

SPHERES IN THE CURVE COMPLEX

SPENCER DOWDALL, MOON DUCHIN, AND HOWARD MASUR

1. INTRODUCTION

The curve graph (or curve complex) $\mathcal{C}(S)$ associated to a surface S of finite type is a locally infinite combinatorial object that encodes topological information about the surface through intersection patterns of simple closed curves. It is known to be δ -hyperbolic [5], a property that is often described by saying that a space is “coarsely a tree.” To be precise, there exists δ such that for any geodesic triangle, each side is in the δ -neighborhood of the union of the other two sides. In this note, we will investigate the finer metric properties of the curve graph by considering the geometry of spheres; specifically, we will study the average distance between pairs of points on $\mathcal{S}_r(\alpha)$, the sphere of radius r centered at α . To make sense of the idea of averaging, we will develop a definition of null and generic sets in §3 that is compatible with the topological structure of the curve graph.

Given a family of probability measures μ_r on the spheres $\mathcal{S}_r(x)$ in a metric space (X, d) , let $E(X) = E(X, x, d, \{\mu_r\})$ be the normalized average distance between points on large spheres:

$$E(X) := \lim_{r \rightarrow \infty} \frac{1}{r} \int_{\mathcal{S}_r(x) \times \mathcal{S}_r(x)} d(y, z) d\mu_r(y) d\mu_r(z),$$

if the limit exists. For finitely generated groups with their Cayley graphs, or more generally for locally finite graphs, we can study averages with respect to counting measure because the spheres are finite sets. It is shown in [2] that non-elementary hyperbolic groups all have $E(G, S) = 2$ for any finite generating set S ; this is also the case in the hyperbolic space \mathbb{H}^n of any dimension endowed with the natural measure on spheres. By contrast, $E(\mathbb{R}^n) < 2$ and $E(\mathbb{Z}^n, S) < 2$ for all n and S , with nontrivial dependence on S . In particular this shows that $E(X)$ varies under quasi-isometry, so it is a fine and not a coarse asymptotic statistic. Note that δ -hyperbolicity itself (without an assumption of homogeneity) does not imply $E = 2$: even for locally finite trees, one can get any value $0 \leq E \leq 2$.

In [1], we show that for Teichmüller space with the Teichmüller metric and various visual measures on spheres, $E(\mathcal{T}(S)) = 2$. Here, we show something stronger for the curve graph.

Theorem. *With respect to our notion of genericity for spheres in the curve graph, almost every pair of points on $\mathcal{S}_r(\alpha)$ is at distance exactly $2r$ apart.*

This holds for every r and is certainly stronger than saying that the average distance is asymptotic to $2r$, so we can write $E(\mathcal{C}(S)) = 2$. That is, suppose we start with α and a pair of curves β and γ such that the shortest path in the curve graph from α to either β or γ has length r . Then, almost surely, there is no more efficient way to connect them with each other than to travel through α . This

result tells us that, even though the space $\mathcal{C}(S)$ is far from uniquely geodesic, the concatenation of two geodesic segments of length r that share a common endpoint is almost always itself geodesic. In this sense the curve graph is “even more hyperbolic than a tree.”

Of course, the meaningfulness of this result depends on the notion of genericity. Lustig and Moriah [4] have introduced a very natural definition for genericity in $\mathcal{C}(S)$ which uses the topology and measure class of $\mathcal{PML}(S)$. We identify the sphere of radius 1 in $\mathcal{C}(S)$ with a lower-complexity curve complex, so that genericity can be defined in the same way. We then extend to spheres of larger radius in a “visual” manner; see Definition 6. While the Lustig–Moriah definition gives content to statements about typical curves on S , our notion of genericity on spheres enables us to talk about typical properties of high-distance curves on S .

2. BACKGROUND

We fix a topological surface $S = S_{g,n}$ with genus g and n punctures, and let $h = 6g - 6 + 2n$. Let \mathcal{S} be the set of homotopy classes of essential nonperipheral simple closed curves on S . From now on, a *curve* will mean an element of \mathcal{S} . Next we define the *curve graph* $\mathcal{C}(S)$: The vertex set of $\mathcal{C}(S)$ is \mathcal{S} . In the case that $h > 2$, two curves are joined by an edge if they are disjointly realizable. In the case of $S_{1,1}$ we join two vertices if the curves intersect once, and in the case of $S_{0,4}$ two vertices are joined by an edge if the curves intersect twice. In each of these cases, $\mathcal{C}(S)$ is endowed with the standard path metric, denoted $d_S(\alpha, \beta)$.

