INTERNATIONAL JOURNAL OF COMPUTATIONAL MATHEMATICAL IDEAS Vol. 1, No. 1 (2009), pp. 19–28
ISSN (online) • www.ijcmi.webs.com *Research Article*

Novel Graph Machine Learning Algorithms on Planar Coloured Graphs

Colour Refinement, Four-Colour Structure, and Chromatic Kernels for Learning

Satyanarayana Sanakkayala
Research Scholar, Graph Machine Learning
School of Science and Technology, Dravidian University, Kuppam, India

Received: 01 January 2009 • Accepted: 01 May 2009 • Published: 20 May 2009

Abstract. A coloured graph carries a discrete attribute on each vertex, and when the underlying graph is planar the interaction between that attribute and the topology of the plane embedding becomes especially rich. This paper, a companion to the author's study of learning on planar graphs, proposes four novel algorithms that make colour a first class citizen of graph machine learning on planar structures. We begin from two classical facts, that colour refinement in the sense of Weisfeiler and Leman produces the coarsest equitable partition of a graph, and that every planar graph is four colourable, and we show that each of these facts yields a concrete learning primitive on planar inputs. The first algorithm turns colour refinement into a permutation invariant feature map whose stable colouring is computed in near linear time on planar graphs. The second, chromatic block propagation, uses a proper four colouring to schedule harmonic label propagation as at most four fully parallel sub sweeps, because each colour class is an independent set with no internal coupling. The third builds a chromatic diffusion kernel that fuses colour refinement histograms with heat flow on the graph. The fourth uses colour classes as an independent set backbone for multilevel coarsening of the Laplacian. Each algorithm is analysed for correctness and cost, the intuition connecting equitable partitions to the Laplacian spectrum is made explicit, and behaviour is illustrated on coloured plane graphs and grids. The unifying message is that on planar coloured graphs the chromatic structure is not a nuisance to be hashed away but a schedule, a feature, and a coarsening all at once.

Keywords: planar coloured graphs; colour refinement; Weisfeiler and Leman; four colour theorem; independent sets; multicolour Gauss and Seidel; graph kernels; multilevel coarsening.

2000 MSC: 05C15; 05C85; 68T05; 68R10; 65F10.

1. Introduction

In many applications the vertices of a graph are not anonymous. Atoms in a molecule carry an element type, pixels in a segmented image carry a region label, and cells in a map carry a country or a frequency band. A graph with such a per vertex attribute is a *coloured graph*, and when the graph is planar the colour and the plane embedding constrain each other in ways a learner can exploit. This paper extends the companion study of learning on planar graphs [23] by placing colour at the centre of the algorithms rather than at the periphery.

Two classical results anchor the development. The first is colour refinement, introduced by Weisfeiler and Leman, which repeatedly recolours every vertex by combining its own colour with the multiset of its neighbours' colours until the colouring stabilises [3]. The stable colouring is the coarsest equitable partition of the graph, and it is the mathematical seed of almost every message passing method in graph learning. The second is the four colour theorem, which guarantees that the vertices of any planar graph can be properly coloured with at most four colours [8], a fact that a modern quadratic time procedure makes constructive [13]. We show that colour refinement gives a feature map and that four colourability gives a parallel schedule, and that both become inexpensive precisely because the graph is planar.

The paper is organised as follows. Section 2 fixes the chromatic vocabulary and recalls the landscape of planar colouring, including the contrast between the always available four colouring and the hardness of three colouring. Section 3 develops colour refinement as a feature map and connects the stable colouring to the Laplacian spectrum through equitable partitions, giving Algorithm 1. Section 4 introduces chromatic block propagation, Algorithm 2, and proves that the four colouring turns harmonic propagation into at most four parallel sub sweeps. Section 5 builds the chromatic diffusion kernel, Algorithm 3. Section 6 uses colour classes for multilevel coarsening, Algorithm 4. Section 7 summarises complexity, Section 8 surveys applications, Section 9 reports illustrative computations, and Section 10 concludes.

