A GEOMETRIC ZERO-ONE LAW

ROBERT H. GILMAN, YURI GUREVICH, AND ALEXEI MIASNIKOV

Abstract. Each relational structure X has an associated Gaifman graph, which endows X with the properties of a graph. If x is an element of X, let $B_n(x)$ be the ball of radius n around x. Suppose that X is infinite, connected and of bounded degree. A first-order sentence ϕ in the language of X is almost surely true (resp. a.s. false) for finite substructures of X if for every $x \in X$, the fraction of substructures of $B_n(x)$ satisfying ϕ approaches 1 (resp. 0) as n approaches infinity. Suppose further that, for every finite substructure, X has a disjoint isomorphic substructure. Then every ϕ is a.s. true or a.s. false for finite substructures of X. This is one form of the geometric zero-one law. We formulate it also in a form that does not mention the ambient infinite structure. In addition, we investigate various questions related to the geometric zero-one law.

1. Introduction

Fix a finite purely relational vocabulary Υ. From now on structures are Υ structures and sentences are first-order Υ sentences by default. By substructure we mean the induced substructure corresponding to a subset of elements. All relationships between the elements are inherited, and other relationships are ignored.

According to the well known zero-one law for first-order predicate logic, a first-order sentence ϕ is either almost surely true or almost surely false on finite structures [7], [9]. In other words if a structure is chosen at random with respect to the uniform distribution on all structures with universe $\{1, 2, \ldots, n\}$, then the probability that ϕ is true approaches either 1 or 0 as n goes to infinity.

There is another version of the zero-one law in which instead of choosing a structure uniformly at random from the set of structures with universe $\{1, 2, \ldots, n\}$ one chooses an isomorphism class of structures uniformly at random from the set of isomorphism classes of structures

Date: June 1, 2007.
2000 *Mathematics Subject Classification.* 03C13.
Key words and phrases. finite structure, zero-one law, percolation.

with universe of size n. This second version is known as the unlabeled zero-one law. The first version, which has received the greater share of attention, is called the labeled zero-one law. It holds for models of parametric axioms, graphs for example, i.e., undirected graphs without loops. For an introduction and surveys see [5], [6, Chapter 3], [10], and [13].

There are many extensions of the zero-one law to different logics and different probability distributions. In this article we consider another kind of extension. We show in Theorem 3 that under certain circumstances there is a zero-one law for the finite substructures of a fixed infinite structure; Theorem 5 gives a variation on this theme which does not refer to the ambient infinite structure. Theorem 6 shows that our results can yield zero-one laws for classes of structures to which neither the labeled nor unlabeled law applies.

Let X be a fixed infinite structure. If X were finite, a natural way to compute the probability that a finite substructure satisfied a sentence ϕ would be to divide the number of substructures of X satisfying ϕ by the total number of substructures of X. As X is infinite, this simple approach does not work; but there is a straightforward extension which does. To explain it we need a few definitions.

Recall that the Gaifman graph [8] of X has the elements of X as its vertices and an undirected edge between any two distinct vertices, x, y, for which there is a relation $R \in \Upsilon$ and elements $z_1, \ldots z_\ell$ in X such that $R(z_1, \ldots, z_\ell)$ is true in X and $x, y \in \{z_1, \ldots z_\ell\}$. Denote the Gaifman graph of X by $[X]$.

If X is a graph, we may identify X with $[X]$. In any case we extend some standard graph-theoretic terminology from $[X]$ to X. The distance, $d(x, y)$, between $x, y \in X$ is the length of the shortest path from x to y in $[X]$ or ∞ if there is no such path. For any $Y \subseteq X$, $d(x, Y)$ is the minimum distance from x to a point in Y, and $B_n(Y)$ is the substructure of X supported by the elements a distance at most n from Y. $B_n(x)$ is an abbreviation of $B_n(\{x\})$. The ambient structure X to which $B_n(Y)$ and $B_n(x)$ refer will be clear from the context.

