}^m, \forall m \in \mathcal{M} \quad (31\text{j})$$

$$\Delta p_{\text{choke}}^m = p_{\text{ds}}^m - \bar{p}_{\text{sep}} \quad \forall m \in \mathcal{M} \quad (31\text{k})$$

$$\left. \begin{array}{l} p_{\text{us}}^c - p_{\text{ds}}^c \geq -M \cdot (1 - y^c) \\ y^c \in \{0, 1\} \end{array} \right\} \quad \forall c \in \mathcal{W} \cup \mathcal{M} \quad (31\text{l})$$

$$\left. \begin{array}{l} t_{\text{GL}}^w \cdot \bar{q}_{\text{inj},\min}^w \leq q_{\text{inj}}^w \leq t_{\text{GL}}^w \cdot \bar{q}_{\text{inj},\max}^w \\ t_{\text{GL}}^w \in \{0, 1\} \end{array} \right\} \quad \forall w \in \mathcal{W}. \quad (31\text{m})$$

where PI^w is the productivity index (PI) of well w , \bar{p}_{res}^w is the reservoir pressure, $\bar{q}_{\text{liq}}^{\max}$ and $\bar{q}_{\text{inj}}^{\max}$ establish the platform's processing capacity, \bar{p}_{sep} is the separator pressure, $\bar{q}_{\text{inj},\min}^w$ and $\bar{q}_{\text{inj},\max}^w$ gas-injection bounds on wells when operating, and M is a sufficiently large constant. These terms are parameters of the model, which are differentiated from variables by a monospaced typography.

The objective aims to maximize the total oil production, which promotes economic gain. All variables are restricted to be non-negative, which is omitted for simplicity. Effectively, the only decision variables are the lift-gas injection rates q_{inj}^w and the pressure drops $\Delta p_{\text{choke}}^w$ across choke valves, as all the others are uniquely determined by assignments of those. Exceptions are the y^c binary variables, which determine whether component c is active (open) or not (closed), and the binary variables t_{GL}^w that impose limits on lift-gas injection when a well is operating.

A brief description of each constraint follows:

- (31a) ensures that the mixture fractions correspond to the flow rates in each component of the network. However, the relations are modeled with bilinear terms to avoid numerical problems involving the division operation, e.g., the relation $wct^w = q_{\text{water}}^w / q_{\text{liq}}^w$ is modeled as $wct^w \cdot q_{\text{liq}}^w = q_{\text{water}}^w$.
- (31b) models the mixture of the production streaming into a manifold from its connected wells without flow splitting, which is allowed in very peculiar conditions [42].
- (31c) defines the VLP curve for each component of the network.
- (31d) is the Inflow Performance Relationship (IPR), modeled here as a linear constraint, but could be treated as another black-box function for increased precision (see, for example, [43]).⁶

⁶Considering the approximation of the black-box function p_{us}^w as a multilinear interpolation model, notice that $p_{\text{us}}^w = 0$ when well w is inactive because the convex combination of well w 's variables add up to y^w .

- (31e-31f) ensure that the manifold is active if, and only if, at least one of the connected wells is active.
- (31g-31h) implement the platform constraints limiting the liquid production (separator capacity) and the total lift-gas injected (compressor capacity).
- (31i-31k) establish the pressure relations between the components through choke valves, ensuring pressure balance.
- (31l) controls the activation ($y^c = 1$) and deactivation ($y^c = 0$) of each component, ensuring feasibility of the model when the choke valves are completely closed.
- (31m) ensures that, if well w is operating with lift-gas injection ($t_{GL} = 1$), the operational limits are respected.

Note that problem (31) can be reformulated to fit the format of problem (1), enabling the use of our relax-fix-and-exclude algorithm to compute a global optimum. The black-box functions are the VLP curves $p_{us}^c(\cdot)$ that model the pressure drop through the components, which we approximate using multilinear interpolation. Although (31) contains linear and bilinear constraints not explicitly present in (1), the former are trivially handled by the algorithm, and the latter can be relaxed in the same way as the multilinear interpolants, as discussed in Section 3.1.

