

The automorphism group of a graphon

LÁSZLÓ LOVÁSZ* and BALÁZS SZEGEDY†

Dedicated to the memory of Ákos Seress

Contents

1	Introduction	2
2	Preliminaries	3
2.1	Graphs and graphons	3
2.2	Graph algebras	4
2.3	Metrics on graphons	5
2.4	Continuity of restricted homomorphism numbers	7
3	Compactness of the automorphism group	10
3.1	Automorphisms of graphons	10
3.2	Compactness	12
4	Spectra	14
4.1	Spectral decomposition	14
4.2	Spectral decomposition of pure graphons	16
4.3	Subdividing edges	18
4.4	Automorphism groups and spectral decomposition	20
5	Orbits of the automorphism group	20
5.1	Characterization of the orbits	20
5.2	Node-transitive graphons	23
5.3	Limits of node-transitive graphons	25
6	Graph algebras of finite rank graphons	27

Abstract

*Institute of Mathematics, Eötvös Loránd University, Budapest, Hungary. Research supported by ERC Advanced Research Grant No. 227701.

†Alfréd Rényi Mathematical Research Institute, Budapest, Hungary. Research supported by ERC Consolidator Grant No. 617747.

We study the automorphism group of graphons (graph limits). We prove that after an appropriate “standardization” of the graphon, the automorphism group is compact. Furthermore, we characterize the orbits of the automorphism group on k -tuples of points. Among applications we study the graph algebras defined by finite rank graphons and the space of node-transitive graphons.

1 Introduction

Graphons have been introduced as limit objects of convergent sequences of dense simple graphs, and many aspects of graphs can be extended to graphons. The goal of this paper is to describe a natural way to extend the notion of graph automorphisms to graphons. Our notion of automorphism group satisfies the natural requirement that it is invariant under weak isomorphism of graphons. (Weakly isomorphic graphons represent the limit objects of the same convergent graph sequences.) Thus our study of the automorphisms of graphons fits well into graph limit theory.

In this paper we heavily use the topological aspects of graph limit theory developed in [7]. It was shown in [7] that every graphon has two “canonical” representations on metric spaces, which we call, informally, the *neighborhood metric* and the *2-neighborhood metric*. These metric spaces depend only on the weak isomorphism class of the graphon. (In [4], these are called the “neighborhood metric” and the “similarity metric”.) The neighborhood metric space is simpler to define and work with, but it is not compact in general; the 2-neighborhood metric space is compact. The automorphism group acts on each of these as a subgroup of isometries. It is a rather straightforward consequence of the compactness of the 2-neighborhood metric that the automorphism group is always a compact topological group (Theorem 10). This fact is also closely related to (and could be derived from) a theorem of Vershik and Haböck [10] on the compactness of isometry groups of multivariate functions. As a consequence we prove that for node-transitive graphons the neighborhood metric is also compact.

The space of graphons (with weakly isomorphic graphons identified) is compact in a natural topology (defined by the “cut distance”). Another result of this paper is that the set of node-transitive graphons is closed, and hence compact, in this topology. As we will see, graph limit theory restricted to this closed set gives rise to a rather interesting limit theory for functions on groups. Such a theory was initiated in [9], and it was a crucial component of the limit approach to higher order Fourier analysis (see [8]).

We give a characterization of the orbits of the automorphism group on k -tuples of points. This generalizes results in [3] from finite graphs to graphons, as well as the characterization of weak isomorphism of graphons by Borgs, Chayes and Lovász [1]. We use this characterization to connect the topic of graph algebras with group theory. As an application, we give a group theoretic description of the graph algebras defined by finite rank graphons.

It follows from our results that the limit of a convergent sequence of finite graphs, each having a node-transitive automorphism group, is a node-transitive graphon. However, the relationship between the automorphism groups of the finite graphs and that

of the limit graphon is more involved.

2 Preliminaries

2.1 Graphs and graphons

A k -labeled graph is a graph (simple or multi) with k of its nodes labeled $1, \dots, k$ ($k \in \{0, 1, 2, \dots\}$). We denote by \mathcal{F}_k the set of k -labeled multigraphs, by \mathcal{G}_k the set of k -labeled simple graphs, and by \mathcal{G}_k^0 the set of k -labeled simple graphs with nonadjacent labeled nodes. In particular, \mathcal{G}_0 is the set of unlabeled simple graphs.

We will need some special k -labeled graphs and multigraphs. We denote by K_2 the (unlabeled) graph with two nodes and one edge, and by C_2 the multigraph consisting of two nodes connected by two edges. We denote by K_2^\bullet and $K_2^{\bullet\bullet}$ the graph K_2 with one and two nodes labeled, respectively; C_2^\bullet and $C_2^{\bullet\bullet}$ are defined analogously. We denote by $P_n^{\bullet\bullet}$ the path with n nodes, with its two endpoints labeled.

For two simple graphs F and G , let $\text{hom}(F, G)$ denote the number of homomorphisms (adjacency-preserving maps) $V(F) \rightarrow V(G)$. We define the *homomorphism density*

$$t(F, G) = \frac{\text{hom}(F, G)}{|V(G)|^{|V(F)|}}.$$

A *graphon* consists of a standard probability space J and a symmetric measurable function $W : J \times J \rightarrow [0, 1]$. To simplify notation, we will omit some letters that may be understood. For the standard probability space J , we let \mathcal{B} denote the underlying sigma-algebra and let π denote the probability measure. Also, we write dx instead of $d\pi(x)$ in integrals if there is only one probability measure considered.

Every graphon (J, W) defines an integral operator T_W on the Hilbert space $L^2(J)$ by

$$(T_W f)(x) = \int_J W(x, y) f(y) dy.$$

We say that W has finite rank if this operator has finite rank (i.e., its range is a finite dimensional subspace of $L^2(J)$).

For every graphon (J, W) and every graph $F = (V, E)$, we define

$$t(F, J, W) = \int_{J^V} \prod_{ij \in E} W(x_i, x_j) \prod_{i \in V} dx_i. \quad (1)$$

We note that the formula makes sense for multigraphs F , but we exclude loops. We write $t(F, W)$ instead of $t(F, J, W)$ if the underlying probability space is clear.

Graphons were introduced to describe limit objects of convergent sequences of dense graphs. A sequence of simple graphs G_n is called *convergent*, if the numerical sequence $t(F, G_n)$ converges for every simple graph F . In this case, there is a graphon (J, W) such that $t(F, G_n)$ converges to $t(F, W)$ for every simple graph F [5].

The limiting graphon is not strictly uniquely determined. Quite often one assumes that $J = [0, 1]$ (with the Lebesgue measure). In this paper, different underlying spaces

will be more useful. We say that two graphons (J, W) and (J', W') are *weakly isomorphic*, if $t(F, W) = t(F, W')$ for every simple graph F . Every graphon is weakly isomorphic to a graphon on $[0, 1]$, but this is not always the most convenient representative of a weak isomorphism class.

Weakly isomorphic pairs of graphons were characterized in [1]. Let J and L be standard probability spaces and let $\varphi : J \rightarrow L$ be a measure preserving map. For any function $U : L \times L \rightarrow \mathbb{R}$, we define the function $U^\varphi : J \times J \rightarrow \mathbb{R}$ by

$$U^\varphi(x, y) = U(\varphi(x), \varphi(y)).$$

It is clear that if (L, U) is a graphon, then so is (J, U^φ) , which we call the *pullback* of (L, U) along φ . It is easy to see that the graphons (L, U) and (J, U^φ) are weakly isomorphic. It follows that all pullbacks of the same graphon are weakly isomorphic. The main result of [1] asserts that two graphons are weakly isomorphic if and only if they are pullbacks of the same graphon.

We can define a sequence of graphons (W_1, W_2, \dots) to be *convergent* with limit graphon W if $t(F, W_n)$ converges to $t(F, W)$ for every simple graph F . (There is a semimetric, called the “cut distance”, on the set of graphons that makes this space compact, and which defines this same notion of convergence. We don’t need the cut distance in this paper, however.)

We can define homomorphism densities of k -labeled graphs in graphons, but these will be k -variable functions $J^k \rightarrow \mathbb{R}$ rather than numbers. These *restricted homomorphism densities* are defined by not integrating the variables corresponding to labeled nodes:

$$t_{x_1, \dots, x_k}(F, W) = \int_{J^{V \setminus [k]}} \prod_{ij \in E} W(x_i, x_j) \prod_{i \in V \setminus [k]} d\pi(x_i). \quad (2)$$

2.2 Graph algebras

Graph algebras are important algebraic structures associated with graph parameters. We give a quick introduction to the subject. For more details see [4].

For two simple graphs $G, H \in \mathcal{G}_k$, the product GH is defined as the graph obtained from G and H by identifying vertices with the same label and by reducing multiple edges. This product defines a commutative semigroup structure on \mathcal{G}_k . Let \mathcal{Q}_k denote the set of formal \mathbb{R} -linear combinations of elements from \mathcal{G}_k . Such linear combinations are usually called *quantum graphs*. The multiplication extends to \mathcal{Q}_k from \mathcal{G}_k using the distributive law and thus \mathcal{Q}_k becomes a commutative algebra. In other words, \mathcal{Q}_k is the semigroup algebra of \mathcal{G}_k .

Let $\llbracket G \rrbracket$ be the graph obtained by removing the labels in the graph G . We extend this notation to quantum graphs by linearity. An arbitrary graph parameter $f : \mathcal{G} \rightarrow \mathbb{R}$ can be extended to quantum graphs by linearity. Similarly, restricted homomorphism densities can be extended to k -labeled quantum graphs by linearity: if $f = \sum_{i=1}^n a_i F_i$, then

$$t_{x_1, \dots, x_k}(f, W) = \sum_{i=1}^n a_i t_{x_1, \dots, x_k}(F_i, W).$$

Every graph parameter gives rise to a symmetric bilinear form on \mathcal{Q}_k defined by $\langle G, H \rangle = f(\llbracket GH \rrbracket)$. Let \mathcal{I}_k be the set of elements Q in \mathcal{Q}_k such that $\langle Q, P \rangle_f = 0$ for every $P \in \mathcal{Q}_k$. Then \mathcal{I}_k is an ideal and $\mathcal{Q}_k/\mathcal{I}_k$ is the *graph algebra* corresponding to f . The infinite matrix $M_k : \mathcal{G}_k \times \mathcal{G}_k \rightarrow \mathbb{R}$ defined by $M_k(G, H) = f(\llbracket GH \rrbracket) = \langle G, H \rangle$ is called the *k-th connection matrix* of f . It is easy to see that the rank of M_k is the dimension of $\mathcal{Q}_k/\mathcal{I}_k$.