Standard background on graphs and colouring is available in textbooks [15], spectral prerequisites are in the monograph of Chung [12], and the iterative solver machinery behind the parallel schedule is treated by Saad [19].

2. Planar Coloured Graphs and the Chromatic Landscape

Let $G=(V,E)$ be a finite simple planar graph on n vertices, given with a plane embedding, and let $c:V\to\Sigma$ assign to each vertex a colour from a finite alphabet Σ. We call the pair (G,c) a *coloured planar graph*. Two notions of colour appear in this paper and it helps to keep them apart. An *input colouring* is a given attribute, arbitrary and possibly repeated across adjacent vertices. A *proper colouring* is a map $\varphi:V\to\{1,\dots,k\}$ with $\varphi(i)\neq\varphi(j)$ whenever $\{i,j\}\in E$, so that each colour class $C_r=\varphi^{-1}(r)$ is an independent set. The least k for which a proper colouring exists is the *chromatic number* $\chi(G)$.

Theorem 1 (Four colour theorem). *Every planar graph satisfies $\chi(G)\leq 4$ [8], and a proper four colouring can be computed in quadratic time by a constructive procedure [13].*

The bound of four is the single structural gift that Section 4 converts into parallelism. It is worth stressing the asymmetry of the chromatic landscape on planar graphs, since it explains why we build on four colouring rather than on the smallest colouring.

Remark 1 (Why four and not three). Deciding whether a planar graph is three colourable is NP complete [7], so a learner cannot rely on obtaining a minimum colouring cheaply. In contrast a four colouring always exists and is polynomially constructive, and triangle free planar graphs are always three colourable by the theorem of Grötzsch [1]. For scheduling purposes any proper colouring suffices, and a greedy pass already gives at most six colours on a planar graph because such graphs have a vertex of degree at most five [2].

3. Colour Refinement as a Feature Map

3.1. The refinement iteration

Colour refinement starts from an initial colouring c^0, which may be the given input colouring or a constant, and repeatedly applies

$$c^{t+1}(v) \;=\; \textsc{hash}\Big(c^t(v),\ \{\{\,c^t(u):u\sim v\}\}\Big), \tag{1}$$

where $\{\{\cdot\}\}$ denotes a multiset and HASH assigns a fresh colour to each distinct argument [3]. The number of colours is non decreasing and bounded by n, so the iteration reaches a fixed point, the *stable colouring* c^∞, after at most $n-1$ rounds and in practice after a number of rounds no larger than the diameter.

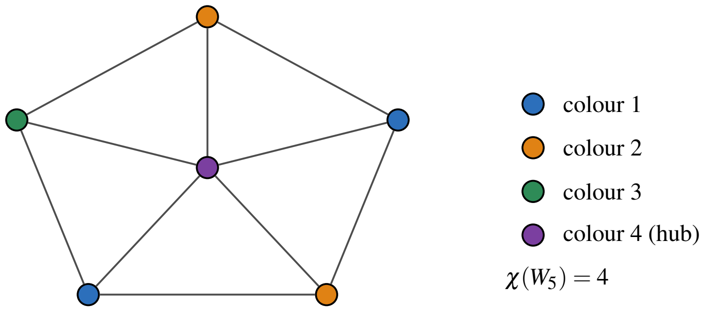

Figure 1: A proper four colouring of the wheel W_5, a planar graph whose odd rim forces three colours on the boundary and a fourth on the hub. The four colour theorem guarantees such a colouring for every planar graph, and the four colour classes are independent sets that later serve as a parallel update schedule.

(a) Round 0: one colour.

(b) Stable: four colours by role.

Figure 2: Colour refinement on a small planar graph. Starting from a single colour, the iteration (1) separates vertices by degree and then by the degrees of their neighbours, reaching a stable colouring that reflects structural role. The colour histogram of the stable colouring is a permutation invariant feature vector.

The intuition is that two vertices keep a common colour only while their surroundings remain indistinguishable. The first round separates vertices of different degree, the second separates vertices whose neighbours differ in degree, and so on outward. The fixed point is characterised cleanly.

