

TAKING THE HIGH-EDGE ROUTE OF RANK-3 OUTER SPACE

DAMARA GAGNIER, CATHERINE PFAFF

ABSTRACT. Principal fully irreducible outer automorphisms were introduced in [AKKP19] to emulate principal pseudo-Anosov surface homeomorphisms, i.e. those whose attracting and repelling invariant foliations have only 3-pronged singularities. It is proved in [AKKP19] that these outer automorphisms semi-mimic principal pseudo-Anosov surface homeomorphisms in important ways. We prove here which rank-3 graphs carry train track representatives of principal fully irreducible outer automorphisms. As a corollary, one obtains which simplices of rank-3 Culler-Vogtmann Outer space principal axes pass through.


1. INTRODUCTION

Let $\text{Out}(F_r)$ denote the outer automorphism group of a fixed free group of rank $r \geq 3$. This group's study has fundamental contributions dating back to Nielsen, starting in 1915, and Whitehead in the 1930's to 1950's. Recent substantial progress has utilized that, just as the mapping class group of a surface encodes all (equivalence classes of) homeomorphisms of the surface, $\text{Out}(F_r)$ encodes all (equivalence classes of) homotopy equivalences of graphs of a fixed rank. In particular, this provides a powerful analogue of Teichmüller space, namely Outer space (denoted CV_r), the deformation space of marked metric graphs. This correspondence further gives rise to related analogues of the train track theory.

Pseudo-Anosov surface homeomorphisms are the generic (in a random walk sense, see [Riv08, KM98, CM15, Mah11, Sis18, Hor16, MS17]) and dynamically minimal elements of a mapping class group. The $\text{Out}(F_r)$ -analogues are fully irreducible outer automorphisms. Various methods for constructing fully irreducible outer automorphisms with an assortment of properties can be found, for example, in [CP10], [Pfa13], [Pfa15a], and [CL15]. We focus on “principal” fully irreducible outer automorphisms, as introduced in [AKKP19] to emulate principal pseudo-Anosov surface homeomorphisms. Principal pseudo-Anosov homeomorphisms were proved generic (in a random-walk sense) in [GM17]. Principal structures in the pseudo-Anosov and Teichmüller space situation play an important role. For each genus, there is a partial ordering on the set of singularity structures, and the 3-pronged (i.e. principal) structure is the unique maximal element. This partial ordering has dynamical significance in that, in the total geodesic flow of the moduli space, the closure of a stratum associated to one singularity structure contains the closure of a lower stratum associated to a second singularity structure if and only if the first structure is greater than the second in the partial ordering. Principal pseudo-Anosovs homeomorphisms play a fundamental role in the work of Masur, for example, in [Mas82] The existence of at least one principal fully irreducible outer automorphism in each rank was proved in [AKKP19].

The authors acknowledge support from NSERC (the Natural Sciences and Engineering Research Council of Canada) and the Queen's University Department of Mathematics and Statistics.

Main Theorem A. The graphs carrying train track representatives of principal fully irreducible outer automorphisms in $Out(F_3)$ are precisely:



That is, each of the graphs of CV_3 where either all vertices are valence-3 or all vertices are valence-3 except that one is valence-4 carries a train track representative of a principal fully irreducible outer automorphism in $Out(F_3)$.¹

A fully irreducible outer automorphism in $Out(F_r)$ (more specifically Stallings fold decompositions of train track representatives of the outer automorphism, please see Definition 7) determines geodesics (called axes) in CV_r . In [AKKP19], it is shown that the axes of principal fully irreducible outer automorphisms share a certain “stability” property with principal pseudo-Anosov axes in Teichmüller space. This stability property is then used in [KMPT22b] and [KMPT22a] to not only prove which outer automorphisms are random walk generic, but to understand properties held by a typical (with respect to the harmonic measure) tree in the boundary of Outer space.

Outer space is a simplicial complex minus some faces: it has an open simplex for each marked graph (obtained by varying the lengths of the edges of that graph). The faces of a simplex are obtained by collapsing the edges of a forest (making their lengths zero). It is known (via [MP16]) that a principal fully irreducible outer automorphism has only a single axis and that this axis is determined by a sequence of “proper full folds.” It is further known that it stays in the highest two dimensions of open simplices. It has not previously been known precisely which of these simplices the principal axes pass through. The main results of this paper also determine these simplices in rank 3.