For $\alpha \in \mathcal{S}$, we write S_α to denote the lower-complexity punctured (possibly disconnected) surface obtained by cutting open S along α . Note that $\mathcal{C}(S_\alpha)$ can be realized as the subgraph of $\mathcal{C}(S)$ consisting of neighbors of α —that is, it is identified with the sphere $\mathcal{S}_1(\alpha) \subset \mathcal{C}(S)$.

Recall that a *measured lamination* on S , given a hyperbolic structure, is a foliation of a closed subset of S by geodesics, together with a measure on transversals that is invariant under holonomy along the leaves of the lamination. We will use $\mathcal{ML}(S)$ to denote the space of *measured laminations* on S . Let $\text{Mod}(S) := \pi_0(\text{Diff}^+(S))$ denote the *mapping class group* of S ; it acts on $\mathcal{C}(S)$ and on $\mathcal{ML}(S)$. The latter has a natural $\text{Mod}(S)$ -invariant measure μ . (This is the Lebesgue measure associated to the piecewise linear structure induced on $\mathcal{ML}(S)$ by train track neighborhoods.) The space of *projective measured laminations* $\mathcal{PML}(S)$ is obtained by identifying laminations whose transverse measures differ only by a scalar multiple; it is endowed with the topology of a sphere of dimension $h - 1$ and it inherits a natural $\text{Mod}(S)$ -invariant measure class, which we will again denote by μ .

A *train track* on S is a finite collection of disjointly embedded arcs, called *branches*, meeting at vertices, called *switches*, such that the branches are C^1 away from switches and have well-defined tangents at the switches. (There are also non-degeneracy conditions for switches and topological conditions on the complement; see [7] for details.) Train tracks exist on every surface with $h > 0$, and we say that a lamination F is *carried* by a train track τ if there is a map $\phi : S \rightarrow S$ isotopic to the identity with $\phi(F) \subset \tau$. Via this carrying relation, measured laminations that are carried by τ correspond to choices of weights on the branches of τ that satisfy switch conditions.

A connected proper subsurface V of S is *essential* if all components of ∂V are essential; i.e., they are homotopically nontrivial and not isotopic to a puncture.

Definition 1. Consider a non-annular essential subsurface V . The *subsurface projection* π_V is a coarsely well-defined map $\mathcal{C}(S) \rightarrow \mathcal{C}(V)$ defined as follows. Realize $\beta \in \mathcal{S}$ and ∂V as geodesics (in any hyperbolic metric on S). If $\beta \subset V$, let $\pi_V(\beta) = \beta$. If β is disjoint from V , then $\pi_V(\beta)$ is undefined. Otherwise, $\beta \cap V$ is a disjoint union of finitely many homotopy classes of arcs with endpoints on ∂V , and we obtain $\pi_V(\beta)$ by choosing any arc and performing a surgery along ∂V to create a simple closed curve contained in V . All possible ways to do this form a non-empty subset of the curve complex $\mathcal{C}(V)$ with uniformly bounded diameter. We can denote by $d_V(\alpha, \beta)$ the diameter in $\mathcal{C}(V)$ of $\pi_V(\alpha) \cup \pi_V(\beta)$. If α and β are disjoint and both intersect V then $d_V(\alpha, \beta) \leq 4$.

There is a well-defined inclusion $\mathcal{S} \hookrightarrow \mathcal{PML}(S)$ whose image is dense and we will identify \mathcal{S} with its image under that map. The *supporting subsurface* of a lamination is the subsurface filled by F . We will denote the geometric intersection number on \mathcal{S} by $i(\alpha, \beta)$, and we recall that it has a well-defined extension to $\mathcal{ML}(S)$. On $\mathcal{PML}(S)$, we can thus talk about whether or not $i(F, G) = 0$. Then $\mathcal{PML}(S_\alpha)$ can be identified with the subset of $\mathcal{PML}(S)$ consisting of those laminations F for which $i(F, \alpha) = 0$. If α is nonseparating, then S_α has complexity $h - 2$; if α is separating, then S_α consists of two surfaces with complexity $h_1 + h_2 = h - 2$. In that case we consider $\mathcal{PML}(S_\alpha)$ as a product of the corresponding spaces for the two components. In either case we see that $\mathcal{PML}(S_\alpha)$ has positive codimension in $\mathcal{PML}(S)$.