We will be interested in the special case when f is defined by $f(G) = t(G, W)$ for some fixed graphon W . In this case, $\mathcal{Q}_k/\mathcal{I}_k$ depends only on the weak isomorphism class of W . It was shown in [5] that in this case the inner product $\langle \cdot, \cdot \rangle$ is positive semidefinite, and hence so are the connection matrices.

There is a concrete representation of $\mathcal{Q}_k/\mathcal{I}_k$ that will be convenient to use and that will create a connection between automorphisms of W and the graph algebras. Let (J, W) be an arbitrary graphon. We define a map $\psi_k : \mathcal{Q}_k \rightarrow L^\infty(J^k)$ by letting $\psi_k(G)$ be the k -variable function $t_{x_1, x_2, \dots, x_k}(G, W) \in L^\infty(J^k)$. We extend this map linearly to general quantum graphs. Note that $L^\infty(J^k)$ is a commutative algebra with pointwise multiplication and addition, and ψ_k is an algebra homomorphism. The kernel of ψ_k is equal to \mathcal{I}_k and thus the range of ψ_k is isomorphic to the k -th graph algebra of W . We denote by $\mathcal{A}_k = \mathcal{A}_k(J, W)$ this subalgebra of $L^\infty(J^k)$.

We will need a subalgebra of \mathcal{A}_k : let \mathcal{A}_k^0 denote the linear span of functions $\psi_k(G)$, where in $G \in \mathcal{G}_k^0$ (so its labeled points are non-adjacent). By definition $\mathcal{A}_1 = \mathcal{A}_1^0$.

2.3 Metrics on graphons

For two points x, y of a graphon (J, W) , we define their *neighborhood distance* by

$$r_W(x, y) = \|W(x, \cdot) - W(y, \cdot)\|_1 = \int_J |W(x, z) - W(y, z)| dz.$$

It may happen that $W(x, \cdot)$ is not measurable for some x ; however, we can always change W on a set of measure 0 to make these one-variable sections of it measurable. We will assume in the sequel that these functions are measurable.

The distance function r_W is not necessarily a metric, only a semimetric, meaning that $r_W(x, y)$ may be 0 for distinct points x and y . Such points are called *twins*. How to merge twins to get a weakly isomorphic graphon for which r_W is a metric, was described in [1] (see also [4]).

As a further step of “purifying” a graphon, we can replace the metric space (J, r_W) by its completion. Furthermore, in this new topology the underlying probability measure may not have full support; we may restrict the graphon to the support of the measure (which is a closed and therefore complete subspace). This procedure is described in [7].

We call a graphon (J, W) *pure*, if (J, r_W) is a complete metric space, and π has full support (i.e., every open set has positive measure). The procedure described above implies that every graphon is weakly isomorphic to a pure graphon. Pure graphons will be crucial in this paper, even in order to define automorphisms.

It will be sometimes convenient to use the L^2 -distance instead of the L^1 -distance:

we consider

$$d_W(x, y) = \left(\int_J (W(x, z) - W(y, z))^2 dz \right)^{1/2}.$$

Since trivially $d_W(x, y)^2 \leq r_W(x, y) \leq d_W(x, y)$, these two metrics define the same topology. In particular, the metric space (J, d_W) associated with a pure graphon (J, W) is also complete and the measure π has full support.

One advantage of d_W is that it can be expressed in terms of restricted homomorphism densities. We consider the 2-labeled quantum graph h in Figure 1. Then it is easy to check that

$$d_W(x, y)^2 = t_{xy}(h, W). \quad (3)$$




Figure 1: The quantum graph h in the representation of the metric d_W .

While the difference between the metrics r_W and d_W is not essential, the 2-neighborhood metric (called the *similarity metric* in [4]) is more substantially different [6, 7]. One way to define it is to introduce the “operator square” of a graphon:

$$(W \circ W)(x, y) = \int_J W(x, z)W(z, y) d\pi(x),$$

and then consider the neighborhood distance of the graphon $(J, W \circ W)$:

$$\bar{r}_W(x, y) = r_{W \circ W}(x, y) = \int_J \left| \int_J (W(x, u) - W(y, u))W(u, z) d\pi(u) \right| d\pi(z).$$

This definition looks artificial, but in fact it has many nice properties. It is easy to see that $\bar{r}(W) \leq r_W$. If (J, W) is a pure graphon, then (J, \bar{r}_W) is a metric space (in particular, the distance between distinct points is positive), which is not necessarily complete, but we can consider its completion (\bar{J}, \bar{r}_W) . We can extend the probability measure to \bar{J} by defining it to be 0 on the set of new points. We can also extend the function W to $\bar{W} : \bar{J} \times \bar{J} \rightarrow [0, 1]$ so that (\bar{J}, \bar{W}) is a graphon, and the metric $\bar{r}_{\bar{W}}$ is equal to the completion of the metric \bar{r}_W (this takes some care). We will not distinguish $\bar{r}_{\bar{W}}$ and \bar{r}_W in the sequel. On the other hand the metric $r_{\bar{W}}$ is quite different: In terms of the \bar{r}_W metric, all open sets have positive measure, while the set $\bar{J} \setminus J$ of new points is closed and has measure 0. On the other hand, in terms of the $r_{\bar{W}}$ metric, the set $\bar{J} \setminus J$ is open (of measure zero).

The main property of this completion, which we will need, is that the space (\bar{J}, \bar{r}_W) is compact ([7]; see also [4], Corollary 13.28). The metric \bar{r}_W has another important property ([4], Theorem 13.27):

Proposition 1 *If (J, W) is pure graphon, then the metric \bar{r}_W defines exactly the weak topology on \bar{J} . In other words, $\bar{r}_W(x_n, x) \rightarrow 0$ ($x, x_n \in \bar{J}$) if and only if*

$$\int_{\bar{J}} \bar{W}(x_n, y) f(y) dy \rightarrow \int_{\bar{J}} \bar{W}(x, y) f(y) dy$$

for every bounded measurable function $f : J \rightarrow \mathbb{R}$.

A further important property of the metric \bar{r}_W , which we don't use in this paper but is worth mentioning, is that a decomposition of \bar{J} into sets with small \bar{r}_W -diameter corresponds to a (weak) regularity partition. We refer to [6, 7, 4] for the exact statement of this correspondence.

2.4 Continuity of restricted homomorphism numbers

We start with citing Lemma 13.19 from [4]:

Lemma 2 *Let (J, W) be a pure graphon and let $F = (V, E)$ be a k -labeled multigraph with nonadjacent labeled nodes. Then*

$$|t_{x_1, \dots, x_k}(F, W) - t_{y_1, \dots, y_k}(F, W)| \leq |E(F)| \max_{i \leq k} r_W(x_i, y_i). \quad (4)$$

for all $x_1, \dots, x_k, y_1, \dots, y_k \in J$.

We need a version of this lemma for the \bar{r}_W -distance instead of the r_W -distance. Some special cases of this were proved in [4], Section 13.4.

Lemma 3 *Let (J, W) be a pure graphon, and let $F = (V, E)$ be a k -labeled simple graph with nonadjacent labeled nodes. Then the restricted homomorphism function $t_{x_1 \dots x_k}(F, W)$ is continuous in each of its variables x_i on the metric space (\bar{J}, \bar{r}_W) .*

Proof. Consider any point $x = (x_1, \dots, x_k) \in \bar{J}^k$, and let $y_1, y_2, \dots \in \bar{J}$ be such that $\bar{r}_W(y_n, x_1) \rightarrow 0$ if $n \rightarrow \infty$. We want to show that

$$t_{y_n x_2 \dots x_k}(F, W) \rightarrow t_{x_1 \dots x_k}(F, W) \quad (m \rightarrow \infty). \quad (5)$$

Let $N(1) = \{k+1, \dots, k+r\}$, and let F' be obtained from F by deleting node 1 and labeling nodes $k+1, \dots, k+r$. Then

$$t_{y_n x_2 \dots x_k}(F, W) = \int_{J^r} \prod_{i=k+1}^{k+r} W(y_n, z_i) t_{x_2 \dots x_k z_{k+1} \dots z_{k+r}}(F', W) dz_{k+1} \dots dz_{k+r}.$$

The condition $\bar{r}_W(y_n, x_1) \rightarrow 0$ implies that $W(y_n, \cdot) \rightarrow W(x_1, \cdot)$ weakly as $n \rightarrow \infty$ ([4], Theorem 13.7). It is easy to see that this implies that $\prod_{i=k+1}^{k+r} W(y_n, z_i) \rightarrow$

$\prod_{i=k+1}^{k+r} W(x_1, z_i)$ (weakly as a function of $(z_{k+1}, \dots, z_{k+r})$), which in turn implies that

$$\begin{aligned} t_{y_n x_2 \dots x_k}(F, W) &\rightarrow \int_{J_r} \prod_{i=k+1}^{k+r} W(x_1, z_i) t_{x_2 \dots x_k z_{k+1} \dots z_{k+r}}(F', W) dz_{k+1} \dots dz_{k+r} \\ &= t_{x_1 x_2 \dots x_k}(F, W), \end{aligned}$$

as claimed. \square

Let us discuss the restrictions in these lemmas. It is obvious that these lemmas do not remain valid if we allow edges between labeled nodes: for example, $W(x, y) = t_{xy}(K_2^{\bullet\bullet}, W)$ itself is not necessarily continuous. Lemma 2 implies that t_{x_1, \dots, x_k} is continuous (even Lipschitz) in the neighborhood distance, simultaneously in all variables. Lemma 3, however, fails to hold in this stronger sense; see Example 4 below (adapted from [4], Example 13.30). This example also shows that in Lemma 3 we have to restrict F to simple graphs. Inequality (4) also shows that, for a fixed F , the difference $|t_x(F, W) - t_y(F, W)|$ can be estimated by $\max_i r_W(x_i, y_i)$, independently of W . Example 5 below shows that, even in the case $k = 1$, no such estimate can be given in terms of $\bar{r}_W(x, y)$.

