



Analyzing RL components for Wagner’s framework via Brouwer’s conjecture

Flora Angileri¹ · Giulia Lombardi² · Andrea Fois³ · Renato Faraone³ · Carlo Metta⁴ · Michele Salvi¹ · Luigi Amedeo Bianchi² · Marco Fantozzi³ · Silvia Giulia Galfrè⁵ · Daniele Pavesi³ · Maurizio Parton⁶ · Francesco Morandin³

Received: 3 April 2025 / Revised: 30 July 2025 / Accepted: 5 September 2025

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2025

Abstract

In this work, we continue the study and systematization started in our previous work (Angileri et al. in: Lecture notes in computer science 15243 LNAI, 2025, pp 325–338. https://doi.org/10.1007/978-3-031-78977-9_21) of Wagner’s Reinforcement Learning framework to investigate graph conjectures. After identifying three main directions that impact the framework’s performance (the environment dynamics, the RL algorithm and the neural network used as a function approximator), we conduct an ablation study to evaluate the effectiveness of each component, analyzing several variations of them. The experiments compare three environment dynamics, implemented as Gym spaces (*Linear*, *Local* and *Global*), two algorithms [PPO and the Cross-Entropy method (Wagner in Constructions in combinatorics via neural networks, 2021, <https://arxiv.org/abs/2104.14516>)], different neural network structures (Multi-Layer Perceptron and Graph Neural Networks) and reward systems. This study was intended not only to test the framework’s capabilities, but also to identify a configuration of environment, algorithm, and neural network that can be effective when exploring graph spaces, even with a complex target. For this reason, all the experiments were executed on Brouwer’s Conjecture. We also present the data collected with the various trained models, as these interesting configurations can be used in the inference process on the problem. Our analysis shows that a proper calibration of the individual components of the framework can significantly improve its performance, suggesting effective settings for addressing complex problems and contributing to the study of Brouwer’s Conjecture. All the codes and data are open source and available at https://github.com/CuriosAI/graph_conjectures.

C. Metta is funded by EU Horizon 2020: G.A. 871042 SoBig-Data++, NextGenEU - PNRR-PEAI (M4C2, investment 1.3) FAIR and “SoBigData.it”. G. Lombardi is funded by FSE REACT-EU, PON Ricerca e Innovazione 2014–2020. M. Salvi is supported by the PRIN project GRAFIA, the MUR Excellence Department Project MatMod@TOV, awarded to the Department of Mathematics, University of Rome Tor Vergata, CUP E83C18000100006. G. Lombardi, M. Salvi, L. A. Bianchi, M. Parton, and F. Morandin are partially funded by INdAM groups GNAMPA and GNSAGA. Computational resources provided by the CLAI laboratory (Amato et al., 2024), University of Chieti-Pescara, Italy.

Editors: Riccardo Guidotti, Anna Monreale, Dino Pedreschi.

Extended author information available on the last page of the article

Keywords Reinforcement learning · Graph theory · Brouwer’s conjecture

1 Introduction

The innovative idea presented in Wagner (2021) by Adam Zolt Wagner falls within an interesting line of research concerning the use of Reinforcement Learning (RL) techniques to investigate conjectures in combinatorics, with a primary focus on graph theory problems. The main vision of this work is that RL can be used as a tool to perform an ‘intelligent’ exploration of huge spaces of configurations, in search of counterexamples or interesting instances.

Wagner’s approach can be applied to conjectures expressed in the normalized form

$$f(S) < 0 \text{ (or } f(S) \leq 0) \quad \forall S \in \mathcal{S},$$

where $\mathcal{S}$ is some configurations’ space, and f is a real-valued function on $\mathcal{S}$. The core of this framework consists of formulating a Reinforcement Learning (RL) problem in which the interaction between the agent and the environment is aimed at constructing a particular configuration $S^* \in \mathcal{S}$ such that $f(S^*) > 0$. To achieve this, the reward system defined should be strictly related to the function f under exam, so that the more a configuration is near to disproving the conjecture, the higher the total reward obtained—by the agent—will be. This idea is brilliant not only when counterexamples actually emerge: in fact, as shown in Wagner’s work, the data collected during the search process can still be used by mathematicians to gain useful insights, potentially guiding them toward a rigorous proof or a confutation.

In our previous work (Angileri et al., 2025), we proposed a systematization of the framework in Wagner (2021), starting to analyze which elements are the most relevant for the agent’s learning, and how to choose and set them in order to improve the search. Among all the factors involved, we identified three pivotal directions to explore further: the environment’s dynamics, the RL algorithm implemented, and the neural network used to approximate policies and value functions. We showed how just changing the algorithm already improved the performances, by exhibiting a new, smaller counterexample to Conjecture 2.1 in Wagner (2021) that was not found by Wagner’s original algorithm. To enhance the research in this direction, we provided a set of tools, available at https://github.com/CuriousAI/graph_conjectures, that can be customized and used to experiment on a wide range of graph conjectures. This set consists of four Gym environments (*Linear*, *Local*, *Global*, *Flip*) that allow for testing different ways of moving through graph spaces, termination conditions, types of actions, and reward systems. Moreover, we remarked on the importance of a proper neural network architecture when used as a policy approximator. In fact, since most graph statements involve only isomorphism-invariant graph functions, it is evident that the agent’s behavior should depend only on the topology of the graph and not be influenced by the labeling of the nodes. This crucial observation moved us towards considering Graph Neural Networks (GNNs) and suggesting them as optimal structures. Moreover, we suggested that pretraining the networks on a supervised learning task related to the problem might be useful to identify the neural network architecture that best extracts relevant features for the task. Once identified, this neural network—whether pretrained or trained from scratch—can then be incorporated into the RL algorithm. We anticipate that this approach will

yield a more accurate search process, as the network can more reliably estimate the total return of a given graph. To this end, we provided a dataset of 1198 graphs with 11 nodes, labeled with their Laplacian spectra, that can be used on problems regarding Laplacian eigenvalues, like Brouwer's Conjecture. Also, the dataset can be found at https://github.com/CuriosAI/graph_conjectures.

In this work, we continue the analysis started in Angileri et al. (2025) by testing several reinforcement learning frameworks on Brouwer's Conjecture, a challenging open problem in graph theory.

