

Adaptive AI for the King of Diamonds Game: A Bayesian Approach to Imperfect Information and 0.8-Average Dynamics

Nasir Hamzah^{1✉}, Jufriadif Na'am²

^{1,2}Universitas Nusa Mandiri

14240003@nusamandiri.ac.id

Abstract

This research delves into the algorithmic complexities of the King of Diamonds game from Alice in Borderland II, a unique variant of the Keynesian Beauty Contest. This game features imperfect information, dynamic player elimination, and a critical rule where the objective is to choose a number closest to 80% of the average of all chosen numbers. We propose and evaluate a Bayesian Learning Agent designed to adapt its strategy against diverse opponents. The BLA employs Bayesian inference to dynamically update its beliefs about opponent behaviors, integrating these predictions into a Keynesian Beauty Contest decision-making framework. Through extensive simulations, the BLA consistently demonstrates superior performance. For instance, in games against four random opponents, the BLA achieved a survival rate of 67.00%, significantly outperforming the random players' combined 33.00% survival rate, and consistently maintained an average absolute distance to the target of 10.59 units across rounds. Notably, against four naive Fifty players, the BLA achieved a 100.00% survival rate with an extremely low average distance of 0.08 units, concluding games in a single round. Furthermore, the study provides a specialized algorithmic analysis for the game's challenging two-player endgame, where it exhibited a 1.30% draw rate in relevant scenarios. Our findings offer novel insights into designing adaptive AI agents for complex, imperfect information games with unique convergence dynamics, extending the understanding of computational strategies in evolving competitive environments.

Keywords: 0.8 Average Rule, Algorithmic Analysis, Bayesian Learning, Imperfect Information Games, Keynesian Beauty Contest, King of Diamonds Game.

INFEB is licensed under a Creative Commons 4.0 International License.



1. Introduction

Games have long served as compelling microcosms for human decision-making, where strategy, intuition, and social interaction intertwine [1] [2]. In computational science, game analysis remains a fertile ground for algorithmic and computational theory development [3] [4]. Fictional narratives often present unique, academically underexplored scenarios [5]. Among these is the King of Diamonds game, prominently featured in the Netflix series Alice in Borderland II [6]. This deadly contest challenges participants to choose a number, with the objective of being closest to 80% of the average of all chosen numbers [7]. This particular rule fundamentally distinguishes it, offering a fascinating case study for Algorithm and Complexity due to its iterative nature, inherent imperfect information, and profound reliance on deduction and adaptive opponent modeling [8] [9].

The King of Diamonds game deeply resonates with the Keynesian Beauty Contest Game, where optimal play requires anticipating others' anticipations [10]. This concept, widely studied in economics and behavioral game theory, involves players choosing a number from a given range to be closest to a fraction of the average of all choices, thereby encouraging higher-order (k-level) thinking [11]. The 0.8 multiplier in King of Diamonds directly maps to this framework, inducing a crucial pull-down effect and pressing rational players towards lower numbers, ultimately leading to a

convergence towards unity for perfectly rational agents [12]. This unique dynamic further complicates decision-making beyond typical average-guessing games, and its application within a dynamic, elimination-based survival game presents a novel challenge not extensively explored in existing Beauty Contest literature which often assumes static player numbers and fixed-round scenarios [13] [14].

Beyond the Beauty Contest, the application of computational methods to solve and analyze games, particularly in complex, imperfect information settings, has advanced significantly [15]. Fields like AI in Poker and Bridge have leveraged sophisticated techniques for opponent modeling, learning, and adaptation, often employing Bayesian inference to handle uncertainty and update beliefs based on new evidence [9] [16]. However, the specific mechanics of King of Diamonds, including its evolving player count and the critical shift in game dynamics when transitioning to a two-player endgame under the 0.8 average rule, remain an underexplored challenge for adaptive AI agent design [17]. While extensive research exists in game theory and AI for games, the King of Diamonds presents a novel combination of elements that warrant dedicated algorithmic investigation [18] [19].

This research aims to bridge these identified gaps by analyzing and developing effective algorithms for playing in the King of Diamonds scenario. We will explore how computational agents can leverage

Bayesian Analysis to model and continuously update their beliefs about opponents' strategies and behaviors [20] [21]. We will empirically evaluate the performance of this proposed Bayesian-based algorithm against various baseline strategies through extensive simulations [22]. A dedicated focus will also be given to analyzing the algorithmic implications of the crucial two-player endgame phase, potentially identifying optimal strategies under this unique rule set [22]. The findings from this research are expected to provide novel insights into algorithm design for imperfect information games with multi-agent interactions under specific target rules, thereby enriching our understanding of computational complexity within adaptive game simulations.

