

Structure of the chromatic polynomial

Radmila Sazdanovic and Daniel Scofield

ABSTRACT. Recently, big data techniques such as machine learning and topological data analysis have made their way to theoretical mathematics. Motivated by the recent work with polynomial invariants for knots, we use manifold learning and topological data analysis techniques to explore the structure and properties of the point cloud consisting of the chromatic polynomials of graphs up to 10 crossings. Although chromatic, as well as the Tutte polynomial fail to distinguish graphs, according to a conjecture by Bollobás, Pebody and Riordan they approximate the space of random graphs. In this work we compare structures in the chromatic data revealed using filtered PCA and Ball Mapper techniques, and relate them with a range of numerical invariants for graphs.

1. Introduction

Graph invariants are properties of graphs that are invariant under graph isomorphism. Their definition implies they can be used to distinguish graphs: if values of the invariant are different the graphs are non-isomorphic, but it is possible to have non-isomorphic graphs with the same invariant. Invariants are abundant in graph theory, as well as in other areas of mathematics such as low-dimensional topology. They may be numerical, polynomials, or more involved algebraic constructions such as homology theories. Examples of such invariants for graphs include numbers of vertices, edges, and cycles, connectivity, vertex degrees, and chromatic number, the chromatic, Tutte and magnitude polynomials, the Stanley chromatic function [56, 54, 35, 51], and the corresponding homology theories categorifying these polynomial invariants [24, 21, 25, 49].

Recent years have brought significant advances of artificial intelligence and machine learning methods in graph theory, which have proven challenging partially due to the heterogeneity of graphs and the wide range of their applications. In this paper we take a different approach motivated by the conjecture of Bollobás, Pebody and Riordan stating that almost every pair of independently chosen random graphs are distinguished by their Tutte and chromatic polynomials. Hence, instead of studying graphs directly, we examine the data derived from their chromatic polynomials. This idea of using invariants instead of objects has successfully been implemented in knot theory [28, 26, 36, 16, 43]. More precisely, we use ideas introduced in the work of Dłotko, Gurnari, Hajij, Levitt, and the first author,

2020 *Mathematics Subject Classification.* Primary 05C31, 55N31.
RS was partially supported by NSF grant DMS-1854705.

applying manifold learning and topological data analysis, and compare the results of these two approaches [36, 43].

In this paper, Principal Component Analysis (PCA), a widely used dimension reduction technique, filtered PCA introduced in [36], and TDA’s Ball Mapper algorithm are applied to chromatic polynomial data for graphs with a fixed number of vertices. The dimensionality of chromatic data for graphs of all orders up to 10 appears to be one, with the first principal component influenced by all non-trivial chromatic coefficients. The corresponding Ball Mapper graph (BMGraph) has remarkable and intriguing resemblance with the PCA structure. Though the evidence is only empirical, the first principal component (PC1) corresponds to the elongated, linear structure of the Ball Mapper graph, and the second principal component (PC2) corresponds, informally, to its “width”.

The structure of the chromatic data is further explored with respect to various graph invariants using Ball Mapper as a tool for visualizing functions between high dimensional spaces. The linear structure of the Ball Mapper, analogous to the first Principal component, is determined by the number of edges in a graph, which is consistent with the prominence of edges E in coefficient formulas and the PCA observations. When colored by the corresponding values of PC2, the linear structure of the BMGraph corresponding to PC1 splits into concurrent/parallel strands determined by increasing values of PC2. Similarly, for a fixed number of edges and vertices, the collection of clusters in the Ball Mapper graph corresponds to a slice of points in the 2-dimensional PCA projection. Clusters on both ends of each slice in the PCA projection are graphs whose chromatic coefficients are extremal in the Rodriguez/Satyanarayana poset [46].

Our analysis facilitates the study of numerical graph invariants through dimensionality reduction and clustering. Recent work on graph compression summarized in Section 2.3 suggests a correlation between the chromatic polynomial and various measures of graph irregularity. While these measures are incomparable, our two-dimension projection of chromatic data, combined with Ball Mapper, reveals clusters of graphs which tend to maximize or minimize these invariants among all graphs with the same numbers of vertices and edges, see Section 4.4.

For example, on one end of the spectrum defined by our second principal component PC2, we find clusters of graphs with smallest chromatic coefficients and maximum irregularity scores, which tend to have more triangles, more cliques (complete subgraphs), and higher differences between vertex degrees; e.g., a complete graph with one pendant edge. The other end of the spectrum features graphs with largest chromatic coefficients and minimum irregularity scores; which tend to have a more regular structure, greater numbers of large cycles, and generally avoid cliques; e.g., Turan graphs as typical examples.

This work leaves a number of natural questions unanswered and raises more. Some of the questions to be addressed in the upcoming work are the following:

Question 1 One of the main requirements in this approach is that the results are stable with respect to the filtration of our data, in this case by the number of vertices. It is quite remarkable that both dimensionality and Ball Mapper structure were stable, in addition to being consistent across methods. While filtered PCA [36] provides a reliable way of dimensionality reduction for data where finding a representative sample is a challenge, the next natural step is to apply these methods to a set of random graphs.

Question 2 Big data techniques, such as machine learning and topological data analysis have only recently been applied to data arising from theoretical mathematics where they could provide valuable insights and conjectures. This was achieved in knot theory [16] where the relation between knot invariants was discovered by AI and later proven [15]. In this paper we provide insights into the missing theoretical characterization of chromatically maximal graphs, or the features corresponding to graph irregularity, which is quite elusive as the most prominent irregularity measures seem to be mutually incomparable [2]. The scope and types of results in graph or knot theory that can be obtained using these approaches is an interesting open question.

Question 3 Methods used in this paper can be readily extended to other graph polynomials. The Tutte polynomial is particularly interesting as it specializes both to the chromatic and the Jones polynomial for knots. Determining the dimensionality of Tutte data is an exciting open problem since it should be at least 3, since that is the dimensionality of the Jones polynomial data [36]. Furthermore, Relational Ball Mapper or Mapper-on-Ball Mapper [43] constructions can be used to explore relations between the Tutte and chromatic, or Tutte and Jones polynomial.

To summarize, this approach provides a new way of vectorizing graphs, and opens up another area of theoretical mathematics to big data techniques. In addition to immediate extensions of this work to other polynomials such as Tutte there are exciting but less clear goals to determining statistical nature of polynomial and other graph invariants and relations between them, as well as potential applications to other sciences.

Acknowledgements

The authors extend their heartfelt gratitude to Pawel Dłotko and Davide Gurnari for many helpful conversations and support with Ball Mapper.

2. Background

In this section we provide necessary background in graph theory, focusing on the chromatic polynomial and irregularity measures.

Let $G = (V, E)$ be a graph, with vertex set V and edge multiset E . If G has an edge between vertices $v_i, v_j \in V$, we write the corresponding element in E as $\{v_i, v_j\}$.

DEFINITION 2.1. The order of a graph G is the number of vertices $n = |V|$. A loop in graph G is an edge of the form $\{x, x\}$ for some $x \in V$. A multiple edge is an element that occurs multiple times in the multiset E . We say that G is a simple graph if it has no loops or multiple edges.