Novel contributions

First, we contribute to the study of Brouwer's Conjecture by providing the most relevant data obtained during the experiments. These data are graphs with a high value of the functions *complete brouwer* and *central brouwer*, as defined in Sect. 4.1. Observing these configurations can help infer which factors are most relevant, in the topology of a graph, for Brouwer's Conjecture, and so it might guide a confutation or a rigorous proof of the statement. To this end, we varied a great portion of the factors involved in the RL process, and aimed to maintain an appropriate level of exploration during the agent's learning to prevent it from getting stuck in local maxima. This contribution aligns with the perspective of introducing a complementary strategy to the well-established lines of research on Brouwer's Conjecture: leveraging reinforcement learning as a tool to support mathematical discovery. Specifically, we explore the potential of RL algorithms to identify rare instances that may either challenge the conjecture or provide structurally interesting examples to inspire deeper mathematical insights.

We also provide two new versions of *Local* and *Global* environments, which are characterized by the monotonicity of the game dynamics.

Another contribution of this work is a wide analysis of several components of this general RL framework. This comparative study allows both to see the impact of each element on the learning process and to offer a range of possible options, evaluated on a difficult problem such as Brouwer's Conjecture. The main aim of this study is to help identify a framework capable of dealing with complex open problems, as this will constitute a valuable tool for mathematical research. All the comparisons and collected data are presented in Sect. 5.

In order to assist the agent in understanding the relevant factors, we considered the use of a Graph Neural Network (GNN). After pondering various architectures and pre-selecting the more promising, the final choice landed on the one that outperformed the task of Laplacian spectrum prediction, on the dataset provided in Angileri et al. (2025). As emerged from the experiments, this method attests its effectiveness by regularizing the learning process, as we will see in Sect. 5.

In Sect. 2 we survey the related works. In Sect. 3, we provide an overview of the fundamental reinforcement learning concepts that represent the theoretical foundation of the paper. Section 4 presents the experimental setup in detail, including updated versions of the *Local* and *Global* environments and a detailed description of the Graph Neural Network used. In Sect. 5, we compare different environments, reward functions, and neural network architectures-offering a comprehensive overview of the main training metrics per epoch, such as the average and maximal reward. Furthermore, the best graphs obtained with the experiments are analyzed in Sect. 5.3, with visual representations also available in Appendix A.

2 Related work

2.1 Reinforcement learning and graph conjectures

Wagner's original paper (2021) has significantly impacted the field, inspiring numerous follow-up studies that leverage AI techniques for mathematical exploration. This line of work fits into the broader context of machine assistance in mathematics, recently surveyed by Terence Tao in his Presidential Address (2025a) and related article (2025b). In his overview of historical and recent developments, Tao speculates on the future roles of machines and cites Wagner's contribution among examples where computational methods address challenging mathematical problems. Wagner's framework, therefore, serves as an instance within the growing field of machine-assisted proof and discovery.

The most recent paper using Wagner's idea is Davies et al. (2025). Here, Davies et al. construct a novel family of geometric objects relevant to the Hirsch Conjecture on polytope diameters. They introduce a specific class of 5-dimensional prisms, referred to as 'drums', and demonstrate that their 'width'—a measure related to the diameter—grows *Linearly* with the number of vertices. This result provides a new infinite family of counterexamples to the Hirsch Conjecture, exhibiting *Linear* growth in excess width and answering an open question previously posed by Matschke, Santos, and Weibel regarding such behavior in low dimensions.

Another paper leveraging Wagner's framework is Swirszcz et al. (2025), where the authors leverage multi-agent RL to generate novel and complex high-dimensional polytopes exhibiting specific extremal properties. Similar to Wagner's original setup, their system is framed as a game where agents collaboratively refine polytope examples. They apply it to challenging geometric problems like the k -neighbourly problem and the search for longest monotone paths, and generate millions of new counterexamples to the 50-year-old Hirsch Conjecture, surpassing previously known human-generated constructions both in quantity and diversity. This highlights the potential of AI as a tool for mathematical discovery in geometry, consistent with Tao's vision.

In Charton et al. (2024), Charton, Ellenberg, Wagner, and Williamson introduce PatternBoost, a hybrid methodology combining classical search algorithms with transformer neural networks. The approach conceptually parallels Wagner's original framework, but substitutes reinforcement learning with transformer-based models to generate promising candidate constructions (the "play" phase), while the policy improvement step of Wagner's cross-entropy method is replaced by classical exhaustive search (the "learning" phase). The method operates iteratively, alternating between a '*Local*' search phase that generates many candidate constructions and a '*Global*' phase where a transformer learns patterns from the best examples found, subsequently seeding the next *Local* search. Successfully applied to various problems in extremal combinatorics, PatternBoost has generated solutions matching or surpassing the best previously known results, including finding a counterexample to a conjecture that stood for 30 years.

In Stevanović et al. (2021), Stevanović, Damnjanović, and Stevanović investigate a conjecture posed by Akbari, Alazemi, and Đelić concerning the relationship between the *energy* of a graph $E(G)$ and its matching number $\mu(G)$. The conjecture proposed that the inequality $E(G) \leq 2\mu(G)$ holds for connected graphs with maximum vertex degree Δ satisfying $2 \leq \Delta \leq 5$, extending a previously proven result for $\Delta \geq 6$. To identify potential

counterexamples, the authors apply Wagner's approach, discovering that likely counterexamples occur at $\Delta = 3$. Subsequently, they perform an exhaustive search among connected graphs with $\Delta \leq 3$ and at most 19 vertices, explicitly enumerating these counterexamples and thus disproving the conjecture. Here, Wagner's approach guides and informs mathematical intuition, in the spirit of Davies et al. (2021).

Wagner's method is not always better than classical graph search algorithms. In this direction, Parczyk et al. (2024) compared Wagner's method with exhaustive search (when possible) and Tabu Search for finding extremal graphs related to Ramsey multiplicity. Based on their comparison, they reported inferior performance for the RL approach in terms of solution quality and computation time for medium and large-sized problems, highlighting the effectiveness of other heuristics in their study.

Adopting Wagner's approach of framing conjecture refutation as a graph construction game, in Vito and Stefanus (2023), Roucairol and Cazenave (2022, 2024), Roucairol, Cazenave, Vito, and Stefanus explored alternative search algorithms for this task, particularly focusing on Monte Carlo Search (MCS) techniques such as Nested Monte Carlo Search, Nested Rollout Policy Adaptation, and Adaptive Monte Carlo Search. This MCS-driven 'game playing' proved effective and highly efficient, rapidly finding counterexamples for several existing spectral and chemical graph theory conjectures, suggesting that more sophisticated RL methods like AlphaZero, which relies on Monte Carlo search, may improve the original Wagner approach in these domains.

