

# Spatial Branch-and-Bound for Computing Multiplayer Nash Equilibrium

Jakub Černý<sup>1</sup>, Shuvomoy Das Gupta<sup>2</sup>, Christian Kroer<sup>1</sup>

<sup>1</sup>Columbia University, USA

<sup>2</sup>Rice University, USA

jakub.cerny@columbia.edu, sd158@rice.edu, christian.kroer@columbia.edu

## Abstract

Equilibria of realistic multiplayer games constitute a key solution concept both in practical applications, such as online advertising auctions and electricity markets, and in analytical frameworks used to study strategic voting in elections or assess policy impacts in integrated assessment models. However, efficiently computing these equilibria requires games to have a carefully designed structure and satisfy numerous restrictions; otherwise, the computational complexity becomes prohibitive. In particular, finding even approximate Nash equilibria in general-sum normal-form games with two or more players is known to be PPAD-complete. Current state-of-the-art algorithms for computing Nash equilibria in multiplayer normal-form games either suffer from poor scalability due to their reliance on non-convex optimization solvers, or lack guarantees of convergence to a true equilibrium. In this paper, we propose a formulation of the Nash equilibrium computation problem as a polynomial complementarity problem and develop a complete and sound spatial branch-and-bound algorithm based on this formulation. We provide a qualitative analysis arguing why one should expect our approach to perform well, and show the relationship between approximate solutions to our formulation and that of computing an approximate Nash equilibrium. Empirical evaluations demonstrate that our algorithm substantially outperforms existing complete methods.

**Code** — [github.com/CoffeeAndConvexity/SBnBForNE](https://github.com/CoffeeAndConvexity/SBnBForNE)

**Extended version** — [arxiv.org/abs/2508.10204](https://arxiv.org/abs/2508.10204)

## 1 Introduction

Computing equilibria in games is a long-standing research topic in artificial intelligence and has served as an important benchmark domain. Advances in equilibrium-finding algorithms enabled breakthroughs in AI for recreational settings such as poker, Stratego, and Diplomacy (Bowling et al. 2015; Brown and Sandholm 2018; Perolat et al. 2022; Bakhtin et al. 2022), and have enabled decision-support systems for security and logistics applications (Pita et al. 2008; Fang, Stone, and Tambe 2015; Tambe 2011; Jain et al. 2010; Černý et al. 2024; Černý, Iyengar, and Kroer

2025). Deployed systems have helped allocate security resources across the eight terminals of Los Angeles International Airport (Pita et al. 2008), support biodiversity protection over large conservation areas (Fang, Stone, and Tambe 2015; Fang et al. 2016; Fang and Nguyen 2016), and assist in screening millions of US air passengers annually (Brown et al. 2016). Related models have also been explored for traffic monitoring, drug interdiction, and cybersecurity (Sinha et al. 2018). These applications depend on algorithms capable of computing game-theoretic equilibria at scale. Achieving such scale in practice typically requires modeling restrictions, most commonly reducing interactions to two players and, in many cases, assuming zero-sum utilities, which allows domain-agnostic methods to solve games efficiently.

While some interactions can be reasonably approximated as two-player, e.g., when dependencies between participants are limited or can be decomposed, many real-world settings inherently involve multiple interacting agents and must be modeled as multiplayer games. This is particularly common in computational economics, where various equilibrium concepts underpin deployed systems. Examples include Bayesian Nash equilibria in online advertising auctions (Edelman, Ostrovsky, and Schwarz 2007), market-clearing equilibria in electricity markets (Kirschen and Strbac 2018), competitive equilibria in matching platforms such as ride-sharing (Banerjee, Riquelme, and Johari 2015; Bimpikis, Candogan, and Saban 2019), and Nash equilibria in spectrum auctions and telecommunications markets (Cramton 1997; Cramton et al. 2006). These scenarios are fundamentally multiplayer in nature. Solving them at scale typically requires even stronger modeling assumptions and significant exploitation of domain-specific structure to design tractable algorithms.

Relaxing these assumptions to develop more generally applicable and domain-agnostic algorithms for multiplayer settings remains a central challenge. This is especially true in contexts beyond traditional economic markets, such as multiplayer card games, network-based games like pursuit-evasion or interdiction, or security scenarios involving more than two self-interested entities, such as coalition/team-based games in organizations like the UN or NATO. Even in one of the most studied models – finite normal-form games – approximating a general-sum Nash equilibrium becomes PPAD-complete with two or more players, while two-player

zero-sum games admit polynomial-time solutions.

Still, computing Nash equilibria in multiplayer normal-form games remains a traditional problem that has been studied for decades, and many classical algorithms exist. Support enumeration methods, such as full support enumeration and vertex enumeration, rely on exhaustively checking combinations of pure strategy supports and solving the resulting best-response constraints, typically expressed as systems of linear inequalities (Porter, Nudelman, and Shoham 2008; Lipton and Markakis 2004). Homotopy-based methods, including the Govindan-Wilson algorithm (Govindan and Wilson 2003), simplicial subdivision (van der Laan, Talman, and Van der Heyden 1987), and extended Lemke-type path-following algorithms (Lemke and Howson 1964; McKelvey and McLennan 1996; Turocy 2005), trace continuous solution paths from a perturbed game to the original game by following parameterized trajectories of fixed points. Polynomial system solvers approach the problem by encoding the equilibrium conditions, such as the Karush-Kuhn-Tucker conditions for each player’s optimization problem, as systems of polynomial equalities and inequalities, then solving them using Grobner bases, homotopy continuation, or algebraic elimination techniques (Datta 2010; McKelvey and McLennan 1996). See Berg and Sandholm (2017) for a more detailed overview.

Classical algorithms often suffer from numerical instability, particularly those based on solving polynomial systems, which can be sensitive to input perturbations and rounding errors. Their computational complexity also grows rapidly with the number of players and strategies, making them difficult to scale and limiting their applicability to relatively small games. Recent complete algorithms instead rely on formulations amenable to highly optimized industrial solvers such as CPLEX and Gurobi. These approaches use techniques such as decomposing multilinear terms into binary variables (Ganzfried 2024) or adding correlation plan constraints to restrict the feasible solution space (Zhang, An, and Subrahmanian 2023). An advantage of this approach is the ability to incorporate objectives into the formulation; for instance, to compute equilibria that maximize social welfare. Another modern approach is the exclusion method, a tree search-based algorithm that iteratively partitions the continuous solution space (Berg and Sandholm 2017). Among incomplete methods, one of the most scalable is the gradient-based ADIDAS algorithm, which combines homotopy methods with iterative polymatrix solvers and can handle games with billions of outcomes (Gemp et al. 2022). However, due to the use of biased gradient descent, ADIDAS lacks convergence guarantees to an exact equilibrium and instead offers local convergence only.