Definition 2. Given a group G that acts on a space X with Borel algebra $\mathcal{B}(X)$, a G -invariant mean on X is a function $\sigma : \mathcal{B}(X) \rightarrow [0, 1]$ such that

- $\sigma(\emptyset) = 0$ and $\sigma(X) = 1$;
- if $B_1, \dots, B_N \in \mathcal{B}(X)$ are pairwise disjoint then $\sigma(\bigsqcup B_i) = \sum \sigma(B_i)$; and
- $\sigma(B) = \sigma(gB)$ for all $B \in \mathcal{B}(X)$ and all $g \in G$.

Note that invariant means are only required to be finitely additive, while measures must be countably additive.

Proposition 3. *There is no $\text{Mod}(S)$ -invariant mean on \mathcal{S} or on \mathcal{PML} .*

Proof. One can choose a finite number $T = T(S)$ of train tracks τ_1, \dots, τ_T such that (1) every curve is carried by τ_i for some i and (2) the set of curves that are carried by τ_i and have positive weights on all branches of τ_i is disjoint from the corresponding set of curves carried by τ_j . Pairs of attracting and repelling fixed points of pseudo-Anosov diffeomorphisms are dense in $\mathcal{PML} \times \mathcal{PML}$.

If there is an invariant mean σ on \mathcal{S} , for some i the set B_i of curves carried by τ_i satisfies $a_i = \sigma(B_i) > 0$. Choose N such that $N > 1/a_i$ for such an index i . Find N pseudo-Anosovs ψ_k with distinct attracting fixed points at laminations carried by τ_i and repelling fixed points carried by some τ_j with $j \neq i$. Disjoint neighborhoods around these attracting fixed points may be chosen such that all curves in each neighborhood are carried by τ_i and have positive weights on all branches of τ_i ; likewise for the repelling fixed points. Raising each ψ_k to a high enough power m_k we can conclude that the curves carried by $\psi_k^{m_k}(\tau_i)$ are disjoint from the set of curves carried by $\psi_l^{m_l}(\tau_i)$ for $l \neq k$. By invariance, each of those sets has σ -mean a_i . Adding N of these we find that $\sigma(\mathcal{S}) > 1$, a contradiction. \square

Corollary 4. *There is no $\text{Mod}(S)$ -invariant Borel probability measure μ on \mathcal{PML} .*

3. GENERICITY IN THE CURVE COMPLEX

In the paper [4], Lustig–Moriah give the following notion of genericity.

Definition 5. Let X be a topological space, provided with a Borel measure or measure class μ . Let $Y \subset X$ be a (possibly countable) subset with $\mu(\overline{Y}) \neq 0$. Then the set $A \subset Y$ is called *generic in Y* (or simply generic, if $Y = X$) if $\mu(\overline{Y \setminus A}) = 0$. (Here closures are taken in X .) On the other hand A is called *null in Y* if $\mu(\overline{A}) = 0$.

We can extend the definitions to products as follows. Given $E \subset Y \times Y$ and $a \in Y$ let $E(a) := \{b \in Y : (a, b) \in E \text{ or } (b, a) \in E\}$.

Definition 6. E is *null in $Y \times Y$* if $\{a \in Y : E(a) \text{ not null in } Y\}$ is null.

This definition for products corresponds to Fubini’s theorem: the set of points with non-null fibers must be null.

We will focus on the case that $X = \mathcal{PML}(S)$ for any surface S and $Y = \mathcal{S}(S)$ is the set of simple closed curves. Several examples and observations can be made immediately to illustrate that this notion is topologically interesting.

- Nullness in \mathcal{S} is preserved by: acting by $\text{Mod}(S)$, passing to subsets, and finite unions. A set is null if and only if its complement is generic.
- The entire set \mathcal{S} is generic in \mathcal{S} , and being generic in \mathcal{S} implies denseness in \mathcal{PML} . (Because if A misses an open set in \mathcal{PML} , then the closure of its complement has positive measure.)
- There are natural subsets of \mathcal{S} that are neither null nor generic. For instance, suppose that $g \geq 2$, so that \mathcal{S} has a nontrivial partition into *separating* and *nonseparating* curves. Each of these subsets is dense in \mathcal{PML} , so neither can be null or generic.
- Our basic example of a null set in \mathcal{S} is the set of all curves disjoint from some α , which is a copy of $\mathcal{S}(S_\alpha)$ sitting inside $\mathcal{S}(S)$. Its closure in $\mathcal{PML}(S)$ consists of those laminations giving zero weight to α , which can be identified with a copy of $\mathcal{PML}(S_\alpha)$; since this has lower dimension, it has measure zero.