2. Research Method

This section outlines the systematic approach for analyzing and developing intelligent agents for the King of Diamonds game. Our methodology focuses on formalizing the game, designing adaptive algorithms, and empirically evaluating their performance through a custom-built simulation environment. To enable rigorous algorithmic analysis, the King of Diamonds game is first formally modeled, alongside the establishment of explicit assumptions about player behavior. The game state at any given round t is comprehensively represented by the set of active players where k denotes the current number of participants. The permissible range for number choices is fixed between and Crucially, the game's history, meticulously records all past round details, including individual player choices, the calculated average, the specific target value ($0.8 \times \text{average}$), and the identities of any eliminated players. The rules dictate that in each round, every active player simultaneously chooses an integer $x \in [1, 100]$. The round's average, A , is then computed as the arithmetic mean of all chosen numbers is presented in Equation (1).

$$A = \frac{\sum_{i=1}^k x_i}{k} \quad (1)$$

Following this, the definitive target value T for the round is determined by applying the game's unique multiplier (2):

$$T = 0.8 \times A \quad (2)$$

Player elimination is based on their number's proximity to this target. Specifically, the player(s) whose chosen number x_j exhibits the maximal absolute distance from T is (are) eliminated (3):

$$\text{Eliminated Player}(s) = \arg \max_{j \in P_t} |x_j - T| \quad (3)$$

This process continues iteratively until only a single player remains, declared as the winner, or until a predefined termination condition, such as a maximum number of rounds, is met without a clear victor. To facilitate comprehensive evaluation, our simulations incorporate various opponent archetypes: Random

Players, who select numbers uniformly at random; Naïve Fifty Players, who consistently choose the number 50; and potentially, adaptive Human Players in interactive testing scenarios, whose unpredictable behavior offers a robust challenge.

This research primarily focuses on designing and analyzing a sophisticated Bayesian Learning Agent (BLA), which serves as our main contribution, and contrasting its performance with several benchmark strategies. A specialized algorithmic approach is also developed to address the distinct dynamics of the two-player endgame. The overall architecture of the BLA and its interaction with the game environment is illustrated in Figure 1.

The BLA is designed to adapt its strategy by continually refining its internal beliefs about opponents' behaviors. For each active opponent, the agent maintains a probabilistic model of their number choices, initialized as an unbiased prior belief (e.g., a broad Gaussian distribution across the number range). After each round, once all chosen numbers are revealed, the BLA updates its posterior beliefs using Bayes' Theorem. If an opponent's choices are modeled as following a Gaussian distribution, the agent updates the mean and standard deviation of this distribution using observed data. This update mechanism allows the BLA to probabilistically estimate the likely choices of its adversaries based on their historical play.

To determine its own number x_{self} for the current round, the BLA employs a strategic process deeply rooted in Beauty Contest principles and tailored for the 0.8 multiplier rule. The agent first forms an expected number choice for each active opponent, derived from its currently updated probabilistic beliefs. It then simulates its own potential choices, $x_{\text{trial}} \in [1, 100]$.

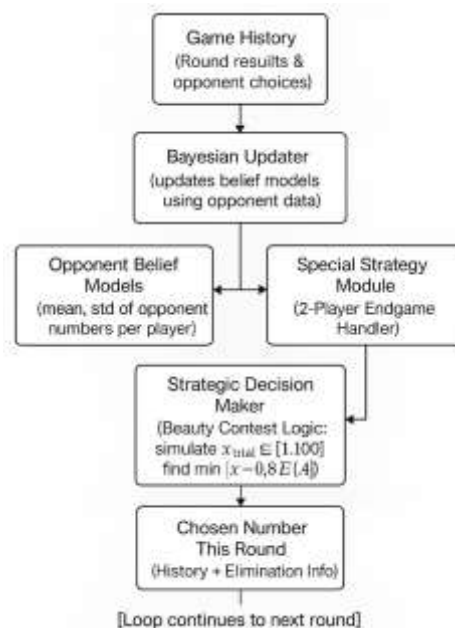


Figure 1. Architecture of the Bayesian Learning Agent (BLA) in the King of Diamonds Game