**Our contributions.** We propose that the Nash equilibrium problem in multiplayer normal-form games should be solved as a polynomial complementarity problem. This formulation allows us to design a sound and complete algorithm based on spatial branch-and-bound. We provide qualitative analysis arguing why the reformulated problem is more tractable, and prove that a suboptimal solution to our reformulation yields an approximate Nash equilibrium, with an explicit bound on

the approximation quality. Empirically, we demonstrate that our algorithm not only outperforms existing complete methods, but also remains competitive with the leading incomplete method on mid-sized games when run with early termination, while offering better-quality solutions. Large games still remain out of reach primarily due to the size of the utility matrix, which cannot be explicitly constructed. However, because our algorithm can serve as a subgame solver within incremental oracle-based methods, we see a path toward improving scalability in future work.

## 2 Preliminaries

We represent a multiplayer normal-form game as a tuple  $G = (N, A, u)$ , where  $N$  denotes a finite set of players with  $n = |N|$ . For each player  $i \in N$ , let  $A_i$  be the finite set of available actions, and define the joint action space as the Cartesian product  $A = \prod_{i \in N} A_i$ . A real-valued utility function  $u_i : A \rightarrow \mathbb{R}$  specifies the payoff for player  $i$  under each action profile  $a \in A$ , and we denote by  $\underline{u}_i$  and  $\bar{u}_i$  the minimal and maximal values of  $u_i$ , respectively. Mixed strategies are denoted by  $\Delta = \prod_{i \in N} \Delta_i$ , where each  $\Delta_i$  is the probability simplex over  $A_i$ . For any tuple indexed over players, we use the subscript  $-i$  to indicate the components corresponding to all players except  $i$ ; for example, for  $a = (a_1, \dots, a_n) \in A$ , we write  $a_{-i} = (a_1, \dots, a_n) \setminus \{a_i\}$ , and for a mixed strategy profile  $\delta \in \Delta$ , we write  $\delta_{-i}(a_{-i}) = \prod_{j \in N, j \neq i} \delta_j(a_j)$ , where  $\delta_j \in \Delta_j$ . The expected utility of player  $i$  under a mixed strategy profile  $\delta \in \Delta$  is given by

$$u_i(\delta) = \sum_{a \in A} \left( \prod_{j \in N} \delta_j(a_j) \right) u_i(a),$$

and the expected utility of player  $i$  when they play a specific action  $a_i \in A_i$  against  $\delta_{-i}$  is

$$u_i(a_i, \delta_{-i}) = \sum_{a_{-i} \in A_{-i}} \delta_{-i}(a_{-i}) u_i(a_i, a_{-i}).$$

A *Nash equilibrium* is a strategy profile  $\delta = (\delta_i)_{i \in N} \in \Delta$  such that no player can improve their expected utility by unilaterally deviating from their strategy: formally,

$$u_i(\delta_i, \delta_{-i}) \geq u_i(\delta'_i, \delta_{-i}) \quad \forall \delta'_i \in \Delta_i. \quad (\text{NE})$$

Finding a Nash equilibrium can be formulated as the following mixed-integer polynomial feasibility problem (Shoham and Leyton-Brown 2008):

$$\left( \begin{array}{ll} v_i = s_i(a_i) + u_i(a_i, \delta_{-i}) & \forall i, a_i \\ \delta_i(a_i) \leq z_i(a_i) & \forall i, a_i \\ s_i(a_i) \leq (1 - z_i(a_i))(\bar{u}_i - \underline{u}_i) & \forall i, a_i \\ \sum_{a_i \in A_i} \delta_i(a_i) = 1 & \forall i \\ 0 \leq \delta_i(a_i) \leq 1 & \forall i, a_i \\ 0 \leq s_i(a_i) \leq \bar{u}_i - \underline{u}_i & \forall i, a_i \\ \underline{u}_i \leq v_i \leq \bar{u}_i & \forall i \\ z_i(a_i) \in \{0, 1\} & \forall i, a_i \end{array} \right) \quad (\mathcal{P})$$

where  $\{\delta_i(a_i)\}_{i \in N, a_i \in A_i}$ ,  $\{v_i\}_{i \in N}$ ,  $\{z_i(a_i)\}_{i \in N, a_i \in A_i}$ ,  $\{s_i(a_i)\}_{i \in N, a_i \in A_i}$  are the decision variables. Moreover, we call a strategy profile  $\delta \in \Delta$  an  $\epsilon$ -equilibrium if for all  $i \in N$

$$u_i(\delta_i, \delta_{-i}) \geq u_i(\delta'_i, \delta_{-i}) - \epsilon \quad \delta'_i \in \Delta_i. \quad (\epsilon\text{-NE})$$

### 3 Our algorithm

This section is organized as follows. First, we propose a continuous formulation of  $(\mathcal{P})$  as a polynomial feasibility problem. Then, we discuss why a spatial branch-and-bound algorithm working with this continuous formulation can be significantly more tractable numerically than the original mixed-integer formulation. Finally, we present a customized spatial branch-and-bound algorithm that exploits the underlying structure of our continuous formulation to solve it more efficiently.

#### 3.1 Equivalent continuous formulation of $(\mathcal{P})$

The equivalent continuous formulation of  $(\mathcal{P})$  is as follows:

$$\left( \begin{array}{ll} v_i - u_i(a_i, \delta_{-i}) \geq 0 & \forall i, a_i \\ \delta_i(a_i) (v_i - u_i(a_i, \delta_{-i})) = 0 & \forall i, a_i \\ \sum_{a_i \in A_i} \delta_i(a_i) = 1 & \forall i \\ 0 \leq \delta_i(a_i) \leq 1 & \forall i, a_i \\ \underline{u}_i \leq v_i \leq \bar{u}_i & \forall i \end{array} \right) \quad (\mathcal{Q})$$

where the decision variables are now  $\{\delta_i(a_i)\}_{i \in N, a_i \in A_i}$ ,  $\{v_i\}_{i \in N}$ , which is a strict subset of the decision variables of  $(\mathcal{P})$ . The second constraint, which states that the product of two nonnegative terms must be zero, is known as a complementarity constraint. We now show that the continuous formulation  $(\mathcal{Q})$  and the mixed-integer formulation  $(\mathcal{P})$  are equivalent to each other.

**Proposition 1.** *The continuous formulation  $(\mathcal{Q})$  and the mixed-integer polynomial formulation  $(\mathcal{P})$  are equivalent.*

*Proof.* First, assume we have a feasible solution  $(\delta, v, z, s)$  for  $(\mathcal{P})$ , then we can show that  $(\delta, v)$  is a feasible solution to  $(\mathcal{Q})$ . The third and fourth constraints  $\sum_{a_i \in A_i} \delta_i(a_i) = 1$ ,  $0 \leq \delta_i(a_i) \leq 1$ , and  $\underline{u}_i \leq v_i \leq \bar{u}_i$  are identical in both formulations, so  $(\delta, v)$  will satisfy those trivially. For the first constraint of  $(\mathcal{Q})$ , note that from  $(\mathcal{P})$ , we have  $s_i(a_i) = v_i - u_i(a_i, \delta_{-i})$ , and  $s_i(a_i) \geq 0$ , which implies  $v_i - u_i(a_i, \delta_{-i}) \geq 0$ . To show that the second constraint of  $(\mathcal{Q})$  will be satisfied by  $(\delta, v)$ , consider two cases: (i)  $\delta_i(a_i) = 0$ , where the second constraint is trivially satisfied (ii)  $\delta_i(a_i) > 0$ , where  $\delta_i(a_i) \leq z_i(a_i)$  from  $(\mathcal{P})$  forces  $z_i(a_i) = 1$  implying  $s_i(a_i) \leq (1 - z_i(a_i))(\bar{u}_i - \underline{u}_i) = 0$ , and because  $s_i(a_i) \geq 0$  in  $(\mathcal{P})$ , we must have  $s_i(a_i) = v_i - u_i(a_i, \delta_{-i}) = 0$ , thus satisfying the second constraint of  $(\mathcal{Q})$  as well. Hence,  $(\delta, v)$  is a feasible solution to  $(\mathcal{Q})$ .

