Book review — L. Lovász, *Large Networks and Graph Limits*

Martin Bálek* Andrew Goodall *

1 Introduction

The book *Large Networks and Graph Limits*, xiv + 475 pp., published in late 2012, comprises five parts, the first an illuminating introduction and the last a tantalizing taste of how the scope of the theory developed in its pages might be extended to other combinatorial structures than graphs. The three central parts treat in depth the topics of graph algebras, limits for sequences of dense graphs (this constitutes the most substantial part, occupying nearly half the book) and limits for sequences of bounded degree graphs. Primarily the book is aimed at graduate students and research mathematicians interested in graph theory and its application to networks (for example, the internet and networks in social science, biology, statistical physics and engineering).

There are 23 chapters and an appendix, the latter conveniently giving necessary background from areas of mathematics outside mainstream graph theory. A bibliography collects together the extensive research in this area up to 2012, and a subject, author and notation index facilitate navigation of the book. The author maintains a webpage[1] for corrections and supplementary material. Indeed, via the author's homepage the reader can freely access the many papers he has written with collaborators on the topic of graph homomorphisms and graph limits. The book synthesizes much of the material in these papers, with some revision in

*IUUK, Charles University, Prague, Czech Republic. Both authors supported by the Center of Excellence-Inst for Theor. Comp. Sci., Prague, P202/12/G061, and the second author by Project ERCCZ LL1201 Cores. Email addresses: andrew@iuuk.mff.cuni.cz, balek@iuuk.mff.cuni.cz

[1] http://www.cs.elte.hu/~lovasz/book/homnotes.html

approach when needed to fit the larger trajectories of ideas that animate the book and make it cohere.

The author is generous in sharing his speculative thoughts as well as his insights: the book charts extensive conceptual territory in a way that simultaneously impresses by its depth and scope and encourages by its invitation to the reader to explore the area too. It is certain that there are still far-reaching discoveries to be made: many will be made by the author and his collaborators, but the territory is open to all those with a firm background in graph theory, linear algebra, probability or analysis. This active area of research – with such importance for applications to large networks – has in this book a timely survey. By following the author the reader is brought right up to the edge of what is currently terra incognita: only, rather than declaring that *hic sunt leones*, the instruments for charting the unmapped regions are put in the reader's hand.

In being brought to the forefront of current research it is naturally the case that the more rounded chapters and structured portions of book are supplemented by those less easily incorporated into an overall picture. For example, in Chapter 14 on the space of graphons, results are reported, but their full import is as yet to be discovered. Sprinkled throughout the book are such phrases as "this is not known" or "this only has a partial answer". However, one must also say that a huge amount *has* been answered – much by the author – and receives a compelling exposition in the shape of this book.

All but a few of the chapters contain many exercises, which appear at the end of each section and help the reader to a deeper understanding of the discussed notions and their properties. Doing the exercises forms an essential part of engaging with the book, for they are complementary to the content of the exposition, although sometimes they seem quite challenging (and only occasionally is one alerted to which are the difficult ones). There is however a plentiful supply of exercises that are not too difficult, or at least not beyond reach given mastery of the chapter material. The book does not contain solutions to these exercises. The compromise of leaving many statements as exercises or without detailed proof is an effective way to keep the book readable, permitting concentration on key ideas while still giving the reader the opportunity to explore the context more widely and deeply. In such a new and rapidly growing field, there is encouragingly large scope for the reader to follow paths as yet untrodden and to discover new results for him/herself too.

Lovász is always punctilious in ascribing credit to others; in this review we refer to many results due not only to the author but to other researchers in the field (many of whom he works with) and which receive an exposition in the book – in the pages of this review we leave out all these citations for reasons of space.

2 Chapter by chapter

2.1 Part 1. Large graphs: an informal introduction

In the first part of the book the author gives to the reader within just over thirty pages an intuitive feel for the principal topics of the book. This is far from just being a list of chapters with brief indications of their content. The author gives a delightfully insightful and broad view of the territory to be explored, in a way that is accessible to the reader who may have an as yet insufficient technical background to understand the finely developed theory that forms the remainder of the book. One purpose of the book is precisely to equip such a reader with the necessary tools to develop such an understanding. Having said that, even experts in the area will appreciate the conspectus Lovász provides of their field of research. In his introduction a clear idea is conveyed of the motivations and key questions that underlie the theory to be elaborated in the core of the book.

Chapter 1 *Very large networks* considers the key questions about large networks that have had such a large impact on the recent evolution of graph theory: how to obtain information about large networks, how networks can be modelled, how given networks can be approximated, and how to run algorithms on them. While so doing the author introduces such fundamental concepts as sampling, partitioning, left and right homomorphisms, random and quasirandom graphs, graph distances, regularity lemmas, graph limits, parameter estimation and property testing. The chapter finishes with a briefer treatment of the case of bounded degree graphs, for the good reason that “the technicalities in the bounded degree case are deeper, and so it is even more difficult to state key results, even informally.” Nevertheless, the author identifies some of the issues that make this area less well understood, and in Part 4, with the help of the appropriate technical apparatus, the current state of the art is presented.

Chapter 2 *Large graphs in mathematics and physics* considers another major motivation for graph limit theory, that of extremal graphs. The theory expounded in Chapter 16 is closely related to Razborov’s flag algebra technique, that has seen recent success in proving long inaccessible results in extremal graph theory. In this introductory chapter the author gives a brief history of the classical theorems on triangle density in graphs, which serves as an exemplar for the class of problems typical in this area. One of the virtues of graph limit theory is that it not only allows quantities vanishing in the limit to actually be let vanish – giving a computational algebra of subgraph densities that allows one to see the wood for the trees – but that it allows one to state in a precise way general questions about the nature of extremal graphs. Examples of such questions are given, such as "Is there always an extremal graph?" – this is vague as just posed here, but is given precision by the author in this introduction, and later in the book given a satisfying answer too. In the remaining section of the chapter there is a brief indication of the application of graph homomorphisms and limits to statistical physics, where the limiting behaviour of graph parameters is a fundamental object of study, although as the author ruefully states this topic is not discussed any further in the book. Both extremal graph theory and statistical physics are motivating forces for the topic of graph algebras, to which the second part of the book is devoted.

2.2 Part 2. The algebra of graph homomorphisms

Necessary preliminaries are given in Chapter 3 *Notation and terminology* (it is a not inconsiderable virtue of this book that the difficulties of unifying notation stemming from disparate fields of mathematics are overcome, and further that notation is kept informal whenever possible to do so). From this chapter we highlight the key notion of a partially labelled graph: a *k-labelled graph* is a graph to which k distinct labels have been assigned to some of its vertices (there is a distinction to be made between the case when a vertex can receive at most one label, and the case when multiple labels are allowed, but this is not important here).