Example 4 For $y \in [0, 1)$, let $y = 0.y_1 y_2 \dots$ be the binary expansion of y . Define $U(x, y) = y_k$ for $0 \leq y \leq 1$ and $2^{-k-1} \leq x \leq 2^{-k}$. Define $U(0, y) = 1/2$ for all y . This function is not symmetric, so we put it together with a reflected copy to get a graphon:

$$W(x, y) = \begin{cases} U(2x, 2y - 1), & \text{if } x \leq 1/2 \text{ and } y \geq 1/2, \\ U(2y, 2x - 1), & \text{if } x \geq 1/2 \text{ and } y \leq 1/2, \\ 0, & \text{otherwise.} \end{cases}$$

Let $u_k \in [2^{-1} + 2^{-k}, 2^{-1} + 2^{-k-1})$, then (as noted in [4]) the sequence (u_1, u_2, \dots) converges to the point 0 in the metric \bar{r}_W . On the other hand, for the 3-node path labeled at both endpoints

$$t_{u_n}(C_2^{\bullet}, W) = t_{u_n, u_n}(P_3^{\bullet\bullet}, W) = \frac{1}{4},$$

but

$$t_0(C_2^{\bullet}, W) = t_{0,0}(P_3^{\bullet\bullet}, W) = \frac{1}{8},$$

showing that $t_x(C_2^{\bullet}, W)$ is not continuous at $x = 0$, and that $t_{x,y}(P_2^{\bullet\bullet}, W)$, as a function of x and y , is not continuous at $(0, 0)$.

Example 5 Consider the weighted graph H given by the matrix of edgeweights

$$A = \begin{pmatrix} 1/4 & 1/2 & 0 & 1 \\ 1/2 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

and the vector of nodeweights

$$b = \begin{pmatrix} 2/3 - \varepsilon \\ 1/3 - \varepsilon \\ \varepsilon \\ \varepsilon \end{pmatrix}.$$

This weighted graph H can be considered as a pure graphon with a 4-point underlying space. Then $H \circ H = H'$ is the weighted graph given by the matrix of edgeweights

$$A' = A \operatorname{diag}(b) A = \begin{pmatrix} 1/8 & 1/4 & 2/9 & 2/9 \\ 1/4 & 1/2 & 1/9 & 2/9 \\ 1/6 & 1/3 & 0 & 0 \\ 1/6 & 1/3 & 0 & 0 \end{pmatrix} + O(\varepsilon),$$

and the same nodeweights as before. Let a and b be the last two nodes, then $\bar{r}_H(a, b) = O(\varepsilon)$, but $t_a(K_2^\bullet, W_H) = 1/3 - \varepsilon$ and $t_b(K_2^\bullet, W_H) = 2/3 - \varepsilon$.

We have seen that excluding edges between the labeled nodes is essential in both previous lemmas. The next lemma expresses W in terms of restricted homomorphism numbers for graphs with nonadjacent labeled nodes, and gives a (rather weak, but still useful) remedy for this restriction. We consider two sequences of quantum graphs f_n and g_n (Figure 2), and define

$$U_n(x, y) = \frac{t_{xy}(g_n, W)}{t_x(f_n, W)}. \quad (6)$$




Figure 2: Quantum graphs f , f_n and g_n . Here f^n is obtained by gluing together n copies of f along the labeled nodes; one of these nodes is then unlabeled (shown in white).

Lemma 6 *Let (J, W) be a pure graphon, and let $U_n : J^2 \rightarrow [0, 1]$ be defined by (6). Then $\|U_n - W\|_1 \rightarrow 0$ ($n \rightarrow \infty$).*

Proof. We have

$$|U_n(x, y) - W(x, y)| = \frac{|t_{xy}(g_n, W) - W(x, y)t_x(f_n, W)|}{t_x(f_n, W)}.$$

Here the numerator can be expressed as

$$\begin{aligned} |t_{xy}(g_n, W) - W(x, y)t_x(f_n, W)| &= \left| \int_J t_{xu}(f, W)^n (W(u, y) - W(x, y)) du \right| \\ &\leq \int_J t_{xu}(f, W)^n |W(u, y) - W(x, y)| du = \int_J (1 - d_W(x, u)^2)^n |W(u, y) - W(x, y)| du, \end{aligned}$$

while the denominator is

$$t_x(f_n, W) = \int_J (1 - d_W(x, y)^2)^n dy.$$

Integrating over y , and using Cauchy–Schwarz, we get

$$\begin{aligned} \int_J |U_n(x, y) - W(x, y)| dy &\leq \frac{\int_J (1 - d_W(x, u)^2)^n r_W(x, y) du}{\int_J (1 - d_W(x, y)^2)^n du} \\ &\leq \left(\frac{\int_J (1 - d_W(x, u)^2)^n d_W(x, u) du}{\int_J (1 - d_W(x, y)^2)^n du} \right)^{1/2} \end{aligned}$$

It is not hard to see that the right hand side tends to 0 as $n \rightarrow \infty$, which implies the lemma (in fact, a little more: $U_n(x, \cdot) \rightarrow W(x, \cdot)$ in the L^1 metric for every x). \square

3 Compactness of the automorphism group

3.1 Automorphisms of graphons

It only makes sense to define automorphisms of pure graphons.

Of course, one could define an “automorphism” of any graphon (J, W) as an invertible measure preserving map $\sigma : J \rightarrow J$ such that $W(x^\sigma, y^\sigma) = W(x, y)$ for almost all $x, y \in J$. However, there is a lot of trouble with this notion: weakly isomorphic graphons will have wildly different automorphism groups. An example with many automorphisms is a stepfunction W : here $\text{Aut}(W)$ contains the group of all invertible measure preserving transformations that leave the steps invariant (in addition to all the automorphisms of the corresponding weighted graph). Note, however, that if we purify a stepfunction, then we get a finite weighted graph, and so the large and “ugly” subgroups consisting of measure preserving transformations of the steps disappear. Another problem would be that any permutation of points of a zero-measure set should be considered an automorphism, so every graphon would have a transitive automorphism group.

Definition 7 Let (J, W) be a pure graphon. A measure preserving bijection $\sigma : J \rightarrow J$ is called an *automorphism* of (J, W) if, for every $x \in J$, the equality $W(x^\sigma, y^\sigma) = W(x, y)$ holds for almost all $y \in J$.

Note the change in the phrasing of the last condition: it is stronger than requiring that $W(x^\sigma, y^\sigma) = W(x, y)$ for almost all $x, y \in J$. This modification will exclude “automorphisms” like interchanging two points.

(The simpler but inadequate definition is given in [4]; the results announced there hold true with the definition given here.)

It is clear that every automorphism preserves the distances r_W and \bar{r}_W , and hence it extends to an automorphism of (\bar{J}, \bar{W}) . The points of $\bar{J} \setminus J$ can be identified in the graphon (\bar{J}, \bar{W}) by the property that every $r_{\bar{W}}$ -neighborhood of them has positive measure. So the automorphism groups of a pure graphon (J, W) and its completion (\bar{J}, \bar{W}) are essentially the same. In this section, we will mostly work with (\bar{J}, \bar{W}) .

We can endow $\text{Aut}(W)$ with a metric (and through this, with a topology) by

$$d(\sigma, \tau) = \sup_{x \in \bar{J}} \bar{r}_W(x^\sigma, x^\tau).$$

Not every isometry of the metric space (J, r_W) (or of the metric space (\bar{J}, \bar{r}_W)) is an automorphism.

Example 8 Let $([0, 1], W)$ be the pure graphon $W(x, y) = xy$, and consider the direct sum $([0, 1], W) \oplus ([0, 1], 1 - W)$. This is pure as well, and interchanging the two components is an isometry but not an automorphism in general.

The following technical lemma shows that a slight apparent weakening of the second condition in the definition of an automorphism leads to the same concept. We will formulate it for the \bar{r}_W -metric; for the r_W -metric the proof is similar (in fact, much simpler).

Lemma 9 *Let (J, W) be a pure graphon, and let $\varphi : \bar{J} \rightarrow \bar{J}$ be a bijective measure preserving map that is an isometry of (\bar{J}, \bar{r}_W) and satisfies $\bar{W}^\varphi = \bar{W}$ almost everywhere. Then φ is an automorphism.*

Proof. Let us call a point $x \in \bar{J}$ *nice*, if $\bar{W}(x, y) = \bar{W}(x^\varphi, y^\varphi)$ for almost all $y \in J$. The condition that $\bar{W}^\varphi = \bar{W}$ almost everywhere implies that almost all points are nice, but we want to show that *all* points are nice.

To this end, let us fix $x \in J$. Since the measure has full support in (\bar{J}, \bar{r}_W) , every neighborhood of x has positive measure, and hence there is a sequence of nice points x_n such that $\bar{r}_W(x_n, x) \rightarrow 0$. This means that

$$\int_J \left| \int_J (\bar{W}(x, y) - \bar{W}(x_n, y)) \bar{W}(y, z) dy \right| dz \rightarrow 0 \quad (n \rightarrow \infty). \quad (7)$$

Also, since φ is an isometry,

$$\int_J \left| \int_J (\bar{W}(x^\varphi, y) - \bar{W}(x_n^\varphi, y)) \bar{W}(y, z) dy \right| dz \rightarrow 0 \quad (n \rightarrow \infty).$$

Since φ is measure preserving, we can replace y by y^φ and z by z^φ in this equation:

$$\int_J \left| \int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x_n^\varphi, y^\varphi)) \overline{W}(y^\varphi, z^\varphi) dy \right| dz \rightarrow 0 \quad (n \rightarrow \infty).$$

Since the points x_n are nice, $\overline{W}(x_n^\varphi, y^\varphi) = \overline{W}(x_n, y)$ for almost all y , and similarly $\overline{W}(y^\varphi, z^\varphi) = \overline{W}(y, z)$ for almost all pairs (y, z) . This implies that

$$\int_J \left| \int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x_n, y)) \overline{W}(y, z) dy \right| dz \rightarrow 0 \quad (n \rightarrow \infty).$$

Comparing with (7), we get

$$\int_J \left| \int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x, y)) \overline{W}(y, z) dy \right| dz \rightarrow 0 \quad (n \rightarrow \infty).$$

The left hand side does not depend on n , and hence it follows that

$$\int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x, y)) \overline{W}(y, z) dy = 0$$

for almost all z . We can choose a sequence z_n for which this holds and for which $\overline{r}_{\overline{W}}(z_n, x) \rightarrow 0$. It is easy to see that this implies that

$$\int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x, y)) \overline{W}(y, x) dy = 0.$$

A similar argument gives

$$\int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x, y)) \overline{W}(y^\varphi, x^\varphi) dy = 0.$$

Subtracting, we get

$$\int_J (\overline{W}(x^\varphi, y^\varphi) - \overline{W}(x, y))^2 dy = 0,$$

which implies that $\overline{W}(x^\varphi, y^\varphi) = \overline{W}(x, y)$ for almost all y . This proves the lemma. \square

3.2 Compactness

The following fact is stated (without proof) in Section 13.5 of [4].