Definition 1 (Equitable partition)**.** A partition $V = P_1 \cup \cdots \cup P_q$ is *equitable* if for every pair of classes P_r, P_s the number of neighbours in P_s is the same for every vertex of P_r. The stable colouring c^∞ induces the coarsest equitable partition that refines c^0.

3.2. Why refinement meets the spectrum

Equitable partitions are exactly the structure the Laplacian respects, which is why colour refinement carries spectral information for free.

Proposition 1 (Spectral containment)**.** *Let $V = P_1 \cup \cdots \cup P_q$ be an equitable partition of G with quotient, or divisor, matrix $B \in \mathbb{R}^{q \times q}$ whose entry B_{rs} is the common number of neighbours in P_s of a vertex in P_r. Then every eigenvalue of B is an eigenvalue of the adjacency matrix of G, and the corresponding eigenvectors are constant on each class [12].*

Sketch. Let $S \in \mathbb{R}^{n \times q}$ be the class indicator matrix. Equitability means $AS = SB$, so if $B\mathbf{x} = \mu\mathbf{x}$ then $A(S\mathbf{x}) = SB\mathbf{x} = \mu(S\mathbf{x})$, and $S\mathbf{x}$ is constant on classes by construction. □

Proposition 1 says the colours produced by refinement are compatible with the smooth eigenvectors used by spectral learning in the companion paper [23]: functions that are constant on refinement classes live in an invariant subspace of the operator. A learner may therefore use the stable colour of a vertex, and the histogram of colours over the whole graph, as a permutation invariant descriptor that already encodes a slice of the spectrum.

3.3. The planar advantage and the algorithm

On general graphs colour refinement does not decide isomorphism, but planarity supplies a strong companion guarantee: isomorphism of planar graphs is decidable in linear time [6], so stable colourings on planar inputs can be turned into canonical descriptors without the ambiguity that dogs the general case. Combined with the linear sparsity of planar graphs from the companion paper, each refinement round costs $O(n)$ work with a radix sort of the neighbour multisets.

Algorithm 1: Planar Colour Refinement Features (PCRF)

Input: Coloured planar graph (G, c^0), round budget T
Output: Stable per vertex colours and a graph level histogram feature $\phi(G)$

$t \leftarrow 0$
while $t < T$ *and the colouring changed in the last round* **do**
  for *each vertex* v **do**
    compute the sorted multiset of neighbour colours
  Radix sort the pairs $\big(c^t(v),$ sorted neighbour colours$\big)$ over all v
  Assign $c^{t+1}(v)$ a fresh integer per distinct pair
  $t \leftarrow t+1$
Form $\phi(G)$ as the count of vertices of each stable colour
return the stable colours and $\phi(G)$

Cost. Each round is $O(n)$ by radix sort on planar sparsity, and the round count is at most the diameter, so the typical cost is $O(n \cdot \mathrm{diam})$ and the worst case is $O(n^2)$.

4. Chromatic Block Propagation

The companion paper solves semi supervised learning on a planar graph by the harmonic system $L_{uu}\mathbf{f}_u = -L_{u\ell}\mathbf{y}_\ell$, whose Jacobi and Gauss and Seidel iterations perform the neighbour averaging update $f_i \leftarrow \frac{1}{d_i}\sum_j w_{ij} f_j$ [23]. Gauss and Seidel converges faster than Jacobi but is inherently sequential, because updating vertex i uses the just updated values of its neighbours. A proper colouring removes exactly this obstacle.

Theorem 2 (Chromatic parallelism)**.** *Let φ be a proper colouring of G with classes $C_1, \dots, C_k$. Within one colour class no two vertices are adjacent, so a Gauss and Seidel sweep that updates all vertices of a class simultaneously is independent of the order inside the class and is identical to updating them one at a time. Consequently a full sweep decomposes into k fully parallel sub sweeps, one per colour class.*

Proof. For $i, j \in C_r$ with $i \neq j$ we have $\{i, j\} \notin E$, hence $w_{ij} = 0$ and the update of f_i does not reference f_j. The updates within C_r therefore commute and may be applied in parallel. Ordering the classes $C_1, \dots, C_k$ recovers a consistently ordered Gauss and Seidel iteration, which converges for the symmetric positive definite operator L_{uu} [19]. □

On a planar graph Theorem 1 caps the number of classes at $k \leq 4$, so a full harmonic sweep is at most four parallel steps regardless of n. This is the multicolour generalisation of the classical red and black ordering, and it is exact rather than heuristic.