5 Computational Experiments

For our computational experiments, we consider an oil production optimization problem on offshore production platforms based on real-world cases from Petróleo Brasileiro S.A. (Petrobras). We compare our proposed algorithm, Relax-Fix-and-Exclude, with the Gurobi optimization software [25], the BARON solver [44], and the SCIP optimization suite [45].

5.1 Problem Instances

We consider nine scenarios (S1-S9) for the oil production optimization problem, each with a different arrangement of wells and manifolds. For each scenario, we generate eight random instances by uniformly sampling the problem parameters from realistic intervals. In all instances of all scenarios, the oil wells have the same geometry. However, because the each well is drilled at different location of the reservoir, each has a unique VLP curve.

The VLP curves are computed through a simulation software, which produces the look-up tables. We consider the GOR and WCT of each well as fixed parameters. Therefore, the resulting VLP curve for each well $w \in \mathcal{W}$ has breakpoints sampled from *three* dimensions: q_{liq}^w , p_{ds}^w , and $iglr^w$. Each manifold also requires a VLP curve to model the pressure drop in the pipeline that connects the manifold to the oil platform (riser), as can be seen in Figure 4. Because the oil-water-gas ratios of the mixture flowing out of the manifold is a function of the production from the connected wells, the VLP curve of the manifold must account for changes in the GOR and WCT inputs. However, we consider that the choke valve downstream of each manifold is always

Consequently, when well w is closed $\implies y^w = 0 \implies p_{us}^w = q_{liq}^w = q_{oil}^w = q_{water}^w = q_{gas}^w = 0$. Therefore, constraint (31d) is satisfied with $q_{liq}^w = 0$.

completely open, such that $p_{\text{ds}}^m = p_{\text{sep}}$ for every manifold m . In other words, the VLP curve for each manifold $m \in \mathcal{M}$ has breakpoints sampled from *four* dimensions: q_{liq}^m , iglr^m , gor^m , and wct^m .

The scenarios are ordered from S1 to S9 by complexity and expected difficulty of the instances, as can be seen in Table 1. The complexity increase comes from adding more components (oil wells and manifolds) to the model, which are equivalent to adding new multilinearly-interpolated black-box functions. In particular, the addition of a manifold significantly increases the difficulty of the instance, even when compared to the addition of an oil well, as the manifold functions have an extra domain dimension.

Table 1: Overview of the MINLP problem sizes in our experiments. For each scenario, eight instances were randomly generated by uniformly sampling the problem parameters. “Nonzeros” is the number of nonzero elements in the constraint matrix, indicating the problem sparsity.

Scenario	$ \mathcal{W} $	$ \mathcal{M} $	$ \xi $	$ \lambda $	$ \mathbf{y} $	Rows	Columns	Nonzeros
S1	1	0	50	4000	2	1889	4267	33505
S2	2	0	100	8000	4	3777	8529	67011
S3	3	0	150	12000	6	5690	12791	100542
S4	4	1	255	46000	9	7549	47520	428042
S5	5	1	305	50000	11	9391	51782	461510
S6	6	1	355	54000	13	11149	56044	494886
S7	7	2	460	88000	16	13296	90773	822666
S8	8	2	510	92000	18	14857	95035	855861
S9	9	2	560	96000	20	16681	99297	889303

5.2 Relax-Fix-and-Exclude Implementation

Our algorithm was implemented in Julia [46], using the JuMP library [47] to manipulate the mathematical programming models. We used Gurobi as the `NLPsSolver` and the `MILPsSolver`. Note that Gurobi provides global optimality guarantees (up to specified tolerance) for non-convex quadratic problems, thus fulfilling our algorithm requirements.

We implement the exclusion cut through a custom SOS2 emulation using auxiliary binary variables, as discussed in Section 3.3. In particular, we emulate the SOS2 constraints using the *lambda model* [48], also known as *convex combination method* [16]. This approach has the advantage of not being invasive in the solver’s branching rules, although it results in larger MILP problems.