Main Theorem B. The open simplices in CV_3 with principal axes passing through them are precisely those whose underlying graph is listed in Main Theorem A, i.e. each of the simplices in the highest 2 dimensions.

It is worth remarking that the code written determining the automata in this paper could in theory be used to determine the automata in other ranks (computation capacity permitting). The arguments used are also likely extendable to prove results in specific higher-rank settings of interest, with computation time again likely the most significant obstacle.

Acknowledgements. This work is inspired by results obtained together with Lee Mosher, Yael Algom-Kfir, and Ilya Kapovich. The second author is grateful to their knowledge and insights shared over the years. We thank Martin Lustig for pointing out an important oversight in a previous version and the referee for their helpful suggestions.

¹The complete automaton can be found written out at <https://mast.queensu.ca/~cpfaff/Automaton>.

2. BACKGROUND

We assume throughout this paper that $r \geq 3$ is an integer, and that F_r is the rank- r free group with a fixed free basis.

2.1. Edge maps & train track maps.

Definition 1 (High-valence graphs). We call a graph Γ *high-valence* if the valence of each vertex is ≥ 3 .

Definition 2 (Edge, vertex, & direction sets $E(\Gamma)$, $E^+(\Gamma)$, $V(\Gamma)$, $\mathcal{D}(\Gamma)$, $\mathcal{D}(\Gamma, v)$, & Turns). Let Γ be a graph with positively oriented edges $\{e_1, e_2, \dots, e_n\}$ and vertices $\{v_1, v_2, \dots, v_m\}$. Then $E(\Gamma) = \{e_1, \bar{e}_1, \dots, e_n, \bar{e}_n\}$ will denote the set of edges of Γ with each orientation listed separately and an overline indicating a reversal of orientation (as it will throughout for both edges and paths). We let $E^+(\Gamma) = \{e_1, e_2, \dots, e_n\}$ denote the set of positively-oriented edges of Γ and $V(\Gamma) = \{v_1, v_2, \dots, v_m\}$ the vertex set of Γ .


A *direction at v* will mean an element of $E(\Gamma)$ with initial vertex v . We denote the set of directions of Γ by $\mathcal{D}(\Gamma)$ and the set of those at v by $\mathcal{D}(\Gamma, v)$.

A *turn at v* will mean an unordered pair $\{d_1, d_2\}$ of directions $d_1, d_2 \in \mathcal{D}(\Gamma, v)$. The turn is *degenerate* if $d_1 = d_2$, and *nondegenerate* otherwise.

Definition 3 (Paths & loops). Let Γ be a graph. A *path ρ in Γ* will mean a finite sequence (a_1, a_2, \dots, a_m) such that

1. $a_j \in E(\Gamma)$ for each j and
2. there exists a sequence $(v_1, v_2, \dots, v_{m+1})$ of vertices in Γ satisfying that, for each $i = 2, 3, \dots, m$, the turn $\{\bar{a}_{i-1}, a_i\}$ is at v_i .

For such a path (a_1, a_2, \dots, a_m) we write $\rho = a_1 a_2 \dots a_m$ and say ρ *contains* the oriented edges a_1, a_2, \dots, a_m and *takes* the turns $\{\bar{a}_1, a_2\}, \{\bar{a}_2, a_3\}, \dots, \{\bar{a}_{m-1}, a_m\}$. If $v_1 = v_{m+1}$, we call ρ a *loop in Γ* . We consider two loops to be equivalent if they are equal up to a cyclic permutation of their edges. Sometimes we ignore orientation in a loop ℓ , considering ℓ and $\bar{\ell}$ the same. In such cases, we call ℓ *unoriented*. Otherwise, we call ℓ *oriented*. If ℓ traverses each edge of Γ in at least one of its orientations at least once, we call ℓ *comprehensive*. We call both a path and loop ℓ *tight* if for each $e \in E(\Gamma)$, the path $e\bar{e}$ does not appear in ℓ . We may *tighten* ℓ by removing, for each $e \in E(\Gamma)$, all instances of $e\bar{e}$ in ℓ .