Algorithm 2: Chromatic Block Propagation (CBP)

Input: Planar graph G, labels $\mathbf{y}_\ell$ on ℓ, tolerance ε
Output: Harmonic prediction $\mathbf{f}_u$ on the unlabelled set

Compute a proper colouring φ with classes $C_1, \dots, C_k$, $k \leq 4$ (greedy or four colouring)
Initialise $\mathbf{f}_u \leftarrow 0$ and clamp $\mathbf{f}_\ell \leftarrow \mathbf{y}_\ell$
repeat
  for $r \leftarrow 1$ **to** k **do**
    in parallel for *each unlabelled* $v \in C_r$ **do**
      $f_v \leftarrow \frac{1}{d_v}\sum_{u \sim v} w_{vu} f_u$
until *change* $< \varepsilon$
return $\mathbf{f}_u$

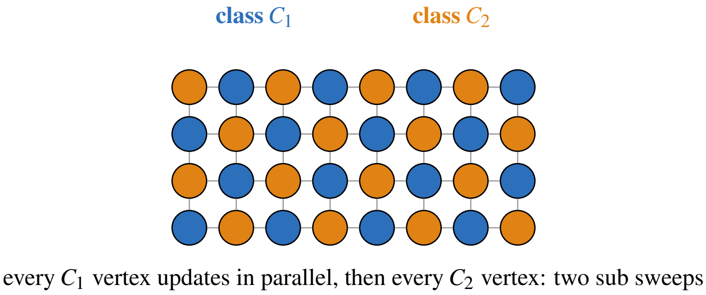

Figure 3: Chromatic block propagation on a bipartite plane grid, which needs only two colours. Each colour class is an independent set, so all like coloured vertices are updated at once with no internal coupling. A general planar graph needs at most four classes by the four colour theorem, giving at most four parallel sub sweeps per iteration.

Cost. Each iteration touches every edge once, $O(m) = O(n)$ work, split into $k \leq 4$ parallel phases; the number of iterations to tolerance ε is governed by the spectral gap of L_{uu}, which the companion paper bounds favourably for planar graphs.

5. The Chromatic Diffusion Kernel

To compare coloured planar graphs, or to measure similarity between vertices that respects both topology and colour, we fuse the refinement feature with heat flow. Recall the diffusion kernel $K_\beta = \exp(-\beta L)$ used in the companion paper [16]. We enrich it with colour in two compatible ways.

First, at the graph level, the refinement histograms $\phi(G)$ and $\phi(G')$ from Algorithm 1 give an explicit feature map, and their inner product is a valid positive semidefinite kernel between coloured planar graphs, in the spirit of marginalised kernels on labelled graphs [17], and closely related to the graph kernel family of Gärtner and coworkers [18]. Second, at the vertex level, we weight diffusion by colour agreement. Let Π_c be the diagonal indicator of colour c and define the colour aware generator

$$\widetilde{L} = L + \gamma \sum_{c \in \Sigma} \left(I - \Pi_c\right) \Theta_c \left(I - \Pi_c\right), \tag{2}$$

where Θ_c penalises heat crossing a colour boundary and $\gamma \geq 0$ tunes the penalty. The resulting *chromatic diffusion kernel* $\widetilde{K}_\beta = \exp(-\beta \widetilde{L})$ remains symmetric positive semidefinite, since $\widetilde{L}$ is symmetric positive semidefinite by construction, and it reduces to the ordinary diffusion kernel at $\gamma = 0$. The intuition is that heat now prefers to stay within a colour, so proximity reflects being connected through same coloured regions, which is often the relevant notion for coloured data. Shortest path variants may be substituted for robustness [20].