Two substructures of X are said to be disjoint if their intersection is empty and there are no edges between them in $[X]$. The disjoint union of structures is defined in the obvious way. Substructures corresponding to the connected components of $[X]$ are called components of X, and substructures which are unions of components are called closed. A structure with just one component is said to be connected. If all vertices of $[X]$ have finite degree, X is locally finite; and if the vertex degrees are uniformly bounded, X has bounded degree. Finally if the vertex degrees of all structures in a collection $\mathcal{C}$ are uniformly bounded, we say that $\mathcal{C}$ has bounded degree.

Definition 1. Suppose X is an infinite, connected, locally finite structure. A sentence is almost surely true for finite substructures of X if for every $x \in X$ the fraction of substructures of $B_n(x)$ for which the sentence is true approaches 1 as n approaches infinity. Likewise a sentence is almost surely false if that fraction approaches 0 as n approaches infinity.

The balls $B_n(x)$ mentioned in Definition 1 are finite because X is locally finite.

Definition 2. A structure X has the duplicate substructure property if for every finite substructure there is a disjoint isomorphic substructure.

Theorem 3. *Let X be an infinite connected structure of bounded degree and possessing the duplicate substructure property. Then any sentence is either almost surely true or almost surely false for finite substructures of X.*

We may think of the structure X from Theorem 3 as inducing a zero-one law on the collection, $\mathcal{C}(X)$, of its finite substructures. Conversely every collection $\mathcal{C}$ of finite substructures which satisfies Hypothesis 4 below obeys a zero one law of this type. (Observe that $C(X)$ satisfies Hypothesis 4.)

Hypothesis 4. The following conditions hold.

(1) $\mathcal{C}$ is closed under taking substructures.
(2) $\mathcal{C}$ has bounded degree bounded.
(3) If F_1 and F_2 are (not necessarily distinct) elements of $\mathcal{C}$, then there exists an element of $\mathcal{C}$ isomorphic to the disjoint union $F_1 \cup F_2$.
(4) $\mathcal{C}$ is *pseudo-connected* in the sense that for every $F \in \mathcal{C}$ there is an embedding of F into a connected member of $\mathcal{C}$.

Theorem 5. *Let $\mathcal{C}$ be a class of finite structures satisfying Hypothesis 4, and let S be the disjoint union of all members of $\mathcal{C}$. We have:*

(1) *There is an infinite structure X, called an* ambient structure *for $\mathcal{C}$, such that X satisfies the hypotheses of Theorem 3, and the finite substructures of X are the same as the elements of $\mathcal{C}$ up to isomorphism.*
(2) *Let X be any ambient structure for $\mathcal{C}$. Then an arbitrary first-order sentence is almost surely true for finite structures of X if*

and only if it holds in S. Consequently all ambient structures give the same same zero-one law on $\mathcal{C}$.

The proof of Theorem 3 proceeds along a well known path. We show that certain axioms are almost surely true for finite substructures of X and that the theory with those axioms is complete. Section 3 contains the proofs of Theorems 3 and 5 including a discussion of the almost sure theory of the finite substructures of X. In Section 4, we consider when the almost sure theory is decidable. In Sections 5 and 6 we show that random substructures of X are elementarily equivalent but not necessarily isomorphic. A random substructure of X is obtained by deleting each element of X with some fixed probability strictly between 0 and 1. The random substructure is the one supported by all the remaining elements. Random sustructures are related to the theory of percolation. See [1, 2].

Now we present some examples. It is straightforward to check that Theorem 3 applies to the following structures.

(1) The Cayley diagram of a finitely generated infinite group. Here Υ consists of one binary relation for each generator.
(2) An infinite connected vertex-transitive graph of finite degree. For example the graph obtained from a Cayley diagram of the type just mentioned by removing all loops and combining all edges between any two distinct vertices joined by an edge into a single undirected edge. See [11] for non-Cayley examples.
(3) The Cayley diagram of a free finitely generated monoid.
(4) The full binary tree; i.e., the tree with one vertex of degree two and all others of degree three. More generally the full k-ary tree for $k \geq 1$.
(5) An infinite connected locally finite and finite dimensional simplicial complex whose automorphism group is transitive on zero-simplices. There is one $n+1$-ary relation for each dimension n.

We conclude this section with an example of a class of structures which satisfies the geometric zero-one law, but for which neither the labeled nor unlabeled law holds. For this purpose a unary forest is defined to be a directed acyclic graph such that each vertex has at most one incoming edge and at most one outgoing edge.