DEFINITION 2.2. Let v be a vertex of a simple graph G . The degree of vertex v , denoted $\deg(v)$, is the number of edges incident to v .

DEFINITION 2.3. A path in graph G is a finite sequence of vertices v_1, v_2, \dots, v_n such that $\{v_i, v_{i+1}\}$ is an edge of G for $i = 1, 2, \dots, n - 1$ and all v_i are distinct.

DEFINITION 2.4. A graph $G = (V, E)$ is connected if any two vertices in G are joined by a path along edges of G . A bridge in a connected graph G is an edge $e \in E$ such that the graph $(V, E - \{e\})$ is disconnected. A pendant edge is an edge incident to a vertex of degree one.

DEFINITION 2.5. The girth of graph G is the length of the shortest cycle in G .

DEFINITION 2.6. Let $H = (V_H, E_H)$ be a subgraph of graph $G = (V, E)$. We say H is an induced subgraph if for every $\{v_i, v_j\} \in E$ with $v_i, v_j \in V_H$, the edge $\{v_i, v_j\}$ is in E_H .

DEFINITION 2.7. The complete graph K_n is the graph which consists of n vertices and $\binom{n}{2}$ edges connecting each distinct pair of vertices. A clique of size m is a subset of vertices of G whose induced subgraph is isomorphic to K_m . The clique number $\omega(G)$ is the maximum size over all cliques in G .

2.1. The chromatic polynomial.

DEFINITION 2.8 ([56, 45, 20]). Let $G = (V, E)$ be a graph. A mapping $f : V \rightarrow \{1, 2, \dots, \lambda\}$ is called a λ -coloring of G if for any pair of vertices $x, y \in V$ such that $\{x, y\} \in E$, $f(x) \neq f(y)$. Two λ -colorings f, g are distinct if $f(x) \neq g(x)$ for some vertex $x \in V$. The chromatic polynomial of the graph G , denoted $P_G(\lambda)$, is equal to the number of distinct λ -colorings of G .

The degree of the chromatic polynomial is equal to the order of G , and thus we write the polynomial with coefficients c_{n-i} as follows:

$$P_G(\lambda) = c_n \lambda^n + c_{n-1} \lambda^{n-1} + c_{n-2} \lambda^{n-2} + \dots + c_1 \lambda + c_0$$

Since no graph has a 0-coloring, the roots of $P_G(\lambda)$ include zero, which forces the free term to be zero $c_0 = 0$ for all graphs.

The chromatic polynomial of a disjoint union of graphs is the product of chromatic polynomials of its components. Additionally, adding multiple edges does not change the chromatic polynomial, while the existence of a loop trivializes it. Therefore, in this paper we consider only simple connected graphs.

THEOREM 2.9. [38] *If G is a graph with m edges, girth $g > 2$, and n_g cycles of length g , then the first g coefficients of the chromatic polynomial are:*

$$c_{n-i} = \begin{cases} (-1)^i \binom{m}{i} & 0 \leq i < g-1 \\ (-1)^{g-1} \left(\binom{m}{g-1} - n_g \right) & i = g-1 \end{cases}$$

For a detailed summary of graph invariants which are determined by the chromatic polynomial see [40].

PROPOSITION 2.10 ([40]). *For a connected graph G , the chromatic polynomial $P_G(\lambda)$ determines the numbers of vertices, edges, triangles, and blocks in G , as well as the girth g and the number of cycles of length g .*

It was long conjectured [45] that the coefficients of $P_G(\lambda)$ are unimodal in the following sense: if $h_i = |c_{n-i}|$, there exists some j with $2 \leq j \leq n-1$ such that

$$h_n \leq h_{n-1} \leq \dots \leq h_{j-1} \leq h_j \geq h_{j+1} \geq \dots \geq h_0.$$

This result has since been proven [27] by showing that the stronger property of log-concavity holds for this sequence: $h_{i-1}h_{i+1} \leq h_i^2$ for all $0 < i < n$.

DEFINITION 2.11. A graph is chromatically unique if no non-isomorphic graph has the same chromatic polynomial.

Cycle graphs, θ -graphs [10] and graphs which consist of wheel graphs with all but three or four adjacent spokes removed [11] are families of chromatically unique graphs.

There are also infinitely many graphs which are not chromatically unique; i.e., non-isomorphic graphs which share a chromatic polynomial. Examples include trees with n vertices, whose chromatic polynomials all have the form $\lambda(\lambda - 1)^{n-1}$ regardless of structure [45, Theorem 13].

For any connected graph G , the first six coefficients of the chromatic polynomial have known formulas computed by counting subgraphs of G .

THEOREM 2.12. [22, 6] *The first four coefficients of the chromatic polynomial $P_G(\lambda)$ are given by the following formulas: $c_n = 1$, $c_{n-1} = -m$, $c_{n-2} = \binom{m}{2} - t_1$, and $c_{n-3} = -\binom{m}{3} + (m-2)t_1 + t_2 - 2t_3$, where m is the number of edges in G , t_1 is the number of triangles, t_2 is the number of induced 4-cycles (pure squares with no diagonal present), and t_3 is the number of complete subgraphs on 4 vertices.*

The 5th and 6th coefficients are given by the following formulas, where the t_i s count induced subgraphs of G on 5 and 6 vertices specified in [6]. For example, t_4 is the number of induced 5-cycles (pure pentagons) and t_8 is the number of induced graphs isomorphic to K_5 .

$$\begin{aligned} c_{n-4} &= \binom{m}{4} - \binom{m-2}{2}t_1 + \binom{t_1}{2} - (m-3)t_2 - (2m-9)t_3 - t_4 + t_5 + 2t_6 + \\ &\quad 3t_7 - 6t_8 \\ c_{n-5} &= -\binom{m}{5} + \binom{m-2}{3}t_1 - (m-4)\binom{t_1}{2} + \binom{m-3}{2}t_2 - (t_2 - 2t_3)t_1 + t_4 \\ &\quad - (m^2 - 10m + 30)t_3 - (m-3)t_5 - 2(m-5)t_6 - 3(m-6)t_7 + 6(m-8)t_8 \\ &\quad + t_9 - t_{10} - 2t_{11} - 2t_{12} - t_{13} + t_{14} - t_{15} - 3t_{16} - 4t_{17} - 4t_{18} + 2t_{19} - 4t_{20} \\ &\quad - t_{21} + 4t_{22} + 3t_{23} + 4t_{24} + 5t_{25} + 4t_{26} + 6t_{27} + 8t_{28} + 16t_{29} + 12t_{30} - 24t_{31} \end{aligned}$$

Note that Whitney's original paper introducing the chromatic polynomial gives an interpretation of each c_{n-i} in terms of spanning subgraphs of G [56]. Gian-Carlo Rota showed that each coefficient can be expressed as a Möbius function on a poset whose elements are subsets of E , and that the sum of the coefficients of $P_G(\lambda)$ is zero [47]. Moreover, Meredith showed that the size of every coefficient is bounded by $|c_{n-i}| \leq \binom{m}{i}$ [38]. These upper bounds are achieved for the first g coefficients of the chromatic polynomial when all cycles of G have length at least g .

In [31] the author defines a Tutte polynomial poset and characterizes the minimal graphs in this poset. With a similar motivation, we define a chromatic polynomial poset as follows.