Theorem 10 *The automorphism group of a pure graphon is compact.*

This theorem is an immediate consequence of the following fact.

Lemma 11 *The automorphisms of a pure graphon (J, W) form a closed subgroup of the isometry group of (\bar{J}, \bar{r}_W) .*

Proof. Clearly every automorphism of (J, W) is an isometry of (\bar{J}, \bar{r}_W) , and these isometries form a subgroup. We want to prove that this subgroup is closed in the topology of pointwise convergence.

Let (φ_n) be a sequence of automorphisms of (J, W) , and assume that they converge to an isometry φ . We want to prove that φ is not only an isometry, but an automorphism. By Lemma 9, it suffices to prove the following claims.

Claim 1 *For every open set X , $\pi(\varphi(X) \Delta \varphi_n(X)) \rightarrow 0$ as $n \rightarrow \infty$.*

Indeed, since $\varphi_n(x) \rightarrow \varphi(x)$ for every $x \in J$, it follows that for every $x \in X$, $\varphi_n(x) \in \varphi(X)$ if n is large enough. This means that every point belongs to a finite number of sets $X \setminus \varphi_n^{-1}(\varphi(X))$ only, which implies that $\pi(X \setminus \varphi_n^{-1}(\varphi(X))) = \pi(\varphi_n(X) \setminus \varphi(X)) \rightarrow 0$. By a similar argument, $\pi(\varphi_n^{-1}(\varphi(X)) \setminus X) = \pi(\varphi(X) \setminus \varphi_n(X)) \rightarrow 0$. This implies the Claim.

Claim 2 *The map φ is measure preserving.*

It suffices to show that φ preserves the measure of any open set $X \subseteq \bar{J}$. By Claim 1, $\pi(\varphi_n(X)) \rightarrow \pi(\varphi(X))$ as $n \rightarrow \infty$. Since φ_n is measure preserving, this implies that $\pi(X) = \pi(\varphi(X))$.

Claim 3 *$\bar{W}^\varphi = \bar{W}$ almost everywhere.*

It suffices to prove that for any two open sets A and B ,

$$\int_{A \times B} \bar{W}(x, y) dx dy = \int_{A \times B} \bar{W}(\varphi(x), \varphi(y)) dx dy. \quad (8)$$

For every $y \in J$,

$$\begin{aligned} & \left| \int_A \bar{W}(\varphi_n(x), \varphi_n(y)) dx - \int_A \bar{W}(\varphi(x), \varphi_n(y)) dx \right| \\ &= \left| \int_{\varphi_n(A)} \bar{W}(x, \varphi_n(y)) dx - \int_{\varphi(A)} \bar{W}(x, \varphi_n(y)) dx \right| \leq \pi(\varphi(A) \Delta \varphi_n(A)). \end{aligned} \quad (9)$$

Using that the maps φ_n are automorphisms,

$$\begin{aligned} \int_{A \times B} \bar{W}(x, y) dx dy &= \int_{A \times B} \bar{W}(\varphi_n(x), \varphi_n(y)) dx dy = \int_{A \times B} \bar{W}(\varphi(x), \varphi_n(y)) dx dy \\ &+ \int_{A \times B} \bar{W}(\varphi_n(x), \varphi_n(y)) dx dy - \int_{A \times B} \bar{W}(\varphi(x), \varphi_n(y)) dx dy \end{aligned}$$

The first term on the right side tends to $\int_{A \times B} \overline{W}(\varphi(x), \varphi(y)) dx dy$ by Proposition 1, and the difference in the last line tends to 0 as $n \rightarrow \infty$ by (9) and Claim 1. This proves Claim 3, and thereby the Lemma. \square

4 Spectra

4.1 Spectral decomposition

Since W is bounded, the operator T_W is Hilbert-Schmidt and hence it has a spectral decomposition

$$W(x, y) \sim \sum_{r=1}^{\infty} \lambda_r f_r(x) f_r(y), \quad (10)$$

where the λ_r are its nonzero eigenvalues and the functions $f_r \in L^2(J)$ are the corresponding eigenfunctions, forming an orthonormal system. Here $\lambda_r \rightarrow 0$. By definition

$$\lambda_r f_r(x) = \int_J W(x, y) f_r(y) dy \quad (11)$$

almost everywhere. We assume that $W(x, \cdot)$ is measurable for every x , and we can change f_r on a set of measure 0 so that (11) holds for every $x \in J$. We note that (11) implies that f_r is bounded:

$$|f_r(x)| = \frac{1}{|\lambda_r|} \left| \int_J W(x, y) f_r(y) dy \right| \leq \frac{1}{|\lambda_r|} \int_J |f_r(y)| dy = \frac{\|f_r\|_1}{|\lambda_r|} \leq \frac{\|f_r\|_2}{|\lambda_r|} = \frac{1}{|\lambda_r|}.$$

We need the following simple observation: for every $x \in J$,

$$\sum_{r=1}^{\infty} \lambda_r^2 f_r(x)^2 = \|W(x, \cdot)\|_2^2 = t_x((K_2^\bullet)^2, W). \quad (12)$$

Indeed, using (11) and the fact that $\{f_r\}$ is an orthonormal system, we get

$$\sum_{r=1}^{\infty} \lambda_r^2 f_r(x)^2 = \sum_{r=1}^{\infty} \left(\int_J W(x, y) f_r(y) dy \right)^2 = \sum_{r=1}^N \langle W(x, \cdot), f_r \rangle^2 = \|W(x, \cdot)\|_2^2.$$

(the last equality follows because even though $\{f_r\}$ may not be a complete orthogonal system, it can be extended by functions in the nullspace of T_W to such a system, and these additional functions contribute 0 terms). The second equality in (12) is trivial by definition. (12) in turn implies that

$$\sum_{r=1}^N \lambda_r^2 f_r(x)^2 \leq 1. \quad (13)$$

Expansion (10) may not hold pointwise, only in L^2 ; but it follows from basic results on Hilbert-Schmidt operators that if we take the inner product with any function $U \in L^2(J \times J)$, then we get an equation:

$$\int_{J \times J} W(x, y) U(x, y) dx dy = \sum_{r=1}^{\infty} \lambda_r \int_{J \times J} f_r(x) f_r(y) U(x, y) dx dy, \quad (14)$$

where the sum on the right side is absolutely convergent. We need the following stronger fact:

Lemma 12 *Let (J, W) be a graphon, and let (22) be its spectral decomposition.*

(a) *For $U \in L^2(J)$ and $y \in J$, the sum*

$$\sum_{r=1}^{\infty} \lambda_r f_r(y) \int_J U(x) f_r(x) dx \quad (15)$$

is absolutely convergent.

(b) *For every bounded measurable function $U : J \times J \rightarrow \mathbb{R}$ and for almost all $y \in J$,*

$$\int_J W(x, y) U(x, y) dx = \sum_{r=1}^{\infty} \lambda_r f_r(y) \int_J U(x, y) f_r(x) dx. \quad (16)$$

Proof. (a) We have

$$\begin{aligned} \sum_{r=N}^{\infty} \left| \lambda_r f_r(y) \int_J U(x) f_r(x) dx \right| &\leq \sum_{r=N}^{\infty} |\lambda_r| |f_r(y)| \left| \int_J U(x) f_r(x) dx \right| \\ &\leq \left(\sum_{r=N}^{\infty} \lambda_r^2 f_r(y)^2 \right)^{1/2} \left(\sum_{r=N}^{\infty} \left(\int_J U(x) f_r(x) dx \right)^2 \right)^{1/2}. \end{aligned} \quad (17)$$

Here the first factor is the tail of a convergent sum by (13), and hence it tends to 0 as $n \rightarrow \infty$. Furthermore, $\{f_r\}$ is an orthonormal system, and hence

$$\begin{aligned} \sum_{r=N}^{\infty} \left(\int_J U(x) f_r(x) dx \right)^2 &\leq \sum_{r=1}^{\infty} \left(\int_J U(x) f_r(x) dx \right)^2 \\ &= \sum_{r=1}^{\infty} \langle U, f_r \rangle^2 \leq \|U\|_2^2, \end{aligned}$$

proving (a).

Let $g_1(y)$ and $g_2(y)$ be the functions on the left and right sides of equation (16). Then for any bounded measurable function $h : J \rightarrow \mathbb{R}$,

$$\begin{aligned} \langle h, g_1 \rangle &= \int_{J \times J} W(x, y) U(x, y) h(y) dx dy \\ &= \sum_{r=1}^{\infty} \lambda_r \int_{J \times J} U(x, y) h(y) f_r(x) f_r(y) dx dy \\ &= \int_J h(y) \sum_{r=1}^{\infty} \lambda_r f_r(y) \int_J U(x, y) f_r(x) dx dy = \langle h, g_2 \rangle \end{aligned}$$

(where we use (14) and the fact that the sum in the third line is absolutely convergent). This proves that $g_1 = g_2$ almost everywhere. \square

4.2 Spectral decomposition of pure graphons

In this chapter we use the topological properties of pure graphons to formulate finer statements about spectral decompositions. First of all, note that if (J, W) is a pure graphon then eigenfunctions of W are continuous functions on \overline{J} in the metric \overline{r}_W ([4], Corollary 13.29). Furthermore, the eigenfunctions separate the points of J :

Lemma 13 *If (J, W) is a pure graphon, then for every pair of distinct points $x, y \in J$ there is an eigenfunction f of W such that $f(x) \neq f(y)$.*

Proof. By way of contradiction, assume that x and y cannot be separated this way. From $x \neq y$ we obtain that $\overline{r}_W(x, y) > 0$, and thus the functions $W \circ W(x, \cdot)$ and $W \circ W(y, \cdot)$ have a positive distance in $L^2(J)$. On the other hand, $W \circ W(z, \cdot) = \sum_{i=1}^{\infty} \lambda_i^2 f_i(z) f_i(\cdot)$ holds for every fixed $z \in J$ where the sum is L^2 -convergent. Applying this formula for $z = x$ and $z = y$ together with our assumption that $f_i(x) = f_i(y)$, we get a contradiction. \square

Lemma 14 *If (J, W) is a pure graphon, then the sum on the left side of (12) converges uniformly for $x \in \overline{J}$.*

Proof. Using continuity of the eigenfunctions we obtain that every term on the left side of (12) is continuous in \overline{r}_W , and so is the right side by Lemma 3. Since every term on the left side is nonnegative, it follows by Dini's Theorem that the convergence is uniform in x . \square

This allows us to get the following stronger version of Lemma 12 for pure graphons:

Lemma 15 (a) *If (J, W) is a pure graphon, then the sum (15) is uniformly absolute convergent for $y \in \overline{J}$.*

(b) *If, in addition, $U(x, y)$ is a continuous function of y for every $x \in J$ in the neighborhood distance, then the expansion (16) holds for every $y \in J$.*

Proof. (a) By Lemma 14,

$$\sup_x \sum_{r=N}^{\infty} \lambda_r^2 f_r(x)^2 \rightarrow 0 \quad (N \rightarrow \infty). \quad (18)$$

Hence the computation in (17) gives an estimate of the tail uniformly for all $y \in \bar{J}$.