Definition 4 (Edge maps, direction map Dg). By an *edge map $g : \Gamma \rightarrow \Gamma'$* we will mean

- a map $\mathcal{V} : V(\Gamma) \rightarrow V(\Gamma')$, where we write $g(v)$ for $\mathcal{V}(v)$, together with
- an assignment, for each $e \in E(\Gamma)$, of a path (denoted $g(e)$) in Γ' such that
 - (1) if the initial vertex of e is v , then the initial vertex of $g(e)$ is $g(v)$, and
 - (2) if $g(e)$ is the edge path $g(e) = a_1 a_2 \dots a_m$, then $g(\bar{e})$ will be the concatenation of edge paths $g(\bar{e}) = \bar{a}_m \dots \bar{a}_2 \bar{a}_1$.

If $\gamma = e_1 e_2 \dots e_n$ is a path in Γ for some $e_1, e_2, \dots, e_n \in E(\Gamma)$, then by $g(\gamma)$ we will mean the concatenation of edge paths $g(\gamma) = g(e_1) g(e_2) \dots g(e_n)$.

We call g *tight* if the image of each edge is a tight path. We often, without explicitly writing it, assume edge maps are tight.

Viewing g as a continuous map of graphs, we say g *represents φ* when $\pi_1(\Gamma)$ has been identified with F_r , i.e. Γ is *marked*, and φ is the induced map of fundamental groups. When a marking is not explicitly given, we mean that “there exists a marking such that.”

To g we associate a *direction map* $Dg : \mathcal{D}(\Gamma) \rightarrow \mathcal{D}(\Gamma')$ such that if $g(e) = a_1 a_2 \dots a_m$, for some $m \geq 1$ and $a_1, a_2, \dots, a_m \in E(\Gamma)$, then $Dg(e) = p_1$ and $Dg(\bar{e}) = \overline{a_m}$. We call a direction e *periodic* if $Dg^k(e) = e$ for some $k > 0$, and *fixed* if $k = 1$.

Remark 1. Let ρ be a path in a graph Γ and g an edge map with domain Γ . Then $g(\rho)$ is tight if and only if g can be viewed as a continuous map that is locally injective on ρ .

Definition 5 (Interior & exterior turns). Suppose $g : \Gamma \rightarrow \Gamma'$ is an edge map, γ a path in Γ , and $\gamma' := g(\gamma)$. The *interior turns of g in Γ'* will be the turns taken by the paths $g(E)$, such that $E \in E^+(\Gamma)$. The *interior turns of g in γ'* will be the turns taken by the paths $g(E)$ such that $E \in E(\Gamma)$ and γ contains E . If $\{a, b\}$ is a turn in γ , then the turn $\{Dg(a), Dg(b)\}$ in γ' is called an *exterior turn of g in γ'* .

Definition 6 (Illegal turns & transparent edge maps). Let $g : \Gamma \rightarrow \Gamma'$ be an edge map. We call a turn $\{d_1, d_2\}$ in Γ *g-prenull* if $\{Dg(d_1), Dg(d_2)\}$ is degenerate. When $g : \Gamma \rightarrow \Gamma$ is a self-map, the turn $\{d_1, d_2\}$ is called an *illegal turn for g* if $\{Dg^k(d_1), Dg^k(d_2)\}$ is degenerate for some k . We call g *transparent* if each illegal turn is prenull, if each turn taken by $g^k(e)$ for some $k \in \mathbb{N}$ and $e \in E(\Gamma)$ is taken by $g(e)$, and if each g -periodic vertex and direction is in fact fixed by g . Notice that every self-map has a transparent power. This is a strengthening of the notion given in [AKKP19] and so in particular implies it.

Lemma 1. Suppose $F : \Gamma \rightarrow \Gamma'$ is an edge map, γ a path in Γ , and $\gamma' := F(\gamma)$. Then

- (1) if γ takes the turn $\{a, b\}$ in Γ , then γ' takes the turn $\{DF(a), DF(b)\}$ in Γ' , and
- (2) each turn in γ' is either an interior or an exterior turn of F in γ' before tightening.

Proof. Suppose that $F : \Gamma \rightarrow \Gamma'$ is as in the lemma statement.