The second part begins in earnest with Chapter 4 *Graph parameters and connection matrices.* Here we meet the notion of a *graph parameter* (a function defined on isomorphism types of multigraphs with loops; a graph property is identified with the graph parameter that is its indicator function) and the fundamental operation of the *gluing product* of k-labelled graphs (disjoint union followed by identification of like-labelled vertices). Graph parameters such as the stability number and matching number have the property that they are additive over disjoint unions; those such as the maximum clique size or chromatic number have the property that they are maxing over disjoint unions; while those such as the number of perfect matchings or number of proper colourings are multiplicative over disjoint unions. It turns out that graph parameters that are multiplicative will be fundamental objects of study (because homomorphism densities have this property). The *connection rank* of a graph parameter is introduced in order to study its properties more finely. This is a function of nonnegative integer k and equal to the rank of the *connection matrix* $M(f, k)$ of the graph parameter f, whose rows and columns are indexed by isomorphism types of k-labelled graphs and entries are the value of the parameter f on the gluing product of the graphs corresponding to the given row and column. The connection rank may of course be infinite, but is often finite: graph parameters of finite connection rank for all k are of particular interest, not least because, as is shown in Chapter 6, they can be computed in polynomial time on graphs of bounded treewidth. As well as the connection rank, another fundamental property of the graph parameter f is that of *reflection positivity*, which the parameter possesses if the matrix $M(f, k)$ is positive semidefinite for all nonnegative integers k. (A graph parameter is multiplicative if and only if $M(f, 0)$ is positive semidefinite and has rank ≤ 1.) The chapter proceeds to give a generous selection of examples of graph parameters and their connection rank. Examples with finite connection rank include the stability number (but not the clique number or chromatic number), number of stable sets and Hamiltonian cycles, the chromatic polynomial and Tutte polynomial (but not the number of Eulerian orientations). The chapter concludes with proofs that any graph property that is minor-closed or definable by a monadic second order formula has finite connection rank.

Chapter 5 *Graph homomorphisms* collects together a diverse set of results about graph homomorphisms (adjacency-preserving mappings between graphs) and its primary purpose is to give the reader a good feel for a key object of study in the book, namely that of a *homomorphism number*. For homomorphisms between simple graphs the homomorphism number is straightforwardly defined as the number of homomorphisms. The homomorphism number is defined more generally for multigraphs, vertex- and edge-weighted graphs (in this case homomorphism numbers are related to partition functions in statistical physics), signed graphs, and partially labelled graphs. It is also sometimes useful to make a restriction to injective homomorphisms, surjective homomorphisms, full homomorphisms (preserving non-adjacency as well as adjacency) or some combination of these (in particular, induced homomorphisms, which are injective full homomorphisms – counting these type of homomorphisms from a graph F into a graph G is to count the number of induced copies of F in G). Useful properties of these homomorphism numbers and the relationship between homomorphism numbers for the various restrictions such as injectivity are recorded. There then follows a discussion on normalizing the homomorphism number to give the *homomorphism density* of a graph F in a graph G, defined in such a way as to give sampling probabilities of seeing F in a large graph G. For dense graphs G, for example, the number of homomorphisms between F and G is divided by the number of mappings from F to G, to give a homomorphism density that is the probability that a mapping from F to G chosen uniformly at random is a homomorphism. There are variants for injective homomorphisms etc. and, for the density to be meaningful, a different normalization is required when considering bounded degree graphs G.

On fixing one of the graphs F and G, homomorphism numbers from F to G include many important graph parameters. For example the number of homomorphisms from F to K_q is the number of proper q-colourings, and the number of homomorphisms from cycles to G determine the spectrum of G. Taken collectively, these graph parameters defined by homomorphism profiles determine a graph G up isomorphism (this holds both for left homomorphisms, where we consider homomorphisms to G, and for right homomorphisms, from G). This is a result the author proved over forty years ago and which was famously used to give a short proof of cancellation for the categorical product. The author reviews this and other applications of homomorphism numbers to problems in graph isomorphism. This is followed by a section in which it is shown that the homomorphism numbers from F to a fixed graph G are linearly independent.

We then move on to several recent characterizations of graph parameters which can be expressed as a homomorphism number from (or to) some finite graph. For example, a graph parameter defined on multigraphs without loops can be expressed as a homomorphism number *to* a finite graph if it is reflection positive and has exponentially bounded rank (and takes value 1 on the empty graph). Thus for instance the number of spanning trees of a graph cannot be expressed as a homomorphism number, for its kth connection rank is the kth Bell number and this is not exponentially bounded. For homomorphisms *from* a fixed graph there is a dually defined connection matrix for which the corresponding conditions must obtain – instead of gluing-products of labelled graphs one has coloured-products of coloured graphs. The proofs are postponed to next chapter where the necessary algebraic tools are developed. The chapter concludes with a brief exposition of two structures, one graphical and the other topological, which can be imposed on the set of homomorphisms from one graph to another (the former significant in the context of statistical physics, the latter in relating topological connectivity to the chromatic number).

Chapter 6 *Graph algebras and homomorphism functions* introduces semigroup algebras on multigraphs and more generally k-labelled multigraphs (labelled multigraphs with gluing product form a commutative semigroup). The elements of these algebras are formal $\mathbb{R}$-linear combinations of multigraphs and are given the name of *quantum graphs*. By linearity, the definitions of graph parameter, graph homomorphisms and homomorphism densities extend to quantum graphs. The tools and results from linear algebra now become available, and will be used in order to derive the characterization of graph parameters expressible as homomorphism numbers stated in the previous chapter. But this is to make sound utilitarian what is in fact a very elegant transposition from graph parameters to algebras, in which properties of graph parameters are reflected closely in the structure of the algebras formed from them – the techniques are novel and deserve the reader's special attention.

Any graph parameter defines an inner product on the vector space of k-labelled quantum graphs to be equal to the evaluation of the graph parameter on the gluing product. This inner product is positive semidefinite if the same is true of the parameter, it satisfies the Frobenius identity (by its definition via the associative gluing product), and its Gram matrix is the connection matrix of the graph parameter used to define it. The dimension of the algebra of quantum graphs factored out by the kernel of the inner product defined by the graph parameter is equal to the connection rank of the parameter. If the graph parameter is reflection positive then the inner product on the quotient space is positive definite, making this an inner product space. If further the graph parameter has finite connection rank then the quotient space is a finite-dimensional commutative Frobenius algebra (structures for which there is a rich theory of duality). In this case the quotient algebra has the simple structure of $\mathbb{R}^m$ with coordinatewise multiplication and the usual inner product (m is the connection rank). The algebra elements corresponding to the standard basis form an idempotent basis for the quotient algebra defined by the graph parameter. As the chapter unfolds the combinatorial significance of these idempotents is revealed.

Another algebra structure can be introduced on the vector space of 2-labelled multigraphs by *concatenation* (gluing a vertex of the first graph labeled by 2 with a vertex of the second graph labeled by 1 and unlabelling this identified vertex). This yields an associative non-commutative algebra with a *-operation (conjugation). The generalization of this construction to more labels has been used in topological quantum field theory but the details of this are not discussed in the book.

After this setting up of the graph algebra apparatus comes the promised construction of a weighted graph H for which the value of a reflection positive graph parameter $f(G)$ of exponentially bounded connection rank equals the number of homomorphisms from G to H. We then move to *contractors* and *connectors* for 2-labelled quantum graphs, which among other things are used in order to give a more explicit description of the special idempotent basis for the graph parameter's quotient algebra. A contractor for a graph parameter f is a quantum 2-labelled graph that upon gluing replicates modulo f the operation of identifying two labelled vertices (other than the 2-labelled graph K_1 with both labels on its single vertex). The deletion-contraction identity for the Tutte polynomial translates to the fact that the Tutte polynomial has as a contractor a linear combination of K_2 and $\overline{K_2}$ (both 2-labelled, one label for each vertex). Connectors are quantum 2-labelled graphs in which the labelled vertices are non-adjacent and which upon gluing replicate modulo f the operation of joining the labelled vertices by an edge. Connectors allow multiple edges to be eliminated, and in this case for each multigraph there is a a simple graph for which the graph parameter takes the same value. Graph parameters f given by a homomorphism number to a weighted graph are shown to have a simple connector and a simple contractor. In fact possession of a contractor substitutes for finiteness of connection rank in another characterization of graph parameters given by homomorphism numbers to a weighted graph H: the requisite conditions are that the graph parameter be reflection positive, multiplicative and have a contractor.