(b) The left side of (16) defines a continuous function of $y \in J$ in the metric r_W . Every term on the right side is also continuous, and the convergence is uniform by the estimate (17), using (18). Hence the limit is a continuous function of $y \in J$. The space J has the property that every nonempty open set has positive measure. If two continuous functions are equal almost everywhere on such a space, then they are equal everywhere. \square

Corollary 16 *If the automorphism group of a pure graphon (J, W) is transitive on J , then (J, W) is compact and $\bar{J} = J$.*

Proof. Let $x \in J$, then the orbit of x is a continuous image of $\text{Aut}(J, W)$, and so it is compact in the metric r_W . If the automorphism group is transitive on J , then this orbit is J , and hence (J, r_W) is compact. Since $\bar{r}_W \leq r_W$, this implies that (J, \bar{r}_W) is compact, and since J is dense in (\bar{J}, \bar{r}_W) , it follows that $J = \bar{J}$. \square

We use our results above about spectra to describe a way, more explicit than convergence in L^2 , of the convergence of the expansion (10). For a graphon (J, W) and $\lambda > 0$, we define the graphon $(J, [W]_\lambda)$ by the following partial sum of (10):

$$[W]_\lambda(x, y) = \sum_{|\lambda_r| \geq \lambda} \lambda_r f_r(x) f_r(y). \quad (19)$$

Note that this sum is finite. If W has multiple eigenvalues, then the terms $\lambda_r f_r(x) f_r(y)$ depend on the basis chosen in the eigenspaces, but $[W]_\lambda$ does not depend on this basis. Let

$$U_\lambda = \bigoplus_{|\lambda_r| \geq \lambda} E_{\lambda_r}, \quad (20)$$

where E_{λ_r} is the eigenspace of W corresponding to λ_r . Let Π_λ denote the orthogonal projection of $L^2(J)$ onto U_λ . Then $T_{[W]_\lambda} = T_W \Pi_\lambda$. From the inequality $\sum_{i=1}^{\infty} \lambda_i^2 \leq 1$, it follows that the rank of $[W]_\lambda$ (the dimension of U_λ) is at most $1/\lambda^2$.

Assume that the eigenvalues are ordered so that $|\lambda_1| \geq |\lambda_2| \geq \dots$. Let μ_λ denote the probability distribution of the vector $(f_1(x), f_2(x), \dots, f_d(x)) \in \mathbb{R}^d$, where $d = \dim(U_\lambda)$ and $x \in J$ is chosen randomly, and let $S_\lambda \subset \mathbb{R}^d$ be the support of μ_λ . Then the purification of $(J, [W]_\lambda)$ can be defined as (S_λ, W'_λ) , where

$$W'_\lambda((x_1, x_2, \dots, x_d), (y_1, y_2, \dots, y_d)) = \sum_{i=1}^d \lambda_i x_i y_i. \quad (21)$$

A coordinate-independent way of describing μ_λ is to consider the dual space of U_λ . For each $x \in J$, we consider the linear functional $f \mapsto (T_W f)(x)$ ($f \in U_\lambda$). If $x \in J$ is

chosen randomly we obtain the probability distribution μ_λ on U_λ^* , and we can define S_λ as its support. We will need the next lemma, which is a direct consequence of the results in the paper [9].

Lemma 17 *Let $\{U_n\}_{n=1}^\infty$ be a convergent sequence of graphons with limit W . Assume that $\lambda > 0$ is not an eigenvalue of W . Then there is subsequence $\{W_n\}_{n=1}^\infty$ in $\{U_n\}_{n=1}^\infty$ and choices of orthonormal eigenvectors for $[W_n]_\lambda$ and $[W_\lambda]$ such that the measures μ_λ^n constructed above for W_n converge to μ_λ weakly. \square*

If $\alpha < \beta$, then the projection

$$P_{\alpha,\beta} : \mathbb{R}^{\dim(U_\alpha)} \rightarrow \mathbb{R}^{\dim(U_\beta)}$$

(by forgetting the last $\dim(U_\alpha) - \dim(U_\beta)$ coordinates) transforms μ_α into μ_β . The map $P_{\alpha,\beta}$ is surjective from S_α to S_β .

Let $\{\alpha_i\}_{i=1}^\infty$ be a decreasing sequence tending to 0. Let S be the inverse limit of the system $\{P_{\alpha_{i+1},\alpha_i} : S_{\alpha_{i+1}} \rightarrow S_{\alpha_i}\}_{i=1}^\infty$. This means that

$$S = \{(s_1, s_2, \dots) : \forall i \in \mathbb{N}, s_i = P_{\alpha_{i+1},\alpha_i} s_{i+1}\} \subset \prod_{i=1}^\infty S_{\alpha_i}.$$

The limit of $\{\mu_{\alpha_i}\}_{i=1}^\infty$ defines a probability measure μ on the compact set S . Let (S, U_{α_i}) be the graphon defined on S using the formula (21) for the i -th coordinate.

Lemma 18 *For every graphon (J, W) there is a measure preserving homeomorphism $\tau : \bar{J} \rightarrow S$ such that $(U_{\alpha_i})^\tau = [\bar{W}]_{\alpha_i}$ holds for every i .*

Proof. Notice that the construction of (S, U_{α_i}) depends only on the weak isomorphism class of W , and so we can assume that (J, W) is pure. The maps $\tau_i : x \mapsto (f_1(x), f_2(x), \dots, f_d(x))$ from \bar{J} to S_i (where $d = \dim(U_{\alpha_i})$) are continuous in the \bar{r}_W metric. Hence the map $\tau = (\tau_1, \tau_2, \dots) : \bar{J} \rightarrow S$ is also continuous. Since τ separates elements in \bar{J} (to see this, apply lemma 13 for $W \circ W$), it is a bijection between \bar{J} and S . The desired property is clear from the definition of τ . \square

4.3 Subdividing edges

As an application of spectral decomposition, we prove the following generalization of Lemma 5.1 in [1] (which will be needed later on).

Lemma 19 *Let (J_1, W_1) and (J_2, W_2) be two pure graphons and let $a \in J_1^k$, $b \in J_2^k$. Let h be a k -labeled quantum multigraph. In every constituent of h , select an edge such that at least one endpoint of it is unlabeled, and let h_m denote the k -labeled quantum multigraph obtained from h by subdividing the selected edge by $m - 1$ new nodes in every constituent. Suppose there exists an $m_0 \geq 2$ such that $t_a(h_m, W_1) = t_b(h_m, W_2)$ for every $m \geq m_0$. Then $t_a(h, W_1) = t_b(h, W_2)$.*

Proof. Let g_i be obtained from h by keeping only those terms in which one endpoint of the selected edge is labeled i ($1 \leq i \leq k$). Let g_0 be the sum of the remaining terms, where the selected edge has no labeled endpoint. Let g'_i be the $(k+1)$ -labeled quantum multigraph obtained from g_i by deleting the selected edge from each constituent and labeling its unlabeled endpoint by $k+1$. Let g'_0 be the $(k+2)$ -labeled quantum multigraph obtained from g_0 by deleting the selected edge from each constituent and labeling its endpoints by $k+1$ and $k+2$. Then

$$\begin{aligned} t_a(h_m, W_1) &= \sum_{i=1}^k \int_{J_1} W_1^{\circ m}(a_i, x) t_{ax}(g'_i, W_1) dx \\ &\quad + \int_{J_1 \times J_1} W_1^{\circ m}(x, y) t_{axy}(g'_0, W_1) dx dy. \end{aligned}$$

We use the spectral decomposition

$$W_1^{\circ m}(x, y) \sim \sum_{r=1}^{\infty} \lambda_r^m f_r(x) f_r(y). \quad (22)$$

This decomposition holds almost everywhere for $m \geq 2$, but for $m = 1$, we can only claim that the sums on the right sides converge to the function on the left in L^2 . Since the graphon is pure, Lemma 15 implies that the expansion

$$\begin{aligned} t_a(h_m, W_1) &= \sum_{i=1}^k \sum_{r=1}^{\infty} \lambda_r^m f_r(a_i) \int_{J_1} t_{ax}(g'_i, W_1) f_r(x) dx \\ &\quad + \sum_{r=1}^{\infty} \lambda_r^m \int_{J_1 \times J_1} f_r(x) f_r(y) t_{axy}(g'_0, W_1) dx dy \end{aligned}$$

holds for all $m \geq 1$. We have an analogous expansion for $t_b(h_m, W_2)$. If these two expressions are equal for every integer $m \geq m_0$, then they are also equal for $m = 1$ (see e.g. [4], Proposition A.21). \square

Corollary 20 *Let (J_1, W_1) and (J_2, W_2) be two pure graphons and let $a \in J_1^k$, $b \in J_2^k$.*

(a) *If*

$$t_a(F, W_1) = t_b(F, W_2) \quad (23)$$

for every k -labeled simple graph F , then (23) holds for every k -labeled multigraph F .

(b) *If (23) holds for every k -labeled simple graph F with nonadjacent labeled nodes, then (23) holds for every k -labeled multigraph F with nonadjacent labeled nodes.*

Proof. (b) follows from Lemma 19 by induction on the number of parallel edges. To prove (a), it suffices to note that $W_1(a_i, a_j) = W_2(b_i, b_j)$ follows by considering the simple graph F with a single edge connecting i and j . \square

4.4 Automorphism groups and spectral decomposition

Let $g: J \rightarrow J$ be an automorphism of a graphon (J, W) . Notice that if f is an eigenfunction of length 1 of W then f^g is also an eigenfunction of length 1 corresponding to the same eigenvalue. As a consequence every automorphism of W acts on the space U_λ defined in (20) as an element in $O_\lambda := \bigoplus_{|\lambda_r| \geq \lambda} O(E_{\lambda_r})$ where $O(E_{\lambda_r})$ is the orthogonal group on E_{λ_r} . The corresponding action on the dual space U_λ^* leaves the measure μ_λ invariant. We will denote by Γ_λ the finite dimensional compact group formed by all elements O_λ that preserve μ_λ . (Note that Γ_λ is the automorphism group of $[W]_\lambda$.)