(1) Let $a_1, \dots, a_n, b_1, \dots, b_m$ be oriented edges of Γ' such that $F(a) := a_1 \dots a_n$ and $F(b) = b_1 \dots b_m$. Suppose further that γ takes $\{a, b\}$. Then γ , up to reorientation, contains $\bar{a}b$. So γ' , also up to reorientation, contains $F(\bar{a})F(b) = (\overline{a_n}) (\overline{a_{n-1}}) \dots (\overline{a_1}) b_1 b_2 \dots b_m$. In particular, γ' takes the turn $\{a_1, b_1\} = \{DF(a), DF(b)\}$ in Γ' . This completes the proof of (1).

(2) Suppose $\gamma = e_1 e_2 \dots e_n$, where $e_1, e_2, \dots, e_n \in E(\Gamma)$, and $F(e_i) = e_{i,1} e_{i,2} \dots e_{i,k_i}$. Then,

$$\gamma' = F(e_1) F(e_2) \dots F(e_n) = e_{1,1} e_{1,2} \dots e_{1,k_1} e_{2,1} e_{2,2} \dots e_{2,k_2} \dots e_{n,1} e_{n,2} \dots e_{n,k_n}.$$


Observe that each turn in γ' is either of the form $\{\overline{e_{i,j}}, e_{i,j+1}\}$ or $\{\overline{e_{i,k_i}}, e_{i+1,1}\}$. A turn of the form $\{\overline{e_{i,j}}, e_{i,j+1}\}$ is in the image $F(e_i)$ and is therefore an interior turn of F in γ' . For a turn of the form $\{\overline{e_{i,k_i}}, e_{i+1,1}\}$, we have $\{\overline{e_{i,k_i}}, e_{i+1,1}\} = \{\overline{DF(e_i)}, DF(e_{i+1})\}$, where $\{\overline{e_i}, e_{i+1}\}$ is a turn taken by γ . Thus, a turn of the form $\{\overline{e_{i,k_i}}, e_{i+1,1}\}$ is an exterior turn of F in γ' . This completes the proof of (2). □

2.2. Folds & (fold-conjugate) Stallings fold decompositions.

Definition 7 ((Different-length) & (permissible) folds). Let Γ be a graph and $\{a, b\}$ a nondegenerate turn in Γ . We define the different-length fold \mathfrak{f} in Γ of $\{a, b\}$ with a longer as follows.


If we subdivide the edge a into a_1 and a_2 , so $a = a_1a_2$, then by the *different-length fold* \mathfrak{f} in Γ of $\{a, b\}$ with a longer we will mean the quotient map of Γ defined by identifying the edges a_1 and b . Note that for the different-length fold \mathfrak{f} of $\{a, b\}$ with a longer, if we label the edges in $\mathfrak{f}(\Gamma)$ so that $\mathfrak{f}(a) = b'a'$, and for all other $e \in E^+(\Gamma)$ with $e \neq a$, we have $\mathfrak{f}(e) = e'$, then the direction map will fix all directions apart from sending a to b' . A different-length fold is sometimes elsewhere in the literature referred to as a “proper full fold,” but we found our terminology more evocative in that, in Culler-Vogtmann Outer space, the different-length fold of $\{a, b\}$ with a longer could only occur if a were longer than b . At times we abuse notation and drop the primes.

Let Γ be a high-valence graph, $\{a, b\}$ a turn in Γ , and ℓ a loop in Γ . We say the different-length fold of $\{a, b\}$ with a longer in Γ is *permissible with respect to ℓ* if ℓ does not take the turn $\{a, b\}$.



Remark 2. Let Γ be a high-valence graph, ℓ a comprehensive loop in Γ , and \mathfrak{f} a permissible fold of Γ with respect to ℓ . Let $\Gamma' := \mathfrak{f}(\Gamma)$ and $\ell' := \mathfrak{f}(\ell)$. Since \mathfrak{f} is a quotient map, and hence surjective, ℓ' is also comprehensive.