Because nullness is not preserved by countable unions, the following proposition is less obvious. It is proved by Lustig–Moriah [4, Cor 5.3] using techniques from handlebodies, but we include a proof for completeness. Our proof is similar to well-known arguments, due to Luo and to Kobayashi [3], showing that the complex of curves has infinite diameter (see the remarks following Proposition 4.6 of [5]).

Proposition 7. *Any bounded-diameter subset of the curve graph $\mathcal{C}(S)$ is a null subset of \mathcal{S} .*

We first prove that chains of disjoint laminations between two curves can be realized with curves.

Definition 8. Given a pair of laminations $F, G \in \mathcal{PML}$, define their *intersection distance* $d_i(F, G)$ to be the smallest n such that there exist $F = G_0, G_1, \dots, G_n = G$ with $i(G_j, G_{j+1}) = 0$.

(We note that $d_i(F, G)$ can be infinite; this occurs if at least one of them is filling and they are topologically distinct.)

Lemma 9. $d_i(\alpha, \beta) = d_S(\alpha, \beta)$ for all $\alpha, \beta \in \mathcal{S}$.

Proof. Since simple closed curves are laminations, it is immediate that $d_i(\alpha, \beta) \leq d_S(\alpha, \beta)$.

For the other direction, we can assume that the laminations are not filling. Then given a lamination $G \in \mathcal{PML}$, let us write Y_G for its supporting subsurface. Now $i(F, G) = 0 \implies i(F, \partial Y_G) = 0 \implies i(\partial Y_F, \partial Y_G) = 0$. But then given a minimal-length disjointness path

$$\alpha - G_1 - G_2 - \cdots - G_{n-1} - \beta,$$

we can realize it by simple closed curves by replacing each G_i with ∂Y_{G_i} . \square

Proof of Proposition 7. First given $\alpha \in \mathcal{S}$, let $\mathcal{S}_r(\alpha) \subset \mathcal{C}$ denote the sphere of radius r centered at α . It is enough to prove the Proposition for each $\mathcal{S}_r(\alpha)$. For the ball of radius 1, the statement follows since each β satisfies $i(\alpha, \beta) = 0$ and, as we saw above, the set of such β has measure 0 closure in $\mathcal{PML}(S)$. Notice that this closure $\overline{\mathcal{S}_1(\alpha)}$, which we identify with a copy of $\mathcal{PML}(S_\alpha)$, is exactly the set of laminations $F \in \mathcal{PML}(S)$ for which $i(\alpha, F) = 0$.

Now we consider the closure of the sphere of radius r and suppose $G \in \overline{\mathcal{S}_r(\alpha)}$. Then $G = \lim_{m \rightarrow \infty} \beta_m^r$ where $\beta_m^r \in \mathcal{S}_r(\alpha)$. Then for each m there is a path $\alpha, \beta_m^1, \dots, \beta_m^r$ in \mathcal{C} , where $\beta_m^j \in \mathcal{S}_j(\alpha)$. Passing to subsequences we can assume that for each j the sequence β_m^j converges to some $G_j \in \overline{\mathcal{S}_j(\alpha)}$ with $G_r = G$. Furthermore since $i(\beta_m^j, \beta_m^{j+1}) = 0$ it follows that

$$i(G_j, G_{j+1}) = 0.$$

Replacing each G_j with ∂Y_{G_j} , as in the proof of Lemma 9, we see that $i(G, \gamma) = 0$ for some $\gamma \in \mathcal{S}_{r-1}(\alpha)$. But then $G \in \overline{\mathcal{S}_1(\gamma)}$. Thus

$$\overline{\mathcal{S}_r(\alpha)} \subset \bigcup_{\gamma \in \mathcal{S}_{r-1}(\alpha)} \overline{\mathcal{S}_1(\gamma)}.$$

Thus $\overline{\mathcal{S}_r(\alpha)}$ is a countable union of measure-zero subsets of $\mathcal{PML}(S)$, hence has zero measure itself. Therefore $\mathcal{S}_r(\alpha)$ is a null set by definition. \square