The group O_α acts on both U_α and U_α^* . Since U_β is an invariant subspace of O_α , the group O_α acts on U_β^* as well. In particular, there is a homomorphism $h_{\alpha,\beta}: \Gamma_\alpha \rightarrow \Gamma_\beta$. We denote by Γ_W the inverse limit of the system $\{h_{\alpha_{i+1}, \alpha_i}\}_{i=1}^\infty$.

We can describe the automorphism group of a compact graphon using representation of a graphon above.

Lemma 21 *For every graphon (J, W) the action of $\text{Aut}(W)$ on \overline{J} can be obtained as $\tau^{-1} \circ \Gamma_W \circ \tau$, where τ is the function in Lemma 18.*

Proof. We may assume that the graphon (J, W) is pure. First we show that $\text{Aut}(W) \subseteq \tau^{-1} \circ \Gamma_W \circ \tau$. Every automorphism of W , restricted to U_{α_i} ($i = 1, 2, \dots$), induces a consistent sequence of elements in $\prod_{i=1}^\infty \Gamma_{\alpha_i}$. It follows that $\tau \circ \text{Aut}(W) \circ \tau^{-1} \subseteq \Gamma_W$. The other containment is a direct consequence of Lemma 18: elements of $\tau^{-1} \circ \Gamma_W \circ \tau$ act on \overline{J} continuously and leave $[\overline{W}]_{\alpha_i}$ invariant for every i . This means that they also fix \overline{W} . \square

5 Orbits of the automorphism group

5.1 Characterization of the orbits

The following theorem characterizes the orbits of the automorphism group of a graphon.

Theorem 22 *Let (J, W) be a pure graphon, and let $a_1, \dots, a_k, b_1, \dots, b_k \in J$. Then there exists an automorphism $\varphi \in \text{Aut}(J, W)$ such that $a_i^\varphi = b_i$ if and only if $t_{a_1 \dots a_k}(F, W) = t_{b_1 \dots b_k}(F, W)$ for every k -labeled simple graph F in which the labeled nodes are independent.*

The following version is more general (at least formally).

Theorem 23 *Let (J_1, W_1) and (J_2, W_2) be two pure graphons and let $\alpha_i \in J_i^k$. Then there exists a measure preserving bijection $\varphi: J_1 \rightarrow J_2$ such that $W_2^\varphi = W_1$ almost everywhere and $\alpha_{1,i}^\varphi = \alpha_{2,i}$ if and only if $t_{\alpha_1}(F, W_1) = t_{\alpha_2}(F, W_2)$ for every k -labeled simple graph F .*

The proof of this theorem is a modification of the proof of the main result of [1], combined with more recent methods involving pure graphons.

First, we note that the condition in the theorem is self-sharpening: by Corollary 20, the condition holds for every k -labeled multigraph F . The following lemma is the main step in the proof.

Lemma 24 *Let (J_1, W_1) and (J_2, W_2) be two graphons and let $a \in J_1^k$, $b \in J_2^k$ such that*

$$t_a(F, W_1) = t_b(F, W_2)$$

for every k -labeled multigraph F . Let π_i denote the probability measure of J_i . Then we can couple π_1 with π_2 so that if (X, Y) is a pair from the coupling distribution, then

$$t_{a_1 \dots a_k X}(F, W_1) = t_{b_1 \dots b_k Y}(F, W_2)$$

almost surely for every $(k+1)$ -labeled multigraph F .

Proof. Consider two random points X from π and Y from π' , and the random variables

$$A = (t_{a_1 \dots a_k X}(F, H) : F \in \mathcal{F}_{k+1}) \quad \text{and} \quad B = (t_{b_1 \dots b_k Y}(F, H) : F \in \mathcal{F}_{k+1})$$

with values in $[0, 1]^{\mathcal{F}_{k+1}}$. We claim that the variables A and B have the same distribution. It suffices to show that A and B have the same mixed moments. If $F_1, \dots, F_m \in \mathcal{F}_{k+1}$, and q_1, \dots, q_m are nonnegative integers, then the corresponding moment of A is

$$\mathbb{E} \left(\prod_{i=1}^m t_{a_1 \dots a_k X}(F_i, H)^{q_i} \right) = \mathbb{E} (t_{a_1 \dots a_k X}(F_1^{q_1} \dots F_m^{q_m}, H)) = t_a(F, H),$$

where the multigraph F is obtained by unlabeled the node labeled $k+1$ in the multigraph $F_1^{q_1} \dots F_m^{q_m}$. Expressing the moments of B in a similar way, we see that they are equal by hypothesis. This proves that A and B have the same distribution.

Using Lemma 6.2 of [1] it follows that we can couple the variables X and Y so that $A = B$ with probability 1. In other words,

$$t_{a_1 \dots a_k X}(F, H) = t_{b_1 \dots b_k Y}(F, H')$$

for every $F \in \mathcal{F}_{k+1}$ with probability 1. □

For an infinite sequence $X \in J^{\mathbb{N}}$, let $X[n]$ denote its prefix of length n .

Lemma 25 *Under the conditions of the previous lemma, we can couple $\pi_1^{\mathbb{N}}$ with $\pi_2^{\mathbb{N}}$ so that if (X, Y) is a pair from the coupling distribution, then for every $n \geq 0$ and every $(k+n)$ -labeled graph F ,*

$$t_{a_1 \dots a_k X[n]}(F, W_1) = t_{b_1 \dots b_k Y[n]}(F, W_2)$$

almost surely.

Proof. By Lemma 24, we can define recursively a coupling κ_n of π_1^n with π_2^n so that $t_X(F, W_1) = t_Y(F, W_2)$ almost surely for every $F \in \mathcal{F}_{k+n}$, and κ_{n+1} , projected to the first n coordinates in both spaces, gives κ_n . The distributions κ_n give a distribution κ on $J_1^{\mathbb{N}} \times J_2^{\mathbb{N}}$, which clearly has the desired properties. \square

The following lemma can be considered as a version of the theorem for infinite sequences.

Lemma 26 *Let (J_1, W_1) and (J_2, W_2) be two pure graphons, and let $a_i = (a_{i,1}, a_{i,2}, \dots) \in J_i^{\mathbb{N}}$ be a sequence whose elements are dense in J_i . Suppose that $t_{a_1}(F, W_1) = t_{a_2}(F, W_2)$ for every partially labeled multigraph F . Then there is a measure preserving bijection $\varphi : J_1 \rightarrow J_2$ such that $W_2^\varphi = W_1$ almost everywhere and $a_{1,j}^\varphi = a_{2,j}$ for all $j \in \mathbb{N}$.*

The notation $t_a(F, W)$, where a is an infinite sequence, means that only those elements of a are considered whose subscript occurs in F as a label.

Proof. We start with noticing that

$$d_{W_1}(a_{1,i}, a_{1,j}) = d_{W_2}(a_{2,i}, a_{2,j}). \quad (24)$$

This follows by (3) and the hypothesis of the lemma.

For $x \in J_1$, take a subsequence $(a_{1,i_1}, a_{1,i_2}, \dots)$ such that $a_{i,i_n} \rightarrow x$. Then $(a_{1,i_1}, a_{1,i_2}, \dots)$ is a Cauchy sequence, and hence, by (24), so is the sequence $(a_{2,i_1}, a_{2,i_2}, \dots)$, and since (J_2, r_{W_2}) is complete, it has a limit x^φ . It is easy to see that this map is well-defined (i.e., it does not depend on the choice of the sequence $(a_{1,i_1}, a_{1,i_2}, \dots)$), and that φ is bijective.

Next, we claim that for every sequence $x_1, \dots, x_k \in J_1$ and every multigraph F with nonadjacent labeled nodes

$$t_{x_1^\varphi, \dots, x_k^\varphi}(F, W_2) = t_{x_1, \dots, x_k}(F, W_1). \quad (25)$$

Indeed, this holds if every x_i is an element of the sequence a_1 by hypothesis, and then it follows for all x_i by the continuity of $t_{x_1, \dots, x_k}(F, W_1)$ (Lemma 2).

Finally, consider the function

$$U_n(x, y) = \frac{t_{x,y}(g_n, W_1)}{t_x(g'_n, W_1)}.$$

By Lemma 6, $\|W_1 - U_n\|_1 \rightarrow 0$ as $n \rightarrow \infty$. Also, by (25),

$$U_n(x, y) = \frac{t_{x^\varphi, y^\varphi}(g_n, W_2)}{t_{x^\varphi}(g'_n, W_2)},$$

and applying Lemma 6 again, $\|W_2^\varphi - U_n\|_1 \rightarrow 0$ as $n \rightarrow \infty$. This implies that $W_1 = W_2^\varphi$ almost everywhere. \square

Now we are ready to prove the main theorem of this section.

Proof of Theorem 23. Let $X_1, X_2 \dots$ be independent random points of J_1 , and let $Y_1, Y_2 \dots$ be independent random points of J_2 . Applying Lemma 25 repeatedly, we can couple $X_1, X_2 \dots$ with $Y_1, Y_2 \dots$ so that, for any $(k+r)$ -labeled graph F ,

$$t_{a_1 \dots a_k X_1 \dots X_r}(F, W_1) = t_{b_1 \dots b_k Y_1 \dots Y_r}(F, W_2). \quad (26)$$

With probability 1, the elements of both sequences $a = (a_1, \dots, a_k, X_1, X_2, \dots)$ and $b = (b_1, \dots, b_k, Y_1, Y_2, \dots)$ are dense in J_1 and J_2 , respectively. Let us fix such a choice, then by Lemma 26 there is a measure preserving bijection $\varphi : J_1 \rightarrow J_2$ such that $W_2^\varphi = W_1$ almost everywhere and $a_i^\varphi = b_i$ for all $i \leq k$. This proves the theorem. \square

Corollary 27 *Let (J, W) be a pure graphon. Then the closure of \mathcal{A}_k^0 in $L^\infty(\bar{J}^k)$ consists of all continuous $\text{Aut}(J, W)$ -invariant functions on $(\bar{J}, \bar{r}_W)^k$.*

Proof. Lemma 2 implies that all functions in \mathcal{A}_k^0 are continuous and clearly they are invariant under automorphisms. The other containment follows from the Stone-Weierstrass theorem, since $(\bar{J}, \bar{r}_W)^k$ is compact, and by Theorem 22 the elements in \mathcal{A}_k^0 separate the orbits of $\text{Aut}(J, W)$. \square

5.2 Node-transitive graphons

Let \mathbb{G} be the automorphism group of the pure graphon (J, W) . We consider the natural action of \mathbb{G} on functions on J defined by $f^g(x) = f(x^g)$. Similarly \mathbb{G} acts diagonally on functions on J^n . For a subset $S \subset L^\infty(J^n)$ we denote by $S^\mathbb{G}$ the set of \mathbb{G} -invariant elements in S . It is clear that restricted homomorphism functions are invariant under the action of \mathbb{G} and thus all the algebras \mathcal{A}_k are \mathbb{G} -invariant.