In Ghebleh et al. (2024a, 2024b), Ghebleh, Al-Yakoob, Kanso, and Stevanović first reimplemented Wagner's approach, and then successfully applied this implementation to generate counterexamples for conjectures concerning the Laplacian spectral radius of graphs and to determine new or improved lower bounds for several small Ramsey numbers.

In Mehrabian et al. (2024), Mehrabian, Wagner et al. searched for graphs of a given size that maximize edge count while avoiding 3-cycles and 4-cycles. They compared Wagner's framework with AlphaZero and Tabu Search; using a curriculum learning strategy (that is, using good graphs found at smaller sizes to initialize the search for larger graphs), both approaches successfully improved several state-of-the-art lower bounds on the maximum number of edges for this problem across various sizes.

In Vott and Lehavi (2023), Vott and Lehavi introduce RamseyRL, a framework using reinforcement learning to guide a best-first search for counterexample graphs for classical Ramsey numbers, $R(s, t)$. Their approach uses a deep neural network as a learned heuristic to estimate a graph's potential to become a counterexample, with the work focusing on presenting and evaluating this RL-guided search methodology rather than finding new record-breaking counterexamples.

In Gladkov et al. (2024), Gladkov, Pak, and Zimin provide a refutation of the long-standing Bunkbed Conjecture. Their successful disproof relies on constructing an explicit, structural counterexample graph based on prior work by Hollom. Interestingly, the authors also report on their separate, unsuccessful attempts to find a counterexample via computational search; for these experiments, they started from Wagner's reinforcement learning approach and used their own reimplementation in Julia, highlighting the practical difficulties in finding such counterexamples computationally despite leveraging AI techniques.

2.2 Brouwer's conjecture

Brouwer's Conjecture was originally formulated in Brouwer and Haemers (2012), as a variation of the Grone-Merris theorem [proven in Bai (2011)], and has since been widely studied, leading to proofs of Conjecture 1 for various graph families: trees (Haemers et al., 2010), unicyclic and bicyclic graphs (Du & Zhou, 2012), regular graphs, split graphs, and cographs (Berndsen & Blokhuis, 2012), to name a few. Although a general proof has not yet been given, the literature is rich in results on very specific cases, like Ganie et al. (2020), Chen (2019), Torres and Trevisan (2024). These results already reduce the search range for counterexamples.

A worth-mentioning paper by Rocha (2020), which gives an asymptotic result about Brouwer's Conjecture. In this work, Rocha proves that the probability for a graph G with n vertices to violate the statement decays exponentially in n . It follows that, from a measure-theoretic standpoint, the set of graphs violating the conjecture forms a negligible subset of the entire graph space and that potential counterexamples, if they exist, become rarer as the size increases. As a result, Rocha (2020) supports a shift in research focus: instead of continuing to identify new families of graphs for which the conjecture holds, future efforts should aim to detect and analyze the exceptional structures that might lead to its violation. Moreover, the decay rate-illustrated in Appendix C-shows that even small increases in the number of nodes lead to a significant reduction in the number of potential counterexamples. For this reason, our analysis focuses on small graphs, with particular attention to the case of 11 nodes.

3 Notation

Here we recall reinforcement learning's basic concepts and notations. An RL environment is described by a Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, \mathcal{R}, p)$ consisting of state space, action space, rewards, and transition model $p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S} \times \mathcal{R})$, where Δ denotes the space of probability distributions. A transition $p(s', r|s, a)$ is the probability of reaching the next state s' with a reward $r \in \mathcal{R} \subset \mathbb{R}$ when the action a is executed in the state s . An agent interacts with the environment by sampling actions from a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. This agent-environment interaction gives rise to a sequence $\mathcal{T} = (S_0, A_0, R_1, S_1, A_1, R_2, \dots)$, called *trajectory*. Here $S_t, A_t, R_{t+1}, S_{t+1}$ denote the state at time t , the action executed at time t sampled from $\pi(\cdot|S_t)$, the reward received at time $t + 1$, and the state reached at time $t + 1$, respectively.

An environment dynamics can be classified in two ways: if there are absorbing states reachable from every state under a uniform policy, it is called *episodic* and interactions will eventually terminate; otherwise, it is called *continuing* and the trajectory will never end, unless an additional termination condition is given.

We refer to the situation in which the only non-zero reward occurs at the end of the episode as a *sparse* reward: $R_t = R(S_t, t) = f(S_T)$ if $t = T$, and 0 otherwise.

4 Methods

The first set of experiments observed the effects of changing the environment by testing the behavior of *Linear*—as presented in Angileri et al. (2025)—and a new implementation of *Local* and *Global*, with the cross-entropy method and an MLP network. We worked with new versions of our *Linear*, *Local*, and *Global* Gym environments. Minor modifications were applied to enable their further use in combination with the GIN architecture and to prevent the agent from visiting the same configuration multiple times. In this section, we briefly recall the logic of each environment and explain these variations. We refer the reader to Angileri et al. (2025) for details. The second batch of experiments is on the neural network’s impact. As highlighted in Angileri et al. (2025), it is reasonable that the agent’s behavior depends only on the graph’s topology. MLPs are unsuitable for this purpose, since this network’s outputs depend on nodes’ labeling.

Starting from this observation, we replaced the MLP with a Graph Isomorphism Network (GIN) (Xu et al., 2019), leaving the algorithm unchanged. GIN networks are an obvious alternative to test for graph problems, as they are capable of deriving useful graph representations—more informative than the simple adjacency matrix—and encoding the graph structure more effectively. In particular, we employed GINEConv, a variant of GIN introduced in Hu et al. (2020) that enables the addition of edge features. The GIN architecture used is described in detail in Sect. 4.1.1.

Given the remarkable results obtained in Angileri et al. (2025) with PPO on Conjecture 2.1 in Wagner (2021), the third batch of experiments is intended to analyze its performance on this complex problem as well, and across all three environments. PPO has been used combined with an MLP, while the integration of the GIN network as a feature extractor poses additional complexity. We are currently working on preliminary experiments.

Before delving into the experimental setup (Sect. 4.1), we briefly recall the different logics of the environments, present their variations, and report all the proven facts on Brouwer’s Conjecture that we considered in designing the experiments.

Linear

Linear consider a fixed ordering of all the possible edges $E = \{(i, j)\}_{i,j=1,\dots,n}^1$ and then at each time t the agent can choose between leaving the edge number t as it is (i.e. passing it), or flipping it. The Edge-flipping operation [as defined in Mehrabian et al. (2024)] changes the state of an edge like a boolean *not* operator, as follows: let $e \in \{0, 1\}$ be the single bit representing the edge, then

$$\text{flip}(e) = \begin{cases} 1 & \text{if } e = 0 \\ 0 & \text{otherwise} \end{cases}$$

The state is given by the graph and the current time t , and the action space is $\{0, 1\}$, where 0 means that the current edge is left unchanged, and 1 that the edge is flipped. The game naturally ends when the agent has taken its decision for every edge in E , i.e. every episode goes on for $T = \frac{n(n-1)}{2}$ rounds, if self-loops are not allowed, and for $T = \frac{n(n+1)}{2}$ otherwise.