5.3 Results and Discussion

In our computational experiments, we measure the performance of the solution methods in terms of the best solution found within a 1-hour budget. We compare Relax-Fix-and-Exclude with three widely-used solvers for mixed-integer nonconvex problems: Gurobi v11.0.3 [25], SCIP v9.0.0 [45], and BARON v24.5.8 [44]. Recall that RFE requires the MILP relaxation to be formulated with the SOS2 constraints emulated

through binary variables. Furthermore, Gurobi and SCIP have their own implementation (native support) of SOS2 constraints, while BARON does not. Thus, to make a fair evaluation, we consider Gurobi and SCIP with the MINLP problem formulated with both the SOS2 constraints and the emulation using binary variables. In our experiments “Gurobi (SOS2)” and “SCIP (SOS2)” indicate the use of the MINLP formulation with the solver’s native implementation of the SOS2 constraints, while “Gurobi (Bin.)” and “SCIP (Bin.)” indicate that the formulation of the SOS2 with binary variables was used. All experiments reported below were performed on a workstation equipped with a 12th Gen Intel(R) Core(TM) i7-12700 processor (20 cores, 40 threads), 64 GB of RAM, and running Ubuntu 24.10 64-bit.

The overall performance of the solution methods in terms of instances solved to optimality within the time limit is illustrated in Figure 5. We notice a significant performance gap between Gurobi (with either formulation) and RFE, in comparison to BARON and SCIP. Furthermore, RFE was the solution method capable of solving most instances within the time limit.




Fig. 5: Total number of instances solved to optimality by each solution method within the time limit.

Table 2 shows further details of the number of instances solved within each scenario. We see that RFE and Gurobi were the only solution methods able to find feasible solutions to all instances in our experiments. In contrast, BARON and SCIP were deeply impacted by the addition of a manifold in scenarios 4 to 9, indicating poor scalability

with the input dimension of the multilinear interpolants, or, equivalently, with the number of variables in the nonlinear constraints. Furthermore, RFE outperforms Gurobi in total number of instances solved to optimality mostly due to its better performance on the most difficult instances (scenarios 7 to 9).

Table 2: Number of instances solved to optimality in each scenario of the oil production optimization problem, within 1 hour. The value in parenthesis indicates the number of instances for which the solution methods found at least one non-trivial feasible solution. All scenarios have the same number of instances (8).

Solver	S1	S2	S3	S4	S5	S6	S7	S8	S9	Total
RFE	8 (8)	8 (8)	8 (8)	8 (8)	8 (8)	8 (8)	6 (8)	3 (8)	2 (8)	59 (72)
Gurobi (SOS2)	8 (8)	8 (8)	8 (8)	8 (8)	6 (8)	8 (8)	2 (8)	1 (8)	0 (8)	49 (72)
Gurobi (Bin.)	8 (8)	8 (8)	8 (8)	7 (8)	6 (8)	5 (8)	1 (8)	0 (8)	1 (8)	44 (72)
SCIP	8 (8)	8 (8)	8 (8)	0 (1)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	24 (25)
BARON	8 (8)	7 (7)	2 (3)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	17 (18)

Suppose we restrict ourselves to the easy instances (without manifolds) from scenarios 1 to 3. In that case, we see mixed results, with Gurobi outperforming RFE most of the time, while SCIP and BARON are significantly worse than either. Table 3 summarizes these results in columns S1 to S3. Notice that the easier the instance, the higher the margin that Gurobi has over RFE, particularly through the problem formulation that uses binary variables. This explains the better performance of Gurobi indicated by Figure 5 for small time budgets (under 5 seconds). However, the conclusion is reversed for the larger instances, as shown in Table 3, columns S4 to S6. In fact, if we take into consideration all instances that both RFE and Gurobi (SOS2) solved to optimality in scenarios 3 to 9, RFE was on average 52% faster than Gurobi (SOS2).