Definition 28 A graphon is called *node-transitive* if the automorphism group of its pure representation (J, W) acts transitively on J .

The next theorem gives an algebraic characterization of node-transitive graphons.

Theorem 29 *Let (J, W) be a graphon. The following statements are equivalent.*

- (i) (J, W) is node-transitive.
- (ii) The functions $t_x(F, W)$ are essentially constant on J for all $F \in \mathcal{G}_1$.
- (iii) $\dim(\mathcal{A}_1) = 1$.
- (iv) The first connection matrix M_1 of W has rank 1.
- (v) $t(\llbracket F^2 \rrbracket, W)t(\llbracket H^2 \rrbracket, W) = t(\llbracket FH \rrbracket, W)^2$ for all $F, H \in \mathcal{G}_1$.

Proof. We may assume that (J, W) is pure. If (i) holds, then \mathbb{G} is transitive on J , and so every function $t_x(F, W)$ is constant on J , which implies (ii). Conversely, (ii) implies by Theorem 22 that \mathbb{G} is transitive on J , so (i) holds. Thus (i) and (ii) are equivalent. Every constant function is in \mathcal{A}_1 , hence $\dim(\mathcal{A}_1) \geq 1$, and so (ii) is

equivalent to (iii). We know that $\text{rk}(M_1) = \dim(\mathcal{A}_1)$, so (iv) is just a re-statement of (iii). Finally, (v) is a re-statement of (iv), since M_k is positive semidefinite. \square

Examples for node-transitive graphons are finite node-transitive graphs. Other examples are graphons defined on compact topological groups.

Definition 30 Let \mathbb{G} be a second countable compact topological group, which, together with its Haar measure, defines a standard probability space. Let $f : \mathbb{G} \rightarrow [0, 1]$ be a measurable function such that $f(x) = f(x^{-1})$. Then the graphon $W : \mathbb{G} \times \mathbb{G} \rightarrow [0, 1]$ defined by $W(x, y) = f(xy^{-1})$ is called a *Cayley graphon*.

Note that the condition $f(x) = f(x^{-1})$ is needed to guarantee that W is symmetric. By omitting this condition we get “directed Cayley graphons”.

Theorem 31 *Cayley graphons are node-transitive. Conversely, every node-transitive graphon is weakly isomorphic to a Cayley graphon.*

Note that a finite node-transitive graph G is not necessarily a Cayley graph (for example, the Petersen graph). However one can obtain a Cayley graph G' from G by replacing every vertex by m vertices and every edge by a complete bipartite graph $K_{m,m}$. The value m is the size of the stabilizer of a vertex in G in the automorphism group. The graph G' is weakly isomorphic to G as a graphon.

Proof. Let (\mathbb{G}, W) be a Cayley graphon on the compact topological group \mathbb{G} . It is clear that \mathbb{G} acts transitively (with multiplication from the right) on this graphon. (However, W might not be pure.) It follows that the restricted homomorphism functions $t_x(F, W)$ are all constant on \mathbb{G} . The third condition in Theorem 29 shows that W is node-transitive.

To prove the second assertion, let (J, π, W) be a node-transitive graphon; we may assume that it is pure. Let \mathbb{G} be its automorphism group. We know that \mathbb{G} is compact, and so it has a normalized Haar measure μ . Let us fix an element $c \in J$, and define the function $U : \mathbb{G} \times \mathbb{G} \rightarrow [0, 1]$ by $U(g, h) = W(c^g, c^h)$. We claim that (\mathbb{G}, μ, U) is a Cayley graphon weakly isomorphic to (J, π, W) .

Claim 4 *The map $\alpha : g \mapsto c^g$ defined on \mathbb{G} is measure preserving.*

The definition of the metric on $\text{Aut}(W)$ implies that $\bar{r}_W(\alpha(g), \alpha(h)) = \bar{r}_W(c^g, c^h) \leq d(g, h)$ for every $g, h \in \mathbb{G}$. This shows that α is continuous and hence, measurable.

Let \mathcal{B} denote the sigma-algebra on \mathbb{G} formed by the sets $\alpha^{-1}(X)$, where X is a Borel set in J , and let ν denote the measure on \mathcal{B} that is the pullback of π . It is clear that ν is \mathbb{G} -invariant. Standard topological group theory shows that ν extends to the Borel sigma-algebra on \mathbb{G} as the normalized Haar measure μ . This proves the Claim.

Since by definition $U = W^\alpha$, the Claim implies that (\mathbb{G}, U) is weakly isomorphic to (J, W) . Let $f : \mathbb{G} \rightarrow [0, 1]$ be defined by $f(g) = W(c^g, c)$. Then $U(g, h) = W(c^g, c^h) = w(c^{gh^{-1}}, c) = f(gh^{-1})$, so (\mathbb{G}, U) is a Cayley graphon. \square

Remark 32 Theorem 31 creates a connection between graph limit theory and an interesting and rich limit theory for functions on groups (see [8], [9]). The idea is the following. Let $\{f_i : \mathbb{G}_i \rightarrow [0, 1]\}_{i=1}^\infty$ be a sequence of measurable functions on compact groups. We say that the sequence f_i is convergent if the corresponding Cayley graphons $\{W_i\}_{i=1}^\infty$ converge. By Proposition 33 and Theorem 31, we prove that the limit of $\{W_i\}_{i=1}^\infty$ is weakly isomorphic to a Cayley graphon defined by a measurable function $f : \mathbb{G} \rightarrow [0, 1]$ on a compact group. We say that f is the *limit object* of the sequence $\{f_i\}_{i=1}^\infty$. It turns out that one can define this limit concept without passing to graphons. This point of view was heavily used in the second author's approach [8] to higher order Fourier analysis.

5.3 Limits of node-transitive graphons

We start with the observation that, if a convergent graph sequence consists of node-transitive graphs, then their limit graphon is node-transitive as well. More generally, we have the following consequence of the fifth condition in Theorem 29.

Corollary 33 *If a sequence of node-transitive graphons is convergent, then their limit graphon is also node-transitive.*

What makes this simple assertion interesting is the fact that the automorphism group of the limit graphon is not determined by the automorphism groups of graphs or graphons in the convergent sequence.

Example 34 Fix any $0 < \alpha < 1$, and define the graph G_n by $V(G_n) = [n]$, where every $i \in [n]$ is connected to the next and previous $\lfloor \alpha n \rfloor$ nodes (modulo n). The automorphism group of G_n is the dihedral group D_n . This sequence tends to the pure graphon on S^1 , with $W(x, y) = \mathbb{1}(\angle(x, y) \leq 1/2)$, whose automorphism group is $O(2)$, the continuous version of the dihedral groups.

No surprise so far. But let us consider the graphs $G_n \times G_{n+1}$. Add edges connecting every node (i, j) to $(i+a, j+a)$, where $\alpha(n+1) < a < n/2$. Let H_n denote the resulting graph.

Identifying node (i, j) with $(i-1)n + j \in [n(n+1)]$, it is not hard to see that $\text{Aut}(H_n) = D_{n(n+1)}$. The limit of this graph sequence is the pure graphon (J, W) , where J is the torus $S^1 \times S^1$, and $W((x_1, x_2), (y_1, y_2)) = \mathbb{1}(\angle(x_1, y_1) \leq 1/2, \angle(x_2, y_2) \leq 1/2)$, whose automorphism group is the wreath product of $O(2)$ with Z_2 , a 2-dimensional group different from $O(2)$.

Example 35 The next example (in a slightly different form) is from the papers [8] and [9]. It shows that even if the underlying group \mathbb{G} is the same for a convergent sequence of Cayley graphons, a transitive action on the limit graphon may need a different, bigger group. Let $\mathbb{G} = \mathbb{R}/\mathbb{Z}$ be the circle group and let $\xi : \mathbb{G} \rightarrow \mathbb{C}$ be the character defined by $\xi(x) = e^{2\pi i x}$. Let f_n be the function $\Im(1 + \xi + \xi^n)/2$ where \Im denotes the imaginary part. It is not hard to see that the limit of the Cayley graphons corresponding to f_n is the Cayley graphon corresponding to the function $f(x, y) = \Im(1 + \xi(x) + \xi(y))/2$ on the torus \mathbb{G}^2 .

In the light of the previous example the next theorem (which we quote from [9]) is somewhat surprising. We need a definition.

Definition 36 Let \mathbb{G} be a compact group with Haar measure μ . Let V_n denote the subspace of $L^2(\mathbb{G}, \mu)$ spanned by the \mathbb{G} -invariant subspaces of dimension at most n . We say that \mathbb{G} is *weakly random* if V_n is finite dimensional for every n .

Theorem 37 Let \mathbb{G} be a weakly random compact group. Let $\{f_n : \mathbb{G} \rightarrow [0, 1]\}_{n=1}^\infty$ be a sequence of measurable functions such that the corresponding Cayley graphons converge. Then the limit graphon is again a Cayley graphon on \mathbb{G} . \square

The best known example for a weakly random group is the orthogonal group $O(3)$. This shows that Cayley graphons on $O(3)$ behave very differently from Cayley graphons on $O(2)$. Cayley graphons on $O(3)$ are closed with respect to graphon convergence, however Cayley graphons on $O(2)$ are not closed.

We cite a related result, which is a consequence of a theorem of Gowers [2], indicating further, more subtle, relations between the automorphism groups of graphs and their limits.