¹ Self-loop edges $\{(i, i)\}_{i=1,\dots,n}$ can be included/excluded in this list by setting an opportune flag to True/False at game’s initialization.

Local

In *Local*, the agent explores the graph space by moving from one node to another. When moving from node i to node j , the agent has the option to flip the edge (i, j) or do nothing. The state is given by the current graph, the current node i where the agent is located, and the current time. In its original version, an action was given by a target node j to move to and a binary value $\{0, 1\}$, with meaning as *Linear*'s actions.

Note that in *Local* the agent can cross a specific edge more times: if this happens, and flipping edges is allowed, as it was in the original version, the edge could be added and removed multiple times (see Fig. 1). This unnecessarily complicates the agent's exploration, making it very easy to encounter previously seen configurations.

To avoid this behavior, the new version of *Local* maintains the logic of moving between the nodes, but allows only one type of action: if the game starts from the empty graph, the only available action will be adding edges; if the starting point is the complete graph, the action will be removing edges; in case the starting graph is a particular configuration given, also the action type (adding or removing) must be specified by the user, otherwise the game cannot start. With this logic, some rounds might be "passed". Consider the case of an agent attempting to cross an edge that has already been removed: the agent will try to remove it again, but since the action type is fixed, it cannot be added back. In this case, the agent's action will have no effect on the graph. We refer to this characteristic as a *monotonicity* of the agent's trajectories, since for every t we will have $G_t \subseteq G_{t+1}$, if edges can be added, or $G_{t+1} \subseteq G_t$, if edges can be removed.

Unlike the *Linear* environment, *Local* does not have a natural stopping time, allowing for a continuous version to be implemented. However, we always tested an episodic version of it by setting a fixed episode duration at the game's initialization.

Global

In *Global*, the agent explores the graph space by acting on any edge across the entire graph. The agent can choose any edge, and as in *Local*, it can decide to flip it or do nothing. As highlighted in the *Local* case, this behavior is not optimal; thus, we created a new version of the *Global* environment, with the same action logic described in the *Local* case. *Global* environment has also been used with an episodic setting.




Fig. 1 During round $t = 0$ the agent moves from node 1 to node 2, and in round $t = 1$ it goes back to node 1. The edge $(1, 2)$ is flipped in both cases, so that $G_0 = G_2$

Conjecture 1 [Brouwer] Let G be a simple² undirected graph with n vertices and $L(G) = D(G) - A(G)$ its Laplacian matrix, where $D(G)$ is the diagonal degree-matrix: $D_{ii} = d(i) \quad \forall i = 1, \dots, n$, and $A(G)$ is the adjacency matrix. Let $(\lambda_i(G))_{i=1, \dots, n}$ be the eigenvalues of L with descending ordering. Then

$$\sum_{i=1}^t \lambda_i(G) \leq m(G) + \binom{t+1}{2} \quad \forall t = 1, \dots, n \quad (1)$$

The inequality 1 is known to be true for any undirected, simple graph G for any t in $\{1, 2, n-1, n\}$. The cases $t \in \{1, n-1, n\}$ are consequences of basic facts on the Laplacian spectrum,³ while the proof of case $t = 2$ is given in Haemers et al. (2010). Moreover, in Brouwer and Haemers (2012), Brouwer states that the conjecture also holds for graphs with at most 10 vertices.

4.1 Experimental setup

Considering the state-of-art on Brouwer's Conjecture (Sect. 2.2), and in particular the asymptotic result (Rocha, 2020), all the experiments have been executed on graphs with a low number of nodes, letting n vary in $\{11, \dots, 15\}$. We recall that all the game modes used are designed to explore some subspace (depending on initial settings) of $\mathcal{G}_n = \{G \mid |V(G)| = n, G \text{ is undirected}\}$, thus a single experiment was needed for every fixed value of n .

The experiments were executed by fixing an arbitrary stopping time T for both *Local* and *Global* environments, while *Linear* dynamics already consists of a finite number of rounds. Therefore, every experiment gave rise to a finite trajectory $S_0, A_0, R_0, \dots, S_T, A_T, R_T$ where the graph G_t constructed by the agent at time t is included in the state information S_t . All the episodes have been started on the complete graph K_n .

Motivation for the two reward functions

We designed two distinct reward strategies to incorporate structural insights about the problem: the *complete reward*

$$\max_{t=3, \dots, n-2} \sum_{i=1}^t \lambda_i(G) - m(G) - \binom{t+1}{2} \quad (2)$$

and the *central reward*

$$\max_{t=\lfloor \frac{n}{4} \rfloor, \dots, \lfloor \frac{3}{4}n \rfloor} \sum_{i=1}^t \lambda_i(G) - m(G) - \binom{t+1}{2}. \quad (3)$$

²Without self-loops and multiple edges.

³Properties of the Laplacian spectrum can be found in Godsil and Royle (2001), Zhang (2011).

The *complete reward* is the most natural function to design for Conjecture 1, since it evaluates the agent's performance over the entire range of spectral indices $t \in \{3, \dots, n-2\}$. While this choice of reward provides a broader signal, the learning's effectiveness is weakened by the presence of graphs that are easy to obtain—starting from K_n —and for which the *complete reward* is too high. These elements represent potential local maxima that prevent the agent from exploring more optimal solutions. In our experimental setting, one of this “tricky” elements is the graph $G = K_n \setminus \{e\}$, where e is an arbitrary edge: in fact

$$\sum_{i=1}^{n-2} \lambda_i(G) = m(G) + \binom{n-1}{2}.$$

and so the *complete reward* of G is 0. The experiments confirmed that the agent struggles more to explore when using the *complete reward*, due to this configuration.

To mitigate this issue and guide the agent's learning on the more informative part of the spectrum, we introduced the *central reward*, which restricts the reward computation to the region of the spectral range with indexes $\lfloor \frac{n}{4} \rfloor \leq t \leq \lfloor \frac{3n}{4} \rfloor$. Even if this score does not cover all the cases to be considered for Conjecture 1, it might encourage the agent to explore, in search of higher rewards.

Both reward functions were used as *sparse* rewards, as recalled in Sect. 3.