Table 3: Average time to solve the instances to optimality in scenarios 1 to 6. The result is absent for SCIP and BARON in columns S4 to S6 because those solvers did not find optimal solutions to any instances in those scenarios. An asterisk (*) indicates that the solver did *not* find an optimal solution for some instances of the respective scenario, and the total time budget was considered for the average in those cases.

Solver	S1	S2	S3	S4	S5	S6
RFE	3.13	3.85	7.86	353.30	335.38	313.21
Gurobi (SOS2)	0.83	3.63	17.51	924.85	1571.51*	835.97
Gurobi (Bin.)	0.49	2.14	15.91	1769.34*	2035.02*	2030.69*
SCIP	21.91	241.43	2409.12	—	—	—
BARON	37.38	1271.03*	3271.66*	—	—	—

In scenarios 7 to 9, for which none of the methods solves all instances to optimality, we evaluate performance in terms of the average primal-dual gap. Because SCIP and

BARON did not find feasible solutions to the instances in those scenarios, they are excluded from the analysis. The results summarized in Table 4 indicate that even when RFE does not find an optimal solution, the returned candidate solution is still of high quality. In fact, in all instances of scenarios 7 to 9, RFE terminated with a gap below 1%, significantly better than Gurobi with both formulations. Furthermore, analyzing all instances in which neither RFE nor Gurobi proved optimality, RFE resulted in a gap *at least* 35% smaller, and on average 91% smaller.

Table 4: Percentual gap of the different solution methods in scenarios 7 to 9. The value shown is the average (maximum) over all instances.

Solver	S7	S8	S9
RFE	0.01 (0.03)	0.12 (0.33)	0.13 (0.57)
Gurobi (SOS2)	3.57 (24.65)	19.27 (49.31)	32.17 (109.65)
Gurobi (Bin.)	9.04 (30.61)	63.28 (229.00)	85.75 (356.28)

We further illustrate this aspect of RFE through two instances from scenarios 7 and 9, for which the performance of RFE and Gurobi is illustrated in Figure 6. In both examples, RFE takes longer than Gurobi to find a non-trivial feasible solution, but the solution it finds is already of very high quality (low gap). Indeed, this behavior was observed in all of our experiments, that is, over all instances, the largest gap observed for the first feasible solution found by RFE was 1.2%.




Fig. 6: Lower bound (objective value) progression during the time limit on two large instances of the oil production optimization problem. Dashed lines indicate upper (dual) bounds. Markers indicate the first feasible solution found and the optimal (if found).

6 Conclusions

This paper introduced a novel algorithm for Mixed-Integer Nonlinear Programming (MINLP) problems involving multilinear interpolations of look-up tables. The algorithm leverages a piecewise-linear relaxation that defines the convex envelope of the multilinear constraints, enabling efficient exploration of the solution space through variable fixing and exclusion strategies. We have provided proof of convergence for the algorithm, demonstrating that it always returns the optimal solution or a valid statement of infeasibility or unboundedness.

Our computational experiments on oil production optimization problems demonstrated the effectiveness of the proposed method, which consistently outperformed all state-of-the-art solvers. The results show that our algorithm, RFE, not only was able to provide optimal solutions faster, but also that it is a better heuristic approach (i.e., when used with a limited time budget) by providing better feasible solutions. On the instances with one manifold, RFE was able to find an optimum for all of them, including those that the alternative solvers could not solve to optimality within the time budget. Further, for the instances with two manifolds, the average gap of RFE was orders of magnitude smaller than the others. We highlight that these results corroborate our initial intuition that the proposed MILP relaxation is of high quality and results in a better enumeration of the combinatorial solution space.

While the results are promising, several avenues for future research could further enhance the performance and applicability of the algorithm. First, the proposed relaxation can be applied to a broader range of nonlinearities, as it is not inherently tailored to multilinear constraints. Second, solving the MILP relaxations remains the primary computational bottleneck, consuming, on average, 85% of the total runtime. To mitigate this, future implementations could explore methods for reusing solver information across iterations. For instance, generating warm-starts using incumbent solutions or recycling parts of the branch-and-bound tree unaffected by exclusion cuts could significantly reduce computational overhead.