We wish to define what it means for a property to be generic for pairs of points in $\mathcal{S}_r(\alpha)$. Although $\mathcal{S}_r(\alpha)$ is contained in \mathcal{S} , the Lustig–Moriah definition of genericity in \mathcal{S} does not apply because the set $\mathcal{S}_r(\alpha)$ is itself null in \mathcal{S} . Nonetheless, $\overline{\mathcal{S}_1(\alpha)} = \mathcal{PML}(S_\alpha)$ is a topological sphere in its own right, and thus has its own natural measure class. Therefore, we may define a subset E of $\mathcal{S}_1(\alpha)$ to be *null* if \overline{E} has measure zero in $\mathcal{PML}(S_\alpha)$; this respects the topology of $\mathcal{PML}(S_\alpha)$ sitting inside of $\mathcal{PML}(S)$. We extend this notion of nullness to subsets $E \subset \mathcal{S}_r(\alpha)$ of larger spheres in a “visual” manner by considering the set of points on the sphere of radius 1 that are metrically between the center and E —these are the points of $\mathcal{S}_1(\alpha)$ that “see” the set E .

Definition 10. E is *null* in $\mathcal{S}_r(\alpha)$ if $E_1 := \{\gamma \in \mathcal{S}_1(\alpha) : E \cap \mathcal{S}_{r-1}(\gamma) \neq \emptyset\}$ is null in $\mathcal{S}_1(\alpha) \hookrightarrow \mathcal{PML}(S_\alpha)$.

Remark 11. The definition given above is the most restrictive notion of nullness that makes use of $\mathcal{S}_1(\alpha)$ as a visual sphere (i.e., that treats the 1–sphere as the sphere of directions). Another possible definition, also natural from the point of view of Fubini’s theorem, would be inductive: suppose nullness has been defined

for spheres of radius $1, \dots, r-1$. Instead of E_1 , the full set of points that see $E \subset \mathcal{S}_r(\alpha)$, we form the smaller set

$$E'_1 = \{\gamma \in \mathcal{S}_1(\alpha) : E \cap \mathcal{S}_{r-1}(\gamma) \text{ is not null in } \mathcal{S}_{r-1}(\gamma)\}.$$

Then we could declare $E \subset \mathcal{S}_r(\alpha)$ to be null in $\mathcal{S}_r(\alpha)$ if E'_1 is null in $\mathcal{S}_1(\alpha)$, completing the inductive definition.

Example 12. To get a feeling for these definitions, consider the examples of \mathbb{R}^2 with the Euclidean metric or ℓ^1 metric, with the Lebesgue measure class on the sphere of radius 1 in each case. To accord with geometric intuition, we expect arcs to be non-null and points to be null.

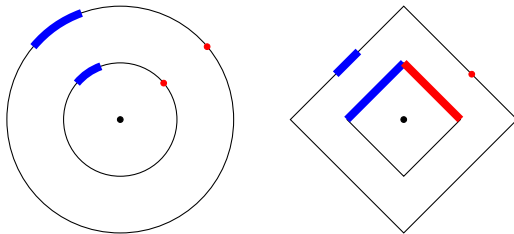
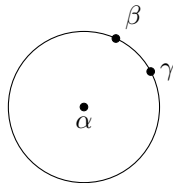


FIGURE 1. In each metric, an arc E and a point E are shown on the sphere of radius two together with the associated E_1 for each.

In the Euclidean metric, if E is an arc on the sphere of radius r , then E_1 is also an arc but E'_1 is empty. If E is a point, then E_1 is a point while E'_1 is again empty. In the ℓ^1 metric, if E is a nontrivial arc on the sphere of radius r , then E_1 is a nontrivial arc, and so is E'_1 . In this setting, however, points in nonaxial directions have a large E_1 but an empty E'_1 .

This means that our visual definition of nullness works intuitively in the ℓ^2 case (points are null but arcs are not), but less so in the ℓ^1 case (where even points are typically non-null). The weaker, inductively defined, notion of nullness makes even arcs null in Euclidean space, but on the other hand behaves intuitively on ℓ^1 . This suggests that the visual definition of nullness is better adapted to capturing the geometry of spheres in certain spaces, while the inductive definition would be better adapted to others. However, being null in our sense implies nullness in the weaker sense.

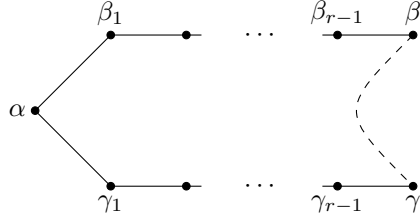
Returning to the curve graph: consider distinct curves $\beta, \gamma \in \mathcal{S}_1(\alpha)$. Clearly $d_S(\beta, \gamma)$ is either 1 or 2. We can easily see that such pairs generically have distance 2 because since α and β are disjoint, the set of $\gamma \in \mathcal{S}_1(\alpha)$ for which $d_S(\beta, \gamma) = 1$ is contained in $\mathcal{PML}(\mathcal{S}_{\alpha, \beta})$, which has codimension at least two in $\mathcal{PML}(\mathcal{S}_\alpha)$.