Next we return to the idempotent basis used in the construction of weighted graph H so that $f(G)$ is the number of homomorphisms from G to H for reflection positive graph parameter with exponentially bounded connection rank. Two questions are addressed: what size this idempotent basis has and how large the constituent graphs need to be in the quantum graphs that represent the basis elements (in the quotient algebra defined by the graph parameter). A useful application of graph algebras follows, namely a proof of the theorem that a graph parameter of finite kth connection rank can be computed in polynomial time for graphs of treewidth at most k. The chapter concludes with description of a method of proving representation theorems for graph parameters, with an application being a proof of a characterization of graph parameters defined on looped multigraphs that are given by homomorphism numbers to an edge-weighted graph.

2.3 Part 3. Limits of dense graph sequences

In the third part of the book we move into a different discourse, namely that of analysis, with the aim being to define limit objects for graph sequences in the dense case, which are called *graphons*. There are similarities here between the completion of $\mathbb{Q}$ to $\mathbb{R}$ and the completion of finite weighted graphs to the space of graphons, with like advantages and insights from working with limiting values rather than the sequences themselves. Before realizing graphons as limits of dense graph sequences they are studied as objects in their own right. A cut norm and cut-distance is defined for graphons, regularity lemmas formulated and proved, and basic properties of sampling from them are established. Just as the rich theory of linear algebra was mined in the second part, so in this third part the resources of the well-developed theory of functional analysis are tapped.

Chapter 7 *Kernels and graphons* introduces these limiting objects for dense graph sequences. Kernels are bounded symmetric measurable functions from the unit square to the reals (the name alludes to the fact that these functions give rise to kernel operators on function spaces on the unit interval) and graphons are those with values in the unit interval only. Kernels are generalization of weighted graphs in the following way. For a vertex- and edge-weighted graph G on n vertices, take a partition of the interval $[0,1]$ into n intervals with lengths equal to the relative weight of the corresponding vertices. The value of the kernel is then set constant equal to the appropriate edge weight on the product of intervals corresponding to endvertices of that edge. In the case of unweighted simple graphs this gives a $\{0,1\}$-valued graphon (which can be seen as a representation of the adjacency matrix as a "pixel picture").

With kernels (and graphons as their normalized version) defined we can transform many basic constructions known for (weighted) graphs into the language of kernels. For example, the *degree function* for a given element $x \in [0,1]$ is defined as the integral of the kernel over $[0,1]$ with one component set equal to x (a straightforward analogue of summing a row/column of an incidence matrix for graphs). A section of the chapter is then devoted to defining homomorphism densities for graphons, analogously to homomorphism densities for finite graphs in which integrals replace summations, and to describing some of their properties such as their reflection positivity and multiplicativity. (Homomorphism numbers between graphons on the other hand seem not to have any sensible way of being defined – it turns out that the parameters defined by homomorphism numbers for finite graphs extend most naturally to graphons by replacing counting by maximization.) An interesting example of a homomorphism density for graphons and a taster of the fruitfulness of the limiting theory to be developed is the graphon for which the density of a finite graph F in it is equal to the number of Eulerian orientations of F. Recall that there was no representation of this graph parameter by a homomorphism number to a finite weighted graph (it has infinite connection rank), although it is the case that for finite graphs of bounded degree the number of Eulerian orientations *can* be represented by a homomorphism number via $\{\pm 1\}$-flows modulo m for sufficiently large m.

Having constructed the limit objects of dense graph sequences (graphons) there arises the question as to when two graphons should be regarded as the same up to isomorphism. One straightforward answer is to declare two kernels to be *isomorphic up to a null set* if there is an invertible measure-preserving mapping under which the kernels are equal almost everywhere. Kernels isomorphic up to a null set have the property that the homomorphism density of any multigraph in either kernel is the same. This turns out to be the defining notion of isomorphism required, and this equivalence with respect to homomorphism densities is called *weak isomorphism* (cf. the Lovász vector or homomorphism profile of a finite graph which determines it up to isomorphism). In the next section sums and products of kernels are defined, beginning with a *direct sum* decomposition of kernels analogous to the decomposition of a graph into connected components and unique up to zero sets. A kernel is then *connected* if it is not isomorphic up to a null set to the direct sum of two kernels; connectivity is invariant under weak isomorphism. Other notions applicable to finite graphs can be easily carried over to kernels. For example, a graphon is bipartite if there is a partition of the unit interval into subsets A and B such that its value on points in $A \times B$ is zero almost everywhere; one can then characterize bipartite graphons as those for which the density of odd cycles is zero. Three product operations are then defined on kernels that will be used later (pointwise, operator, and tensor products) and to finish the chapter kernel operators are used to derive spectral properties of kernels, using the rich theory of Hilbert-Schmidt operators. This allows the expression of many homomorphism densities in kernels in terms of its eigenvalues (for example, the density of a k-cycle gives a sum of kth powers of eigenvalues entirely analogous to the finite case). Indeed, there is a general spectral formula for the homomorphism density of a multigraph in a kernel which is an analogue of the expression of a homomorphism number to a finite weighted graph as an edge-colouring model (the latter notion is briefly discussed in the appendix).

In Chapter 8 *The cut distance* we return to finite graphs with a view to defining a structural similarity measure between graphs that survives the passage to the limit in a dense graph sequence to give a distance measure on graphons. The latter is in fact a more tractable quantity to work with: this is a phenomenon typical of working with graphons, where all the extraneous detritus prevalent when arguing about finite graphs has been pruned away. The cut distance is defined at first for two graphs on the same vertex set of size n. Take a pair of subsets of the common vertex set and calculate the number of edges between them in either graph, and then take the difference of these two quantities and divide by n^2 (normalization is not by a number dependent on the size of the vertex subsets, which has the effect of biasing the measure away from small vertex subsets in favour of larger ones). To obtain the *cut distance* between the two graphs maximize this quantity over all possible pairs of vertex subsets. For two different vertex sets of the same cardinality, take the minimum of the cut distance over all possible labellings of the vertex sets by $1, 2, \ldots, n$ (this refractory measure is but an auxiliary stepping stone to the final definition for arbitrary pairs of graphs). For graphs G and G' on vertex sets of sizes n and n' by replacing each vertex of G by kn' twin vertices and each vertex of G' by kn twin vertices (two new vertices being adjacent if and only if their progenitors are), for any postive integer k, we obtain a pair of graphs on the same number knn' vertices. The cut distance between G and G' is defined to be the limit of the cut distances as k tends to infinity in this construction. The cut distance is in fact a pseudometric, but it will turn out that a zero cut distance will be the same as indistinguishable with respect to weak isomorphism of the limiting graphon of a dense sequence. There is also a finite construction of cut distance which uses *fractional overlays*, which, roughly speaking, for probability distributions with uniform marginals on the vertices of G and vertices of G' measures correlations between uv and $u'v'$ being an edge when (u, u') and (v, v') are chosen from the distribution: the cut distance is obtained for the distribution that minimizes this correlation.

The corresponding construction of cut distance for kernels proceeds by first defining the *cut norm* on the linear space of kernels as the supremum of absolute value of integrals of the kernel over all possible pairs of measurable subsets of $[0, 1]$. (This is a direct parallel to the matrix cut norm, which defines the cut distance for finite graphs.) Properties of the cut norm are then described, for example its equivalence with the operator norm, and that the supremum of its definition is in fact attained (i.e., it is a maximum). Likewise for the cut distance, defined from the cut norm: the infimum that defines it (in a similar way to fractional overlays of finite graphs) is in fact a minimum. The chapter concludes with some relationships between the cut norm and the L_1-norm that need to be recorded for the sequel.