Additionally, we expect the runtime to be considerably reduced by a more efficient implementation of the exclusion cuts over the SOS2 constraints. For example, implementing the cuts through modifications of the branching rules, exploiting the classical approach of Beale [49], could result in a more efficient formulation of the MILP relaxation, as it would have a significantly reduced number of integer variables.

Another approach would be to explore the formulations for SOS2 constraints. The literature presents multiple alternatives [16, 17] that might suit combinatorial cuts better than the convex combination, used in our experiments. Furthermore, the entire problem could admit different formulations based on the choice of representation of the SOS2 constraints using binary variables. Even using the convex combination, as in our experiments, additional choice constraints could be used to significantly reduce the number of continuous variables in our MILP relaxation. More precisely, one could formulate the problem using only n continuous ξ variables, and 2^n continuous λ variables. That would introduce products between those and the SOS2 binary variables, which could be handled through McCormick linearization in an extended formulation.

Acknowledgments. This research was funded in part by Petróleo Brasileiro S.A. (Petrobras) under grant SAP N° 4600677797, CAPES under grant 88887.827578/2023-00, and CNPq under grants 403092/2024-8, 402099/2023-0, and 308624/2021-1.

Code and Data Availability. Our complete implementation of the Relax-Fix-and-Exclude algorithm as well as all necessary data to reproduce our results is available at <https://github.com/gos-ufsc/pwnl-oil>.

References

- [1] Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* **13**(4), 455–492 (1998) <https://doi.org/10.1023/A:1008306431147>
- [2] Bischl, B., Binder, M., Lang, M., Pielok, T., Richter, J., Coors, S., Thomas, J., Ullmann, T., Becker, M., Boulesteix, A.-L., Deng, D., Lindauer, M.: Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *WIREs Data Mining and Knowledge Discovery* **13**(2), 1484 (2023) <https://doi.org/10.1002/widm.1484>
- [3] Shields, B.J., Stevens, J., Li, J., Parasram, M., Damani, F., Alvarado, J.I.M., Janey, J.M., Adams, R.P., Doyle, A.G.: Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **590**(7844), 89–96 (2021) <https://doi.org/10.1038/s41586-021-03213-y>
- [4] Rao, S.S.: *Engineering Optimization: Theory and Practice*, 1st edn. Wiley, Hoboken, NJ (2019). <https://doi.org/10.1002/9781119454816>
- [5] Epelle, E.I., Gerogiorgis, D.I.: A computational performance comparison of MILP vs. MINLP formulations for oil production optimisation. *Computers & Chemical Engineering* **140**, 106903 (2020) <https://doi.org/10.1016/j.compchemeng.2020.106903>
- [6] Grimstad, B., Foss, B., Heddle, R., Woodman, M.: Global optimization of multiphase flow networks using spline surrogate models. *Computers & Chemical Engineering* **84**, 237–254 (2016) <https://doi.org/10.1016/j.compchemeng.2015.08.022>
- [7] Gastellu-Etchegorry, J.P., Gascon, F., Estève, P.: An interpolation procedure for generalizing a look-up table inversion method. *Remote Sensing of Environment* **87**(1), 55–71 (2003) [https://doi.org/10.1016/S0034-4257\(03\)00146-9](https://doi.org/10.1016/S0034-4257(03)00146-9)
- [8] Pini, M., Spinelli, A., Persico, G., Rebay, S.: Consistent look-up table interpolation method for real-gas flow simulations. *Computers & Fluids* **107**, 178–188 (2015) <https://doi.org/10.1016/j.compfluid.2014.11.001>
- [9] Nelles, O., Fink, A.: Grid-Based Look-Up Table Optimization Toolbox. IFAC