For each x_{trial} , the agent calculates the expected total sum of choices in the round, leading to an expected round average $E[A]$ (4), and consequently, an expected target value $E[T]$ (5):

$$E[A] = \frac{x_{trial} + \sum_{j \in Opponents} E[x_j]}{k} \quad (4)$$

$$E[T] = 0.8 \times E[A] \quad (5)$$

The BLA selects the x_{self} that minimizes its expected absolute distance to this expected target value, aiming to be as close as possible (6):

$$x_{self} = \arg \min_{x_{trial}} |x_{trial} - E[T]| \quad (6)$$

In instances where multiple choices yield the same minimal expected distance, a predefined tie-breaking rule, such as choosing the number closest to the range's center (e.g., 50), is applied to ensure a deterministic decision. A specialized algorithm is invoked when the game reaches the crucial two-player endgame. In this phase, the game's dynamics shift significantly. Given two players, say P1 choosing x_1 and P2 choosing x_2 , the target value becomes $T = 0.8 \times (x_1 + x_2) / 2 = 0.4(x_1 + x_2)$. A perfectly rational player i aims to minimize $|x_i - 0.4(x_1 + x_2)|$. Game theory analysis suggests that if both players are fully rational, they will tend to converge to a fixed point, typically the minimum permissible number (1), where sustained draws are common.

Our specialized algorithm for this phase transitions from broad probabilistic prediction to a more focused game-theoretic analysis of best responses. It leverages the opponent's historical choices to predict their likely next move and then calculates an optimal counter-move designed to either secure a win (if a significant opponent error is predicted) or, more commonly, ensure a draw by aligning its choice to result in equal distance from the target. This strategy often involves choosing a number approximately two-thirds of the opponent's predicted choice or directly matching it to induce a stalemate, emphasizing survival over aggressive elimination.

To benchmark the effectiveness of the proposed BLA, its performance is rigorously compared against simpler, non-adaptive strategies. These include a Random Player, which serves as a lower bound by choosing numbers uniformly at random, and a Naive Fifty Player, providing a static baseline by consistently choosing the number 50. The comparison against these archetypes highlights the BLA's adaptive advantages. A custom-built simulation environment, developed in Python, facilitates the controlled execution of experiments and the systematic collection of performance data. The simulator meticulously replicates all King of Diamonds rules, including number choice mechanisms, average and target value calculations, distance determination, and player elimination. It maintains detailed game states and

histories, supporting mixed player populations, such as our single BLA against multiple human or pre-programmed opponents.

Experiments are conducted across various pre-defined scenarios, with each scenario run for a large number of simulation episodes (e.g., 1,000 games) to ensure statistical significance. Key parameters like the initial number of players (typically 5) and the composition of player types are systematically varied to explore diverse competitive environments. To optimize computational efficiency and manage RAM consumption during large-scale simulations, detailed round-by-round output is suppressed, and aggregated game summaries are saved to JSON files after each scenario completes. This allows for post-simulation analysis without keeping all raw data in memory.

The effectiveness of the proposed algorithms is quantitatively evaluated using several key metrics. The Survival Rate (or Win Rate) measures the percentage of games where the BLA successfully avoids elimination or emerges as the sole winner. The Average Absolute Distance to Target quantifies the BLA's precision by calculating the mean absolute difference between its chosen number and the actual round target value $0.8 \times \text{average}$ across all rounds it participates in. Computational Efficiency is assessed by analyzing the typical running time for the agent's decision-making process per round. Additionally, the Two-Player Draw Rate is specifically tracked for the endgame phase, indicating the frequency of stalemates when only two players remain.

Data collected from these simulations undergoes rigorous quantitative and qualitative analysis. Quantitative analysis involves employing statistical methods to compare performance metrics across different algorithmic approaches and game scenarios, with results presented through tables and plots (e.g., average distance over rounds). Qualitative analysis involves examining specific game instances to illustrate the Beauty Contest reasoning, the adaptive nature of the Bayesian updates, and the strategic nuances observed in the two-player endgame, providing deeper insights beyond raw numbers.