Chapter 9 *Szemerédi partitions* surveys various forms of the Szemerédi Regularity Lemma and extends it to graphons, in a streamlined formulation already observed to be typical of the passage from convergent sequences of dense graphs to the limiting graphon. Beginning with the original formulation, a reformulation in the language of graph distances follows and after this it is generalized to kernels. We can find there both the weak and strong version of the regularity lemma. The most general form of regularity lemma is then derived from compactness of the metric space of graphons defined by the cut distance. The next section concerns the technically useful fact that the limiting cut distance between graphs on vertex sets of equal size defined by replacing each vertex by a number of twin vertices is bounded by a certain function of their actual cut distance (as defined by a min-max formula over all possible labellings of vertices), with the consequence that the property of Cauchy convergence of a sequence under either distance coincides. The final section of the chapter is devoted to a proof of the uniqueness of strong regularity partitions.

In Chapter 10 *Sampling* we come to techniques for the analysis of large dense graphs by sampling from a limiting graphon. For graphon W, a W-random graph on n vertices can be extracted by joining distinct vertices i and j with probability $W(x_i, x_j)$, where $(x_1, \ldots, x_n)$ is an ordered tuple of n points chosen uniformly at random from $[0, 1]$. When W is constant equal to p this is the classical Erdős–Rényi random graph $G_{n,p}$.

Another sampling method is to extract a weighted graph for given tuple $(x_1, \ldots, x_n)$ where the edge joining i and j is assigned weight $W(x_i, x_j)$. This sampling method is useful when studying parameters defined on edge-weighted graphs, or in order to discuss nearness in cut distance of samples to the graphon from which they have been sampled (a W-random graph can be generated using the edge-weights as probabilities). A graph parameter such as a homomorphism density changes in value slowly between graphs on the same vertex set that differ only in the edges incident with a single vertex. For such graph parameters there is concentration of its value on random induced subgraphs of a given large graph G around its expected value, and similar behaviour is observed for its value on a W-random graph. In fact, as is shown later in the chapter, the samples themselves are concentrated in the cut-distance, not only the values of graph parameters on them. Before this though it is shown that the cut distance between two large graphs on the same vertex set can be estimated by sampling; this extends to kernels, in the sense that the cut norm of a random sample from the kernel is close to the cut norm of the kernel. A corollary is that when sampling from a pair of graphons the cut distance between the samples is closely matched by the cut distance between the graphons. The broad message is that the metric space of graphons encapsulates information about random sampling from large dense graphs, and in a form that permits easier analysis of random samples and graph parameters defined on them. After this section we then have, with another rather technically involved proof, the result to which we have just alluded, that samples are close in cut distance to the original large graph or graphon. Homomorphism densities are equivalent to sample distributions, and a generalization of the classical Counting Lemma (the density of a fixed graph in a quasirandom graph is roughly what it should be) for them is proved: the density of a simple graph F in a given simple graph is a Lipschitz-continuous function under the cut distance metric, and the same is true of the density of F in graphons. This is followed by a converse Inverse Counting Lemma, which states that if two large dense graphs or graphons are close in the sense of homomorphism densities then they are close in cut distance. Finally, the Counting Lemma and Inverse Counting Lemma are used to tighten the elementary observation that weak isomorphism of graphons is equivalent to being indistinguishable by sampling (all homomorphism densities equal). By way of topological equivalence of sampling and cut distance on the space of graphons assured by the Counting Lemma and its converse, two graphons are weakly isomorphic if and only if their cut distance is zero (this statement easily extends to kernels more generally, and as the author points out can be proved more elementarily than via the heavyweight Counting Lemma and Inverse Counting Lemma). The chapter closes with other characterizations of weak isomorphism that enrich its intuitive meaning.

After the technically demanding exertions of the previous couple of chapters, in Chapter 11 *Convergence of dense graph sequences* a consolidating approach is adopted in order to bring together the advances thus far to bear on the topic of convergent graph sequences. Convergence is defined with reference to a sampling method. In the case of dense graphs subgraph sampling is used, i.e. take an induced subgraph on a k-element vertex set chosen uniformly at random. Equivalently and sometimes more conveniently, homomorphism densities or injective homomorphism densities can be used to define convergence rather than the injective full homomorphism densities that are induced subgraph densities. A sequence of graphs (G_n), $n \to \infty$ (where n is number of vertices), is then *(left-) convergent* if for each finite graph F the induced subgraph density (or homomorphism density, or subgraph density) of F in G_n converges. For example, given fixed $0 \leq p \leq 1$, any sequence of random graphs G_n taken from the Erdős–Rényi model $G_{n,p}$ is convergent with probability 1 (as a random induced subgraph on k vertices is very close in distribution to $G_{k,p}$ for almost all choices of G_n). The choice of cut distance as the natural metric to take on graphs (and graphons) is vindicated by the result that a sequence of simple graphs of increasing order is convergent if and only if it is a Cauchy sequence in the cut distance metric (and a similar result holds for sequences of graphons). So far the message has been that, for the seeker of limits of dense graph sequences, "graphons are what you are looking for". The author pauses to consider the alternatives, and during the course of the chapter works to a summary of "cryptomorphic" objects that all can play the role of a graph limit in an equivalent way to graphons. These include consistent and local random graph models, local random countable graph models, multiplicative normalized simple graph parameters nonnegative on signed graphs, and points in the completion of the space of finite graphs with the cut distance metric.

A random graph model is sequence of random variables indexed by positive integer k whose values are simple graphs on k vertices in which isomorphic graphs have the same probability. A random graph model is consistent if deleting a vertex from the kth random variable yields a random variable with the same distribution as that given by the model on

$k-1$ vertices; it is local if for any pair of disjoint subsets of vertices the subgraphs induced on these subsets are independent as random variables. It is proved that a convergent sequence of graphs (G_n) gives rise to a consistent and local random graph model by taking the limiting distribution of induced homomorphism densities on k-vertex graphs, and that conversely every consistent and local random graph model arises in this way. It then turns out that a consistent and local random graph model can be represented by sampling k vertices from a graphon.

The countable random graph model is a probability distribution on countable graphs on vertex set $\mathbb{N}^*$ (positive integers) invariant under permutations of $\mathbb{N}^*$. The correspondence to consistent random graph models is uncomplicated and preserves the notion of locality: in the countable model, locality means that the induced subgraphs on disjoint finite subsets of positive integers are independent as random variables. The Rado graph in which two positive integers are joined with probability $1/2$ is a classical example of a local countable random graph model, whose corresponding consistent and local random graph model is the Erdős–Rényi model $G_{n,\frac{1}{2}}$. Given a graphon W, selecting a sequence of points $x_1, x_2, \ldots$ independently uniformly at random and joining i and j with probability $W(x_i, x_j)$ we obtain a local countable random graph model, and it turns out all such models can be constructed in this way.

Having strayed from graphons momentarily, we are now led back to consider them, and to see they are more explicit than the alternatives, in the sense that any convergent sequence of graphs has homomorphism densities convergent to a homomorphism density of a fixed graphon W (unique up to weak isomorphism). Conversely, it is shown that any measurable function W occurs as the limit object of some convergent sequence of simple graphs. Three different proofs are given, the first based on techniques presented in the book (compactness of the space of graphons, and the Counting Lemma). The second proof is based on model theory using ultraproducts and ultralimits, the relevant notions receiving a brief exposition here, supplemented by the two-page crash course in the appendix, while the details of the proof can be found in the recent paper that is cited. The third proof by exchangeable random variables uses representation theorems of Aldous and Hoover, and receives a very brief outline with references given to other sources for the details.