- [10] Nelles, O.: Linear, polynomial, and look-up table models. In: Nelles, O. (ed.) *Nonlinear System Identification: From Classical Approaches to Neural Networks, Fuzzy Models, and Gaussian Processes*, pp. 249–278. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-47439-3_10
- [11] Bohn, C., Stober, P., Magnor, O.: An optimization-based approach for the calibration of lookup tables in electronic engine control. In: *IEEE Conference on Computer Aided Control System Design, IEEE International Conference on Control Applications, IEEE International Symposium on Intelligent Control*, pp. 2315–2320 (2006). <https://doi.org/10.1109/CACSD-CCA-ISIC.2006.4777001>
- [12] Gupta, M., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., Esbroeck, A.: Monotonic calibrated interpolated look-up tables. *Journal of Machine Learning Research* **17**(109), 1–47 (2016) <https://doi.org/10.48550/arXiv.1505.06378>
- [13] Furlan, F.F., de Andrade Lino, A.R., Matugi, K., Cruz, A.J.G., Secchi, A.R., de Campos Giordano, R.: A simple approach to improve the robustness of equation-oriented simulators: Multilinear look-up table interpolators. *Computers & Chemical Engineering* **86**, 1–4 (2016) <https://doi.org/10.1016/j.compchemeng.2015.12.014>
- [14] Martins, C.d.O., Furlan, F.F., Giordano, R.d.C.: Implementation of multilinear look-up tables as surrogate models for the ethyl transesterification reactor in equation-oriented simulator EMSO. *The Journal of Engineering and Exact Sciences* **6**(4), 0467–0473 (2020) <https://doi.org/10.18540/jcecvl6iss4pp0467-0473>
- [15] Beale, E.M.L.: A mathematical programming model for the long term development of an off-shore gas field. *Discrete Applied Mathematics* **5**(1), 1–9 (1983) [https://doi.org/10.1016/0166-218X\(83\)90011-2](https://doi.org/10.1016/0166-218X(83)90011-2)
- [16] Vielma, J.P., Ahmed, S., Nemhauser, G.L.: Mixed-integer models for nonseparable piecewise linear optimization: Unifying framework and extensions. *Operations Research* **58**(2), 303–315 (2010) <https://doi.org/10.1287/opre.1090.0721>
- [17] Silva, T.L., Camponogara, E.: A computational analysis of multidimensional piecewise-linear models with applications to oil production optimization. *European Journal of Operational Research* **232**(3), 630–642 (2014) <https://doi.org/10.1016/j.ejor.2013.07.040>
- [18] Huchette, J., Vielma, J.P.: Nonconvex piecewise linear functions: Advanced formulations and simple modeling tools. *Operations Research* **71**(5), 1835–1856 (2022) <https://doi.org/10.1287/opre.2019.1973>