Next, assume we have a feasible solution  $(\delta, v)$  for the continuous problem  $(\mathcal{Q})$ . Then we can construct some  $(z, s)$  such that  $(\delta, v, z, s)$  is a feasible solution to  $(\mathcal{P})$ , which we can show as follows. Define  $s_i(a_i) := v_i - u_i(a_i, \delta_{-i})$  (thus automatically satisfying the first constraint of  $(\mathcal{P})$ ) and  $z_i(a_i) := 1$  if  $\delta_i(a_i) > 0$  and  $z_i(a_i) := 0$  if  $\delta_i(a_i) = 0$

(which satisfies the last constraint of  $(\mathcal{P})$  via construction). For the second constraint of  $(\mathcal{P})$ , we consider two cases: (i) if  $\delta_i(a_i) > 0$ , then  $z_i(a_i) = 1$  by construction and because  $\delta_i(a_i) \leq 1$ , it implies  $\delta_i(a_i) \leq z_i(a_i)$  holds, (ii) If  $\delta_i(a_i) = 0$ , then  $z_i(a_i) = 0$  by construction satisfying  $\delta_i(a_i) \leq z_i(a_i)$ . The constraint  $s_i(a_i) \geq 0$  is guaranteed by our construction and  $s_i(a_i) \leq \bar{u}_i - \underline{u}_i$  is always true as  $v_i \leq \bar{u}_i$  and  $u_i(a_i, \delta_{-i}) \geq \underline{u}_i$ . For  $s_i(a_i) \leq (1 - z_i(a_i))(\bar{u}_i - \underline{u}_i)$ , we consider two cases again: (i) if  $\delta_i(a_i) > 0$ , then  $z_i(a_i) = 1$  by construction and the second constraint of  $(\mathcal{Q})$  implies  $v_i - u_i(a_i, \delta_{-i}) = 0$  in this case so  $s_i(a_i) = 0$  thus  $0 = s_i(a_i) \leq (1 - z_i(a_i))(\bar{u}_i - \underline{u}_i) = 0$ , (ii) If  $\delta_i(a_i) = 0$ , then  $z_i(a_i) = 0$  by construction, leading to  $s_i(a_i) \leq (1 - z_i(a_i))(\bar{u}_i - \underline{u}_i) = \bar{u}_i - \underline{u}_i$ , and we showed earlier  $s_i(a_i) \leq \bar{u}_i - \underline{u}_i$  is true. Hence, any solution to  $(\mathcal{Q})$  implies a solution to  $(\mathcal{P})$ .  $\square$

To solve  $(\mathcal{Q})$  using the spatial branch-and-bound algorithm effectively, we work with the following penalized version of  $(\mathcal{Q})$ , where the complementarity constraint of  $(\mathcal{Q})$  is penalized through the  $\infty$ -norm via the epigraph approach (Boyd and Vandenberghe 2004, page 134) as:

$$\left( \begin{array}{ll} \min & \varpi \\ \text{s.t.} & \varpi \geq \delta_i(a_i) (v_i - u_i(a_i, \delta_{-i})) \quad \forall i, a_i \\ & \varpi \geq -\delta_i(a_i) (v_i - u_i(a_i, \delta_{-i})) \quad \forall i, a_i \\ & v_i - u_i(a_i, \delta_{-i}) \geq 0 \quad \forall i, a_i \\ & \sum_{a_i \in A_i} \delta_i(a_i) = 1 \quad \forall i \\ & 0 \leq \delta_i(a_i) \leq 1 \quad \forall i, a_i \\ & \underline{u}_i \leq v_i \leq \bar{u}_i \quad \forall i \end{array} \right) \quad (\mathcal{R})$$

where the decision variables are  $\{\delta_i(a_i)\}_{i \in N, a_i \in A_i}$ ,  $\{v_i\}_{i \in N}$  and the scalar epigraph variable  $\varpi$ . We now show that any solution to  $(\mathcal{Q})$  is also a solution to  $(\mathcal{R})$  and vice versa.

**Proposition 2.** *Problems  $(\mathcal{Q})$  and  $(\mathcal{R})$  are equivalent.*

*Proof.* Note that the constraints of  $(\mathcal{R})$  are the same as  $(\mathcal{Q})$ , except the complementarity constraint, which is penalized via the epigraph approach (Boyd and Vandenberghe 2004, page 134) with objective effectively minimizing the nonnegative objective  $\max_{i \in N, a_i \in A_i} |\delta_i(a_i) (v_i - u_i(a_i, \delta_{-i}))|$ . Because every finite game admits at least one Nash equilibrium (Nash 1951), there exists an optimal solution to  $(\mathcal{Q})$ , and by construction, such a Nash equilibrium will be a feasible solution to  $(\mathcal{R})$  with objective value 0, hence it will also be the globally optimal solution to  $(\mathcal{R})$  because 0 is the lowest possible value the objective of  $(\mathcal{R})$ .

Similarly, due to (Nash 1951), an optimal solution to  $(\mathcal{R})$  will have objective value 0 and such a solution will satisfy all the constraints shared by  $(\mathcal{R})$  and  $(\mathcal{Q})$ . Because the objective  $\max_{i \in N, a_i \in A_i} |\delta_i(a_i) (v_i - u_i(a_i, \delta_{-i}))|$  being zero is equivalent to each complementarity constraint being satisfied, a solution to  $(\mathcal{R})$  thus satisfies every constraints of  $(\mathcal{Q})$ .  $\square$

Proposition 2 leads to the following corollary:

**Corollary 1.** *The optimal objective value of  $(\mathcal{R})$  is zero and any optimal solution to  $(\mathcal{R})$  is an exact Nash equilibrium.*

Moreover, if any algorithm solving  $(\mathcal{R})$  terminates with a feasible solution having a non-zero objective value, we have the following guarantee regarding the solution being an  $\epsilon$ -Nash equilibrium.