Definition 8 ((Fold-conjugate) Stallings fold decompositions). Stallings introduced a version of a fold in [Sta83]. Let $g: \Gamma \rightarrow \Gamma'$ be a tight graph homotopy equivalence. Let $e'_1 \subset e_1$ and $e'_2 \subset e_2$ be maximal, initial, nontrivial subsegments of edges e_1 and e_2 emanating from a common vertex and satisfying: $g(e'_1) = g(e'_2)$ as edge paths and the terminal endpoints of e'_1 and e'_2 are distinct points in $g^{-1}(V(\Gamma'))$. Redefining Γ having vertices at the endpoints of e'_1 and e'_2 if necessary, one can obtain a graph Γ_1 by identifying the points of e'_1 and e'_2 that have the same image under g . We call such a fold a *Stallings fold*. Stallings [Sta83] also showed that if $g: \Gamma \rightarrow \Gamma'$ is a tight homotopy equivalence, then g factors as a composition of Stallings folds and a final homeomorphism. We call such a decomposition a *Stallings fold decomposition*. It can be obtained as follows: At an illegal turn for $g: \Gamma \rightarrow \Gamma'$, one can fold two maximal initial segments having the same image in Γ' to obtain a map $\mathfrak{g}_1: \Gamma_1 \rightarrow \Gamma'$ of the quotient graph Γ_1 . The process can be repeated for \mathfrak{g}_1 and recursively. If some $\mathfrak{g}_k: \Gamma_{k-1} \rightarrow \Gamma'$ has no illegal turn, then \mathfrak{g}_k will be a homeomorphism and the fold sequence is complete.



Notice that choices of illegal turns are made in this process and that different choices lead to different Stallings fold decompositions of the same homotopy equivalence.

A *subdivided fold* will mean a fold written as a composition of 2 folds. Suppose $\Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \dots \xrightarrow{g_{k-1}} \Gamma_{k-1} \xrightarrow{g_k} \Gamma_k$ and $\Gamma'_0 \xrightarrow{h_1} \Gamma'_1 \xrightarrow{h_2} \dots \xrightarrow{h_{n-1}} \Gamma'_{n-1} \xrightarrow{h_n} \Gamma'_n$ are Stallings fold decompositions. We say the decompositions are *fold-conjugate* if, possibly after a fold subdivision

of g_k into $\Gamma_{k-1} \xrightarrow{g'_k} \Gamma_{k'} \xrightarrow{g_{k+1}} \Gamma_k$ or h_n into $\Gamma'_{n-1} \xrightarrow{h'_n} \Gamma'_{n'} \xrightarrow{h_{n+1}} \Gamma'_n$, we have for some j :

$\Gamma_j \xrightarrow{g_{j+1}} \dots \xrightarrow{g_k} \Gamma_k = \Gamma_0 \xrightarrow{g_1} \dots \xrightarrow{g_j} \Gamma_j$ is $\Gamma'_0 \xrightarrow{h_1} \Gamma'_1 \xrightarrow{h_2} \dots \Gamma'_{n-1} \xrightarrow{h'_n} \Gamma'_{n'} \xrightarrow{h_{n+1}} \Gamma'_n$ or

$\Gamma_j \xrightarrow{g_{j+1}} \dots \xrightarrow{g_k} \Gamma_k = \Gamma_0 \xrightarrow{g_1} \dots \xrightarrow{g_j} \Gamma_j$ is $\Gamma'_0 \xrightarrow{h_1} \Gamma'_1 \xrightarrow{h_2} \dots \xrightarrow{h_{n-1}} \Gamma'_{n-1} \xrightarrow{h_n} \Gamma'_n$ or

$\Gamma_0 \xrightarrow{g_1} \dots \xrightarrow{g_{k-1}} \Gamma_{k-1} \xrightarrow{g'_k} \Gamma_{k'} \xrightarrow{g_{k+1}} \Gamma_k$ is $\Gamma'_j \xrightarrow{h_{j+1}} \dots \xrightarrow{h_n} \Gamma'_n = \Gamma'_0 \xrightarrow{h_1} \dots \xrightarrow{h_j} \Gamma'_j$.

The following lemma will be used in Section 3.2 in understanding the structure of the local Whitehead graph for a train track representative of a principal fully irreducible outer automorphism.

Lemma 2. *Let Γ be a graph, \bar{a} and b directions at a common vertex of Γ , and ℓ a comprehensive loop in Γ . Let \mathfrak{f} be the different-length fold of $\{\bar{a}, b\}$ with b longer. We use the notation for the image of \mathfrak{f} established in Definition 7: i.e. $\mathfrak{f}(b) = \bar{a}'b'$, and for all other $e \in E^+(\Gamma)$ with $e \neq b$, we have $\mathfrak{f}(e) = e'$. Let Γ' and ℓ' denote the \mathfrak{f} -images of Γ and ℓ respectively. Then, before tightening, ℓ' takes the turn $\{a', b'\}$ and ℓ' takes no other turns containing b' .*