DEFINITION 2.13. Let $G(n, m)$ denote the set of simple connected graphs with n vertices and m edges. The chromatic polynomial poset is defined on $G(n, m)$ by $H \leq G$ if and only if $P_G(\lambda) - P_H(\lambda) = r(x)$ for some polynomial $r(x)$ with non-negative coefficients. In other words, if we denote chromatic polynomial coefficients by c_i , then $H \leq G$ if and only if $|c_i(H)| \leq |c_i(G)|$ for all $1 \leq i \leq n$.

Rodriguez and Satyanarayana [46] give a complete characterization of the graphs which are minimal in the chromatic polynomial poset.

DEFINITION 2.14 ([20]). Let $J(n, m)$ be the family of graphs $G \in G(n, m)$ such that all blocks of G are complete graphs, including possibly triangles K_3 and double edges K_2 . Let $L(n, m)$ be the family of graphs $G \in G(n, m)$ such that one block B of G has clique number $\omega(B) \geq |V(B)| - 1$ and all other blocks are K_2 s.

Note that B is a complete graph or contains a complete subgraph containing all but one of the vertices of B , while all other blocks are bridges or pendant edges.

PROPOSITION 2.15 ([46]). Any pair of graphs $G, H \in L(n, m) \cup J(n, m)$, have the same chromatic polynomials $P_G(\lambda) = P_H(\lambda)$.

THEOREM 2.16 ([46]). Any graph $H \in L(n, m) \cup J(n, m)$ is a minimal element of the poset $G(n, m)$.

In Section 4.4, we describe extremal graphs in our data set using the following definition based on the chromatic polynomial poset.

DEFINITION 2.17. A chromatically minimal (resp., chromatically maximal) graph with n vertices and m edges is one which is a minimal (resp. maximal) element of the poset $G(n, m)$.

2.2. Irregularity measures for graphs. A graph G is k -regular if $\deg(v) = k$ for each vertex $v \in V(G)$. Numerous invariants have been proposed to describe how far a given graph is from being regular, i.e., how irregular it is. A generally accepted criterion for an irregularity measure is that it should return a value of zero if and only if a graph is regular. The difference between the maximum and minimum vertex degrees in a graph is one example of an irregularity measure. Two other commonly cited measures are spectral irregularity and variance irregularity of a graph.

DEFINITION 2.18 (Spectral irregularity [55]). Let $\lambda_1(G)$ be the largest eigenvalue of the adjacency matrix of graph G and let $\bar{d}(G)$ be the average degree of the vertices of G . The spectral irregularity measure of G is $\epsilon(G) = \lambda_1(G) - \bar{d}(G)$, which is zero if G is regular and positive otherwise.

DEFINITION 2.19 (Variance irregularity [5]). Suppose G has order n and the vertices are labeled, with $d_i = \deg(v_i)$. The variance irregularity measure of G is

$$\sigma(G) = \frac{1}{n} \sum_{i=1}^n d_i^2 - \frac{1}{n^2} \left(\sum_{i=1}^n d_i \right)^2.$$

For other proposed irregularity measures see [4, 39, 1] or a survey [41, 17] that also talks about their comparative strengths.

The following classes of graphs are the most irregular under spectral and variance irregularity.

DEFINITION 2.20 ([3, 41]). Let G be a connected graph with n vertices and m edges. Let d, t be the unique integer values with $2 \leq d$ and $0 \leq t < d$ such that $m = \binom{d}{2} + t$. Then G is the quasi-complete graph $QC(n, m)$ if it is the unique graph such that:

- (1) G has a clique C of size $n - 1$;
- (2) The single vertex v' not included in C is adjacent to exactly t of the other vertices.

DEFINITION 2.21 ([3, 41]). Let G be a connected graph with n vertices and m edges. Let d, t be the unique integer values with $2 \leq d$ and $0 \leq t < d$ such that $m = \binom{n}{2} - \binom{d}{2} - t$. Then G is the quasi-star graph $QS(n, m)$ if it is the unique graph that:

- (1) contains a set S of $n - d - 1$ vertices, each of which is adjacent to all other vertices of G ;
- (2) contains one vertex adjacent to exactly $n - t$ vertices of S ;
- (3) the only edges of G are the ones already specified.

PROPOSITION 2.22 (Proposition 1, [5]). *Among all graphs in $G(n, m)$, the maximum spectral irregularity is achieved by quasi-complete $QC(n, m)$, and the maximum variance irregularity is achieved by $QC(n, m)$ for $m > \frac{1}{2}\binom{n}{2} + \frac{n}{2}$ and by quasi-star graphs $QS(n, m)$ such that $m < \frac{1}{2}\binom{n}{2} - \frac{n}{2}$.*

On the other hand, we have graphs which are as close to regular as possible, such as Turan graphs.

DEFINITION 2.23. A Turan graph $T(n, r)$ is the graph formed by partitioning a set of n vertices into r groups, with sizes as equal as possible, and connecting two vertices by an edge if and only if the vertices belong to separate groups. If we let $t, p \in \mathbb{Z}$ such that $n = tr + p$, each vertex of $T(n, r)$ has degree either $n - t$ or $n - t - 1$.

Turan graphs are a family of extremal graphs whose degree sequences are constant or nearly constant and they are characterized by the following theorem.

THEOREM 2.24 (Turan's Theorem, [53]). *Turan graph on n vertices with r groups $T(n, r)$ has the maximal number of edges among all order n graphs that contain no $(r + 1)$ -clique.*

2.3. Threshold graphs and compression. Compression of a graph is an operation first introduced in [32] and known by a variety of names, including the Kelmans transformation [13] and swing surgery [48]. See the introduction of [31] for details.

DEFINITION 2.25. The neighborhood of vertex v in graph G , denoted $N_G(v)$, is the set of all vertices which are connected to v by an edge. We say that vertex v dominates vertex u if $N_G(u) \subset N_G[v]$, where $N_G[v] = N_G(v) \cup \{v\}$.

DEFINITION 2.26. A graph G is threshold if for all vertices $u, v \in V(G)$, either u dominates v or v dominates u .

For other characterizations of threshold graphs, see [13, 33].

DEFINITION 2.27. Let $G = (V, E)$ be a graph and let $u, v \in V$. The compression of G from vertex u to vertex v is an operation defined as follows: for each vertex $x \in N_G(u) - N_G(v) - \{v\}$ we delete all edges of the form $\{x, u\}$ from G and replace them with corresponding edges of the form $\{x, v\}$. The new graph produced by the compression of G from u to vertex v is denoted $G_{u \rightarrow v}$.

Compression preserves the number of vertices and edges of a graph, and the result is independent of the order of in which vertices u and v are used, up to isomorphism (i.e., $G_{u \rightarrow v}$ and $G_{v \rightarrow u}$ are isomorphic graphs). Compression of a graph monotonically increases a number of graph invariants, such as the spectral radius

and the number of independent sets of order k , while decreasing other invariants such as the number of spanning trees and the vertex connectivity [13, 30, 31].

Any connected graph can be transformed into a connected threshold graph by repeated applications of graph compression [7, 48]. Thus, threshold graphs are often extremal with respect to these types of invariants.