**Proposition 3.** *A candidate solution  $\delta$  that satisfies all the constraints in  $(\mathcal{R})$  with objective value  $\varpi$  constitutes an  $\epsilon$ -Nash equilibrium with  $\epsilon = (\varpi \times \max_{i \in N} |A_i|)$ .*

*Proof.* For each player  $i$ , let  $v_i^*$  be their best-response value against the profile  $\delta_{-i}$ , i.e.,  $v_i^* = \max_{a_i \in A_i} u_i(a_i, \delta_{-i})$ . The utility deviation improvement  $\epsilon_i$  is then equal to

$$\begin{aligned} \epsilon_i &= v_i^* - u_i(\delta) = v_i^* - \sum_{a_i \in A_i} \delta_i(a_i) u_i(a_i, \delta_{-i}) \\ &\leq v_i - \sum_{a_i \in A_i} \delta_i(a_i) u_i(a_i, \delta_{-i}) \\ &= \sum_{a_i \in A_i} \delta_i(a_i) (v_i - u_i(a_i, \delta_{-i})) \leq \varpi |A_i|, \end{aligned}$$

where  $v_i^* \leq v_i$  follows from the third constraint in  $(\mathcal{R})$ , and the final inequality uses the first two constraints in  $(\mathcal{R})$ . Moreover, we use the fact that  $\sum_{a_i \in A_i} \delta_i(a_i) = 1$ . The approximation parameter of the equilibrium  $\epsilon$  is then equal to  $\epsilon = \max_{i \in N} \epsilon_i = \varpi \max_{i \in N} |A_i|$ .  $\square$

In the next section, we discuss why for an effective spatial branch-and-bound implementation, it is more convenient to work with  $(\mathcal{R})$  over  $(\mathcal{P})$ .

### 3.2 Tractability of $(\mathcal{R})$ over $(\mathcal{P})$

For a spatial branch-and-bound algorithm, working with the penalized continuous formulation  $(\mathcal{R})$  can be significantly more tractable than with the mixed-integer formulation  $(\mathcal{P})$ . This advantage stems from three primary factors: the dimensionality of the search space, the complexity of the branching scheme, and the ease of generating high-quality feasible solutions. To explain these advantages clearly, we first provide a brief overview of how a standard spatial branch-and-bound algorithm works, for more details we refer the interested readers to (Liberti 2008; Horst and Tuy 2013; Locatelli and Schoen 2013; Gurobi Optimization, LLC 2020).

**Brief review of the standard spatial branch-and-bound algorithm.** Consider a nonconvex mixed-integer polynomial minimization problem where the objective and constraint functions are polynomials, and the decision variables are a mix of continuous and binary types. We assume that the problem's feasible region is nonempty and bounded and that it admits a finite optimal value  $p^*$ , as is the case for our problem. The spatial branch-and-bound algorithm employs a divide-and-conquer strategy to solve the problem by recursively partitioning the domain to create a search tree. The process begins with a *presolve* phase, where a linear relaxation of the original problem is solved to obtain valid bounds on the continuous variables. Following this, the algorithm dynamically partitions the feasible region into a finite collection of subregions, each corresponding to a subproblem.

By definition, the smallest among the optimal values of these subproblems is the optimal value of the original problem. The speed of the algorithm is essentially controlled by the efficiency of solving these subproblems and, if necessary, branching them further into smaller subproblems, which we discuss next.

While solving a particular subproblem—which is also a nonconvex polynomial minimization problem, albeit over a smaller region—can be as hard as solving the original problem, computing lower and upper bounds for that subproblem via linear or convex relaxations is much easier. At the root node of the spatial branch-and-bound tree, a linear relaxation of the original problem is constructed and solved to yield a lower bound on  $p^*$ , denoted by  $\underline{p}^*$ . The tighter this relaxation, the closer  $\underline{p}^*$  is to  $p^*$ . In the absence of any special structure, the initial upper bound on  $p^*$  is set to  $\bar{p}^* = \infty$ . However, a user can *warm-start* the algorithm by providing a known feasible solution, which can establish a nontrivial initial  $\bar{p}^*$ . Efficient warm-starting procedures that exploit the problem structure can massively accelerate branch-and-bound solvers, a technique we will leverage for problem  $(\mathcal{R})$  (described in the next section). Note that a polynomial feasibility problem, such as finding a Nash Equilibrium in  $(\mathcal{P})$ , cannot be warm-started, as providing a feasible solution is equivalent to having already solved it. Throughout its execution, the algorithm updates  $\underline{p}^*$  and  $\bar{p}^*$  to progressively reduce the optimality gap,  $\bar{p}^* - \underline{p}^*$  and the algorithm terminates when  $\underline{p}^*$  matches  $\bar{p}^*$ , having found a globally optimal solution.

The algorithm reduces this gap by iteratively processing subproblems. It selects an active subregion from the list of unprocessed nodes in the search tree (typically guided by heuristics) and constructs two convex (usually linear) minimization subproblems. These convex formulations are created using convex lower and upper envelopes (McCormick 1976) for the nonlinear terms, whereas the linear constraints and terms remain unchanged. Additionally, valid linear inequalities (cuts) can be added to strengthen these convex relaxations (Sherali and Adams 1990; Padberg 1989; Sherali and Fraticelli 2002). Solving these convex subproblems provides valid lower and upper bounds on the optimal value for the active subregion, leading to one of the following three possibilities:

- **Possibility 1: pruning.** If the convex optimization problem associated with the lower bound is infeasible or its objective value is greater than the current global upper bound  $\bar{p}^*$ , then the subregion cannot contain the global optimum. Consequently, this subregion is discarded, or *pruned*, from the search tree.
- **Possibility 2: fathoming by optimality.** For the subproblem, if its lower bound matches its upper bound, we have found this subproblem's optimal solution without solving the original nonconvex problem on that subregion. This local solution is also a feasible solution to the main problem. If its objective value is better than the current  $\bar{p}^*$ , it becomes the new incumbent solution, and  $\bar{p}^*$  is updated accordingly. The subregion is now fully explored, or *fathomed*, and is not branched further.

- **Possibility 3: branching.** If neither of the first two possibilities occurs, the current subregion may still contain the global optimum and must be explored further. It is partitioned into smaller subregions by performing *branching*. The algorithm typically prioritizes *combinatorial branching* on a binary variable that has a fractional value in the convex relaxation’s solution. If all binary variables are integer-valued, then it performs *spatial branching* by selecting a continuous variable involved in a violated nonlinear constraint. These new, smaller subregions are then added to the list of active nodes to be processed.

This process is repeated, continually updating the global bounds. The global upper bound,  $\bar{p}^*$ , is always the objective value of the best-known feasible solution i.e., the incumbent. The global lower bound,  $\underline{p}^*$  is updated by taking the minimum of the best objective values of all the leaf nodes associated with the subproblems that are pruned or fathomed. As the algorithm explores the search space, the gap  $\bar{p}^* - \underline{p}^*$  progressively narrows, and when the lower bound  $\underline{p}^*$  matches the upper bound  $\bar{p}^*$ , the algorithm terminates, returning the incumbent solution as a global optimum for the original problem.