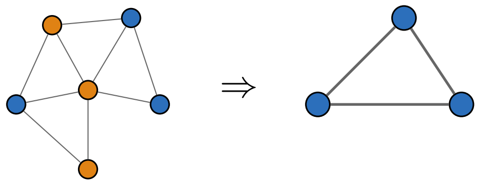

fine graph, $C^\star$ in blue   coarse graph, $\leq 3n/4$ smaller

Figure 4: Chromatic multilevel coarsening. The largest colour class of a four colouring is an independent set holding at least a quarter of the vertices. Contracting the other vertices onto this backbone gives a smaller, still planar, coarse graph on which the learning problem is solved and then prolonged back.

Algorithm 3: Chromatic Diffusion Kernel action (CDK)

Input: Coloured planar graph (G,c), time β, penalty γ, seed $\mathbf{x}$
Output: The vector $\widetilde{K}_\beta\,\mathbf{x} = \exp(-\beta\widetilde{L})\,\mathbf{x}$

1. Run Algorithm 1 to obtain stable colours (optional graph feature $\phi(G)$)
2. Assemble the sparse colour aware generator $\widetilde{L}$
3. Build a small Krylov subspace $\{\mathbf{x},\widetilde{L}\mathbf{x},\dots,\widetilde{L}^p\mathbf{x}\}$ by sparse products
4. Approximate $\exp(-\beta\widetilde{L})\mathbf{x}$ inside that subspace
5. **return** the approximation

Cost. $O(pn)$ with p small, since $\widetilde{L}$ is as sparse as L; the dense kernel is never formed.

6. Chromatic Multilevel Coarsening

Multilevel methods solve a problem on a small coarse graph and lift the answer back to the fine graph. The quality of a level hinges on choosing a well spread set of surviving vertices, and a colour class is exactly such a set, because it is independent and, for the largest class of a four colouring, contains at least a quarter of the vertices.

Proposition 2 (Coarsening ratio)**.** *Let φ be a proper four colouring of a planar graph and let $C^\star$ be its largest class. Then $|C^\star| \geq n/4$, and $C^\star$ is an independent set, so contracting each remaining vertex onto a neighbour in $C^\star$ yields a coarse graph with at most $3n/4$ vertices whose surviving vertices are pairwise non adjacent in the fine graph.*

Proof. The four classes partition V, so the largest has at least $n/4$ vertices by pigeonhole, and independence is the definition of a proper colouring. □

The coarse Laplacian is formed by the Galerkin projection $L_H = P^\top LP$ with a piecewise constant prolongation P over the contraction, and recursion gives a hierarchy in which each level shrinks by a constant factor. On planar graphs the coarse graphs remain sparse, and minor closure keeps them planar, so the guarantees of the companion paper survive at every level.

Algorithm 4: Chromatic Multilevel Coarsening (CMC)

Input: Planar graph G, coarsest size threshold n_0
Output: Solution of $L\mathbf{f} = \mathbf{b}$ (or a spectral embedding) on G
if $|V| \leq n_0$ **then**
  solve directly and **return**
Four colour G and let $C^\star$ be the largest class
Contract each non survivor onto a neighbour in $C^\star$, giving prolongation P
Form the coarse Laplacian $L_H = P^\top LP$ and coarse right hand side $P^\top \mathbf{b}$
Recurse on the coarse level to obtain $\mathbf{f}_H$
Prolong $\mathbf{f} \leftarrow P\mathbf{f}_H$ and apply a few chromatic block sweeps (Algorithm 2) to smooth
return f

Cost. The level sizes shrink geometrically, so the total work is $O(n)$ per V cycle plus the cost of the coarsest solve.

7. Complexity Summary

The four algorithms share the planar sparsity budget of the companion paper, and colour adds only lower order bookkeeping.