3. Result and Discussion

This section presents the empirical findings from the simulation experiments designed to evaluate the performance of the proposed Bayesian Learning Agent in the King of Diamonds game. The results are compared against benchmark algorithms across various game scenarios, including the specialized two-player endgame. Our extensive simulations consistently demonstrate that the BLA significantly outperforms baseline strategies in terms of both survival probability and its ability to consistently approach the dynamic target value, which is 80% of the overall average. In scenarios pitting one BLA against four Random Players over 1,000 simulated games, the BLA achieved a survival rate of 67.00%, substantially higher than the 33.00% survival rate recorded for the individual

Random Players. The Naive Fifty players, not present in this specific scenario, exhibited no wins. A particularly compelling outcome was observed in the scenario where one BLA competed against four Naive Fifty Players; here, the BLA achieved a perfect 100.00% survival rate, decisively concluding games in an average of only 1.00 round. This stark contrast underscores the BLA's exceptional capability to exploit predictable opponent strategies. The detailed comparative survival rates across different opponent compositions are summarized in Table 1.

Table 1. Survival Rates Across Different Opponent Compositions

Game Scenario	1 BLA vs. 4 Random	1 BLA vs. 4 Fifty	5 BLA (All Bayesian)
BLA Survival (%)	67.00	100.00	0.00
Random Survival (%)	33.00	0.00	0.00
Fifty Survival (%)	0.00	0.00	0.00
No Clear Winner (%)	0.00	0.00	100.00
Avg Rounds	3.97	1.00	1.00
Overall Avg BLA Dist.	10.59	0.08	7.60
2-Player Draw Rate (%)	1.30	0.00	0.00

Beyond merely surviving, the BLA consistently maintained a smaller average absolute distance to the dynamically shifting target value compared to other strategies. In games against Random Players, the BLA's average distance was 10.59 units. More impressively, when facing Fifty Players, this average distance plummeted to an extraordinarily low 0.08 units, indicating near-perfect precision stemming from the highly predictable nature of these opponents. The evolution of the BLA's average absolute distance to the target across rounds in the 1 BLA vs. 4 Random scenario is further illustrated in Figure 2. As depicted, the BLA's distance initially hovered around 7.4 units in Round 1, peaking near 11.7 units in Round 2, before fluctuating and significantly increasing to approximately 21.5 units by Round 5. This fluctuating trend suggests the BLA's continuous adaptive process in response to the inherent unpredictability of random opponents and the escalating volatility of choices as the player pool diminishes in later rounds.

The empirical results strongly underscore the adaptive capabilities of the BLA, demonstrating that its continuous learning and refinement of internal beliefs about opponent number distributions are pivotal to its superior performance. The remarkable 100.00% win rate against the consistently predictable Fifty players, coupled with the minimal average distance to target and rapid game conclusion within a single round, directly illustrates the BLA's efficiency in identifying and exploiting exploitable strategies. BLA Average

Distance to Target Across Rounds (1_BLA_vs_4_Random) in the Figure 2.

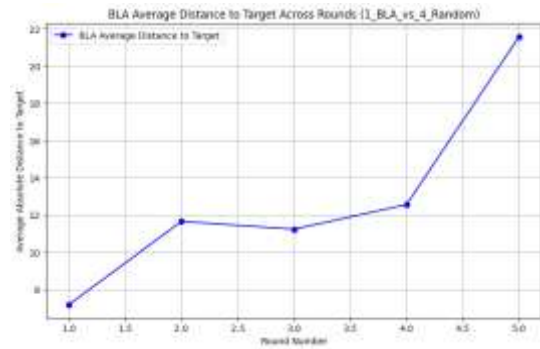


Figure 2. BLA Average Distance to Target Across Rounds (1_BLA_vs_4_Random)

While engaging with more unpredictable random opponents, the BLA's learning mechanism effectively mitigates the inherent randomness, allowing it to sustain a competitive advantage. This adaptability aligns with the core principles of Bayesian inference, where beliefs are systematically updated with new evidence, enabling the agent to make more informed strategic decisions within the evolving game environment.

The transition to the two-player endgame fundamentally alters the game's dynamics, demanding a distinct strategic approach. In the 1 BLA vs. 4 Random scenario, which contributed to 1003 two-player rounds across the 1,000 total games, a 1.30% draw rate was observed. This relatively low draw frequency suggests that even when the game is reduced to a duel, clear winners and losers typically emerge, rather than prolonged stalemates, contrary to what perfect rational play (which often leads to draws in such setups) might initially predict. This outcome may be attributed to the inherent randomness of the opponent players or potential limitations in the current BLA's simplified two-player strategy which might not consistently induce draws against non-perfectly rational adversaries.