4.1.1 GIN network

To select a GNN for Conjecture 1, we tested GCN and GIN on the supervised task of predicting Laplacian eigenvalues $\lambda_1(G), \dots, \lambda_n(G)$. Training was performed on the $n = 11$ Laplacian dataset from Angileri et al. (2025). Both networks use convolutional layers to extract graph topology features, and were implemented with PyTorch Geometric (Fey & Lenssen, 2019).

The GCN uses three GCNConv layers (Kipf & Welling, 2016) with channels (1, 32), (32, 32), (32, 32), followed by Jumping Knowledge(max) (Xu et al., 2018) and a two-layer ReLU MLP.

The GIN uses three GINEConv layers (Hu et al., 2020), each with a different 4-layer MLP of 32 units. Edge features are used; aggregation via Jumping Knowledge(max).

Both networks have a single input channel with every value set to the constant 1 on the first layer (because we do not have features on the nodes). The graph structure used is the complete graph. The rationale for the choice of the complete graph as the underlying structure is due to the fact that if the current observation is used and the graph is not connected, then the information of the different connected components can not cross the barrier comprised of the missing edges. The actual information about the current state of the graph is passed as edge information to the convolutional layers. This is the necessity which made us choose the GINEConv layer instead of the GINConv layer for the GIN net, and decide to put the information on the structure of the graph into the edge information parameter instead.

The training has been for 5 epochs, using Mean-Square loss, Adam optimizer, and 0.001 learning rate. The dataset has been split into 70% training, 15% validation, and 15% test.

Figure 2 shows the comparison between the average loss obtained on the validation set in every epoch, while the final mean test loss is given in Table 1. The results clearly show




Fig. 2 Boxplots of the average validation loss in 5 training epochs of the GIN and GCN (from left to right) on the Laplacian dataset

Table 1 Final average test loss of the two networks on approximately 1.797 elements

	GCN	GIN
Mean test loss	0.0561782897%	0.0098509728%

how the GINEConv layers are more accurate in capturing information about the graph's structure. This leads to the use of the GIN architecture in our experiments.

Remark 1 Since our work focuses on Brouwer's Conjecture, for which the asymptotic result in Rocha (2020) holds, we do not investigate scalability, as it is not expected to yield particularly insightful observations regarding Brouwer's Conjecture. Moreover, due to the pre-training on graphs of size $n = 11$, we focus on this dimension. A systematic analysis of the method's scalability in the face of the combinatorial explosion of the graph space will be the subject of future work.

4.1.2 PPO setup

This subsection describes the experimental setup for training Proximal Policy Optimization (PPO) in our reinforcement learning framework. We tested three custom environments-*Linear*, *Local*, and *Global*-for graph sizes $n = 11$ – 15 , using two sparse reward functions (*complete score* and *central score*), each provided only at episode end, to analyze the effect of reward design on PPO learning.

Network and hyperparameters The policy network is a two-layer MLP (128 units each). Fixed hyperparameters for all runs:

- Learning rate: 0.0003
- Discount factor (γ): 0.99
- GAE parameter (λ): 0.95
- Entropy coefficient: 0.01
- Batch size: 2048
- Max steps: 200,000

The discount factor $\gamma = 0.99$ prioritizes long-term rewards, which is important since the final graph structure depends on a sequence of actions. GAE with $\lambda = 0.95$ helps stabilize learning by balancing bias and variance, resulting in smoother updates—especially valuable when rewards are delayed and sparse. The remaining hyperparameters use standard PPO defaults for stable training. All values were kept constant to ensure fair and consistent comparisons across experiments.

Early stopping Training stops when the variance of average episode rewards over five rollouts drops below $5 \cdot 10^{-4}$ and the last reward deviates from the rolling mean by less than 10^{-3} .

Remark 2 A key goal for future work is to use PPO with a GIN architecture. However, this integration is not a straightforward network replacement. A robust PPO+GIN agent requires substantial re-engineering of the PPO's data handling pipeline to properly interface with the GIN encoder. This will be the subject of a future paper.

4.1.3 Cross-entropy setup

The code used for implementing the cross-entropy method is presented from Ghebleh et al. in (2024a) and available at <https://github.com/dragance106/cema-for-graphs>. Each experiment has produced 3500 generations of 1000 matches.

Hyperparameters The cross-entropy method counts several degrees of freedom. The most characteristic are

- `percent learn`: the top percents of graphs used for training the agent's neural network
- `percent survive`: the top percents of graphs that are kept in the next generation
- `act rndness`, `act rndness wait`, `act rndness mult`, `act rndness max`: these parameters regulate the exploration. In each round, `act rndness` actions of the 1000 total are chosen randomly. This value grows by an `act rndness mult` factor if after `act rndness wait` generations the maximum reward did not increase. `act rndness max` is the maximum exploration factor allowed. `percent learn`, `percent survive` has been left to default, while the randomness parameters have been increased to encourage exploration. However, `act rndness max` was set to 0.05.

This algorithm has been tested with the GIN pre-trained architecture and an MLP, to compare the different structures. The learning rate was set to 0.003.

5 Results

In this section, we present and analyze the results obtained from our framework. The analysis is divided into three parts: the first examines the reward behavior across different settings during training (Sect. 5.1), the second provides an overview of the performance of the various frameworks (Sect. 5.2), and the third investigates the structural properties of the best graphs identified during testing (Sect. 5.3).

5.1 Training phase

Figure 3 provides a comprehensive comparison of training dynamics for all models applied to graphs of size 11. In these plots, as in the rest of Sect. 5, the cross-entropy method is referred to as the “Wagner algorithm” to highlight its connection with the framework proposed in Wagner (2021). Note that the curves *Learnt Reward* and *Survived Reward* represent the minimum reward observed in the training batch and in the survived batch, respectively. Several key observations can be drawn from this comparison, shedding light on how different model architectures and training schemes influence learning behavior.

High variance in cross-entropy The cross-entropy method exhibits high variance throughout training. This behavior is primarily due to its inherently exploratory nature, which is preserved across all generations. However, the instability is also exacerbated by the limitations of the MLP network architecture, which struggles to capture and extract graph features effectively. In contrast, GIN-based models show more stable learning curves—particularly in the average reward—even while maintaining strong exploratory behavior.

Impact on learning speed This high variance has a direct effect on the learning dynamics, as it introduces significant oscillations in the learnt reward curve, ultimately slowing down convergence.