3. Methods and data

3.1. Principal component analysis. Principal component analysis (PCA) is a common technique for exploring multivariate data. The algorithm finds a set of directions within the feature space which explain most of the variance in the data. The number of these directions is often far smaller than the number of original features, and so PCA can be used as a method of dimensionality reduction when visualizing and interpreting a high-dimensional data set.

Let $X \in \mathbb{R}^{n \times p}$ be a matrix representing a data set of n individuals, each with p features. PCA takes the sample covariance matrix S for the data set and finds an orthonormal eigensystem $\{(\lambda_j, a_j)\}_{j=1}^p$ for S with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$. The eigenvector a_1 is the unit vector in \mathbb{R}^p such that the entries of Xa_1 have maximal variance (up to an overall sign change for a_1). The second eigenvector a_2 maximizes variance subject to the constraint that a_2 is orthogonal to a_1 , and all successive vectors a_j maximize variance among directions orthogonal to $\{a_1, \dots, a_{j-1}\}$.

DEFINITION 3.1. [29] Given a data set $X \subset \mathbb{R}^{n \times p}$ and PCA eigensystem $\{(\lambda_j, a_j)\}_{j=1}^p$, the j th principal component of X is the product Xa_j . For the j th principal component of X , the p elements of the vector a_j are called the PC loadings for the j th principal component, and the n elements of Xa_j are called the PC scores.

Note that if a_j is an eigenvector of S with eigenvalue λ_j , then so is its opposite $-a_j$. Thus the signs of PC loadings and PC scores are arbitrary and may be reversed depending on the implementation of PCA in a given instance.

DEFINITION 3.2. [29] The explained variance associated with the j th principal component of the data set X is the eigenvalue λ_j , and the normalized explained variance $\bar{\lambda}_j$ is defined by $\bar{\lambda}_j = \frac{1}{\sum_{i=1}^p \lambda_i} \lambda_j$.

The proportion of variance explained by the first j principal components of X will be denoted $S_j = \sum_{i=1}^j \bar{\lambda}_i$.

The eigenvector a_1 is a unit vector that determines a line in the feature space \mathbb{R}^p . This line indicates the axis in \mathbb{R}^p along which the data set represented by X has the greatest variation. Similarly, a_2 gives the direction orthogonal to a_1 for which the data set has the second-greatest variation, and so on for $2 < j \leq p$.

Suppose that $x \in \mathbb{R}^p$ is a column of X representing one individual in the data set. The PC score for the j th principal component represents the length of the projection of x onto the line corresponding to eigenvector a_j . The implementation of PCA in the Python `scikit-learn` package [44] centers the data by moving the mean to the origin. Thus a PC score of zero in the j th component would indicate that x lies at the center of the distribution with respect to the j th direction of variation. A positive or negative score indicates that x lies further from the center on either side along this direction.




FIGURE 1. Ball Mapper construction: Illustration. The input point cloud is shown in (A), ϵ -net (B), Ball Mapper cover (C), corresponding Ball Mapper graph (D).

3.2. Ball Mapper. Ball Mapper [19, 42] is a recent Mapper-type algorithm for visualizing and analyzing the local and global structure of a data set. Like the widely-used Mapper algorithm [50], Ball Mapper takes a data set in \mathbb{R}^n as an input and returns a graph which describes the proximity of points within the data. While Mapper is highly dependent on a choice of filter function for the data set and a choice of clustering method, Ball Mapper’s approach to obtaining the cover is simpler and requires only one input parameter, a radius ϵ .

Given a set of points $X \subset \mathbb{R}^d$ and $\epsilon > 0$, see Figure 1(A), Ball Mapper finds a collection $C \subset X$, called an ϵ -net Figure 1(B), such that no point in X is located further than distance ϵ from an element of C . This collection determines an overlapping cover $B(X, \epsilon) = \{B_1, B_2, \dots, B_i, \dots\}$: a set of d -dimensional balls with radius ϵ that contain all of X , Figure 1(C). The algorithm then constructs the nerve of $B(X, \epsilon)$, called Ball Mapper graph (BMGraph) Figure 1(D). In a Ball Mapper graph, vertices v_i correspond to the balls B_i in the cover, and two vertices v_i, v_j are connected by an edge if and only if B_i and B_j contain shared points in X .

The Ball Mapper parameter ϵ determines the scale of the features highlighted in the output. By varying ϵ , we create a linear sequence of images which highlight structures at successive scales, inspired by the filtration in persistent homology (but in this case we do not have the stability result). Additionally, we can choose to filter our data according to a range of values of a given criterion and obtain an output that can be displayed as a sequence of graphs, see Figure 4.

Results in this paper are obtained using the R implementation of Ball Mapper [18]. This version of Ball Mapper supports additional visualization features; e.g. the disks corresponding to vertices of the Ball Mapper graph are scaled according to the number of points they contain. Thus larger disks in the visualization correspond to larger clusters of points within radius ϵ of each other. Another useful feature is coloring the Ball Mapper graph using values of some function $f : X \rightarrow \mathbb{R}$ defined on the input data. The color of each Ball Mapper disc is the average value of the functions f on the points of X in that cluster. For convenience, we include the colored vertical bar on the right of each image which describes a uniform scale illustrating the distribution of colors from the minimum to maximum values of f within the data set.

3.3. Graph data. The focus of this paper is on the chromatic polynomials and our data contains chromatic polynomials for all simple connected graphs up to ten vertices for a total of 11,716,571 graphs. Our database contains many additional

G	Chromatic polynomial $P_G(\lambda)$	Vector $Q(G)$
C_3	$\lambda^3 - 3\lambda^2 + 2\lambda$	(1, 3, 2, 0, 0, 0, 0, 0, 0)
K_5	$\lambda^5 - 10\lambda^4 + 35\lambda^3 - 50\lambda^2 + 24\lambda$	(1, 10, 35, 50, 24, 0, 0, 0, 0)
C_{10}	$\lambda^{10} - 10\lambda^9 + 45\lambda^8 - 120\lambda^7 + 210\lambda^6 - 252\lambda^5 + 210\lambda^4 - 120\lambda^3 + 45\lambda^2 - 9\lambda$	(1, 10, 45, 120, 210, 252, 210, 120, 45, 9)

TABLE 1. Chromatic point cloud examples: chromatic coefficient vectors (absolute values) for the triangle, complete graph on 5 vertices, the cycle graph with all ten non-zero entries.

properties including numbers of vertices and edges, spectral irregularity, variance irregularity, numbers of subgraphs including cycles and K_4 s, maximal and minimal vertex degrees, and identifications of all graphs which belong to classes including trees, Turan graphs, or bipartite graphs.

Chromatic polynomials are computed using the Sage `GraphGenerators` class implemented in Python 3 [52, 23]. The point cloud consists of the set of all chromatic polynomial coefficients obtained in the following way. Given graph G with chromatic polynomial $P_G(\lambda) = c_n\lambda^n + c_{n-1}\lambda^{n-1} + \dots + c_1\lambda$, we represent $P_G(\lambda)$ by its coefficient vector $Q(G) = (|c_n|, |c_{n-1}|, \dots, |c_1|, 0, \dots, 0)$. For simplicity and computational efficiency, we take absolute values of each coefficient, since the sign is given by $(-1)^k$ for each coefficient c_{n-k} and thus remains constant within each feature of our data set. Our data set consists of coefficient vectors with at most 10 non-zero entries, so we standardize length by adding zero entries for graphs with less than 10 vertices. Table 1 contains sample vectors for the triangle C_3 , the complete graph K_5 , and the cycle C_{10} with 10 edges.