A unary tree is a connected unary forest; that is, a directed graph consisting of a single finite or infinite directed path. $\mathcal{C}$ is the class of finite unary forests with edges labeled by 0 and 1; Υ consists of two binary relations, one for each edge label. $\mathcal{C}$ is closed under isomorphism, disjoint union, and restriction to components.

Theorem 6. *$\mathcal{C}$, the class of finite unary forests with edges labeled by 0 and 1, obeys the geometric zero-one law but does not obey either the labeled or unlabeled law.*

Proof. Pick an infinite labeled unary tree, X, such that all finite sequences of 0's and 1's appear as the labels of subtrees of X; observe that X satisfies the hypotheses of Theorem 3. Thus $\mathcal{C}$ obeys the geometric zero-one law.

To show that $\mathcal{C}$ does not satisfy the labeled or unlabeled law, we apply [4, Theorem 5.9]. Let $\mathcal{A}_n$ be the set of structures in $\mathcal{C}$ with universe $\{1, 2, \ldots, n\}$, and $\mathcal{B}_n$ a set of representatives for the isomorphism classes of structures in $\mathcal{A}_n$. The cardinalities of $\mathcal{A}_n$ and $\mathcal{B}_n$ are denoted a_n and b_n respectively. It follows immediately from [4, Theorem 5.9] that if $\sum_{n=1}^{\infty} \frac{a_n}{n!} t^n$ has finite positive radius of convergence, then $\mathcal{C}$ does not obey the labeled zero-one law. Likewise if $\sum_{n=1}^{\infty} b_n t^n$ has radius of convergence strictly between 0 and 1, then $\mathcal{C}$ does not obey the unlabeled zero-one law.

Consider a single unary tree with n vertices. The 2^{n-1} different ways of labeling the edges of this tree yield pairwise non-isomorphic labeled trees; and for each labeled tree, the $n!$ different ways of labeling the vertices yield different structures on $\{1, 2, \ldots, n\}$. Thus $2^{n-1} \leq b_n$ and $2^{n-1} n! \leq a_n$. On the other hand each unary forest of size n is isomorphic to a structure obtained by labeling the edges of a unary tree of size n with letters from the alphabet $\{0, 1, 2\}$ and then deleting all edges with label 2. It follows that $2^{n-1} \leq b_n \leq 3^{n-1}$ and $2^{n-1} n! \leq a_n \leq 3^{n-1} n!$. By the results mentioned above neither the labeled nor unlabeled zero-one law holds for $\mathcal{C}$. ☐

Acknowledgment. We thank Andreas Blass for useful discussions related to Section 6.

2. A Sufficient Condition for Elementary Equivalence

The main result of this section is that two structures which satisfy the following condition are elementarily equivalent.

Definition 7. Two structures satisfy the disjoint ball extension condition if whenever either structure contains a ball $B_n(x)$ disjoint from a finite substructure F, and the other structure has a substructure F' isomorphic to F, then the other structure also contains $B_n(y)$ disjoint from F' isomorphic to $B_n(x)$ by an isomorphism matching x to y.

Lemma 8. *Let X and X' be structures and Y a substructure of X. If α is an isomorphism of $B_n(Y)$ to a substructure of X', the following conditions hold.*

(1) *If $x_1 \in B_{n-1}(Y)$ and $x_2 \in B_n(Y)$ are joined by an edge in $[X]$, then $\alpha(x_1)$ and $\alpha(x_2)$ are joined by an edge in $[X']$.*
(2) *For any $x \in B_n(Y)$, $d(x, Y) \geq d(\alpha(x), \alpha(Y))$.*
(3) *$\alpha(B_n(Y)) \subseteq B_n(\alpha(Y))$.*
(4) *If α maps$(B_n(Y))$ onto $B_n(\alpha(Y))$, then for any $x \in B_n(Y)$, $d(x, Y) = d(\alpha(x), \alpha(Y))$.*

Proof. If x_1, x_2 are as above, then $R(t_1, \ldots, t_k)$ is true for some relation $R \in \Upsilon$ and elements $t_1, \ldots, t_k \in X$ with $x_1, x_2 \in \{t_1, \ldots, t_k\}$. It follows that $d(x_1, t_i) \leq 1$ for all i, which implies $\{t_1, \ldots, t_k\} \subseteq B_n(Y)$. As α is an isomorphism, $R(\alpha(t_1), \ldots, \alpha(t_k))$ holds in X'. Thus the first part is proved. The first part implies the next two, and the last one holds by symmetry. □