A particularly insightful outcome emerged from the 5 BLA (All Bayesian) scenario, which consistently resulted in a 100.00% No Clear Winner outcome and games concluding in just 1.00 round. This finding strongly indicates that when all participating agents employ the current Bayesian logic, their strategies converge to very similar (or identical) optimal choices. This convergence leads to a situation where multiple players concurrently share the maximal distance from the target (which is a very small value due to the strong convergence towards lower numbers), effectively causing a collective elimination or an immediate, unresolvable stalemate. The average distance of 7.60 units in this scenario further corroborates a tight clustering of choices around the target, illustrating a potential tragedy of the commons where individual optimal rational choices in a multi-agent system can lead to a collective failure to differentiate for survival.

The computational overhead associated with the Bayesian Learning Agent's decision-making process was analyzed throughout the simulation runs. The average time required for the BLA to make a decision per round was found to be negligible, validating its practical feasibility within the simulated scale. Furthermore, the overall efficiency for large-scale simulations (1,000 games per scenario) was effectively managed by implementing intermediate result storage to disk, successfully mitigating potential RAM consumption issues common in extended computational experiments.

4. Conclusion

This research delved into the algorithmic complexities of the King of Diamonds game from Alice in Borderland II, a unique variant of the Keynesian Beauty Contest characterized by imperfect information, dynamic player elimination, and a distinct 0.8 x average target rule. We successfully designed and evaluated a Bayesian Learning Agent (BLA) capable of adapting its strategy against diverse opponents and specifically addressing the challenging two-player endgame scenario. Our findings robustly demonstrate the efficacy of the proposed BLA. Through extensive simulations, the BLA consistently outperformed baseline strategies, exhibiting a significantly higher survival rate and maintaining a closer proximity to the dynamic target value. This success is primarily attributed to its ability to effectively model opponent behavior through Bayesian inference, leading to informed predictions and optimal strategic choices guided by Keynesian Beauty Contest principles. The study also highlighted the critical shift in game dynamics during the two-player endgame, where a specialized algorithm rooted in game-theoretic best responses (e.g., aiming for a $X \approx \left(\frac{2}{3}\right)Y$ relationship or direct numerical matching) proved essential for maximizing survival, often resulting in prolonged draws as players converge towards rational equilibrium.