**Why  $(\mathcal{R})$  is more tractable than  $(\mathcal{P})$ .** We now discuss three key reasons why the penalty formulation  $(\mathcal{R})$  is more tractable than the mixed-integer formulation  $(\mathcal{P})$  within a spatial branch-and-bound framework.

First, the set of decision variables in  $(\mathcal{R})$ ,  $\{\delta, v\}$ , is a strict subset of the variables in  $(\mathcal{P})$ , which additionally includes the auxiliary binary variables  $z$  and slack variables  $s$ . A spatial branch-and-bound algorithm’s performance is highly sensitive to the number of variables on which it must branch. By operating in a lower-dimensional space, the formulation  $(\mathcal{R})$  inherently defines a smaller and less complex search tree, mitigating the curse of dimensionality and reducing the overall computational effort.

Second, in  $(\mathcal{R})$ , the algorithm only needs to perform spatial branching, which occurs exclusively on the continuous variables  $\delta$  and  $v$  to resolve violations in the polynomial complementarity constraints. In contrast, solving  $(\mathcal{P})$  requires a hybrid branching scheme involving both combinatorial branching on the binary variables and spatial branching on the continuous variables. This dual-purpose branching complicates the node selection logic and increases the number of potential branches at each node of the tree.

Third, the formulation  $(\mathcal{R})$  is exceptionally well-suited for computing high-quality feasible solutions using nonlinear interior-point methods (Byrd, Hribar, and Nocedal 1999; Waltz et al. 2006; Wächter 2002). In theory, when warm-started with a good initial point, such algorithms converge quadratically to a locally optimal solution for nonlinear programs, where all functions are twice continuously differentiable (Wächter and Biegler 2005a,b; Byrd, Gilbert, and Nocedal 2000); empirically, they often perform well even when the warm-start point is not high-quality (Das Gupta, Van Parys, and Ryu 2024). Every function in  $(\mathcal{R})$  is twice continuously differentiable, and a basic feasible solution can be computed trivially. Furthermore, because minimizing the objective of  $(\mathcal{R})$  to zero corresponds to satisfying the un-

derlying complementarity conditions, modern interior-point solvers that include special interfaces for complementarity constraints can handle the problem structure with particular efficiency (Artelys 2025; Byrd, Nocedal, and Waltz 2006). The ability to quickly find a locally optimal solution is a significant advantage, as this solution provides a strong upper bound that allows the branch-and-bound algorithm to prune large portions of the search tree, thereby drastically improving its overall performance. Furthermore, such a locally optimal solution is likely to have a small objective value, thus corresponding to an  $\epsilon$ -Nash equilibrium by Proposition 3.

### 3.3 Customized spatial branch-and-bound

The computational tractability of  $(\mathcal{R})$  allows us to apply the following customized two-stage spatial branch-and-bound algorithm that exploits the problem structure:

- **Stage 1.** We first compute a locally optimal solution to  $(\mathcal{R})$  using a nonlinear interior point method, as discussed in the previous section. From Proposition 3, we know that the objective value of  $(\mathcal{R})$  for a feasible solution tells us how close the found solution is to a Nash equilibrium. Because of this, a locally optimal solution with a small objective value acts as an excellent warm-starting point for the spatial branch and bound algorithm.
- **Stage 2.** We then feed this locally optimal solution as a warm-starting point to the spatial branch-and-bound algorithm. Since many sub-regions of the branch-and-bound tree will likely not contain any Nash equilibrium, their convex relaxations will often yield objective values strictly greater than zero. Hence a small upper bound provided by a locally optimal solution ensures these sub-regions are quickly identified and eliminated, thus aggressively pruning the vast majority of the search space, drastically reducing the number of nodes to explore in turn and leading to a significant reduction in computation time.

Our discussion above results in the following proposition:

**Proposition 4.** *The two-stage spatial branch-and-bound algorithm is a sound and complete algorithm for finding a Nash equilibrium in multiplayer games. If we choose to terminate the algorithm early, it returns an approximate Nash equilibrium along with its corresponding approximation parameter.*

*Proof.* The algorithm’s soundness follows from Corollary 1 in the case of exact solutions, and from Proposition 3 when terminated early. To see the completeness of our spatial branch-and-bound algorithm (i.e. that it computes a solution if one exists), consider the worst-case scenario, where the stage 1 is unable to provide a point to warm-start the spatial branch-and-bound algorithm. In that case, our customized spatial branch-and-bound algorithm reverts to the standard spatial branch-and-bound algorithm, which is complete (Horst and Tuy 2013).  $\square$

## 4 Empirical Evaluation

We empirically analyze the performance and scalability of our proposed branch-and-bound method, focusing on (i)