Lemma 9. *Let X and X' be structures. Suppose that for some $n \geq 1$ and substructures $Y \subseteq X$, $Y' \subseteq X'$ there is an isomorphism $\alpha : B_n(Y) \to B_n(Y')$ with $\alpha(Y) = Y'$. Then for any substructure Z of X with $B_m(Z) \subseteq B_{n-1}(Y)$, α maps $B_m(Z)$ isomorphically to $B_m(\alpha(Z))$.*

Proof. First suppose that $B_m(\alpha(Z)) \subseteq B_{n-1}(Y')$. Lemma 8(3) applied to α and α^{-1} yields $\alpha(B_m(Z)) \subseteq B_m(\alpha(Z))$ and $\alpha^{-1}(B_m(\alpha(Z))) \subseteq B_m(Z)$. It follows immediately that α maps $B_m(Z)$ isomorphically to $B_m(\alpha(Z))$ as desired.

Thus it suffices to show that $B_m(\alpha(Z)) \subseteq B_{n-1}(Y')$. Assume not. As $\alpha(Z) \subseteq B_{n-1}(Y')$, there must be an element $\alpha(x) \in B_n(Y') - B_{n-1}(Y')$ with $d(\alpha(x), \alpha(Z)) = k \leq m$. Consequently there is a path in $[X']$ from some $\alpha(z) \in \alpha(Z)$ to $\alpha(x)$ of length at most m and with all vertices of the path in $B_m(\alpha(Z))$. Without loss of generality assume that $\alpha(x)$ is the first point on that path not in $B_{n-1}(Y')$. But then Lemma 8 implies $x \in B_m(Z) - B_{n-1}(Y)$ contrary to hypothesis. □

Theorem 10. *If two locally finite structures satisfy the disjoint ball extension condition, then they are elementarily equivalent.*

Proof. Let X and X' be the two structures. We show that for each n the duplicator can win the n-step Ehrenfeucht game by constructing isomorphisms α_i from a substructure $F_i \subseteq X$ to a substructure $F'_i \subseteq X'$, where F_i and F'_i consist of the elements chosen by the spoiler and the duplicator in the first i steps. Each α_i will be the restriction of an isomorphism, also called α_i, from $B_{5^{n-i}}(F_i)$ to $B_{5^{n-i}}(F'_i)$.

We argue by induction on i. Suppose $i = 1$. By symmetry we may suppose that the spoiler picks $x \in X$. By hypothesis there is an isomorphism $\alpha_1 : B_{5^{n-1}}(x) \to B_{5^{n-1}}(x') \subseteq X'$ with $x' = \alpha_1(x)$. The duplicator chooses x'.

Assume $\alpha_i : B_{5^{n-i}}(F_i) \to B_{5^{n-i}}(F_i')$ is an isomorphism for some $i < n$. Again by symmetry the spoiler picks $x \in X$. We have $F_{i+1} = F_i \cup \{x\}$. If $B_{5^{n-i-1}}(x) \subseteq B_{5^{n-i}-1}(F_i)$, then we take α_{i+1} to be the restriction of α_i to $B_{5^{n-i-1}}(F_{i+1})$ and set $x' = \alpha_i(x)$, $F'_{i+1} = F'_i \cup \{x'\}$. By Lemma 9, α_{i+1} maps $B_{5^{n-i-1}}(F_{i+1})$ onto $B_{5^{n-i-1}}(F'_{i+1})$.

Otherwise $B_{5^{n-i-1}}(x)$ is not a subset of $B_{5^{n-i}-1}(F_i)$. Some $y \in B_{5^{n-i-1}}(x)$ must be a distance at least 5^{n-i} from F_i. Thus the distance of every vertex $z \in B_{5^{n-i-1}}(x)$ from F_i is at least $5^{n-i} - d(y,z) \geq 5^{n-i} - 2(5^{n-i-1}) \geq 3(5^{n-i-1})$ from F_i. It follows that $B_{5^{n-i-1}}(x)$ and $B_{5^{n-i-1}}(F_i)$ are a distance at least $3(5^{n-i-1}) - 5^{n-i-1} \geq 2(5^{n-i-1}) \geq 2(5^0) = 2$. Thus $B_{5^{n-i-1}}(x)$ and $B_{5^{n-i-1}}(F_i)$ are disjoint.