PPO efficiency and stability Compared to the cross-entropy method, PPO demonstrates both faster convergence and greater stability. Each PPO epoch consists of 2048 timesteps, which corresponds to approximately 37 matches on $\mathcal{G}_{11}$. All PPO-based models achieved satisfactory performance before epoch 100 (i.e., around 3700 matches), whereas the cross-entropy method required up to 500,000 matches to stabilize. The average reward curve confirms this enhanced consistency.

Influence of the environment The structure of the environment has a notable impact on training. In particular, the *Linear* environment is highly sensitive to the order in which edges are visited. As a result, the agent must coordinate its actions precisely to construct a specific graph. In contrast, environments with greater flexibility in edge visitation accelerate the discovery of promising configurations.

Role of the reward function The reward function proves fundamental in shaping the agent’s learning trajectory. The *central reward* leads to a slower and more gradual increase in performance compared to the *complete reward*. Comparing the two may help identify structural factors more relevant to Conjecture 1.

5.2 Testing phase

PPO with MLP setup Notably, the highest-scoring configurations are achieved by PPO combined with the MLP architecture, consistently producing graphs with scores close to




Fig. 3 Comparison of training performance across architectures, reward functions, and environments for graphs with 11 nodes. Each pair of rows represents a different environment (*Linear*, *Local*, *Global*), comparing central (top) and complete (bottom) rewards. Columns show results for PPO with MLP (left), and Wagner with MLP (center) or GIN (right). PPO reports average and max rewards per epoch; Wagner includes additional metrics such as survived and learnt rewards

zero, apart from negligible variations attributable to numerical precision. Among these, the solution found in the Linear-complete reward setting stands out for its sparsity, contrasting with the denser block-structured adjacency matrices observed in the other cases. Analyzing the structure of the adjacency matrices, it can be observed that for both Local and Global, the reward functions drive edge elimination in opposite directions. This leads to larger zero

blocks in the north-west corners for Local-central and Global-complete combinations, and in the south-east corners for Local-complete and Global-central combinations.

Wagner with MLP setup The Wagner-MLP setup tends to produce denser graphs compared to the previous configurations, with the Linear-complete setting reaching a score of zero under the densest condition. Overall, this setup exhibits a less refined exploratory capability in uncovering sparser structures, as reflected by the significantly lower scores observed in the Local and Global environments.

Wagner with GNN setup The Wagner-GNN setup reveals the most diverse range of configurations with scores approaching zero, with both the Local-central and Global-central settings achieving the best performance with identical scores. Notably, the Linear-complete configuration, although not reaching zero, results in a highly sparse graph structure. Overall, this setup outperforms the Wagner-MLP combination, demonstrating that the inclusion of a GNN significantly enhances the exploratory capability of the Wagner algorithm.

Building upon the encouraging outcomes observed with PPO, we extended our experiments to larger graph sizes, training agents on instances with $n \in \{12, 13, 14, 15\}$. This extension serves to explore how structural patterns evolve across different configurations. The results are given in Appendix B.

5.3 Structural properties

All best-performing graphs obtained after the train phase were subsequently tested to determine whether they belong to graph classes already known to satisfy Brouwer's Conjecture. The results are summarized in Fig. 4, which reports the structural properties analyzed-




Fig. 4 Plot showing the presence (1) or absence (0) of structural properties across best graph configurations for 11 nodes. Each row corresponds to a specific model configuration (environment, reward function, algorithm, and network type). The properties include being a Tree, Unicyclic, Bicyclic, Regular, Split, or Cograph. To the right of the structural heatmap, the numerical values of the final score and the number of edges are displayed. Configurations are approximately sorted in descending order, as values of order $1e-14$ or smaller may result from numerical errors

namely, trees, unicyclic and bicyclic graphs, regular graphs, split graphs, and cographs. Interestingly, all graphs achieving higher scores (with the sole exception of the Linear–Complete–Wagner–GNN configuration) fall exclusively within the split and cograph families. In contrast, graphs with lower scores do not belong to any of the known categories for which Brouwer’s Conjecture has been proven. This observation raises the question of whether training such models from an initially empty graph could yield different outcomes, particularly since identifying split or cograph structures may pose a greater challenge. Moreover, our experimental findings—at least in the case of 11-node graphs—provide supporting evidence that Brouwer’s Conjecture may still hold. We believe this represents a meaningful contribution that could support more in-depth, mathematically grounded investigations.

Visualizations of these graphs can be found in Appendix A.

6 Future work and conclusions

A natural next step will be to integrate the GIN architecture (Xu et al., 2019), which has been proven to enhance stability and learning by capturing graph topological properties when used in conjunction with the Cross-Entropy method (Ghebleh et al., 2024a), into the PPO algorithm (Schulman et al., 2017). Considering the greater stability and efficiency of PPO compared to the Cross-Entropy method observed in our experiments (albeit with MLP architectures), we expect that the PPO+GIN combination could further improve the framework’s exploratory capabilities and performance in identifying relevant graph configurations for the conjecture.

In this paper we have tested two GNN flavours, GCN and GIN. One could investigate other graph-based architectures, like for instance GAT (Velickovic et al., 2018) or GraphSAGE (Hamilton et al., 2017), augmented with GloNet (Di Cecco et al., 2024), Switch-Path (Di Cecco et al., 2025a, 2025b), DAC (Metta et al., 2024, 2025).

Furthermore, the best-performing graphs generated during the experiments constitute a valuable dataset. We intend to subject these results to a more in-depth structural analysis. We will employ various graph metrics—such as the number of connected components, expansion properties, node degree distribution, and various centrality measures—to identify potential recurring topological patterns or salient features. This analysis aims to understand which structural properties of graphs are most correlated with high scores in Brouwer’s Conjecture formulation, potentially guiding future theoretical investigations.

Appendix A Best 11-node graphs

Figure 5 presents the best-performing graphs and their corresponding adjacency matrices obtained by evaluating all trained models—PPO with MLP, Wagner with MLP, and Wagner with GNN—under both reward functions (*central* and *complete*).




Fig. 5 Comparison of the best graphs generated for 11 nodes under different configurations: the first plot corresponds to PPO with MLP, the second to Wagner with MLP, and the third to Wagner with GNN. Columns represent the two reward functions: central (left) and complete (right)

Appendix B Extensive PPO results

Due to its promising results, the PPO configuration was examined in greater detail. PPO agents were trained on graphs of various sizes, ranging from 12 to 15 nodes. Figures 6, 7 shows the progression of the average reward per epoch throughout the training process for each graph size.

Furthermore, Figs. 8, 9, 10, 11 display the best graphs obtained for each configuration, along with the corresponding score evolution over time that led to their final structures.