Theorem 38 (Gowers) Let G_n be a Cayley graph of a group Γ_n ($n = 1, 2, \dots$), where the edge-density of G_n tends to a limit $0 \leq c \leq 1$, and the minimum dimension in which Γ_n has a nontrivial representation tends to infinity. Then the sequence $(G_n)_{n=1}^\infty$ is quasirandom, i.e., it tends to a pure graphon (J, W) where J has a single point. \square

Our goal is to determine the automorphism group of the limit of a sequence of node transitive graphs. Using lemma 21 one can reduce the problem of computing the automorphism group of W to the same problem about bounded rank graphons. To demonstrate this principle we show the next theorem. Recall that a compact group Γ is *abelian by pro-finite* if it has a closed abelian normal subgroup A such that Γ/A is the inverse limit of finite groups.

Theorem 39 Let $\{G_n\}_{n=1}^\infty$ be a sequence of node-transitive graphs converging to a graphon (J, W) . Then (J, W) is weakly isomorphic to a Cayley graphon on an abelian by pro-finite group.

Proof. We want to show that $G = \text{Aut}(W)$ has a closed, abelian by pro-finite subgroup that acts transitively on J . Let $\{\alpha_i\}_{i=1}^\infty$ be a decreasing sequence of real numbers with $\lim_{i \rightarrow \infty} \alpha_i = 0$ that contains no eigenvalue of W . We can assume that (J, W) is pure. Since W is node transitive, theorem 10 implies that J is compact and $\bar{J} = J$. We will use the notation from chapter 4.2.

For every W_j let $\mu_{\alpha_i}^j$ denote the measure defined above for W_j in the explicit coordinate system \mathbb{R}^{d_i} where $d_i = \dim(U_{\alpha_i})$. For finitely many values of j the measure $\mu_{\alpha_i}^j$ may exist in a different dimension but we ignore those values. By choosing a subsequence we can assume without loss of generality that the conditions of the lemma 17 hold for every i .

Let $G_i^j \subset O(d_i)$ denote the automorphism group of $\mu_{\alpha_i}^j$ and let H_i denote the closed subgroup in $O(d_i)$ whose elements are ultra-limits (for some fixed ultrafilter ω)

of sequences (g_1, g_2, \dots) where $g_j \in G_i^j$. It is clear that elements of H_i preserve ν_i and it acts transitively on S_i .

We claim that H_i is abelian by finite. A classical theorem by Camille Jordan [11] states that there is a function $f(n)$ such that any finite subgroup of $GL(n, \mathbb{C})$ contains an abelian group of index at most $f(n)$. Using this theorem, we see that each G_i^j has an abelian subgroup of index at most $f(d_i)$. It is a standard technique to show that this property is inherited by the ultralimit H_i . If the groups G_i^j are all abelian, then the continuity of the commutator word shows that H_i is abelian. For the general case, choose $f(d_i)$ coset representatives $g_{i,j,k}$ in each group G_i^j for the abelian subgroup where $1 \leq k \leq f(d_i)$. Their limits as $j \rightarrow \infty$ will be coset representatives for the limiting abelian group.

To finish the proof, let H be the inverse limit of the groups H_i with respect to the homomorphisms $P_{\alpha_{i+1}, \alpha_i}$. Then $H \subseteq \Gamma_W$ and H acts transitively on S . By lemma 21 we obtain that $\tau^{-1} \circ H \circ \tau \subseteq \text{Aut}(W)$ is transitive on J . \square

6 Graph algebras of finite rank graphons

We conclude with an application of our results on automorphisms of graphons to characterize graph algebras of graphons that have finite rank as integral kernel operators. Let (J, W) be a pure graphon with finite rank. The spectral decomposition (10) takes the simpler form

$$W(x, y) = \sum_{i=1}^t \lambda_i f_i(x) f_i(y). \quad (27)$$

For any sufficiently small $\lambda > 0$, we have $[W]_\lambda = W$, and so the considerations in Section 4.2 imply that (J, r_W) is compact.

Let $\mathbb{G} = \text{Aut}(W)$, and let S be the function algebra generated by the eigenfunctions of W . We denote by S_n the space of homogeneous polynomials of degree n in the eigenfunctions of W , so that $S = \bigoplus_n S_n$. Substituting (27) in the definition (2) of restricted homomorphism numbers, we see that $\mathcal{A}_1 \subseteq S$. Since the functions in \mathcal{A}_1 are \mathbb{G} -invariant, it follows that $\mathcal{A}_1 \subseteq S^{\mathbb{G}}$. Our main goal is to prove that equality holds here.

For $h \in L^\infty(J^n)$, we define

$$r(h, x) = \int_{x_1, x_2, \dots, x_n} h(x_1, x_2, \dots, x_n) \prod_{i=1}^n W(x, x_i). \quad (28)$$

The following lemma states some elementary properties of this function.

Lemma 40 (a) *If $h \in L^\infty(J^n)$ then $r(h, x) \in S_n$ (as a function of $x \in J$).*

(b) *If $h \in \mathcal{A}_n$, then $r(h, x) \in \mathcal{A}_1$.*

(c) *$r(h^g, x) = r(h, x)^g$ for every $g \in \mathbb{G}$.*

Proof. Assertion (a) follows by substituting formula (27) in (28). To prove (b), let $h(x_1, \dots, x_n) = t_{x_1 \dots x_n}(s, W)$, and let $s' \in \mathcal{Q}_1$ denote the one-labeled quantum graph obtained from s by connecting a new node with label 1 to all the labeled nodes and then we removing the original labels. Then $r(h, x) = t_x(s', W)$. Finally, (c) follows by replacing $W(x, x_i)$ by $W(x^g, x_i^g)$ in the formula for $r(h^g, x)$. Since the action of \mathbb{G} is measure preserving, the integration over $(x_1^g, x_2^g, \dots, x_n^g)$ is equivalent to the integration over (x_1, x_2, \dots, x_n) . \square

Lemma 41 *Every function $f \in S_n$ can be expressed as $f = r(h, x)$ for some function $h \in L^\infty(J^n)$.*

Proof. If $h(x_1, x_2, \dots, x_n) = f_{i_1}(x_1)f_{i_2}(x_2) \dots f_{i_n}(x_n)$, then

$$r(h, x) = \lambda_{i_1} \lambda_{i_2} \dots \lambda_{i_n} f_{i_1}(x) f_{i_2}(x) \dots f_{i_n}(x).$$

Every function $f \in S_n$ can be expressed as a linear combination of functions such as that on the right side of the previous formula. Since $r(h, x)$ is linear in h , this completes the proof. \square

Lemma 42 *Every function $f \in S_n^{\mathbb{G}}$ can be expressed as $f = r(h, x)$ for some \mathbb{G} -invariant function $h \in L^\infty(J^n)$.*

Proof. By Lemma 41, $f(x) = r(q, x)$ for a suitable $q \in L^\infty(J^n)$. Let $h = \int_{\mathbb{G}} q^g$. It is clear that h is \mathbb{G} -invariant. By lemma 40 and the linearity of $r(\cdot, \cdot)$ in the first variable it follows that $f(x) = r(h, x)$. \square

Lemma 43 $S_n^{\mathbb{G}} = \mathcal{A}_1 \cap S_n$.

Proof. Trivially $S_n^{\mathbb{G}} \supseteq \mathcal{A}_1 \cap S_n$. To prove the reverse, let $f \in S_n^{\mathbb{G}}$. Trivially $f \in S_n$, so it suffices to prove that $f \in \mathcal{A}_1$. Lemma 42 shows that $f(x) = r(h, x)$ for some \mathbb{G} -invariant function $h \in L^\infty(J^n)$. Using Corollary 27, there is a sequence of functions $q_k \in \mathcal{A}_n^0$ such that $q_k \rightarrow h$ ($k \rightarrow \infty$) uniformly in x . By Lemma 40, $r(q_k, x) \in \mathcal{A}_1$ (as a function of $x \in J$), and clearly $r(q_k, x) \rightarrow f = r(h, x)$ uniformly in x . This implies that $f \in \mathcal{A}_1$. \square

Theorem 44 $\mathcal{A}_1 = S^{\mathbb{G}}$.

Proof. We have seen that $\mathcal{A}_1 \subseteq S^{\mathbb{G}}$. To prove the reverse, we note that every function $f \in S$ is a finite sum of functions $\sum_n f_n$, where $f_n \in S_n$, and if f is \mathbb{G} -invariant, then so are the terms f_n . Hence $S^{\mathbb{G}}$ is the linear span of the spaces $S_n^{\mathbb{G}}$. By the previous lemma we get that $S^{\mathbb{G}} \subseteq \mathcal{A}_1$. \square

Corollary 45 \mathcal{A}_1 is finitely generated.

Proof. The algebra S is a finitely generated commutative algebra and the compact group \mathbb{G} acts on S via automorphisms. Hilbert's theorem on \mathbb{G} -invariant rings implies that $\mathbb{A}_1 = S^{\mathbb{G}}$ is finitely generated. \square

References

- [1] C. Borgs, J. Chayes and L. Lovász: Moments of Two-Variable Functions and the Uniqueness of Graph Limits, *Geom. Func. Anal.* **19** (2010), 1597–1619.
- [2] W.T. Gowers: Quasirandom groups, *Combin. Probab. Comput.* **17** (2008), 363–387.
- [3] L. Lovász: The rank of connection matrices and the dimension of graph algebras, *Europ. J. Combin.* **27** (2006), 962–970.
- [4] L. Lovász: *Large networks and graph limits*, Amer. Math. Soc., Providence, R.I. (2012).
- [5] L. Lovász and B. Szegedy: Limits of dense graph sequences, *J. Combin. Theory B* **96** (2006), 933–957.
- [6] L. Lovász and B. Szegedy: Szemerédi's Lemma for the analyst, *Geom. Func. Anal.* **17** (2007), 252–270.
- [7] L. Lovász and B. Szegedy: Regularity partitions and the topology of graphons, in: *An Irregular Mind, Szemerédi is 70*, J. Bolyai Math. Soc. and Springer-Verlag (2010), 415–446.
- [8] B. Szegedy: On higher order Fourier analysis, <http://arxiv.org/abs/1203.2260>
- [9] B. Szegedy: Limits of kernel operators and the spectral regularity lemma, *Europ. J. Combin.* **32** (2011), 1156–1167.
- [10] A.M. Vershik and U. Haböck: Compactness of the congruence group of measurable functions in several variables, *Zapiski Nauchnykh Sem. POMI* **334** (2006), 57–67; English translation: *J. Math. Sci.* (N.Y.) **141** (2007), 1601–1607.
- [11] C. Curtis and I. Reiner: Representation Theory of Finite Groups and Associated Algebras, *John Wiley & Sons.* (1962), 258–262.