By hypothesis there is an isomorphism $\beta : B_{5^{n-i-1}}(x) \to B_{5^{n-i-1}}(x')$ with $\beta(x) = x'$ and $B_{5^{n-i-1}}(x')$ disjoint from $\alpha_i(B_{5^{n-i-1}}(F_i))$. Combining the restriction of α_i to $B_{5^{n-i-1}}(F_i)$ with β, we obtain α_{i+1}. $\square$

3. The Almost Sure Theory

Fix an infinite connected structure X of bounded degree satisfying the duplicate substructure property. Let $\mathcal{C}$ be the collection of all structures isomorphic to finite substructures of X. By construction $\mathcal{C}$ is closed under passage to substructures. By the duplicate substructure property of X, $\mathcal{C}$ is closed under disjoint union.

Let $\mathcal{A}$ be a set of representatives for the isomorphism classes of all finite structures, and define sentences σ_F, $F \in \mathcal{A}$, as follows. For $F \in \mathcal{A} \cap \mathcal{C}$, σ_F says that there there is a closed substructure isomorphic to F; for $F \in \mathcal{A} - \mathcal{C}$, σ_F says that there is no substructure isomorphic to F. Define T to be the theory with axioms $\{\sigma_F\}$.

Observe that the disjoint union of $\{F \mid F \in \mathcal{A} \cap \mathcal{C}\}$ is a model of T.

Lemma 11. *The following conditions hold for any model Y of T.*

(1) *Every finite substructure of Y is isomorphic to a closed substructure.*
(2) *For any two finite substructures, there is a finite substructure isomorphic to their disjoint union.*
(3) *The union of all finite closed substructures of Y is a model of T and consists of infinitely many disjoint copies of each finite substructure of X.*

Proof. Item (1) and the first part of (3) hold by construction of T. For (2) observe that as $\mathcal{C}$ is closed under disjoint union, for any $F_1, F_2 \in \mathcal{A} \cap \mathcal{C}$ there is an $F_3 \in \mathcal{A} \cap \mathcal{C}$ isomorphic to the disjoint union of F_1 and F_2. Finally the last part of (3) follows from (1) and (2). □

Lemma 12. *T is complete.*

Proof. It suffices to show that any two models of T are elementarily equivalent. Up to isomorphism the finite substructures of any model of T are the same as those of X. Thus models of T have bounded degree. By Theorem 10 it suffices to show that any two models Y, Y' of T satisfy the disjoint ball extension condition.

Suppose that F is a finite substructure of Y and $B_n(y) \subseteq Y$ is disjoint from F, and F is isomorphic to $F' \subseteq Y'$. $B_n(y)$ is a finite substructure of Y and hence isomorphic to a finite closed substructure $Z' \subseteq Y'$. By Lemma 11 we may assume Z' is disjoint from F'. Let y' be the image of y under this isomorphism mapping $B_n(y)$ to Z'. By Lemma 8, $Z' \subseteq B_n(y')$. As Z' is closed, it follows that $Z' = B_n(y')$. □

Lemma 13. *Each axiom σ_F is almost surely true for finite substructures of X.*

Proof. If σ_F says there is no substructure isomorphic to F, then F is not isomorphic to any substructure of X. Hence σ_F holds for all substructures of every ball in X. In the remaining case σ_F says that there is a closed substructure isomorphic to F. It follows that F is isomorphic to a substructure F_1 of X.

Choose F_1 such that $G_1 = B_1(F_1)$ has maximum possible size, k. This is possible because the vertex degree of $[X]$ is bounded. G_1 has 2^k subsets, one of which supports F_1. Further our choice of F_1 guarantees that if G' is any substructure isomorphic to G_1, then $G' = B_1(F')$ for some substructure F' isomorphic to F. By Lemma 11 there are denumerably many substructures $G_2, G_3, \ldots$ isomorphic to G_1 and disjoint from G_1 and each other. Each G_i is $B_1(F_i)$ for a substructure F_i of G_i isomorphic to F.