The chapter now moves on to the concrete problem of how to show a given sequence of graphs converges and determining its limit graphon. Disconcerting behaviour is possible here: a randomly growing sequence of graphs may be convergent with probability 1 and yet its limit itself be a random variable. But on the whole convergence is with probability 1 to a single well-defined limit graphon, and moreover is often $\{0,1\}$-valued. The latter property of a graphon, called *random-free*, has the consequence that convergence is not only in cut distance but also L_1 (edit distance), which means that while the graph sequence may grow randomly, there is high concentration, with two instances of the nth term of the random sequence differing in only $o(n^2)$ edges when overlaid properly.

Several well-known examples of convergent sequences are given to illustrate methods of establishing convergence and the limit graphon (there is by no means a unified approach to take here). These include simple threshold graphs (for a graph on n vertices, vertices i, j are connected if and only if $i + j \leq n$), (multitype) quasirandom graphs, growing uniform attachment graphs (in the n-th step we add a vertex and every nonadjacent vertices are connected with probability $1/n$; the corresponding limit function is $1 - \max(x, y)$), prefix attachment graphs (a newly added vertex is connected to all those created earlier than a given randomly chosen existing vertex). The growing preferential attachment graph, where a new vertex is joined to an old vertex with probability proportional to its current degree, with probability 1 converges to a constant function (i.e., is quasirandom), only the limit constant may differ on running the process again (the distribution of the limit is unknown).

Earlier, in Chapter 7, it was shown how graphons, considered as kernel operators, have a discrete spectrum such that any neighbourhood of 0 contains all but a finite number of eigenvalues. Eigenvalues occur with finite multiplicity and are contained in $[-1, 1]$. It is shown here that if a sequence of simple graphs is convergent to graphon W, then the (suitably ordered and normalized) eigenvalues of their adjacency matrices converge to the spectrum of W. In fact the stronger result is proved that a sequence of graphons convergent to W in the cut distance have spectra convergent to the spectrum of W.

The chapter concludes with two applications of graph limits to proving results about finite graphs: the Chung–Graham–Wilson characterization of quasirandom graphs by edge and C_4 densities alone, and the Removal Lemma (hinting at deeper connections between the Regularity Lemma and measure theory).

Having explored at length convergence of graph sequences (G_n) with respect to left homomorphism densities (from a fixed graph to G_n), the dual question of counting right homomorphisms (from G_n to a fixed graph) is the subject of Chapter 12 *Convergence from the right.* The problem considered is whether left convergence can be characterized in terms of right homomorphism densities, and not to define a new metric that might in a similar way to the cut distance give a condition for Cauchy convergence equivalent to convergence of right homomorphism densities (i.e., the structural similarity between graphs measured so well by the cut distance is taken as a paramount). That there are difficulties characterizing convergence of a graph sequence in the cut distance metric space (which is equivalent to convergence of left homomorphism densities) by means of right homomorphisms is already suggested by the fact that any sequence of graphs (G_n) such that the chromatic number of G_n tends to infinity will eventually have zero homomorphism density to any given fixed simple graph: here convergence of right homomorphism densities gives no information about the convergence or otherwise of the sequence (G_n). It turns out that there is more than one suitable definition of right convergence in terms of homomorphisms from G_n to a fixed small graph that can be adopted to surmount the problem that even taking weighted target graphs is not sufficient to characterize graph sequence convergence with respect to cut distance. Two routes are taken in this chapter: a modification of counting homomorphisms, or replacement of counting by maximization and introducing restricted multicuts.

For a fixed weighted graph H on q vertices the number of homomorphisms from G with n vertices to H grows exponentially with n^2; the *homomorphism entropy* is defined as the logarithm of the number of homomorphisms from G to H divided by n^2. The homomorphism entropy can be tightly bounded by the *maximum weighted multicut density* of G with respect to positively edge-weighted graph H. This is defined as the maximum over all possible partitions of the vertices of G into at most q parts of normalized sums of edge densities between pairs of subsets of the partition (cuts) weighted by the logarithm of the edge-weight of H. Using the fact that weighted multicut density is invariant under blowing up of vertices and the fact that the multicut density with respect to H is Lipschitz-continuous in the cut distance metric (a simpler variant of the Counting Lemma) a necessary condition for graph convergence emerges, namely that for every weighted graph with positive edge-weights, the sequence of homomorphism entropies is convergent. This condition is not sufficient, suggested already by the fact the any vertex weights of H are ignored, and confirmed by the example of interleaved sequences of quasirandom graphs of density p and $2p$, which has convergent homomorphism entropies but is not a left convergent sequence. This brings us to the notion of *typical homomorphisms*, which are those whose inverse images have size proportional to the vertex weights of the target graph H. Counting only with those we obtain the *typical homomorphism entropy* and likewise the *maximum restricted multicut density*.

The generalization of maximum weighted multicut density to graphons (using supremum in place of maximum and integration instead of summation) is called the *overlay functional*. This has many good properties, such as invariance under measure-preserving transformations of the two kernels that are its arguments, symmetry, and behaviour close to an inner product (although it is not, for instance, bilinear, only subadditive in each variable). Finally, on the way to proving a characterization of convergence in terms of right homomorphisms, the *quotient set* of a kernel W is defined as the set of all its quotient graphs, where the quotient graph of a kernel with respect to a partition of $[0, 1]$ into q measurable subsets $S_i, i \in \{1, \ldots, q\}$, is defined analogously to the template graph with respect to a partition of the vertices of a graph into q subsets. Specifically, the quotient graph is a weighted graph on $\{1, \ldots, q\}$ where vertex i has weight the measure of S_i and edge ij has weight the integral of W over $S_i \times S_j$ normalized by the measure of $S_i \times S_j$.

Convergence of a graphon sequence from the right is shown to be equivalent to (1) convergence of restricted multicut densities for all simple graphs, (2) convergence of overlay functional values for every kernel, or (3) that, for each q, the quotient sets of terms of the sequence by q-partitions form a Cauchy sequence in the Hausdorff metric defined via the cut distance. The remainder of the chapter then translates this characterization of convergent graphon sequences by behaviour of right homomorphisms to the setting of graph sequences, the complication lying in the fact that quantities such as multicuts associated with the graphon defined from a finite graph G are only approximations of the analogous combinatorial quantities for G. The chapter finishes with a question: what other convergence results are there for right-homomorphism parameters of a left convergent graph sequence? In the negative direction, convergence of the normalized logarithm of right homomorphism densities (from a term of the convergent sequence (G_n) to fixed graphon W) does not follow from left convergence of (G_n), as shown by an example given in the exercises.

A leitmotiv of part 3 is that graphons bring a leanness and clarity to reasoning about sequences of large graphs and convergence to a limit. In Chapter 13 *On the structure of graphons* this is taken a step further, with a generalization of kernels and graphons to arbitrary probability spaces (when the space is finite one obtains a weighted graph with normalized vertex weights). While this may seem redundant, given that graphons on $[0, 1]$ with the Lebesgue measure account for all limit objects of convergent graph sequences, a felicitous choice of underlying probability spaces can radically simplify finding the limits of such sequences as those of prefix attachment graphs and interval graphs. As far as subgraph densities and multicuts etc. are concerned the underlying probability space is a matter of indifference, permitting one to choose the most convenient at hand.