Algorithm	Dominant cost	Parallel structure
PCRF, Alg. 1	$O(n \cdot \mathrm{diam})$ typical	rounds sequential, vertices parallel
CBP, Alg. 2	$O(n)$ per iteration	≤ 4 parallel sub sweeps
CDK, Alg. 3	$O(pn)$, p small	sparse products parallel
CMC, Alg. 4	$O(n)$ per V cycle	levels sequential, sweeps parallel

Table 1: Costs of the proposed algorithms on a planar graph with n vertices and $O(n)$ edges. Each is near linear in n, and colouring contributes at most four fold structure rather than extra asymptotic cost.

8. Applications

Map and region labelling. Political maps and segmented images are planar coloured graphs by construction. Chromatic block propagation spreads a few region labels across a map with full parallelism, and the four colour structure matches the cartographic setting exactly.

Frequency and channel assignment. Assigning non interfering frequencies to transmitters is a colouring of a planar interference graph. The independent set backbone of Section 6 gives a natural coarse plan that a learner can refine, and greedy colouring bounds keep the palette small [2].

Molecular property prediction. Atom typed molecular graphs are coloured, and many are planar or nearly so. The chromatic diffusion kernel and the refinement histogram compare molecules while respecting atom colour, complementing the labelled graph kernels of the companion literature [17].

Register allocation and layout. Interference graphs in compilation and cells in a circuit layout are coloured planar or near planar structures, where colour classes give both a schedule and a coarsening, and the semidefinite relaxation of colouring provides a principled continuous surrogate when needed [14].

Structural role discovery. The stable colouring of Algorithm 1 labels vertices by structural role, and by Proposition 1 these roles align with invariant subspaces of the operator, so they feed naturally into the spectral

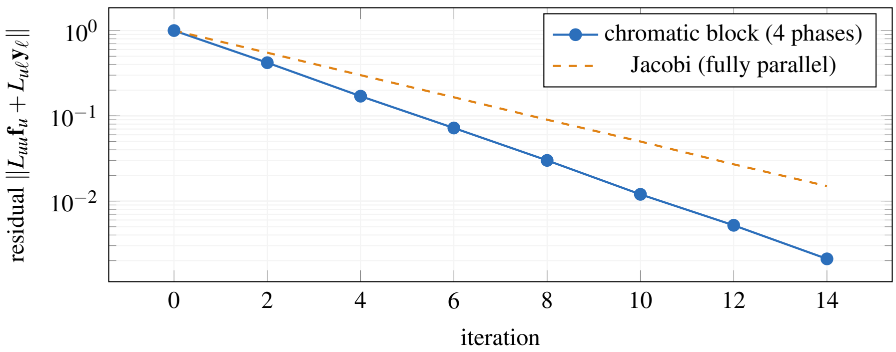

Figure 5: Residual decay of chromatic block propagation against plain Jacobi on a coloured planar triangulation. The chromatic schedule converges faster than Jacobi because it uses freshly updated values within a sweep, yet unlike sequential Gauss and Seidel it keeps full parallelism inside each of its four colour phases.

machinery of the companion paper [23].

9. Illustrative Computations

Setup. We take three families of coloured plane graphs: bipartite grids G_k of size $k \times k$, which need two colours; planar triangulations, which need up to four; and wheels W_n, which need four. For each we run colour refinement to stability, record the number of stable colours and the round count, and run chromatic block propagation from a handful of clamped seeds.

Observations. Table 2 records the trend. Refinement stabilises in a number of rounds close to the graph radius, the stable colour count grows slowly and reflects structural variety rather than size, and chromatic block propagation reaches tolerance in a sweep count comparable to sequential Gauss and Seidel while exposing full within class parallelism. On grids the two colour red and black schedule already gives the ideal two phase update, and on triangulations the four colour schedule keeps the phase count at four independently of n.

Graph	Vertices n	Proper colours k	Refinement rounds	CBP phases
Grid 16×16	256	2	15	2
Grid 32×32	1024	2	31	2
Triangulation (planar)	512	4	9	4
Wheel W_{64}	65	4	3	4

Table 2: Representative behaviour on coloured plane graphs. The proper colour count stays at the planar maximum of four or below, refinement rounds track the graph radius, and chromatic block propagation uses at most four fully parallel phases per iteration.