Our main result is that given any limit on their length, paths connecting points on the sphere “almost surely” pass through the center.

Theorem 13 (Avoiding the center). *For a surface S with $h \geq 4$, consider a point $\alpha \in \mathcal{C}(S)$. For $K > 0$, let $P_r^K(\alpha) \subseteq \mathcal{S}_r(\alpha) \times \mathcal{S}_r(\alpha)$ consist of those pairs (β, γ) that are connected by some path of length $\leq K$ that does not go through α . Then for any K and r , the set $P_r^K(\alpha)$ is null.*

Proof. For any pair $(\beta, \gamma) \in P_r^K(\alpha)$, there is a path $\beta = \delta_0, \delta_1, \dots, \delta_k = \gamma$ in $\mathcal{C}(S)$ with $k \leq K$, and $\delta_i \neq \alpha$ for each i . Two successive curves δ_i and δ_{i+1} , since they are disjoint and intersect S_α , project to nonempty sets in $\mathcal{C}(S_\alpha)$ whose distance from each other is at most 4; thus $d_{S_\alpha}(\beta, \gamma) \leq 4K$.



Let β_1, γ_1 be any closest points on $\mathcal{S}_1(\alpha)$ to β, γ , respectively. Since we can join γ, γ_1 by a geodesic in $\mathcal{C}(S)$ that misses α , there is a constant $M = M(S)$ coming from Masur-Minsky [6, Thm 3.1] such that $d_{S_\alpha}(\gamma, \gamma_1) \leq M$. By the triangle inequality,

$$d_{S_\alpha}(\beta, \gamma_1) \leq M + 4K.$$

For each $\beta \in \mathcal{S}_r(\alpha)$, let $E(\beta) = \{\gamma : (\beta, \gamma) \in P_r^K(\alpha)\}$ and then consider the corresponding $E_1(\beta) = \{\gamma_1 \in \mathcal{S}_1(\alpha) : E(\beta) \cap \mathcal{S}_{r-1}(\gamma_1) \neq \emptyset\}$.

We have shown that $E_1(\beta)$ has diameter at most $2M + 8K$ by the triangle inequality and is therefore null in $\mathcal{S}_1(\alpha)$ by Proposition 7. Thus $E(\beta)$ is null for all β , so $P_r^K(\alpha)$ is null. \square

Corollary 14 (Statistical hyperbolicity). *Consider pairs of points $\beta, \gamma \in \mathcal{S}_r(\alpha)$ such that $d_S(\beta, \gamma) < 2r$. The set of all such pairs is null. So for a generic pair of points on $\mathcal{S}_r(\alpha)$, the distance is exactly $2r$.*

In closing, we note that there is a coarsely well-defined, coarsely Lipschitz map from Teichmüller space to the curve complex given by taking a short marking at every point. With respect to natural measures on $\mathcal{T}(S)$, generic geodesics make definite progress in the curve graph, as we show in [1]. Thus we can see a very loose heuristic for calculating average distances in $\mathcal{T}(S)$ by projecting to $\mathcal{C}(S)$ and appealing to the phenomena that we have just demonstrated in the curve graph. Obtaining the precise asymptotics needed for statistical hyperbolicity, as we do in [1], takes much more work.

REFERENCES

[1] S. Dowdall, M. Duchin, and H. Masur, *Statistical hyperbolicity in Teichmüller space*, arXiv:1108.5416
 [2] M. Duchin, S. Lelièvre, and C. Mooney, *Statistical hyperbolicity in groups*, Alg. Geom. Top. 12 (2012), 1–18.
 [3] T. Kobayashi, *Heights of simple loops and pseudo-Anosov homeomorphisms*. Braids (Santa Cruz, CA, 1986) 327–338. Amer. Math. Soc., Providence, RI, 1988.
 [4] M. Lustig and Y. Moriah, *Horizontal Dehn surgery and genericity in the curve complex*. Journal of Topology 3 (2010), 691–712.

- [5] H. Masur and Y. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. 138 (1999), no. 1, 103–149.
- [6] H. Masur and Y. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*. GAFA 10 (2000), no. 4, 902–974.
- [7] R.C. Penner with J.L. Harer, *Combinatorics of train tracks*. Annals of Mathematics Studies, 125. Princeton University Press, Princeton, NJ, 1992.