Fig. 6 PPO training performance for graph sizes 12 and 13. Each block shows average and maximum rewards over training epochs. Columns represent the three environments (*Linear*, *Local*, *Global*), and rows compare the two reward functions (central vs. complete)




Fig. 7 PPO training performance for graph sizes 14 and 15

Appendix C Theorems

Definition 1 [Condition 1, Rocha (2020)]. Let $\{G_1, G_2, G_3 \dots\}$ be independent random weighted graphs, where weights $\xi_{ij}^{(n)}$ are bounded r.v.s on the same probability space and independent for each n , with

$$\begin{aligned}\xi_{ij}^{(n)} &= \xi_{ji}^{(n)} \\ \mathbb{E}[\xi_{ij}^{(n)}] &= \mu_n \\ \text{Var}[\xi_{ij}^{(n)}] &= \sigma_n^2 \\ \sup_{i,j,n} \mathbb{E}[|(\xi_{ij}^{(n)} - \mu_n)/\sigma_n|^p] &< \infty \quad \text{for some } p > 6\end{aligned}$$

Theorem 3 [1 and 2, Rocha (2020)] Assume $\{G_1, G_2, G_3 \dots\}$ are independent weighted random graphs as in Definition 1. If $\mu_n \in (0, 1 - \gamma]$ for some $\gamma \in (0, 1)$ and $\frac{\mu_n}{\sigma_n} \left(\frac{n}{\log n}\right)^{1/2} \rightarrow \infty$ as $n \rightarrow \infty$, then for n large enough, we have that




Fig. 8 Best-performing graphs (12 nodes) generated by trained PPO agents with MLP policy, along with score evolution over the episode. Rows correspond to environments (*Linear*, *Local*, *Global*); columns to reward functions (central vs. complete)




Fig. 9 Best graphs obtained by trained PPO agents for each environment and reward function (13 nodes)




Fig. 10 Best graphs obtained by trained PPO agents for each environment and reward function (14 nodes)




Fig. 11 Best graphs obtained by trained PPO agents for each environment and reward function (15 nodes)

$$\mathbb{P} \left[\sum_{i=1}^k \lambda_i(G_n) \leq m(G_n) + \binom{k+1}{2} \right] \geq 1 - \exp(-C\mu_n n) \quad (C1)$$

for some constant $C > 0$ that depends only on γ . Otherwise, if $\frac{\mu_n}{\sigma_n} \left(\frac{n}{\log n} \right)^{1/2} \rightarrow 0$ and $\sigma_n^2 \frac{\log n}{\mu_n n} \rightarrow 0$ as $n \rightarrow \infty$, then for n large enough, we have that

$$\mathbb{P} \left[\sum_{i=1}^k \lambda_i(G_n) \leq m(G_n) + \binom{k+1}{2} \right] \geq 1 - \exp(-C\mu_n n) \quad (C2)$$

for some constant $C > 0$.

Author contributions All authors contributed equally

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

- Amato, G., Amelio, A., Caroprese, L., Chiacchiaretta, P., Fioravanti, F., Ippoliti, L., Meo, M. C., Monaco, G., Morbidoni, C., Moscardelli, L., Parton, M., Scozzari, F. (2024). AI for sustainability: Research at Ud'A node. In *CEUR workshop proceedings* (Vol. 3762, pp. 494–498).
- Angileri, F., Lombardi, G., Fois, A., Faraone, R., Metta, C., Salvi, M., Bianchi, L. A., Fantozzi, M., Galfrè, S. G., Pavesi, D., Parton, M., Morandin, F. (2025). A systematization of the Wagner framework: Graph theory conjectures and reinforcement learning. In *Lecture notes in computer science 15243 LNAI* (pp. 325–338). https://doi.org/10.1007/978-3-031-78977-9_21
- Bai, H. (2011). The grone-merris conjecture. *Transactions of the American Mathematical Society*, 363(8), 4463–4474. <https://doi.org/10.1090/S0002-9947-2011-05393-6>
- Berndsen, J., & Blokhuis, A. (2012). *Three problems in algebraic combinatorics*. Dissertação de Mestrado: Eindhoven University of Technology.
- Brouwer, A. E., & Haemers, W. H. (2012). *Spectra of graphs*. <https://doi.org/10.1007/978-1-4614-1939-6>
- Charton, F., Ellenberg, J. S., Wagner, A. Z., & Williamson, G. (2024). PatternBoost: Constructions in mathematics with a little help from AI. <https://arxiv.org/abs/2411.00566>
- Chen, X. (2019). On Brouwer's conjecture for the sum of k largest Laplacian eigenvalues of graphs. *Linear Algebra and its Applications*, 578, 402–410.
- Davies, A., Gupta, P., Racaniere, S., Swirszcz, G., Wagner, A. Z., Weber, T., & Williamson, G. (2025). Drums of high width. <https://arxiv.org/abs/2503.09919>
- Davies, A., Velickovic, P., Buesing, L., Blackwell, S., Zheng, D., Tomašev, N., Tanburn, R., Battaglia, P., Blundell, C., Juhász, A., Lackenby, M., Williamson, G., Hassabis, D., Kohli, P. (2021). Advancing mathematics by guiding human intuition with AI. *Nature*, 600(7887), 70–74. <https://doi.org/10.1038/s41586-021-04086-x>
- Di Cecco, A., Metta, C., Fantozzi, M., Morandin, F., Parton, M. (2024). GloNets: Globally connected neural networks. In *IDA 2024. Lecture notes in computer science* (Vol. 14641, pp. 53–64). https://doi.org/10.1007/978-3-031-58547-0_5
- Di Cecco, A., Papini, A., Metta, C., Fantozzi, M., Galfrè, S. G., Morandin, F., Parton, M. (2025a). Switch-Path: Enhancing exploration in neural networks learning dynamics. In *Lecture notes in computer science 15243 LNAI* (pp. 275–291). https://doi.org/10.1007/978-3-031-78977-9_18