FIGURE 2. A PCA projection of the chromatic polynomial data to 2 dimensions determined by the first two principal components and colored by number of edges for graphs with eight (A) and nine vertices (B).

In the sections that follow, we visualize data sets corresponding to graphs with a fixed number of vertices. Within each data set, the varying sizes of chromatic coefficients make it difficult to see important features during visualization. We apply a log transformation $f(x) = \ln(1 + x)$ to the data set derived from graphs with n vertices and then separately normalize each feature to the range $[0, 1]$ with min-max scaling over this particular data set.




FIGURE 3. The PCA projection from Figure 2(B) with clusters containing some well known graphs circled and labeled. In addition to multiple non-isomorphic graphs the following clusters contain A: all 9-vertex trees, B: complete graph K_5 with 4 pendant edges, C: complete graph K_6 with 3 pendant edges, D: complete graph K_7 with 2 pendant edges, E: complete graph K_8 with 1 pendant edge, F: complete graph K_8 with triangle attached along 1 edge, G: complete graph K_9 , and K_9 with 1 edge removed.

4. Chromatic polynomial data: structure and dimensionality

The main results of this paper rely on introducing big data analysis to graph theory in a novel way inspired by applications to knot theory [43, 36]. In these approaches a graph is represented by a unique vector, consisting of coefficients of the chromatic polynomial. In this paper we focus on filtered PCA and Ball Mapper, but this graph representation lends itself to any of the big data techniques, including machine learning and other artificial intelligence tools.

Input data consists of chromatic polynomials for graphs up to ten crossings as described in Section 3.3. This point cloud naturally embeds in \mathbb{R}^9 and we are using Euclidean distance as our metric for this initial analysis.

Note that Theorem 2.12 provides an interpretation of the first few chromatic coefficients in terms of number of vertices and edges, as well as the number of special subgraphs such as cycles, complete graphs, and many others. Moreover, Bollobás, Pebody and Riordan conjecture that almost every pair of independently chosen random graphs are distinguished by their Tutte and chromatic polynomials.

CONJECTURE 4.1 ([8]). *Let $G(n, 1/2)$ be the space of random graphs on a fixed set of n vertices in which each pair of vertices is joined independently with probability $1/2$. Let $T(G), p(G)$ denote the Tutte polynomial and chromatic polynomial of G respectively. Then almost every graph $G \in G(n, 1/2)$ is such that $T(G') = T(G)$ implies G' is isomorphic to G . In addition, almost every graph $G \in G(n, 1/2)$ is such that $p(G') = p(G)$ implies G' is isomorphic to G .*

While the conjecture remains open, it has been shown that the probability that two independently chosen random graphs have the same Tutte polynomial is of order $O(1/\log n)$ [37].

4.1. Dimension of the chromatic data. Inspired by the big data approach to the Jones polynomial, a polynomial invariant for knots [36], we apply Principal Component Analysis (PCA) to the chromatic data for graphs with increasing number of vertices, up to ten maximum. The *dimension* of the Jones polynomial

is defined to be the smallest value \mathbf{d} for which the normalized explained variance of the first \mathbf{d} principal components sums to more than 95% across all considered orders [36].

According to this definition and PCA computations, chromatic polynomial data is essentially one-dimensional, and this dimensionality is stable with respect to increasing the order of a graphs, see Table 2. This table presents a snapshot of PCA computations for chromatic polynomial data for graphs of orders $n = 8, 9, 10$ and the first three principal components for each data set. In particular the last two rows in Table 2 contain the normalized (NEV) and cumulative explained variance (CEV) for $n = 8, 9, 10$ indicating that almost all variation is explained by the first principal component. More precisely, the cumulative explained variance of PC1, S_1 , is greater than 0.99 for all orders up to 10. In Section 4.3 we revisit this question since the structure of the Ball Mapper graph is linear, see comparison with the PCA in Figure 6.

	$n = 8$			$n = 9$			$n = 10$		
	PC1	PC2	PC3	PC1	PC2	PC3	PC1	PC2	PC3
NEV	0.99188	0.00792	0.00018	0.99178	0.00806	0.00015	0.99189	0.00796	0.00015
CEV	0.99188	0.99980	0.99998	0.99178	0.99984	0.99999	0.99189	0.99985	0.99999

TABLE 2. The normalized explained variance $\overline{\lambda}_j(n)$ (NEV) and cumulative normalized explained variance $S_j(n)$ (CEV) for the first three principal components of the chromatic polynomial data of graphs with orders 8, 9, and 10 respectively.

4.2. Graph-theoretic interpretation of the principal components. To get an insight into the graph-theoretic relevance of the principal components and the dimension reduction of our data from 10 to one (or two), we provide visualisations of the first two principal directions and further insights obtained using Ball Mapper; see Figure 6.

Focusing on the two most significant principal components on chromatic polynomials of graphs with fixed order, we visualize the structure of each data set by plotting each graph as a coordinate pair (x_1, x_2) pair where x_i represents the i th PC score of that graph, see Figure 2. This figure illustrates the persistent structure found for graphs with eight and nine vertices and the stability of the dimension reduction with respect to the number of vertices of a graph. The first principal component corresponds to the horizontal axis in Figure 2, which appears to be roughly correlated with an increase in the number of edges. Graphs with a fixed number of edges are found in clusters separated by thin regions of empty space in this projection. Figure 3 identifies the location of certain well-known graphs, such as trees and complete graphs, at the boundary of the projection. This boundary also contains maximum coefficient graphs; see Figure 12.

The data set of chromatic polynomials for graphs of 9 vertices has 8 dimensions, corresponding to the eight coefficients which are not trivially equal to 0 or 1. The loadings for PC1 and PC2 in Table 3 are the entries of the eigenvectors that give the directions of greatest variation in this eight-dimensional data set. We observe that all loadings for PC1 are positive, so that movement along this axis corresponds to an overall increase or decrease in coefficient size. This agrees with our previous

observation that PC1 is correlated with the number of edges in a graph. For PC2, the loadings for the last four coefficients have the opposite sign from the first four, indicating that an increase in PC2 score correlates with an decrease in relative size of the last few coefficients.

	c_8	c_7	c_6	c_5	c_4	c_3	c_2	c_1
PC1	0.34694	0.34683	0.34761	0.34876	0.35005	0.35172	0.35579	0.379545
PC2	0.54568	0.39243	0.24013	0.08614	-0.06850	-0.22145	-0.37150	-0.53983

TABLE 3. PC loadings for two most significant principal components, PC1 and PC2, for the chromatic polynomial data for all graphs of order 9, with chromatic polynomial determined by eight coefficients $P_G(\lambda) = \lambda^9 + c_8\lambda^8 + \dots + c_1\lambda$

PC loadings provide a way to think about principal components in terms of the coefficients of the chromatic polynomial, and in turn, based on the formulas in Theorem 2.12, features of a graph. We hypothesize that PC1 measures the average size of the coefficient vector as a whole, and that PC2 detects the relative difference between the first and last coefficients.