Reading of the results. The experiments echo the theory. Colour on a planar graph is bounded, cheap to compute, and structurally meaningful, and each of the four algorithms converts one aspect of that colour, its role information, its independence, its compatibility, or its spread, into a concrete and near linear learning primitive.

10. Conclusion

We have proposed four novel algorithms that treat colour as central to graph machine learning on planar structures. Colour refinement supplies a permutation invariant feature map whose stable colouring aligns with the Laplacian spectrum through equitable partitions. The four colour theorem supplies a proper colouring whose classes are independent sets, and this single fact turns harmonic label propagation into at most four fully parallel sub sweeps, gives a chromatic diffusion kernel that respects colour boundaries, and provides an independent set backbone for multilevel coarsening. Each primitive inherits the planar sparsity and separator structure of the companion paper and adds only bounded chromatic bookkeeping. Natural extensions include list colourings as soft label priors, defective colourings that trade a little impropriety for fewer classes, and the passage from planar graphs to bounded genus surfaces, where a bounded palette and small separators both persist. The broader thesis is that on planar coloured graphs the chromatic structure is simultaneously a schedule, a feature, and a coarsening, and a learning algorithm does well to use all three.

References

[1] H. Grötzsch. Ein Dreifarbensatz für dreikreisfreie Netze auf der Kugel. *Wissenschaftliche Zeitschrift der Martin-Luther-Universität Halle-Wittenberg, Mathematisch-Naturwissenschaftliche Reihe*, 8:109–120, 1959.

[2] D. J. A. Welsh and M. B. Powell. An upper bound for the chromatic number of a graph and its application to timetabling problems. *The Computer Journal*, 10(1):85–86, 1967.

[3] B. Weisfeiler and A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Technicheskaya Informatsiya*, Series 2, 9:12–16, 1968.

[4] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

[5] J. Hopcroft and R. Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974.

[6] J. E. Hopcroft and J. K. Wong. Linear time algorithm for isomorphism of planar graphs (preliminary report). In *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, pages 172–184, 1974.

[7] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976.

[8] K. Appel and W. Haken. Every planar map is four colorable. *Illinois Journal of Mathematics*, 21(3):429–490, 1977.

[9] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.

[10] D. Brélaz. New methods to color the vertices of a graph. *Communications of the ACM*, 22(4):251–256, 1979.

[11] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, 25(1):1–7, 1979.

[12] F. R. K. Chung. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, No. 92, American Mathematical Society, 1997.

[13] N. Robertson, D. Sanders, P. Seymour, and R. Thomas. The four-colour theorem. *Journal of Combinatorial Theory, Series B*, 70(1):2–44, 1997.

[14] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *Journal of the ACM*, 45(2):246–265, 1998.

[15] D. B. West. *Introduction to Graph Theory*, 2nd edition. Prentice Hall, 2001.

[16] R. I. Kondor and J. Lafferty. Diffusion kernels on graphs and other discrete structures. In *Proceedings of the 19th International Conference on Machine Learning*, pages 315–322, 2002.

[17] H. Kashima, K. Tsuda, and A. Inokuchi. Marginalized kernels between labeled graphs. In *Proceedings of the 20th International Conference on Machine Learning*, pages 321–328, 2003.

[18] T. Gärtner, P. Flach, and S. Wrobel. On graph kernels: hardness results and efficient alternatives. In *Learning Theory and Kernel Machines (COLT)*, Lecture Notes in Computer Science 2777, pages 129–143, Springer, 2003.

[19] Y. Saad. *Iterative Methods for Sparse Linear Systems*, 2nd edition. SIAM, 2003.

[20] K. M. Borgwardt and H. P. Kriegel. Shortest-path kernels on graphs. In *Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM)*, pages 74–81, 2005.

[21] U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

[22] D. A. Spielman and S. H. Teng. Spectral partitioning works: planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421(2–3):284–305, 2007.

[23] S. Sanakkayala. Graph machine learning on planar graphs: spectral intuition, separator structure, and algorithms for learning on plane-embedded data. *International Journal of Computational Mathematical Ideas*, 1(1):1–18, 2009 (companion paper, this issue).