- Di Cecco, A., Papini, A., Metta, C., Fantozzi, M., Galfrè, S. G., Morandin, F., Parton, M. (2025b). Exploration and generalization in deeplearning with SwitchPath activations. *Machine Learning*, 114(9), 200. <https://doi.org/10.1007/s10994-025-06840-y>
- Du, Z., & Zhou, B. (2012). Upper bounds for the sum of laplacian eigenvalues of graphs. *Linear Algebra and its Applications*, 436(9), 3672–3683. <https://doi.org/10.1016/j.laa.2012.01.007>
- Fey, M., & Lenssen, J. E. (2019). Fast graph representation learning with PyTorch geometric. *arXiv Preprint arXiv:1903.02428*
- Ganie, H. A., Pirzada, S., Rather, B. A., & Trevisan, V. (2020). Further developments on brouwer's conjecture for the sum of laplacian eigenvalues of graphs. *Linear Algebra and its Applications*, 588, 1–18.
- Ghebleh, M., Al-Yakoob, S., Kanso, A., & Stevanović, D. (2024a). Reinforcement learning for graph theory, I. Reimplementation of Wagner's approach. <https://arxiv.org/abs/2403.18429>
- Ghebleh, M., Al-Yakoob, S., Kanso, A., & Stevanović, D. (2024b). Reinforcement learning for graph theory, II. Small Ramsey numbers. <https://arxiv.org/abs/2403.20055>
- Gladkov, N., Pak, I., & Zimin, A. (2024). The bunkbed conjecture is false. <https://arxiv.org/abs/2410.02545>
- Godsil, C., & Royle, G. F. (2001). *Algebraic graph theory. Graduate texts in mathematics* (Vol. 207). <https://doi.org/10.1007/978-1-4613-0163-9>
- Haemers, W. H., Mohammadian, A., & Tayfeh-Rezaie, B. (2010). On the sum of laplacian eigenvalues of graphs. *Linear Algebra and its Applications*, 432(9), 2214–2221. <https://doi.org/10.1016/j.laa.2009.03.038>
- Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *NeurIPS* (pp. 1024–1034).
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., Leskovec, J. (2020). Strategies for pre-training graph neural networks. In *International conference on learning representations*. <https://openreview.net/forum?id=HJIWWJSFDH>
- Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv Preprint arXiv:1609.02907*
- Mehrabian, A., Anand, A., Kim, H., Sonnerat, N., Balog, M., Comanici, G., Berariu, T., Lee, A., Ruoss, A., Bulanova, A., Toyama, D., Blackwell, S., Paredes, B. R., Veličković, P., Orseau, L., Lee, J., Naredla, A. M., Precup, D., Wagner, A. Z. (2024). Finding increasingly large extremal graphs with AlphaZero and Tabu search. In *IJCAI-24* (pp. 6985–6993). <https://doi.org/10.24963/ijcai.2024/772>
- Metta, C., Fantozzi, M., Papini, A., Amato, G., Bergamaschi, M., Galfrè, S. G., Marchetti, A., Vegliò, M., Parton, M., Morandin, F. (2024). Increasing biases can be more efficient than increasing weights. In *IEEE/CVF, winter conference on applications of computer vision WACV* (pp. 2798–2807). <https://doi.org/10.1109/WACV57701.2024.00279>
- Metta, C., Fantozzi, M., Papini, A., Amato, G., Bergamaschi, M., Andrea, F., Galfrè, S. G., Marchetti, A., Vegliò, M., Parton, M., Morandin, F. (2025). Increasing biases can be more efficient than increasing weights. *Advances in Data Analysis and Classification*. <https://doi.org/10.1007/s11634-025-00649-2>
- Parczyk, O., Pokutta, S., Spiegel, C., & Szabó, T. (2024). *New Ramsey multiplicity bounds and search heuristics*. Foundations of Computational Mathematics.
- Rocha, I. (2020). Brouwer's conjecture holds asymptotically almost surely. *Linear Algebra and its Applications*, 597, 198–205. <https://doi.org/10.1016/j.laa.2020.03.019>
- Roucairol, M., & Cazenave, T. (2022). Refutation of spectral graph theory conjectures with Monte Carlo Search. In *Computing and combinatorics* (pp. 162–176).
- Roucairol, M., & Cazenave, T. (2024). Refutation of spectral graph theory conjectures with search algorithms. <https://arxiv.org/abs/2409.18626>
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *CoRR arXiv:1707.06347*
- Stevanović, Đ., Damnjanović, I., & Stevanović, D. (2021). Finding counterexamples for a conjecture of Akbari, Alazemi and Andelić. <https://arxiv.org/abs/2111.15303>
- Swirszcz, G., & Wagner, A. Z., Williamson, G., Blackwell, S., Georgiev, B., Davies, A., Eslami, A., Racaniere, S., Weber, T., Kohli, P. (2025). Advancing geometry with AI: Multi-agent generation of polytopes. <https://arxiv.org/abs/2502.05199>
- Tao, T. (2025a). Machine-assisted proofs. Simons Foundation / American Mathematical Society. AMS Presidential Address. <https://www.simonsfoundation.org/event/machine-assisted-proofs/>
- Tao, T. (2025b). Machine-assisted proof. In *Notices of the American Mathematical Society* (Vol. 72) (pp 6–13) <https://www.ams.org/journals/notices/202501/noti3041/noti3041.html>
- Torres, G. S., & Trevisan, V. (2024). Brouwer's conjecture for the cartesian product of graphs. *Linear Algebra and its Applications*, 685, 66–76.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In *ICLR*

- Vito, V., & Stefanus, L. Y. (2023). Adaptive Monte Carlo search for conjecture refutation in graph theory. <https://arxiv.org/abs/2306.07956>
- Vott, S., & Lehavi, A. M. (2023). RamseyRL: A framework for intelligent Ramsey number counterexample searching. <https://arxiv.org/abs/2308.11943>
- Wagner, A. Z. (2021). Constructions in combinatorics via neural networks. <https://arxiv.org/abs/2104.14516>
- Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How powerful are graph neural networks? In *ICLR*. <https://openreview.net/forum?id=ryGs6iA5Km>
- Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K., Jegelka, S. (2018). Representation learning on graphs with jumping knowledge networks. In *ICML-18* (pp. 5453–5462). <https://proceedings.mlr.press/v80/xu18c.html>
- Zhang, X. D. (2011). The Laplacian eigenvalues of a graph: A survey. <https://arxiv.org/abs/1111.2897>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Authors and Affiliations

Flora Angileri¹ · Giulia Lombardi² · Andrea Fois³ · Renato Faraone³ · Carlo Metta⁴ · Michele Salvi¹ · Luigi Amedeo Bianchi² · Marco Fantozzi³ · Silvia Giulia Galfrè⁵ · Daniele Pavesi³ · Maurizio Parton⁶ · Francesco Morandin³

✉ Maurizio Parton
curiosailab@gmail.com

¹ Tor Vergata University of Rome, Rome, Italy

² University of Trento, Trento, Italy

³ University of Parma, Parma, Italy

⁴ ISTI-CNR, Pisa, Italy

⁵ University of Pisa, Pisa, Italy

⁶ University of Chieti-Pescara, Pescara, Italy