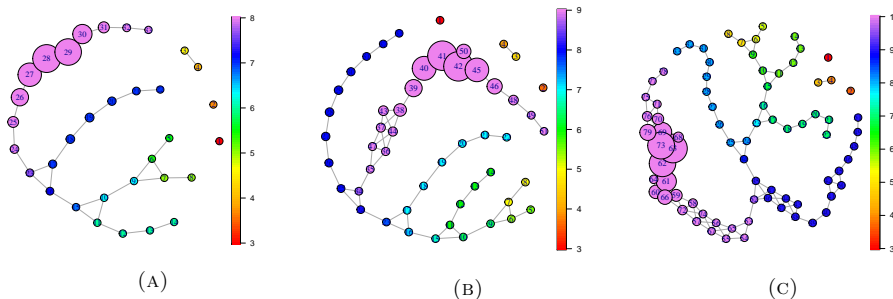


FIGURE 4. Stability of BMGraphs for chromatic data with respect to the number of vertices with graphs of order at most 8 (A), nine (B) and ten (C) with the Ball Mapper parameter $\epsilon = 0.1$, colored by the number of vertices.

4.3. Chromatic data: Ball Mapper structure. In this section we apply Ball Mapper to the chromatic point cloud consisting of coefficient vectors of all graphs up to $n = 10$ vertices.

First, we provide the evidence for the stability of the Ball Mapper approach to the chromatic polynomial data with respect to the order of a graph. In Figures 4 and 5 we provide evidence for a persistent structure across data sets of graphs with order equal to eight, nine, and ten respectively. Graphs of each given order are clustered in long branch-like features, with neighboring branches of adjacent orders across all three BMGraphs. These branch-like features are also reflected in the two-dimensional PCA plot for all graphs up to 10 vertices combined, see Figure 5. This structure invites and justifies exploring the data corresponding to graphs of a given order.




FIGURE 5. A two-dimensional PCA projection of the chromatic polynomial data for all graphs with 3-10 vertices. The projection is based on the first two principal components for normalized data, colored by number of vertices.

Figure 7 and Figure 8 provide a detailed structure of branches consisting only of graphs with exactly 8 and 9 vertices respectively. Both Ball Mapper graphs have linear structure consistent with the PCA analysis in Section 4.1. To confirm the informal correspondence between the 1-dimensionality of the data obtained by the PCA approach and the “linear” structure of the BMGraph, we color each cluster in the BMGraph by the average PC scores of the graphs it contains, Figure 6. The remarkable agreement between the two methods suggests that the first principal direction corresponds to the linear feature of the Ball Mapper structure.




FIGURE 6. Clusters 9-12 in BMGraph of the Chromatic polynomial data shown in Figure 8 at radius $\epsilon = 0.06$ colored by the average PC1 score (A), and the average PC2 score (B) of graphs contained in each cluster.

The linear direction within BMGraphs corresponds to the number of edges in a graph, see Figures 7(B) and 8(B), as well as the L^2 norm of the chromatic coefficient vectors, see Figure 7(C), 8(C). This result can be explained by the fact that the

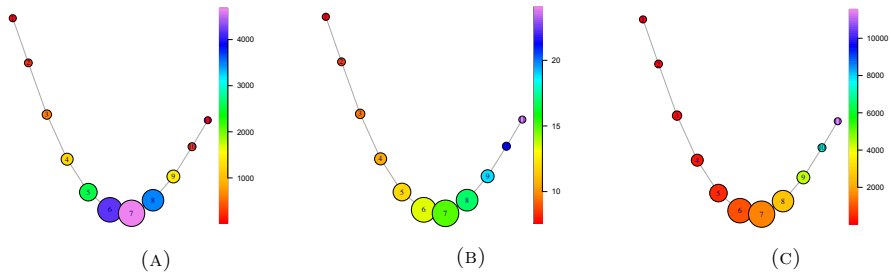


FIGURE 7. BMGraph of the chromatic polynomial data of all graphs with exactly 8 vertices and the Ball Mapper parameter $\epsilon = 0.22$ colored by cluster size (A), number of edges (B), and the norm of the chromatic polynomial (C).

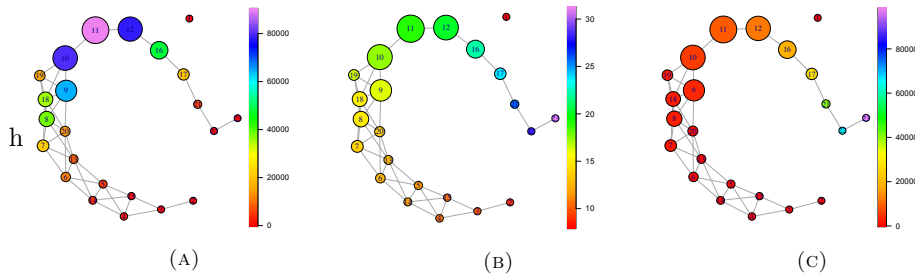


FIGURE 8. BMGraph of the chromatic polynomial data of all graphs with exactly 9 vertices and the Ball Mapper parameter $\epsilon = 0.15$ colored by cluster size (A), number of edges (B), and the norm of the chromatic polynomial (C).

number of edges E contributes to the largest terms of each formula in Theorem 2.12.

Based on the BMGraph in Figure 8(A) and associated data, four largest clusters contain more than 50,000 out of total of 261079 points/graphs: cluster 9 (64442), cluster 10 (80017), cluster 11 (90054), and cluster 12 (78878). Therefore, we analyze the collection of these 4 clusters using a smaller value of radius ϵ , see BMGraphs on Figure 9. As expected, the major PC1 direction of this particular embedding on the BMGraph is correlated with the number of edges Figure 9(A), and the PC2 direction is roughly consistent with the L^2 norm of the chromatic coefficients, Figure 9(B). The average number of triangles in a graph contained in any given cluster, as well as the average degree distance seem to be varying in the PC2 direction, complementary to the main linear structure, from north-west to south-east in Figures 9(C-D).

4.4. Extremal chromatic polynomials and irregularity. Section 2.2 provides a few details on regular graphs and irregularity measures, but is by no means comprehensive. Note that there is a lot of ambiguity when it comes to irregularity measures. It was recently shown that no two of the following irregularity measures are mutually comparable: spectral irregularity, variance irregularity, Albertson index, and total irregularity [2].




FIGURE 9. Clusters 9-12 in BMGraph of the chromatic polynomial data shown in Figure 8 at radius $\epsilon = 0.06$ colored by: number of edges (A) and by L^2 norm of chromatic coefficients (B), average number of triangles in a graph (C), average difference between the maximal and minimal degree (D).

Section 2.3 offers a brief survey of results on graph compression, which uniformly increases or decreases values of a number of graph invariants related to the chromatic and Tutte polynomials.