Consider balls $B_n(x)$ for some x. It follows from the connectedness of X that for any m, $B = B_n(x)$ will contain at least m of the G_i's if n is large enough. For each $G_i \subseteq B$, the fraction of substructures of B whose restriction to that G_i is not F_i is at most $1 - 2^{-k}$. Thus the fraction whose restriction to some G_i in $B_n(x)$ equals F_i is at least $1 - (1 - 2^{-k})^m$, which is arbitrarily small when m is large enough and hence when n is large enough. Further when the restriction of a substructure of B to G_i is F_i, then because the substructure does not contain any points of $B_1(F_i) - F_i$, F_i is closed in the substructure. □

Now we complete the proofs of Theorems 3 and 5. Let σ be an arbitrary first-order sentence. Since T is complete, it follows that either σ or $\neg\sigma$ is derivable from a finite set of axioms of T. Clearly the conjunction of this finite set of almost surely true sentences is almost surely true for finite substructures of X. It follows that σ or $\neg\sigma$, whichever one is derivable from T, is almost surely true for finite substructures of X. The proof of Theorem 3 is complete.

To prove the first part of Theorem 5 construct X as follows. Let $\{F_i \mid i = 1, 2, \ldots\}$ be a set of representatives of the isomorphism classes of elements of $\mathcal{C}$. It follows from Part (4) of Hypothesis 4 that F_1 lies in a connected structure G_1 which is isomorphic to an element of $\mathcal{C}$. By Parts (3) and (4), for each $i \geq 2$ the disjoint union $F_i \cup G_{i-1}$ lies in a structure G_i isomorphic to an element of $\mathcal{C}$. Define X to be the union of ascending series $G_1 \subset G_2 \subset \cdots$.

Since each G_i is connected, so is X. By Part (2), X has uniformly bounded degree. The construction of X together with Part (1) implies that the finite substructures of X are exactly those in $\mathcal{C}$ up to isomorphism. The duplicate substructure property for X follows from Part 3. Thus X satisfies the conditions of Theorem 3.

Let X be any ambient structure for $\mathcal{C}$. The axioms for the almost sure theory T defined above with respect to X assert that every finite substructure is isomorphic to a substructure of X, and that every finite substructure of X is isomorphic to a closed substructure. As the finite substructures of X are the same as the elements of $\mathcal{C}$ up to isomorphism, the definition of S implies that S is a model of T. Hence the second assertion of Theorem 5 holds.

4. Decidability

In this and subsequent sections we develop our theme further. From now on X is any structure satisfying the hypotheses of Theorem 3 and T is the almost sure theory for finite substructures of X.

Definition 14. X is locally computable if for every natural number n one can effectively find a set of representatives of the isomorphism classes of balls of radius n.

Notice that by hypothesis X is of bounded degree. Thus for any n there are up to isomorphism only a finite number of balls of radius n.

Lemma 15. *T is decidable if and only if X is locally computable.*

Proof. Assume X is locally computable. To prove that T is decidable, it suffices to show that the axioms for T are computable. Indeed if the axioms are computable, then T is recursively enumerable; and because T is complete, enumeration of T produces either σ or $\neg\sigma$ for every sentence σ. Thus T is decidable.

The axioms of T are computable if we can decide for any finite structure F whether or not F is isomorphic to a substructure of X. If $[F]$ is connected, then any isomorphic substructure F_1 of X must lie in some ball of radius at most equal to the size of F. By hypothesis we can examine the finitely many representatives of the isomorphism classes of these balls to check if F is isomorphic to a substructure of X.

If $[F]$ is not connected, we can check as above if its connected substructures are isomorphic to substructures of X. If some connected substructure fails the test, then F cannot be a substructure of X. If they all pass, then by the duplicate substructure property they can be embedded into X in such a way that they are a distance at least 2 from each other. It follows that their union is isomorphic to F.