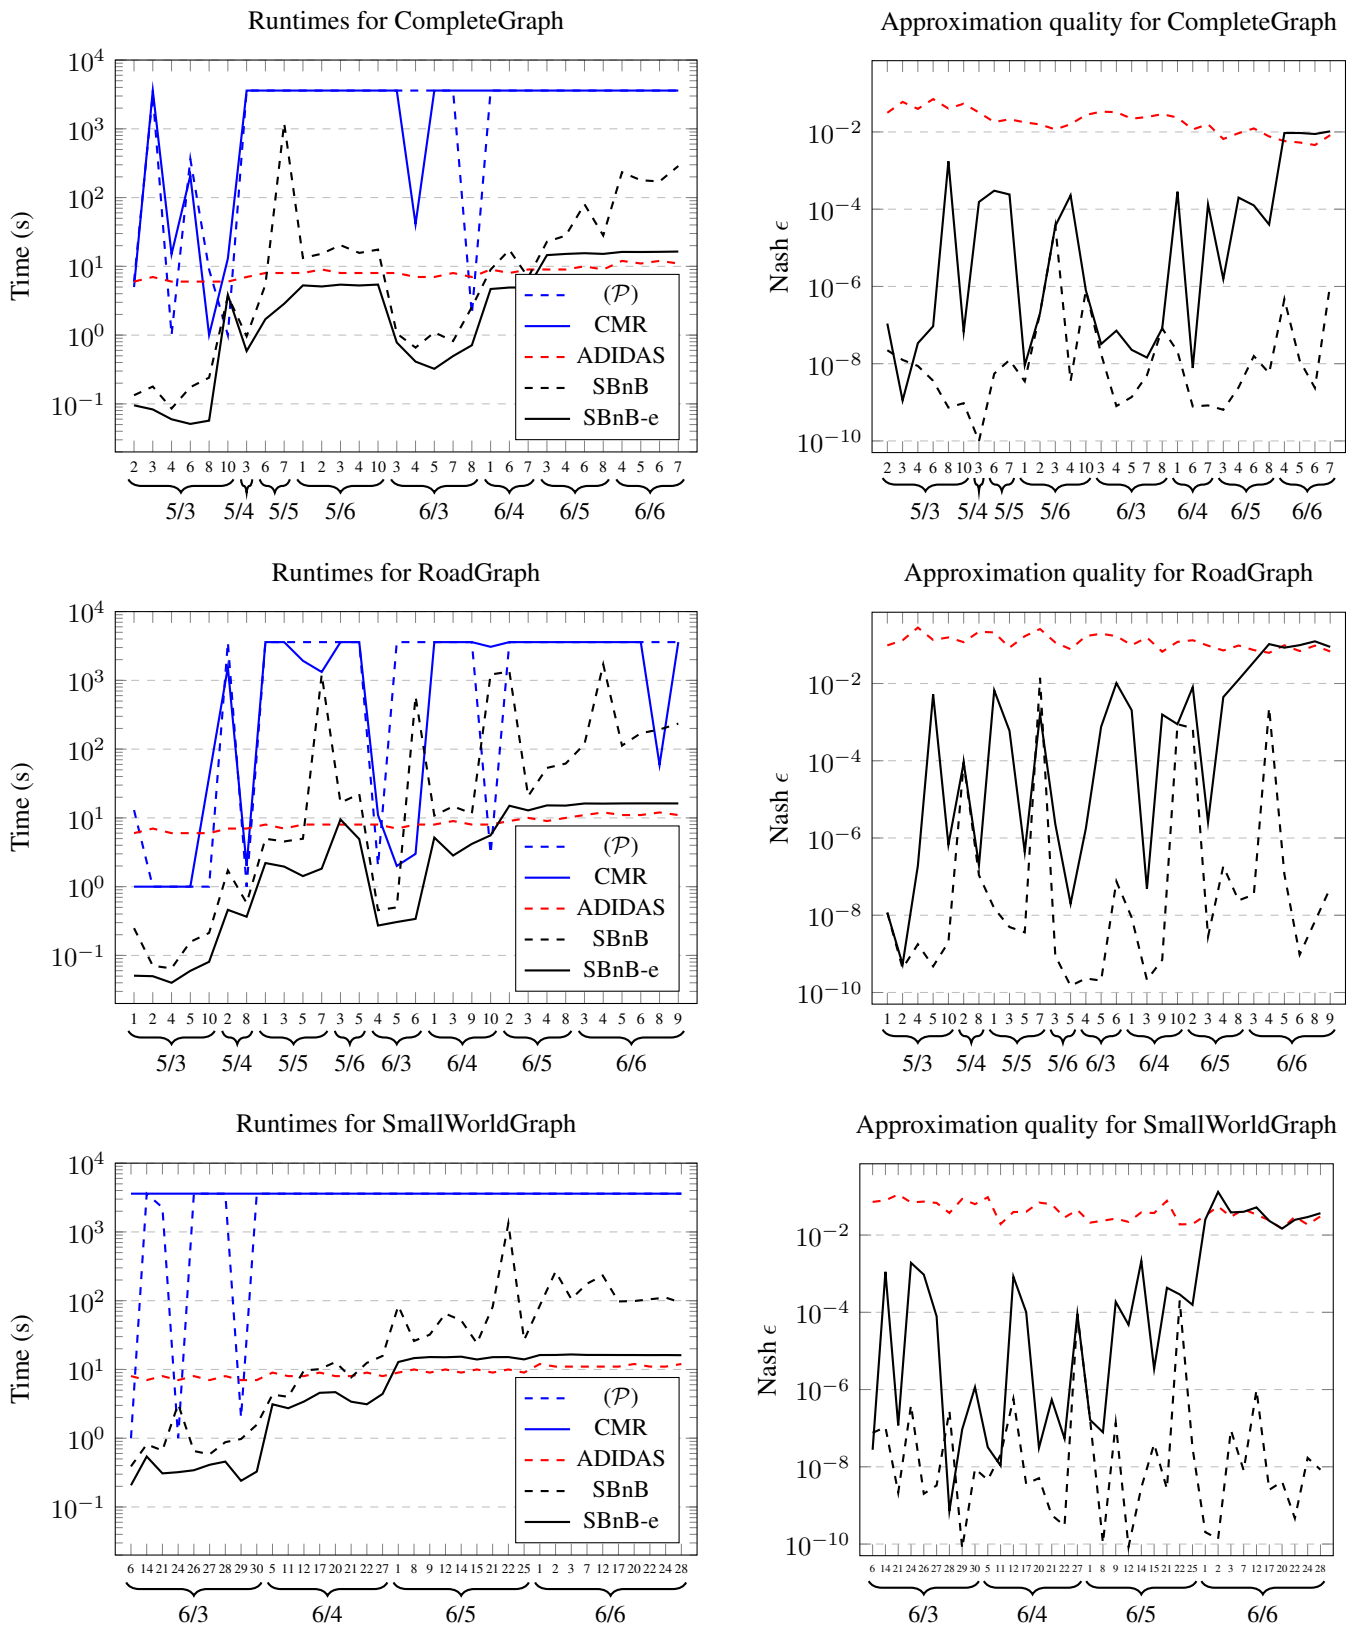


Figure 1: Runtime and approximation quality results for Random Graphical Games on different underlying graphs. Each tick corresponds to a seed, the numbers below indicate the number of players / number of actions per player.

wall-clock computational time, (ii) equilibrium approximation quality. While prior work has evaluated algorithms on fully random games (Ganzfried 2024; Zhang, An, and Subrahmanian 2023; Berg and Sandholm 2017), we do so only in the appendix, as we believe the resulting performance reports are not sufficiently informative. Random games frequently admit sparse equilibria, which may render them more tractable than instances with more structured or adversarial properties. For example, although computing a Nash equilibrium in two-player general-sum games is PPA-hard (Chen and Deng 2006), games with appropriately chosen random payoff distributions can often be solved in expected polynomial time (Bárány, Vempala, and Vetta 2007). More broadly, in multi-player games with i.i.d. payoffs, the existence of pure Nash equilibria is almost certain, with the probability approaching  $1 - 1/e$  as action set sizes increase (Goldberg, Goldman, and Newman 1968; Dresher 1970). Similarly, two-player Gaussian or uniform games typically admit equilibria supported on just two actions with high probability, specifically,  $1 - O(1/\log n)$  (Bárány, Vempala, and Vetta 2007).

Interestingly, we observed that similar characteristics appear even in games generated by the widely used GAMUT library (Nudelman et al. 2004), despite these games incorporating various structural constraints rather than being entirely random. To investigate this, we implemented a simple vectorized method in NumPy to identify pure strategy profiles with minimal  $\epsilon$  and examined the following GAMUT game classes: BidirectionalLEG, CollaborationGame, CovariantGame, RandomGraphicalGame, and RandomLEG. For games with up to 6 players and 6 actions per player (where applicable), we generated 30 random instances for each setting. Our method found pure profiles with least exploitability for all games within 2 seconds. Among the tested classes, only RandomGraphicalGame consistently generated instances with non-trivial equilibria. For this reason, in this paper we chose to focus our experimental evaluation on this class, using three types of underlying graphs: CompleteGraph, RoadGraph, and SmallWorldGraph, with the graph parameter setting `-reflex ok 0`. All utilities were normalized to the  $[0, 1]$  range.