PROPOSITION 4.2. *Suppose that G' is obtained from G by compression. Then*

- (1) *The coefficients of the chromatic polynomial of G' are smaller than those of G : $|c_i(G')| \leq |c_i(G)|$ for all $1 \leq i \leq n$.*
- (2) *Graph G' has greater variance irregularity than G : $\sigma(G') \geq \sigma(G)$.*
- (3) *Graph G' has greater spectral irregularity than G : $\epsilon(G') \geq \epsilon(G)$.*

PROOF. Part (1) is a restatement of Theorem 6.3 of [13]. Part (2) derived from Lemma 2.3 of [14] with Definition 2.19 in mind. Part (3) follows from Theorem 2.1 of [12] with respect to Definition 2.18. \square

As a consequence, if we consider graphs with a fixed number of vertices and edges, then part (1) claims that threshold graphs and their co-chromatic graphs are among the chromatically minimal and part (2) says that threshold graphs maximize variance irregularity.

The rest of this section provides visualizations of these observations and a large-scale comparison between the spectral and variance irregularity measures, highlighting the relation between threshold graphs (the extremal graphs with respect to compression) and the coefficients of the chromatic polynomial. For example, the graphs with extremal PCA2 scores and minimal coefficient graphs, as in Figure 12(C), are either threshold or co-chromatic with threshold graphs, see Figure 10.

The BMGraph 11(A) is colored by spectral irregularity and (B) by variance irregularity emphasizing that, in these particular embeddings of BMGraphs, clusters on the left- and right-hand sides contain graphs with extreme irregularity measures, with the average irregularity monotonically increasing from left to right. The same phenomenon appears in Figure 9(D) when the coloring function is the difference between maximal and minimal vertex degrees.

These observations based on Ball Mapper graphs are matched by visualizations of the chromatic data in the 2-dimensional PCA projection, see Figure 12. In both projections, higher PC2 score corresponds to higher values of graph irregularity measures within the subset consisting of graphs with equal number of edges. However, the overall distributions of spectral Figures 12(A), and variance irregularity




FIGURE 10. A PCA projection of the chromatic polynomial data for all graphs of order 9 into 2-dimensions determined by the two most significant principal directions, PC1 and PC2, colored by presence of threshold graphs denoted by blue color.




FIGURE 11. Clusters 9-12 in GMGraph of the Chromatic polynomial data shown in Figure 8 at radius $\epsilon = 0.06$ colored by average measure of spectral irregularity in each cluster (A), variance irregularity in each cluster (B), presence of chromatically minimal (purple color) (C) and presence of chromatically maximal graphs (D)

Figures 12(B) are different, with the variance regularity being lower on the middle left section.

The location of chromatically minimal and maximal graphs of order 9 (see Definition 2.17) in the BMGraph and 2-dimensional PCA projection is shown on Figures 11(C), 11(D) and 12(C), 12(D), respectively. Both methods show that chromatically minimal graphs are found in regions of highest irregularity. This is expected since the families of minimal graphs characterized by Definition 2.14 include members of the high-irregularity connected quasi-complete graphs and connected quasi-stars (Section 2.2). For example, it follows from Definition 2.20 that the connected quasi-complete graph $QC(n, m)$ is an element of the family $L(n, m)$.

While, to the best of our knowledge, there exists no similar characterization for chromatically maximal graphs, their location among the most regular graphs is supported by results from extremal graph theory. Lazebnik and others prove partial characterizations for graphs which maximize evaluations of the chromatic polynomial $P_G(\lambda)$ at given values of λ [34]. These classes of graphs include Turan graphs (Definition 2.23) and are generally similar to Turan graphs in structure.




FIGURE 12. A PCA projection of the chromatic polynomial data for all graphs of order 9 into 2-dimensions determined by two most significant principal directions, PC1 and PC2, colored by spectral (A) and variance irregularity measure (B), the presence of minimum-coefficient (C) and maximum-coefficient graphs (D) denoted by red color.




FIGURE 13. Extremal graphs in $G(9, 11)$ with coefficient vectors equal to: $(1, 11, 51, 131, 205, 201, 121, 41, 6)$ (A), and $(1, 11, 55, 165, 328, 446, 406, 224, 56)$ (B).

Turan graphs are extremal graphs which are chromatically unique [9] and close to regular, see Section 2.2. The location of all Turan graphs of order 9 within the 2-dimensional PCA projection, among other chromatically maximal graphs, is highlighted in red in 12(D).

Computations for graphs up to nine vertices imply that:

- The minimal and maximal elements of the chromatic polynomial poset $G(n, m)$ achieve the extremal PC2 scores among all graphs with n vertices and m edges when the principal components are computed over the data set of all n -vertex graphs.
- The Turan graph $T(n, m)$ is maximal in the poset $G(n, m)$.

Further investigation of chromatically maximal graphs in our data suggests that there is a relationship between the “tail” of the polynomial (the last few coefficients) and graph irregularity. For example, consider the graphs in Figure 13 with $n = 9$ vertices and $E = 11$ edges. The graph on the left is one of three co-chromatic non-isomorphic graphs which are minimal in $G(9, 11)$, with higher irregularity measures and significantly smaller tail than the graph in Figure 13(B), which is maximal in $G(9, 11)$.