Atoms in a probability space are singleton subsets with positive measure; kernels can be made atom-free by a procedure analogous to the construction of a kernel from a weighted graph, replacing atoms by intervals of their measure. Two points in the probability space underlying a kernel are called twins if the marginal of the kernel with respect to one point is almost everywhere equal to the marginal with respect to the other point. Twins can be removed to leave a twin-free kernel that is related to the orginal kernel by a measure-preserving map, and moreover based on a standard probability space (completion of a Borel space) if the original kernel is. With these preparations, it is proved that two weakly isomorphic standard twin-free kernels are in fact isomorphic up to a null set. From this follows a clutch of characterizations of weak isomorphism of standard kernels extending those given in Chapter 10. Since null sets are usually happily ignored it would seem that with atom-free and twin-free kernels we have reached our ultimate destination. However, purifying kernels of their null sets proves useful in deriving further results about homomorphism densities. A *neighbourhood distance* is defined on the probability space underlying the kernel W by setting the distance between points x and y equal to the L_1-distance of the marginals of W along x and along y. A kernel is *pure* if with the neighbourhood distance it is based on a complete separable metric space and every open set has positive measure. Every twin-free kernel is isomorphic up to a null set to a pure kernel. After using pure kernels to prove some results about homomorphism densities, such as the fact that as a multigraph parameter the density in a kernel is contractible, the topology of the underlying space of a pure graphon is examined for correspondences with combinatorial properties of sequences that converge to it. In order to extract these correspondences a different metric than neighbourhood distance is introduced, namely the *similarity distance* which is defined like the neighbourhood distance except with the replacement of W by the operator product of W with itself (like a convolution, or taking the square of the adjacency matrix in the case of finite graphs). The similarity distance defines a Hausdorff space, but not necessarily a complete metric space. The differences in topologies defined by the neighbourhood and similarity distances appear to have combinatorial significance, in particular those graphons for which the finer space defined by neighbourhood distance is also compact. A compelling reason for introducing the similarity distance is the correspondence of the subsets in a weak regularity partition with the Voronoi cells of an average ϵ-net in the metric space defined by similarity distance. Further, the Minkowski dimension of this metric space has implications for the number of classes in a weak regularity partition of the graphon, and the Vapnik–Chervonenkis dimension has implications for densities of signed bipartite graphs.

The chapter closes with a brief look at the automorphism group of a graphon, which for a pure graphon when given the topology of pointwise convergence under the similarity distance is compact. After proving that a graphon must be a step-function when the connection matrix of homomorphism densities in W for 2-labelled graphs has finite rank, the corollary is deduced that a reflection positive, multiplicative and normalized simple graph parameter f either has infinite kth connection rank for k greater than 1 or there is a twin-free weighted graph H such that $f(G)$ is the normalized homomorphism number from G to H, in which case the kth connection rank is finite and its kth root tends to $|V(H)|$.

Chapter 14 *The space of graphons* consists of disparate technical results about graphons, difficult to summarize in any detail here. It begins by considering alternative norms than the cut norm, such as those defined by homomorphism densities. Then the interrelationships between different norms that have combinatorial interpretations are explored. For example, the relation between the cut norm and L_2 norm features in the proof of the Regularity Lemma and the relationship between the cut norm and the L_1-norm will be seen in the succeeding two chapters to be significant in property testing and the stability theory of extremal graphs. The chapter then moves on to consider closures of graphon properties and derives some characterizations of hereditary properties from it, and those that are random-free (the closure of the property is $\{0, 1\}$-valued; a graph property is random-free if it does not contain large quasirandom bipartite graphs).

The next section studies graphon varieties, defined by equations specifying linear dependence between subgraph densities (linear dependence is equivalent to algebraic dependence due to multiplicativity of subgraph densities), and gives many examples to illustrate the theory. This is followed by a section on random graphons, or more precisely probability distributions on equivalence classes of graphons under weak isomorphism, which form a compact metric space whose sigma-algebra of Borel subsets allows some headway in analysis. Random graphons give a representation of graph parameters more general than those represented by graphons. Random graphon models are shown to be cryptomorphic to countable random graph models, consistent random graph models and normalized simple graph parameters indifferent to isolated vertices and with nonnegative upper Möbius inverse. There follows a number of equivalent descriptions of isolate-indifferent normalized simple graph parameters that are reflection positive and that are multiplicative.

The final section of the chapter summarizes some recent work by other authors about *exponential random graph models* on the compact space of graphons. These give a way to understand the structure of a random graph conditioned on the value of a given graph parameter being small (for example, having triangle density much less than 1/8) by imposing a probability distribution weighted in favour of those graphs for which the value of the parameter is small, and then studying the limiting distribution on graphons.

The next two chapters treat two major areas of success for the application of the abstract theory developed hitherto.

Chapter 15 *Algorithms for large graphs and graphons* concerns parameter estimation, property testing and distinguishing between two properties, and computation of structures such as maximum cuts.

A parameter is *estimable* if it can be approximated with high probability by taking a sample of sufficiently large size (dependent only on the desired error bound). Estimability of a graph parameter f is equivalent to saying that f preserves left convergence of graph sequences (in particular all left homomorphism densities are estimable by definition). The opening result of the chapter is a useful set of criteria that a parameter must satisfy in order to be testable, which are derived using facts about the cut distance. This is followed by showing that the density of maximum cuts and the free energy of a states model (from statistical physics, and related to the right homomorphism density to an edge-weighted graph) are both estimable.

Before considering the more complicated task of testing for a single property the more tractable problem of distinguishing between two properties $\mathcal{P}_1$ and $\mathcal{P}_2$ by sampling is addressed. Here an auxiliary test property

$\mathcal{Q}$ is queried of the sample, and the answer determines which of the two properties to plump for: the properties $\mathcal{P}_1$ and $\mathcal{P}_2$ are said to be distinguishable by sampling if for every sufficiently large sample size k the probability of the test property $\mathcal{Q}$ being satisfied on a sample of k vertices given $\mathcal{P}_1$ is at least $2/3$ while that given $\mathcal{P}_2$ is at most $1/3$ (it is then seen to suffice to require that the former be strictly larger than the latter). Then it is shown that this is equivalent to either of two conditions involving the cut distance and the total variation distance between two graphs on sufficiently large number of vertices one of which has property $\mathcal{P}_1$ and the other $\mathcal{P}_2$.

Next the widely studied problem of tesing for a single property $\mathcal{P}$ is considered in its aspect relating to graph limit theory. First testability of a graphon property $\mathcal{P}$ is defined in terms of a test property that for a graph of any sample size from a graphon satisfying $\mathcal{P}$ holds with probability at least $2/3$, and such that, for a given postive ϵ, a sample graph of sufficiently large size (dependent on ϵ) from a graphon that has edit distance from $\mathcal{P}$ at least ϵ does not have the test property with probability at least $2/3$. A test property for $\mathcal{P}$ is then constructed in terms of graphs of sufficiently small cut distance from $\mathcal{P}$ and a condition for testability of a property framed in terms of distinguishability from the property of having edit distance from $\mathcal{P}$ of at least a given ϵ. A number of corollaries are deduced giving equivalent formulations of testability that permit more examples of testable properties to be offered, such as subgraph density.

From graphon properties we then move to testing graph properties, with an analogous definition of testability as that for graphon properties (this notion of testability for finite graphs is called *oblivious testing* as it does not assume any information about the size of graphs involved). That being triangle-free is testable is equivalent to the Removal Lemma. Testable graphon properties are shown to be precisely the closures of testable graph properties, and a graph property testable if and only if it has a property called robustness and its closure is testable.

Next the theorem that every hereditary property is testable is proved using the machinery developed in the chapter. A characterization of testability relating it to estimability is given: $\mathcal{P}$ is testable if and only if the normalized edit distance from $\mathcal{P}$ is an estimable parameter.