- [19] Misener, R., Floudas, C.A.: Piecewise-linear approximations of multidimensional functions. *Journal of Optimization Theory and Applications* **145**(1), 120–147 (2010) <https://doi.org/10.1007/s10957-009-9626-0>
- [20] Rikun, A.D.: A convex envelope formula for multilinear functions. *Journal of Global Optimization* **10**(4), 425–437 (1997) <https://doi.org/10.1023/A:1008217604285>
- [21] Keha, A.B., de Farias, I.R., Nemhauser, G.L.: Models for representing piecewise linear cost functions. *Operations Research Letters* **32**(1), 44–48 (2004) [https://doi.org/10.1016/S0167-6377\(03\)00059-2](https://doi.org/10.1016/S0167-6377(03)00059-2)
- [22] McCormick, G.P.: Computability of global solutions to factorable nonconvex programs: Part I – convex underestimating problems. *Mathematical Programming* **10**(1), 147–175 (1976) <https://doi.org/10.1007/bf01580665>
- [23] Sundar, K., Nagarajan, H., Linderoth, J., Wang, S., Bent, R.: Piecewise polyhedral formulations for a multilinear term. *Operations Research Letters* **49**(1), 144–149 (2021) <https://doi.org/10.1016/j.orl.2020.12.002>
- [24] Kim, J., Richard, J.-P.P., Tawarmalani, M.: Piecewise polyhedral relaxations of multilinear optimization. *SIAM Journal on Optimization*, 3167–3193 (2024) <https://doi.org/10.1137/22M1507486>
- [25] Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2024). <https://www.gurobi.com>
- [26] Gleixner, A.M.: Exact and Fast Algorithms for Mixed-Integer Nonlinear Programming. Logos Verlag Berlin, Berlin, Germany (2015)
- [27] Weiser, A., Zarantonello, S.E.: A note on piecewise linear and multilinear table interpolation in many dimensions. *Mathematics of Computation* **50**(181), 189–196 (1988) <https://doi.org/10.1090/S0025-5718-1988-0917826-0>
- [28] Sahinidis, N.V.: BARON: A general purpose global optimization software package. *Journal of Global Optimization* **8**(2), 201–205 (1996) <https://doi.org/10.1007/BF00138693>
- [29] Zhang, Y., Sahinidis, N.V.: Solving continuous and discrete nonlinear programs with BARON. *Computational Optimization and Applications* (2024) <https://doi.org/10.1007/s10589-024-00633-0>
- [30] Bolusani, S., Besançon, M., Bestuzheva, K., Chmiela, A., Dionísio, J., Donkiewicz, T., Doornmalen, J., Eifler, L., Ghannam, M., Gleixner, A., Graczyk, C., Halbig, K., Hedtke, I., Hoen, A., Hojny, C., Hulst, R., Kamp, D., Koch, T., Kofler, K., Lentz, J., Manns, J., Mexi, G., Mühmer, E., Pfetsch, M.E., Schlösser, F., Serrano, F., Shinano, Y., Turner, M., Vigerske, S., Weninger, D., Xu, L.: The SCIP Optimization Suite 9.0. Technical report, Optimization Online (February

- 2024). <https://optimization-online.org/2024/02/the-scip-optimization-suite-9-0/>
- [31] Bolusani, S., Besançon, M., Bestuzheva, K., Chmiela, A., Dionísio, J., Donkiewicz, T., Doornmalen, J., Eifler, L., Ghannam, M., Gleixner, A., Graczyk, C., Halbig, K., Hedtke, I., Hoen, A., Hojny, C., Hulst, R., Kamp, D., Koch, T., Kofler, K., Lentz, J., Manns, J., Mexi, G., Mühmer, E., Pfetsch, M.E., Schlösser, F., Serrano, F., Shinano, Y., Turner, M., Vigerske, S., Weninger, D., Xu, L.: The SCIP Optimization Suite 9.0. ZIB-Report 24-02-29, Zuse Institute Berlin (February 2024). <https://nbn-resolving.org/urn:nbn:de:0297-zib-95528>
 - [32] Biegler, L.T., Grossmann, I.E.: Retrospective on optimization. *Computers & Chemical Engineering* **28**(8), 1169–1192 (2004) <https://doi.org/10.1016/j.compchemeng.2003.11.003>
 - [33] Beale, E., Tomlin, J.: Special facilities in a general mathematical programming system for nonconvex problems using ordered sets of variables. In: *Proceedings of the Fifth International Conference on Operational Research*, pp. 447–454 (1970)
 - [34] Castro, P.M.: Tightening piecewise McCormick relaxations for bilinear problems. *Computers & Chemical Engineering* **72**, 300–311 (2015) <https://doi.org/10.1016/j.compchemeng.2014.03.025>
 - [35] Al-Khayyal, F.A., Falk, J.E.: Jointly constrained biconvex programming. *Mathematics of Operations Research* **8**(2), 273–286 (1983) <https://doi.org/10.1287/moor.8.2.273>
 - [36] D’Ambrosio, C., Frangioni, A., Liberti, L., Lodi, A.: On interval-subgradient and no-good cuts. *Operations Research Letters* **38**(5), 341–345 (2010) <https://doi.org/10.1016/j.orl.2010.05.010>
 - [37] Balas, E., Jeroslow, R.: Canonical cuts on the unit hypercube. *SIAM Journal on Applied Mathematics* **23**(1), 61–69 (1972) <https://doi.org/10.1137/0123007>
 - [38] Puranik, Y., Sahinidis, N.V.: Domain reduction techniques for global NLP and MINLP optimization. *Constraints* **22**(3), 338–376 (2017) <https://doi.org/10.1007/s10601-016-9267-5>
 - [39] Foss, B., Jenson, J.P.: Performance analysis for closed-loop reservoir management. *SPE Journal* **16**(01), 183–190 (2010) <https://doi.org/10.2118/138891-PA>
 - [40] Müller, E.R., Camponogara, E., Seman, L.O., Hülse, E.O., Vieira, B.F., Miyatake, L.K., Teixeira, A.F.: Short-term steady-state production optimization of offshore oil platforms: wells with dual completion (gas-lift and ESP) and flow assurance. *TOP* **30**(1), 152–180 (2022) <https://doi.org/10.1007/s11750-021-00604-2>
 - [41] Codas, A., Camponogara, E.: Mixed-integer linear optimization for optimal lift-gas allocation with well-separator routing. *European Journal of Operational Research*