To prove the converse suppose that T is decidable. For any finite structure F one can write down a formula which says that there is an element u for which the ball of radius n around u is isomorphic to F. Hence one can decide whether or not F is isomorphic to a ball of radius n in X. As X has bounded degree, only finitely many F's have to be checked in order to generate a complete list of isomorphism types of balls of radius n in X. □

Corollary 16. *If X is the Cayley diagram of a finitely generated group G, then T is decidable if and only if X has solvable word problem*

Proof. Recall that there is one binary predicate for each generator of G. If the word problem is decidable, one can construct the ball of radius n around the identity. Since all balls of radius n are isomorphic, X is locally computable. Conversely if X is locally computable, T is decidable by Lemma 15. For any word w in the generators of G, the binary relation $R_w(x, y)$ which holds when there is a path with label w from x to y in X is definable. Thus we can decide if $\exists x\, R_w(x, x)$ is true, i.e., if w defines the identity in G. □

5. Random Substructures

Let X be a structure satisfying the hypotheses of Theorem 3. For a fixed p, $0 < p < 1$, we may imagine generating a random substructure of X by deleting each element of X with probability $1-p$. The random substructure is the one supported by all the remaining elements. We will show that almost all random substructures are elementarily equivalent but not necessarily isomorphic.

A more precise definition of random substructures of X is obtained by first defining a measure on cones. For each pair, S, T, of disjoint finite subsets of elements of X, the corresponding cone consists of all subsets of elements which include S and avoid T. The measure of this cone is defined to be $p^{|S|}q^{|T|}$, where $|S|$ and $|T|$ are the cardinalities of S and T respectively, and $q = 1-p$. By a well known theorem of Carathéodory the measure on cones extends uniquely to a probability measure, μ, on the σ–algebra generated by the cones.

Lemma 17. *Let F be a finite substructure of X. With probability 1 a random substructure of X contains a closed substructure isomorphic to F.*

Proof. The proof is just a modification of the proof of Lemma 13. Fix F, and pick a substructure F_1 of X which is isomorphic to F and for which $B_1(F_1)$ is maximal. By the duplicate substructure property X has denumerably many pairwise disjoint and isomorphic substructures $H_1 = B_1(F_1), H_2, H_3, \ldots$. For any i there is an isomorphism $\alpha_i : H_1 \to H_i$ carrying F_1 to $F_i = \alpha(F_1)$. By Lemma 8 $H_i \subseteq B_1(F_i)$. By maximality of $B_1(F_1)$ we have $H_i = B_1(F_i)$.

Let Y be a random substructure of X. If $Y \cap B_1(F_i) = F_i$, then Y contains F_i as a closed substructure. By disjointness the denumerably many events $Y \cap B_1(F_i) \neq F_i$ are independent. As each of these events has the same probability, and that probability is less than 1, we conclude that the probability of a random graph containing at least one of the F_i's as a closed substructure is 1. □

Now define X^* to be the structure consisting of the disjoint union of a denumerable number of copies of each finite substructure of X. It is clear that X^* is a model of T.

Lemma 18. *With probability 1 a random substructure of X contains a closed substructure isomorphic to X^*.*

Proof. The duplicate substructure property and Lemma 17 together guarantee that the set of substructures with the desired property is the intersection of a countable number of sets of measure 1. □

Theorem 19. *With probability* 1 *a random substructure of* X *is a model of* T*. In particular, almost all random substructures of* X *are elementarily equivalent.*

Proof. By Lemma 18 it suffices to show that if a substructure X_0 of X contains a union of connected components isomorphic to X^*, then X_0 is elementarily equivalent to X^*. The argument used in the proof of Lemma 12 applies. □

6. Random Subgraphs of Trees

In this section we sharpen the results of the preceding section in the case of random subtrees of trees. Let Γ_k, $k \geq 1$, be the full k-ary tree, that is, the tree with one vertex, the root, of degree k and all others of degree $k+1$. As we noted earlier, Theorem 3 applies to Γ_k. We maintain the following notation from Section 5: p is a number strictly between 0 and 1, $q = 1 - p$, and μ is the corresponding measure on subgraphs of Γ_k.

A descending path in Γ_k is one which starts at any vertex and continues away from the root. Let p_n be the probability that a random subgraph admits no descending path of length n starting at a fixed vertex v. A moment's thought shows that $p_0 = q$, and $p_{n+1} = q + pp_n^k$. In particular p_n is independent of the choice of v. The probability that a random subtree contains an infinite descending path starting at a particular vertex v is $1 - \lim_{n\to\infty} p_n$.