We conducted all the experiments on a computer running Ubuntu 20.04.6 LTS on an Intel(R) Core(TM) i9-12900 CPU with 16 cores and 24 threads with 64 GB of RAM. We used JuMP (Lubin et al. 2023), a domain specific modeling language embedded in the open-source programming language Julia (Bezanson et al. 2017), to model our optimization problems, and our customized two-stage spatial branch-and-bound algorithm uses the solvers KNITRO 14.2 (Byrd, Nocedal, and Waltz 2006) and Gurobi 12.0 (Gurobi Optimization, LLC 2025). We compared our approach against three baselines: (1) solving formulation ( $\mathcal{P}$ ) directly using Gurobi, which in version 12 applies similar transformations as those described by Ganzfried (2024); (2) the CMR algorithm proposed and open-sourced by Zhang, An, and Subrahmanian (2023); and (3) the incomplete ADIDAS algorithm introduced by Gemp et al. (2022) and implemented in the OpenSpiel library (Lanctot et al. 2019) which neither of the complete methods previously

compared to. Since the ( $\mathcal{P}$ ) formulation solved with Gurobi and the CMR algorithm have already been shown to outperform classical methods such as those implemented in the Gambit library (Savani and Turocy 2025) and the exclusion method (Berg and Sandholm 2017), we focused our comparisons on these more competitive baselines.

The computational results are shown in the left-hand plots of Figure 1, where the vertical axis is presented on a logarithmic scale. All computations were terminated after one hour if not completed, and such cases are indicated with this time limit in the plots. We refer to our spatial branch-and-bound algorithm as SBnB, and to its early-terminated variant, which stops once the approximation quality of ADIDAS is approximately matched or exceeded, as SBnB-e. Out of the 96 tested instances, formulation ( $\mathcal{P}$ ) timed out in 77 cases, and CMR in 76 cases, solving 19.8% and 20.8% of the testing instances, respectively. In contrast, both variants of our algorithm, as well as ADIDAS, successfully computed an equilibrium for all instances. With the exception of several instances in the RoadGraph class, where CMR occasionally produced better results, SBnB and its early-terminated variant consistently outperformed the other complete methods. Notably, SBnB-e achieved not only faster runtimes than the incomplete ADIDAS algorithm on smaller instances but also comparable runtimes on larger games with 6 players and 5 or 6 actions.

The right-hand plots in Figure 1 show the equilibrium approximation quality, measured as the  $\epsilon$  of the computed solutions – a minimal value satisfying Equations ( $\epsilon$ -NE).<sup>1</sup> For Gurobi-based methods (formulation ( $\mathcal{P}$ ) and CMR) the resulting  $\epsilon$  values are on the order of Gurobi’s numerical precision, typically around  $10^{-5}$  to  $10^{-6}$ . SBnB consistently produces highly accurate solutions, in some cases achieving several orders of magnitude greater precision than alternative methods. Interestingly, SBnB-e performs comparably to SBnB on smaller instances, offering significant improvements over ADIDAS in both accuracy and runtime, and matches its solution quality on larger games.

We also considered the observation by Sandholm, Gilpin, and Conitzer (2005) that adding a welfare-maximizing or support-minimizing objective can accelerate solver-based methods. While we observed a speedup on random games, it had no consistent effect on the instances reported here.

## 5 Conclusion

We proposed solving Nash equilibria in multiplayer normal-form games via a polynomial complementarity problem, with a sound and complete spatial branch-and-bound algorithm. Our method offers a principled, scalable, and robust alternative to classical approaches, avoiding numerical instability while providing formal solution guarantees. Though challenges remain for very large games, our results indicate that general-purpose equilibrium computation is more tractable than previously thought, and our algorithm shows promise for integration into broader oracle-based solvers.

<sup>1</sup>Note that the  $\epsilon$  reported by default in the OpenSpiel implementation of ADIDAS corresponds to the average player-wise  $\epsilon$ , rather than the maximum deviation from best response.

## Acknowledgments

This research was supported by the Office of Naval Research awards N00014-22-1-2530 and N00014-23-1-2374, and the National Science Foundation awards IIS-2147361 and IIS-2238960. We thank Moïse Blanchard for his assistance with pure Nash equilibrium computation vectorization.

## References

- Artelys. 2025. Knitro special interface for complementarity constraints. [https://www.artelys.com/app/docs/knitro/2\\_userGuide/complementarity.html](https://www.artelys.com/app/docs/knitro/2_userGuide/complementarity.html). Accessed: 2025-07-10.
- Bakhtin, A.; Brown, N.; Dinan, E.; Farina, G.; Flaherty, C.; Fried, D.; Goff, A.; Gray, J.; Hu, H.; et al. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624).
- Banerjee, S.; Riquelme, C.; and Johari, R. 2015. Pricing in ride-share platforms: A queueing-theoretic approach. *Available at SSRN 2568258*.
- Bárány, I.; Vempala, S.; and Vetta, A. 2007. Nash equilibria in random games. *Random Structures & Algorithms*, 31(4): 391–405.
- Berg, K.; and Sandholm, T. 2017. Exclusion method for finding Nash equilibrium in multiplayer games. In *Proceedings of the AAAI conference on artificial intelligence*, 1.
- Bezanson, J.; Edelman, A.; Karpinski, S.; and Shah, V. B. 2017. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1): 65–98.
- Bimpikis, K.; Candogan, O.; and Saban, D. 2019. Spatial pricing in ride-sharing networks. *Operations Research*, 67(3): 744–769.
- Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up limit hold'em poker is solved. *Science*, 347(6218): 145–149.
- Boyd, S. P.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Brown, M.; Sinha, A.; Schlenker, A.; and Tambe, M. 2016. One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.
- Byrd, R. H.; Gilbert, J. C.; and Nocedal, J. 2000. A trust region method based on interior point techniques for nonlinear programming. *Mathematical programming*, 89(1): 149–185.
- Byrd, R. H.; Hribar, M. E.; and Nocedal, J. 1999. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization*, 9(4): 877–900.
- Byrd, R. H.; Nocedal, J.; and Waltz, R. A. 2006. KNITRO: An integrated package for nonlinear optimization. In *Large-scale nonlinear optimization*, 35–59. Springer.
- Černý, J.; Iyengar, G.; and Kroer, C. 2025. Contested Route Planning. In *International Conference on Decision and Game Theory for Security*. Springer.
- Černý, J.; Ling, C. K.; Chakrabarti, D.; Zhang, J.; Farina, G.; Kroer, C.; and Iyengar, G. 2024. Contested Logistics: A Game Theoretic Approach. In *International Conference on Decision and Game Theory for Security*. Springer.
- Chen, X.; and Deng, X. 2006. Settling the Complexity of Two-Player Nash Equilibrium. In *FOCS*, volume 6.
- Cramton, P. 1997. The FCC spectrum auctions: An early assessment. *Journal of Economics & Management Strategy*, 6(3): 431–495.
- Cramton, P. C.; Shoham, Y.; Steinberg, R.; and Smith, V. L. 2006. *Combinatorial auctions*, volume 1. MIT press Cambridge.
- Das Gupta, S.; Van Parys, B. P.; and Ryu, E. K. 2024. Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. *Mathematical Programming*, 204(1): 567–639.
- Datta, R. S. 2010. Finding all Nash equilibria of a finite game using polynomial algebra. *Economic Theory*, 42(1).
- Dresher, M. 1970. Probability of a pure equilibrium point in n-person games. *Journal of Combinatorial Theory*, 8(1).
- Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1): 242–259.
- Fang, F.; Nguyen, T.; Pickles, R.; Lam, W.; Clements, G.; An, B.; Singh, A.; Tambe, M.; and Lemieux, A. 2016. Deploying paws: Field optimization of the protection assistant for wildlife security. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 3966–3973.
- Fang, F.; and Nguyen, T. H. 2016. Green security games: Apply game theory to addressing green security challenges. *ACM SIGecom Exchanges*, 15(1): 78–83.
- Fang, F.; Stone, P.; and Tambe, M. 2015. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Ganzfried, S. 2024. Fast complete algorithm for multiplayer Nash equilibrium. In *International Conference on Decision and Game Theory for Security*, 109–123. Springer.
- Gemp, I.; Savani, R.; Lanctot, M.; Bachrach, Y.; Anthony, T.; Everett, R.; Tacchetti, A.; Eccles, T.; and Kramár, J. 2022. Sample-based Approximation of Nash in Large Many-Player Games via Gradient Descent. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 507–515.
- Goldberg, K.; Goldman, A.; and Newman, M. 1968. The probability of an equilibrium point. *Journal of Research of the National Bureau of Standards*, 72(2): 93–101.
- Govindan, S.; and Wilson, R. 2003. A global Newton method to compute Nash equilibria. *Journal of Economic Theory*, 110(1): 65–86.
- Gurobi Optimization, LLC. 2020. Non-Convex Quadratic Optimization in Gurobi 9.0. [https://www.gurobi.com/wp-content/uploads/2020-01-14\\_Non-Convex-Quadratic-Optimization-in-Gurobi-9.0-Webinar.pdf?x41151](https://www.gurobi.com/wp-content/uploads/2020-01-14_Non-Convex-Quadratic-Optimization-in-Gurobi-9.0-Webinar.pdf?x41151).