- 217**(1), 222–231 (2012) <https://doi.org/10.1016/j.ejor.2011.08.027>
- [42] Camponogara, E., Seman, L.O., Muller, E.R., Gaspari, L.K. Eduardo F. Miyatake, Vieira, B.F.: A ReLU-based linearization approach for maximizing oil production in subsea platforms: An application to flow splitting. *Chemical Engineering Science* **295**, 120165 (2024) <https://doi.org/10.1016/j.ces.2024.120165>
 - [43] Vogel, J.V.: Inflow performance relationships for solution-gas drive wells. *Journal of Petroleum Technology* **20**(01), 83–92 (1968) <https://doi.org/10.2118/1476-PA>
 - [44] Kılınç, M.R., Sahinidis, N.V.: Exploiting integrality in the global optimization of mixed-integer nonlinear programming problems with BARON. *Optimization Methods & Software* **33**(3), 540–562 (2018) <https://doi.org/10.1080/10556788.2017.1350178>
 - [45] Bolusani, S., Besançon, M., Bestuzheva, K., Chmiela, A., Dionísio, J., Donkiewicz, T., Doornmalen, J., Eifler, L., Ghannam, M., Gleixner, A., Graczyk, C., Halbig, K., Hedtke, I., Hoen, A., Hojny, C., Hulst, R., Kamp, D., Koch, T., Kofler, K., Lentz, J., Manns, J., Mexi, G., Mühmer, E., Pfetsch, M.E., Schlösser, F., Serrano, F., Shinano, Y., Turner, M., Vigerske, S., Weninger, D., Xu, L.: The SCIP Optimization Suite 9.0. Technical report, Optimization Online (February 2024). <https://optimization-online.org/2024/02/the-scip-optimization-suite-9-0/>
 - [46] Bezanson, J., Edelman, A., Karpinski, S., Shah, V.B.: Julia: A fresh approach to numerical computing. *SIAM Review* **59**(1), 65–98 (2017) <https://doi.org/10.1137/141000671>
 - [47] Lubin, M., Dowson, O., Dias Garcia, J., Huchette, J., Legat, B., Vielma, J.P.: JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation* (2023) <https://doi.org/10.1007/s12532-023-00239-3>
 - [48] Dantzig, G.B.: On the significance of solving linear programming problems with some integer variables. *Econometrica* **28**(1), 30–44 (1960) <https://doi.org/10.2307/1905292>
 - [49] Beale, E.M.L.: Branch and bound methods for mathematical programming systems. *Annals of Discrete Mathematics* **5**, 201–219 (1979) [https://doi.org/10.1016/S0167-5060\(08\)70351-0](https://doi.org/10.1016/S0167-5060(08)70351-0)