Lemma 20. *The probability that a random subtree contains an infinite descending path starting at a particular vertex* v *is* 0 *if* $p \leq 1/k$ *and strictly between* 0 *and* 1 *otherwise.*

Proof. Define $f(x) = q + px^k$. Observe that $f(0) = q = p_0$, $f(f(0)) = p_1$, etc. Further f maps the unit interval to itself and is strictly increasing on that interval. Thus $p_0, p_1, p_2, \dots$ is an increasing bounded sequence which converges to a fixed point of f. When $k = 1$, f is linear with a single fixed point (on the unit interval) at $x = 1$. Otherwise f is concave up and has a single fixed point at $x = 1$ if $p \leq 1/k$ and two fixed points if $p > 1/k$. Let x_0 be the least fixed point of f on the unit interval. Since f is increasing, $0 \leq x_0$ implies that every point in the forward orbit of 0 under f is no greater than x_0. Thus $p_0, p_1, p_2, \dots$ converges to x_0. As $0 < q \leq x_0$, we are done. □

We observe that the statement that there is an infinite descending path starting at the root of a full k-ary tree can be formulated in monadic second-order logic, in fact in existential monadic second-order logic. Thus we have evidence that Theorem 19 does not extend to this more powerful logic.

References

[1] I. Benjamini and O. Schramm, Percolation beyond Z^d, many questions and a few answers. Electron. Comm. Probab., 1 (no. 8) , 71-82 (electronic), 1996.

[2] ———, Recent progress on percolation beyond Z^d, `http://research.microsoft.com/~schramm/pyondrep/`.

[3] A. Blass and Y. Gurevich, Zero-One Laws: Thesauri and Parametric Conditions. Bull. Euro. Assoc. for Theor. Computer Science, # 91 (February 2007).

[4] K. J. Compton, A logical approach to asymptotic combinatorics I. First order properties, Advances in Math., Volume 65, 1987, 65-96.

[5] K. J. Compton, 0–1 laws in logic and combinatorics, in I. Rival (ed.) NATO Advanced Study Institute on Algorithms and Order, D. Reidel, 1989, 353–383.

[6] H-D Ebbinghaus and J. Flum, *Finite Model Theory*, Springer, 1995.

[7] R. Fagin, Probabilities on finite models, Journal of Symbolic Logic, Volume 41, 1976, 50-58, 1976.

[8] H. Gaifman, On local and nonlocal properties, Proceedings of the Herbrand symposium (Marseilles, 1981), 105–135, Stud. Logic Found. Math., Volume 107, North-Holland, Amsterdam, 1982.

[9] Y. Glebski, V. Kogan, M. I. Liogonkij and V.A. Talanov, The extent and degree of satisfiability of formulas of the restricted predicate calculus, Kibernetika, Volume 2, 1969, 17–27.

[10] Y. Gurevich *Zero-One laws*, The logic in computer science column, Current Trends in Theoretical Computer Science", Eds. G. Rozenberg and A. Salomaa, World Scientific, Series in Computer Science, Volume 40, 1993.

[11] R. Jajcay and J. Siráň, A construction of vertex-transitive non-Cayley graphs, Austalas. J. Combin., Volume 10, 1994, 105–114.

[12] Ph. G. Kolaitis, H. J. Promel, and B. L. Rotschild, K_{l+1}–free graphs: asymptotic structure and a $0-1$ law, Trans. Amer. Math. Soc., Volume 303, 1987, 637-671.

[13] P. Winkler, Random structures and zero-one laws, in N. W. Sauer et al. (eds.), *Finite and Infinite Combinatorics in Sets and Logic*, NATO Advanced Science Institutes Series, Kluver, 1993, 399–420.

Department of Mathematical Sciences, Stevens Institute of Technology, Hoboken, NJ 07030

E-mail address: `rgilman@stevens.edu`

Microsoft Research, One Microsoft Way, Redmond, WA 98052

E-mail address: `gurevich@microsoft.com`

Department of Mathematics and Statistics, McGill University, Montreal, Quebec H3A 2K6

E-mail address: `alexeim@math.mcgill.ca`