- Gurobi Optimization, LLC. 2025. Gurobi Optimizer Reference Manual, Version 12.
- Horst, R.; and Tuy, H. 2013. *Global optimization: Deterministic approaches*. Springer Science & Business Media.
- Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rath, S.; Tambe, M.; and Ordóñez, F. 2010. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4): 267–290.
- Kirschen, D. S.; and Strbac, G. 2018. *Fundamentals of power system economics*. John Wiley & Sons.
- Lancot, M.; Lockhart, E.; Lespiau, J.-B.; Zambaldi, V.; Upadhyay, S.; Pérolat, J.; Srinivasan, S.; Timbers, F.; Tuyls, K.; Omidshafiei, S.; Hennes, D.; Morrill, D.; Muller, P.; Ewalds, T.; Faulkner, R.; Kramár, J.; Vylder, B. D.; Saeta, B.; Bradbury, J.; Ding, D.; Borgeaud, S.; Lai, M.; Schrittwieser, J.; Anthony, T.; Hughes, E.; Danihelka, I.; and Ryan-Davis, J. 2019. OpenSpiel: A Framework for Reinforcement Learning in Games. *CoRR*, abs/1908.09453.
- Lemke, C. E.; and Howson, J. T., Jr. 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2): 413–423.
- Liberti, L. 2008. Introduction to global optimization. *Ecole Polytechnique*.
- Lipton, R. J.; and Markakis, E. 2004. Nash equilibria via polynomial equations. In *Latin American Symposium on Theoretical Informatics*, 413–422. Springer.
- Locatelli, M.; and Schoen, F. 2013. *Global optimization: theory, algorithms, and applications*. SIAM.
- Lubin, M.; Dowson, O.; Dias Garcia, J.; Huchette, J.; Legat, B.; and Vielma, J. P. 2023. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*.
- McCormick, G. P. 1976. Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems. *Mathematical programming*, 10(1): 147–175.
- McKelvey, R. D.; and McLennan, A. 1996. Computation of equilibria in finite games. *Handbook of computational economics*, 1: 87–142.
- Nash, J. F. 1951. Non-cooperative games. *Annals of mathematics*, 54(2): 286–295.
- Nudelman, E.; Wortman, J.; Shoham, Y.; and Leyton-Brown, K. 2004. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *AAMAS*, volume 4, 880–887.
- Padberg, M. 1989. The boolean quadric polytope: some characteristics, facets and relatives. *Mathematical programming*, 45(1): 139–172.
- Pérolat, J.; De Vylder, B.; Hennes, D.; Tarassov, E.; Strub, F.; de Boer, V.; Muller, P.; Connor, J. T.; Burch, N.; Anthony, T.; et al. 2022. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science*, 378(6623): 990–996.
- Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, 125–132.
- Porter, R.; Nudelman, E.; and Shoham, Y. 2008. Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior*, 63(2): 642–662.
- Sandholm, T.; Gilpin, A.; and Conitzer, V. 2005. Mixed-integer programming methods for finding Nash equilibria. In *AAAI*, 495–501.
- Savani, R.; and Turocy, T. L. 2025. Gambit: The package for computation in game theory, Version 16.3.0. <https://www.gambit-project.org>.
- Sherali, H. D.; and Adams, W. P. 1990. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3): 411–430.
- Sherali, H. D.; and Fraticelli, B. M. 2002. Enhancing RLT relaxations via a new class of semidefinite cuts. *Journal of Global Optimization*, 22(1): 233–261.
- Shoham, Y.; and Leyton-Brown, K. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.
- Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg security games: Looking beyond a decade of success. *IJCAI*.
- Tambe, M. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge university press.
- Turocy, T. L. 2005. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2): 243–263.
- van der Laan, G.; Talman, A.; and Van der Heyden, L. 1987. Simplicial variable dimension algorithms for solving the nonlinear complementarity problem on a product of unit simplices using a general labelling. *Mathematics of operations research*, 12(3): 377–397.
- Wächter, A. 2002. *An interior point algorithm for large-scale nonlinear optimization with applications in process engineering*. Ph.D. thesis, Carnegie Mellon University.
- Wächter, A.; and Biegler, L. T. 2005a. Line search filter methods for nonlinear programming: Local convergence. *SIAM Journal on Optimization*, 16(1): 32–48.
- Wächter, A.; and Biegler, L. T. 2005b. Line search filter methods for nonlinear programming: Motivation and global convergence. *SIAM Journal on Optimization*, 16(1): 1–31.
- Waltz, R. A.; Morales, J. L.; Nocedal, J.; and Orban, D. 2006. An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical programming*, 107(3): 391–408.
- Zhang, Y.; An, B.; and Subrahmanian, V. 2023. Computing optimal Nash equilibria in multiplayer games. *Advances in Neural Information Processing Systems*, 36: 21230–21254.