References

1. Hosam Abdo, Stephan Brandt, and Darko Dimitrov, *The total irregularity of a graph*, Discrete Mathematics and Theoretical Computer Science **16** (2014), no. 1, 201–206.
2. Hosam Abdo, Darko Dimitrov, and Ivan Gutman, *Graph irregularity and its measures*, Applied Mathematics and Computation **357** (2019), 317–324.
3. Rudolf Ahlswede and Gyula OH Katona, *Graphs with maximal number of adjacent pairs of edges*, Acta Mathematica Academiae Scientiarum Hungarica **32** (1978), no. 1-2, 97–120.
4. Michael O Albertson, *The irregularity of a graph*, Ars Combinatoria **46** (1997), 219–225.
5. Francis K Bell, *A note on the irregularity of graphs*, Linear Algebra and its Applications **161** (1992), 45–54.
6. Halina Bielak, *A new method for counting chromatic coefficients*, Annales UMCS Informatica AI **3** (2005), no. 1, 179–189.
7. Zbigniew Ryszard Bogdanowicz, *Spanning trees in undirected simple graphs*, Stevens Institute of Technology, 1985.
8. Béla Bollobás, Luke Pebody, and Oliver Riordan, *Contraction–deletion invariants for graphs*, Journal of Combinatorial Theory, Series B **80** (2000), no. 2, 320–345.
9. Chong-Yun Chao and George A Novacky Jr, *On maximally saturated graphs*, Discrete Mathematics **41** (1982), no. 2, 139–143.
10. Chong-Yun Chao and Earl Whitehead Jr., *On chromatic equivalence of graphs*, Theory and Applications of Graphs (Yousef Alavi and Don Lick, eds.), Lecture Notes in Mathematics, Springer, 1976, pp. 121–131.
11. ———, *Chromatically unique graphs*, Discrete Mathematics **27** (1979), no. 2, 171–177.
12. Péter Csikvári, *On a conjecture of V. Nikiforov*, Discrete Mathematics **309** (2009), no. 13, 4522–4526.
13. ———, *Applications of the Kelmans transformation: extremality of the threshold graphs*, The Electronic Journal of Combinatorics (2011), P182–P182.
14. Jonathan Cutler and AJ Radcliffe, *Extremal graphs for homomorphisms*, Journal of Graph Theory **67** (2011), no. 4, 261–284.
15. Alex Davies, András Juhász, Marc Lackenby, and Nenad Tomasev, *The signature and cusp geometry of hyperbolic knots*, arXiv preprint arXiv:2111.15323 to appear in Geometry & Topology (2021).
16. Alex Davies, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, et al., *Advancing mathematics by guiding human intuition with AI*, Nature **600** (2021), no. 7887, 70–74.
17. Darko Dimitrov and Riste Škrekovski, *Comparing the irregularity and the total irregularity of graphs*, Ars Mathematica Contemporanea **9** (2014), no. 1, 45–50.
18. P Dlotko, *Ballmapper: Create a ball mapper graph of the input data*, R package version 0.1.0 (2019).
19. Paweł Dłotko, *Ball mapper: a shape summary for topological data analysis*, arXiv preprint arXiv:1901.07410 (2019).
20. Fengming Dong, Khee Meng Koh, and Kie Leong Teo, *Chromatic Polynomials and Chromaticity of Graphs*, World Scientific, Singapore, 2005.
21. Michael Eastwood and Stephen Huggett, *Euler characteristics and chromatic polynomials*, European Journal of Combinatorics **28** (2007), no. 6, 1553–1560.
22. Edward Farrell, *On chromatic coefficients*, Discrete Mathematics **29** (1980), 257–264.
23. Conda forge community, *The conda-forge Project: Community-based Software Distribution Built on the conda Package Format and Ecosystem*, July 2015.
24. Laure Helme-Guizon and Yongwu Rong, *A categorification for the chromatic polynomial*, Algebraic & Geometric Topology **5** (2005), no. 4, 1365–1388.
25. Richard Hepworth and Simon Willerton, *Categorifying the magnitude of a graph*, Homology, Homotopy and Applications (2017), 31–60.
26. Mark C Hughes, *A neural network approach to predicting and computing knot invariants*, Journal of Knot Theory and Its Ramifications **29** (2020), no. 03, 2050005.
27. June Huh, *Milnor numbers of projective hypersurfaces and the chromatic polynomial of graphs*, Journal of the American Mathematical Society **25** (2012), no. 3, 907–927.
28. Vishnu Jejjala, Arjun Kar, and Onkar Parrikar, *Deep learning the hyperbolic volume of a knot*, Physics Letters B **799** (2019), 135033.

29. Ian T Jolliffe and Jorge Cadima, *Principal component analysis: a review and recent developments*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **374** (2016), no. 2065, 20150202.
30. Nathan Kahl, *Graph vulnerability parameters, compression, and threshold graphs*, Discrete Applied Mathematics **292** (2021), 108–116.
31. ———, *Extremal graphs for the Tutte polynomial*, Journal of Combinatorial Theory, Series B **152** (2022), 121–152.
32. Alexander K Kelmans, *On graphs with randomly deleted edges*, Acta Mathematica Academiae Scientiarum Hungarica **37** (1981), no. 1, 77–88.
33. Lauren Keough and Andrew J Radcliffe, *Graphs with the fewest matchings*, Combinatorica **36** (2016), no. 6, 703–723.
34. Felix Lazebnik, *The maximum number of colorings of graphs of given order and size: A survey*, Discrete Mathematics **342** (2019), no. 10, 2783–2791.
35. Tom Leinster, *The magnitude of metric spaces*, Documenta Mathematica **18** (2013), 857–905.
36. Jesse SF Levitt, Mustafa Hajij, and Radmila Sazdanovic, *Big data approaches to knot theory: Understanding the structure of the Jones polynomial*, arXiv:1912.10086 (2019).
37. Martin Loebl, Jiří Matoušek, and Ondřej Pangrác, *Triangles in random graphs*, Discrete mathematics **289** (2004), no. 1-3, 181–185.
38. Guy Meredith, *Coefficients of chromatic polynomials*, Journal of Combinatorial Theory, Series B **13** (1972), 14–17.
39. Vladimir Nikiforov, *Eigenvalues and degree deviation in graphs*, Linear Algebra and its Applications **414** (2006), no. 1, 347–360.
40. Marc Noy, *Graphs determined by polynomial invariants*, Theoretical Computer Science **307** (2003), no. 2, 365–384.
41. Joelma Ananias de Oliveira, Carla Silva Oliveira, Claudia Justel, and Nair Maria Maia de Abreu, *Measures of irregularity of graphs*, Pesquisa Operacional **33** (2013), no. 3, 383–398.
42. Dłotko Paweł, *Ball mapper: a shape summary for topological data analysis*, arXiv:1901.07410 (2019).
43. Dłotko Paweł, Davide Gurnari, and Radmila Sazdanovic, *Knot invariants and their relations: a topological perspective*, arXiv preprint arXiv:2109.00831 (2021).
44. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Scikit-learn: Machine learning in Python*, Journal of Machine Learning Research **12** (2011), 2825–2830.
45. Ronald Read, *An introduction to chromatic polynomials*, Journal of Combinatorial Theory **4** (1968), 52–71.
46. J Rodriguez and A Satyanarayana, *Chromatic polynomials with least coefficients*, Discrete Mathematics **172** (1997), no. 1-3, 115–119.
47. Gian-Carlo Rota, *On the foundations of combinatorial theory I. Theory of Möbius functions*, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete **2** (1964), no. 4, 340–368.
48. Appajosyula Satyanarayana, L Schoppmann, and Charles L Suffel, *A reliability-improving graph transformation with applications to network reliability*, Networks **22** (1992), no. 2, 209–216.
49. Radmila Sazdanović and Martha Yip, *A categorification of the chromatic symmetric function*, Journal of Combinatorial Theory, Series A **154** (2018), 218–246.
50. Gurjeet Singh, Facundo Mémoli, and Gunnar Carlsson, *Topological methods for the analysis of high dimensional data sets and 3D object recognition*, SPBG **91** (2007), 100.
51. Richard P Stanley, *A symmetric function generalization of the chromatic polynomial of a graph*, Advances in Mathematics **111** (1995), no. 1, 166–194.
52. W. A. Stein et al., *Sage Mathematics Software (Version 9.1)*, The Sage Development Team, 2020, <http://www.sagemath.org>.
53. Paul Turán, *Eine extremalaufgabe aus der graphentheorie*, Mat. Fiz. Lapok **48** (1941), 436–452.
54. William Thomas Tutte, *A contribution to the theory of chromatic polynomials*, Canadian Journal of Mathematics **6** (1954), 80–91.
55. Lothar Von Collatz and Ulrich Sinogowitz, *Spektren endlicher grafen*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, vol. 21(1), Springer, 1957, pp. 63–77.

56. Hassler Whitney, *A logical expansion in mathematics*, Bulletin of the American Mathematical Society **38** (1932), no. 8, 572–579.

DEPARTMENT OF MATHEMATICS, NORTH CAROLINA STATE UNIVERSITY
Email address: rsazdanovic@math.ncsu.edu

DEPARTMENT OF MATHEMATICS, FRANCIS MARION UNIVERSITY
Email address: daniel.scofield@fmarion.edu