Finally in this chapter the problem of computing structures in large graphs is treated, with the examples of representative sets (under a definition of similarity distance), regularity partitions, and maximum cuts.

Graph limits have seen their richest application in extremal graph theory, and this is taken up in Chapter 16 *Extremal theory of dense graphs*. The first section relates reflection positivity to inequalities between subgraph densities, and the second section develops techniques of variational calculus of graphons (which have the advantage of being able to be continuously deformed, as opposed to finite graphs). We are now ready to see applications to the problem of describing relationships between densities of complete graphs, which of course has a long history in combinatorics, and here is surveyed with formulations and proofs using the theory developed thus far.

Another advantage of graphons is exploited, namely that local optima can be defined and the tools of analysis used to study them. Three examples are given: that for every hereditary property a random graph of appropriate density is asymptotically furthest from it in edit distance (here the local optimum is global, giving a short proof), Sidorenko's Conjecture relating bipartite subgraph densities to edge densities, and "common graphs" (a prototypical example being triangles, for which the sum of their density in a given graph and the density in its complement is at least 1/4).

After this survey of historically important examples in extremal graph theory comes the general question of deciding inequalities between subgraph densities, beginning with the recent result that it is algorithmically undecidable to decide whether a quantum graph with rational coefficients is nonnegative, which is proved by a reduction to Matiyasevich's solution of Hilbert's Tenth Problem. By a nonnegative quantum graph is meant one whose density (extended linearly from graphs to quantum graphs) in any graphon is nonnegative. The similar result is true, and older and easier to prove, for homomorphism numbers into finite simple graphs (rather than graphons). Then the related question of a Positivstellensatz for graphs is raised: is there a nonnegative quantum graph that is not a sum of squares? Well, yes, one has been recently constructed, but a weaker result is true: a nonnegative quantum graph is arbitrarily close in edit distance to a sum of squares.

The rest of the chapter is devoted to the interesting question of describing the structure of extremal graphs for an extremal problem. For classical extremal graph theory this has been answered in the dense case by the Erdős–Simonovits–Stone theory of extremal graphs. In the asymptotic sense, the only extremal graphs are the Turán graphs. A more general type of extremal problem on simple graphs is considered: maximize the density of a given quantum graph in a graphon W subject to a finite set of constraints giving the densities of other quantum graphs in W. The question is whether (as for the classical extremal problem, where the constraints all stipulate density zero and the given quantum graph whose density is to be maximized is the single edge) there is a special family of graphons (like the graphons obtained from complete graphs in the classical case) such that every extremal graph problem has a solution from this family?

A class of graphons necessary to cater for these sorts of extremal problems is then described, that is conjectured to be sufficient too, namely the *finitely forcible graphons.* These are defined as those graphons for which knowing the densities of a finite number of simple graphs is sufficient to determine them up to weak isomorphism. The Chung–Graham–Wilson characterization of quasirandom graphs by edge and 4-cycle densities alone is equivalent to saying that every graphon given by a constant function is finitely forcible. It is shown that "most" graphons (in the Baire category sense) are not finitely forcible, but that there are interesting families of finitely forcible graphons. The author conjectures that every extremal problem of the type described has a finitely forcible optimum. Finite forcible graphons include step-functions (almost all classical extremal problems have a solution whose "template" is a step-function), threshold graphons (defined by an inequality on a real symmetric bivariate polynomial), and the fractal-like complement reducible graphons. We are then shown that there are not too many finitely forcible graphons, and after discussion of a more technical notion of infinitesimally finitely forcible graphons, the chapter concludes with some speculation about the nature of finitely forcible graphons.

The final Chapter 17 *Multigraphs and decorated graphs* of Part 3 adumbrates extensions of the theory developed in detail for sequences of simple graphs to multigraphs; it turns out that there is (modulo the differences in how homomorphisms between multigraphs may be defined) when edge multiplicities are uniformly bounded a not too unstraightforward transportation of results to not only multigraphs but *decorated graphs*, where edges are assigned elements from a compact Hausdorff space (for multigraphs one can work with the compactification of $\mathbb{N}$ by adjoining ∞). After an initial discussion of the variations one might consider, namely sequences of multigraphs with (un)bounded edge multiplicities, and densities of (multi)graphs in terms of this sequence, the case of bounded edge multiplicities is subsumed under the more general consideration of decorated graphs. These include simple graphs, (discrete space on two elements “edge” and “non-edge”), coloured graphs (finite discrete space), multigraphs of edge multiplicity at most m (discrete space $\{0, 1, \ldots .d\}$), and weighted graphs (closed bounded interval). The definition of homomorphism number is not so straightforward, and a definition is offered for homomorphisms between decorated graphs where the source graph is decorated with the space of continuous real-valued functions on the compact space decorating the target graph (thus a map between vertices gets a real weight according to the evaluation of the source edge on its target value). Defining convergence is unproblematic since samples define a distribution on a compact space, and the requisite notion of weak convergence in distribution suffices. A characterization of convergence of a graph sequence in terms of homomorphism numbers from the left analogous to the simple case is then proved.

Limits of graphs decorated by a compact Hausdorff space K are bivariate functions taking probability measures as values rather than values in K. After defining the appropriate limit objects (K*-graphons*), it is proved that a convergent sequence of K-decorated graphs has left homorphism densities of a graph decorated by continuous real-valued functions on K convergent to the corresponding density in a K-graphon. The section concludes with a proof deferred of the theorem from Chapter 5 that a multigraph parameter is equal to a homomorphism number to a random weighted graph if and only if it is multiplicative, reflection positive and has finite second connection rank.

The final section sketches the problems surrounding sequences of multigraphs of unbounded edge multplicity, with references to as yet unpublished work for further reading.

2.4 Part 4. Limits of bounded degree graphs

This part is devoted to developing a theory of graph limits for graphs of bounded degree (all degrees are bounded by a fixed constant D). The limit objects that will play a role analogous to that of graphons are infinite graphs generalizing finite bounded degree graphs, namely *Borel graphs*, and these are the subject of Chapter 18 *Graphings*. There is a close connection to the construction of limit objects for bounded degree sequences by Benjamini and Schramm, and there will be some appeal to their construction, in particular to extend the notion of weak isomorphism to graphings.

We are given a Borel sigma-algebra $(\Omega, \mathcal{B})$. A graph on vertex set Ω is a *Borel graph* if its edge set is a Borel set in $\mathcal{B} \times \mathcal{B}$; the tacit assumption that degrees are bounded by D is maintained. Two basic examples on the unit interval and unit circle are given as illustrations: given a fixed $a \in (0, 1)$ we connect by an edge points at distance a. In the case of an interval we thus obtain a union of finite paths (for a greater than $1/2$ just a matching and isolated vertices). For a circle we obtain cycles if a is rational and two-way infinite paths if a is irrational.

The neighbourhood of a vertex in a Borel graph is obtained by projection onto a coordinate, and this is again a Borel set by a classical theorem of Lusin. It is shown that the converse holds: a graph on a Borel space is Borel if and only if vertex neighbourhoods are Borel. The theory of Borel graphs is already well developed, and the author marshals just those results that are needed in the sequel. For example, a proper colouring of the vertices of a Borel graph is a *Borel colouring* if each colour class is a Borel set: every Borel graph has a Borel colouring with $D + 1$ colours (extending Brooks' Theorem) and a Borel edge colouring with $2D - 1$ colours (a bit weaker than Shannon's Theorem for multigraphs that $3D/2$ colours suffices, and Vizing's Theorem for simple graphs that $D + 1$ suffices).