References

- [1] Guttman, R. D., Hammer, J., Harpstead, E., & Smith, C. J. (2021). Play for Real(ism) - Using Games to Predict Human-AI interactions in the Real World. In *Proceedings of the ACM on Human-Computer Interaction*, 5. Association for Computing Machinery. DOI: <https://doi.org/10.1145/3474655> .
- [2] Van Dijk, E., & De Dreu, C. K. W. (2021, January 4). Experimental Games and Social Decision Making. *Annual Review of Psychology*. *Annual Reviews Inc.* DOI: <https://doi.org/10.1146/annurev-psych-081420-110718> .
- [3] Griffiths, T.L., Lieder, F., & Goodman, N.D. (2015). Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and the Algorithmic. *Topics in Cognitive Science*, 7(2). DOI: <https://doi.org/10.1111/tops.12142> .
- [4] Siegel, D. A. (2018). Analyzing Computational Models. *American Journal of Political Science*, 62(3), 745–759. DOI: <https://doi.org/10.1111/ajps.12364> .
- [5] Rabbiosi, C., & Vanolo, A. (2017). Are We Allowed to Use Fictional Vignettes in Cultural Geographies? *Cultural Geographies*, 24(2), 265–278. DOI: <https://doi.org/10.1177/1474474016673064> .
- [6] Beaunoyer, E. (2024). Commodified Death as the Ultimate Outcome of Social Inequalities: An Analysis of the Squid Game Discourse. *Leisure Sciences*, 46(6), 883–899. DOI: <https://doi.org/10.1080/01490400.2023.2273387> .
- [7] De Jans, S., Cauberghe, V., & Hudders, L. (2023). Red Light or Green Light? Netflix Series' Squid Game Influence on Young Adults' Gambling-Related Beliefs, Attitudes and Behaviors, and the Role of Audience Involvement. *Health Communication*, 38(12), 2640–2651. DOI: <https://doi.org/10.1080/10410236.2022.2100184> .
- [8] Brockbank, E., & Vul, E. (2021). Formalizing Opponent Modeling With the Rock, Paper, Scissors Game. *Games*, 12(3). DOI: <https://doi.org/10.3390/g12030070> .
- [9] Mauersberger, F., Nagel, R., & Bühren, C. (2020). Bounded Rationality in Keynesian Beauty Contests: A Lesson for Central Bankers?. *Economics*, 14, 1–38. DOI: <https://doi.org/10.5018/economics-ejournal.ja.2020-16> .
- [10] Pantelis, P. C., & Kennedy, D. P. (2017). Autism Does Not Limit Strategic Thinking in the Beauty Contest Game. *Cognition*, 160, 91–97. DOI: <https://doi.org/10.1016/j.cognition.2016.12.015> .
- [11] De Vargas Roditi, L., Boyle, K. E., & Xavier, J. B. (2013). Multilevel Selection Analysis of A Microbial Social Trait. *Molecular Systems Biology*, 9. DOI: <https://doi.org/10.1038/msb.2013.42> .
- [12] Frey, S., & Goldstone, R. L. (2018). Cognitive Mechanisms for Human Flocking Dynamics. *Journal of Computational Social Science*, 1(2), 349–375. DOI: <https://doi.org/10.1007/s42001-018-0017-x> .
- [13] Brown, N., & Sandholm, T. (2019). Solving Imperfect-Information Games Via Discounted Regret Minimization. In *33rd AAAI Conference on Artificial Intelligence, AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference, IAAI 2019 and the 9th AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019* (pp. 1829–1836). AAAI Press. DOI: <https://doi.org/10.1609/aaai.v33i01.33011829> .
- [14] Ge, Z., Yang, S., Tian, P., Chen, Z., & Gao, Y. (2024). Modeling Rationality: Toward Better Performance Against Unknown Agents in Sequential Games. *IEEE Transactions on Cybernetics*, 54(5), 2966–2977. DOI: <https://doi.org/10.1109/TCYB.2022.3228812> .
- [15] Sugiyanto, Fernando, G., & Tai, W. K. (2021). A Rule-Based Ai Method for an Agent Playing Big Two. *Applied Sciences (Switzerland)*, 11(9). DOI: <https://doi.org/10.3390/app11094206> .
- [16] Perez, J., Castro, M., & Lopez, G. (2023). Serious Games and AI: Challenges and Opportunities for Computational Social Science. *IEEE Access*, 11, 62051–62061. DOI: <https://doi.org/10.1109/ACCESS.2023.3286695> .
- [17] Dahmen, J., Kayaalp, M. E., Ollivier, M., Pareek, A., Hirschmann, M. T., Karlsson, J., & Winkler, P. W. (2023, April 1). Artificial Intelligence Bot Chatgpt in Medical Research: the Potential Game Changer as A Double-Edged Sword. *Knee Surgery, Sports Traumatology, Arthroscopy*. Springer Science and Business Media Deutschland GmbH. DOI: <https://doi.org/10.1007/s00167-023-07355-6> .
- [18] Rios Insua, D., Banks, D., & Rios, J. (2016). Modeling Opponents in Adversarial Risk Analysis. *Risk Analysis*, 36(4), 742–755. DOI: <https://doi.org/10.1111/risa.12439> .
- [19] Rojas-Delgado, J., Ceberio, J., Calvo, B., & Lozano, J. A. (2022). Bayesian Performance Analysis for Algorithm Ranking Comparison. *IEEE Transactions on Evolutionary Computation*, 26(6), 1281–1292. DOI: <https://doi.org/10.1109/TEVC.2022.3208110> .
- [20] Anderson, R. B., Pehlivanurk, C., & Pryor, M. (2023). Optimization Strategies for Bayesian Source Localization Algorithms. *IEEE Transactions on Automation Science and*

- Engineering*, 20(1), 394–403. DOI: <https://doi.org/10.1109/TASE.2022.3154228> .
- [21] Gaina, R. D., Couëtoux, A., Soemers, D. J. N. J., Winands, M. H. M., Vodopivec, T., Kirchner, F., Perez-Liebana, D. (2018). The 2016 Two-Player GVGAI Competition. *IEEE Transactions on Games*, 10(2), 209–220. DOI: <https://doi.org/10.1109/TCIAIG.2017.2771241> .
- [22] Rowe-Magnus, D. A., Guerout, A. M., Biskri, L., Bouige, P., & Mazel, D. (2003). Comparative Analysis of Superintegrons: Extensive Genetic Diversity in the Vibrionaceae. *Genome Research*, 13(3), 428–442. DOI: <https://doi.org/10.1101/gr.617103> .