The notion of graphing as an analogue of graphon in the dense case is now introduced. We take a probability measure on $(\Omega, \mathcal{B})$ and call a graph with vertex set Ω a *graphing* if measuring (counting) edges between any two measurable sets of vertices is the same if we measure from the point of view of either of them. For example, the previously mentioned example on the unit circle with uniform probability measure is a graphing. There follows a section on how to verify that a distribution on a Borel graph makes it into a graphing, which includes introducing the notion of *measure-preserving family* of graphs, which can be thought of as a graphing whose edges have been coloured and oriented so that each colour defines an invertible measure-preserving map. (The connection of these families with finitely generated groups is indicated right at the end of the chapter, but this aspect, which makes limits of bounded degree graphs of such interest to group theorists, is beyond the scope of the book.)

There is no space in this review to describe all the technical machinery needed (concerning *involution-invariant random rooted graphs*) for the main goal of this chapter, which is, in a similar manner to Chapter 5, to introduce homomorphism densities and *local equivalence* (i.e., same homomorphism densities, equivalently, same r-neighbourhood distributions) for graphings.

Chapter 19 *Convergence of bounded degree graphs* continues the journey towards defining a limit process for bounded degree graphs. For this a distance has to be introduced, but unfortunately there is no analogue to cut distance and we are thus left with a *sampling distance* construction, which is locally defined as opposed to the globally defined cut distance, and so is less powerful an instrument for comparing graphs. The sampling distance leads to the notion of *local convergence* of graph sequences (r-neighbourhood densities converge for every r and given r-ball, equivalently, convergence of homomorphism densities for every connected graph) and a description of their limits (*involution-invariant distributions*), followed by an interesting selection of examples as illustrations, including the representation of the limit of Penrose tilings as a graphing. The description of the limit of a locally convergent sequence by means of graphings is not unique; the sigma-algebra of the limit object carries combinatorial information, which is in contrast to the graphon limits of convergent sequences of dense graphs. On the other hand, limits of locally convergent sequences *are* uniquely described by an involution-invariant distribution; however, under a stronger notion of convergence (*local–global*), graphings are precisely what are required to describe the limits. Local-global convergence uses the *nondeterministic sampling distance*, in which r-neighbourhoods of two graphs are compared in sampling distance over all possible pairs of k-colourings of the graphs. Local-global convergence enables combinatorial information to be passed from the the limiting graphing back down to the graphs in the sequence (for example, whether they are expanders).

Chapter 20 *Right convergence of bounded degree graphs* may be understood as a direct analogue of Chapter 12 (which treats dense graphs). A strong motivation comes from statistical physics, where for example the Ising model is most commonly studied on regular lattices, but again the bounded degree case brings many difficulties, which we do not have the space to elaborate on here. In particular, specific conditions need to be fulfilled in order to enable the construction of random homomorphisms. In case the source graph has small maximum degree and the edge weights of the destination graph are close to 1 this is ensured by the Dobrushin Uniqueness Theorem, which receives here a proof in its combinatorial version. Convergence from the right can be then characterized by the convergence of normalized homomorphism numbers to certain fixed graphs. The proof of this result is included in the second half of this chapter.

Chapter 21 *On the structure of graphings* contains a detailed description of a specific type of graphing, namely *hyperfinite graphings*. Hyperfiniteness of a graph means that we obtain bounded size connected components after deleting a small fraction of edges. Many families used in practice have this property (trees, grids, planar graphs, random graphs, expanders) and thus an understanding of the limiting behaviour in the hyperfiniteness case would be very welcome. This definition of hyperfiniteness generalizes naturally to graphings. Any two locally equivalent hyperfinite atom-free graphings are locally–globally equivalent, which reduces the complexity of the bounded degree case and the results are in such cases direct analogues to those of dense case.

It is shown that a sequence of graphs converges locally (not necessarily local-globally) to a hyperfinite graphing if and only if the family of graphs in the given sequence is hyperfinite. Moreover a graphing is hyperfinite only if it is a limit of a hyperfinite graph sequence.

After a short section concerning the question of how far a general bounded degree graph can be simplified by removing a small fraction of its edges, the chapter concludes with brief discussion about what a Regularity Lemma might look like for non-dense graphs.

Chapter 22 *Algorithms for bounded degree graphs* steps into territory that has as yet to submit to a unified theoretical approach. The absence of a Regularity Lemma and a metric as discerning as the cut distance are sorely felt. The chapter tries as far as it can to illustrate previous theoretical results in the context of algorithms, addressing as for dense graphs (see Chapter 15) the problems of parameter estimation, property distinction, property testing and computation of a structure. A selection of recent non-trivial results are presented on these topics. An algorithmic theory of bounded-degree graphs of corresponding power to that of dense graphs is still in the future.

2.5 Part 5. Extensions: a brief survey

The last part consists of one chapter and the appendix. Chapter 23 *Other combinatorial structures* briefly surveys some possible extensions of the theory developed in parts 2, 3 and 4. First the author discusses the limits of sparse (but not too sparse) graphs, and a Markov chain on a graphon is defined. But it is an open problem whether it can be used for construction of a limit object for suitable graph sequences (neither bounded degree nor dense).

Next come edge-connection matrices as analogues to connection matrices, where gluing is done not on vertices but on edges. An analogous theorem to the vertex-gluing version for homomorphism numbers characterizes those graph parameters that can be expressed as edge-colouring models as those that are multiplicative and have positive semidefinite edge-connection matrices. The connection between edge-colouring models and multilinear algebra (tensor networks) is explained in some detail.

Another generalization goes in the direction of hypergraphs, introducing homomorphism density, constructing a limit object *hypergraphon*, and concluding with the formulation of a Strong Hypergraph Regularity Lemma.

The categorical way of looking at mathematical structures is then adopted, and a categorical version of the Regularity Lemma formulated. Many questions arise concerning the interpretation of statements in categorical language in concrete instances.

Lastly, a short concluding section surveys other discrete structures for which convergence and limits might be defined, including directed graphs, posets, permutations, metric spaces, and functions on Abelian groups. While partial results have been obtained in these areas, and there are resonances with classical mathematics of the last century such as von Neumann's theory of continuous geometries and the number theoretical study of sequences of integers, it is at this point that a cloud of unknowing descends. There remains much to be discovered.

The *Appendix* contains definitions and technical details on the following topics: Möbius functions, the Tutte polynomial, basics of probability and measure theory, moments and moment problem, ultraproducts and ultralimits, Artin's theorem on nonnegative polynomials, and basics of category theory.

3 Concluding remarks

All in all this is a very interesting book at the frontier of research, giving a detailed exposition of a well developed theory (especially for sequences of dense graphs) and at the same time describing many open problems for further research. It is informative and the technical passages are relieved by an engaging informal written style and the regular interspersal of paragraphs signposting where we have been and where we are going. There are only a few inconsistencies in notation to watch out for and subjects omitted can be pursued by consulting the references the author gives. The author occasionally refers to supposedly well-known terms and techniques from other disciplines that would perhaps only be immediately understandable to an expert from that field, and may even be used in a sense that differs a little from its source due to nature of its application here. For instance, in Chapter 8 he defines ℓ_1 and ℓ_2 matrix norms which are normalized versions of the usual definitions (to fit in with the cut norm). Also, the term "graphing" has a meaning of some twenty years' vintage, but, given that this meaning seems labile, the author adjusts its meaning to suit its role as a counterpart to graphons. These shifts are explicitly remarked by the author when the terms are introduced, but perhaps could be overlooked by the nonlinear reader.

The book is highly recommended to general mathematicians and computer scientists, both students and professionals. Every reader will find inspiration from some of its pages.