Proof. Since ℓ is a comprehensive loop, ℓ contains b . So ℓ' contains $\mathfrak{f}(b) = \bar{a}'b'$ and thus, pre-tightening, ℓ' takes the turn $\{a', b'\}$. Now suppose $\{d_1, d_2\}$ is a turn taken by ℓ' pre-tightening. By Lemma 1, we have that $\{d_1, d_2\}$ is either an interior or exterior turn of \mathfrak{f} in ℓ' . If $\{d_1, d_2\}$ is an interior turn, it must be $\{a', b'\}$, as the quotient map fixes all edges but b . So we assume $\{d_1, d_2\}$ is an exterior turn. Since there is no $e \in E(\Gamma)$ such that $D\mathfrak{f}(e) = b'$, we then have that $\{d_1, d_2\}$ cannot contain the direction b' , completing the proof. \square

3. THE RANK-3 LONELY DIRECTION TRAIN TRACK AUTOMATON AND MAP F

In this section we define the automaton satisfying that each Stallings fold decomposition of each train track representative of each principal fully irreducible outer automorphism is represented by a directed loop in the automaton (see §4 and §5).

3.1. Sequences of permissible folds and the automaton $\mathcal{A}(\Gamma, \ell)$ generated by (Γ, ℓ) .

Definition 9 (Sequence of permissible folds). Let $\Gamma, \Gamma_1, \Gamma_2, \dots, \Gamma_n$ be rank- r high-valence graphs and let ℓ be a comprehensive loop in Γ . Let $\mathfrak{f}_1 : \Gamma \rightarrow \Gamma_1, \mathfrak{f}_2 : \Gamma_1 \rightarrow \Gamma_2, \dots, \mathfrak{f}_n : \Gamma_{n-1} \rightarrow \Gamma_n$ be folds and $\ell_1 := \mathfrak{f}_1(\ell), \ell_2 := \mathfrak{f}_2(\ell_1), \dots, \ell_n := \mathfrak{f}_n(\ell_{n-1})$. Suppose \mathfrak{f}_i is a permissible fold in Γ_{i-1} with respect to ℓ_i for each i . We then call $\mathfrak{f} = (\mathfrak{f}_1, \mathfrak{f}_2, \dots, \mathfrak{f}_n)$ a *sequence of permissible folds in Γ* , from which we obtain Γ_n . Abusing notation, we let \mathfrak{f} also denote $\mathfrak{f}_n \circ \mathfrak{f}_{n-1} \circ \dots \circ \mathfrak{f}_1$.

Lemma 3. *Let Γ and Γ' be high-valence graphs, ℓ and ℓ' comprehensive loops in Γ and Γ' respectively, and $\mathfrak{f} : \Gamma \rightarrow \Gamma'$ a permissible different-length fold with respect to ℓ such that $\mathfrak{f}(\ell) = \ell'$. Then if ℓ is tight, ℓ' is tight.*

Proof. Denote by $\{a, \bar{b}\}$ the turn folded by \mathfrak{f} and assume, without loss of generality, that a is longer. We use the notation for the image of \mathfrak{f} established in Definition 7. That is $\mathfrak{f}(a) = \bar{b}'a'$, and for each other $e \in E^+(\Gamma)$ with $e \neq b$, we have $\mathfrak{f}(e) = e'$. Let $\{d'_1, d'_2\}$ be an arbitrary turn taken by ℓ' . For ℓ' to be tight it suffices to show $d'_1 \neq d'_2$. By Lemma 1, $\{d'_1, d'_2\}$ is either an interior or exterior turn of \mathfrak{f} . We examine the cases separately:

Case 1 ($\{d'_1, d'_2\}$ is an interior turn): Since all edges in $E^+(\Gamma)$ except a are fixed by \mathfrak{f} (but for the addition of primes), the interior turn $\{d'_1, d'_2\}$ must be taken by $\mathfrak{f}(a) = \bar{b}'a'$. Thus, $\{d'_1, d'_2\} = \{a', b'\}$. Thus, it suffices to show $a' \neq b'$.

Suppose for the sake of contradiction that $a' = b'$. Subdivide a into a_1a_2 with $\mathfrak{f}(a_1) = \bar{b}'$ and $\mathfrak{f}(a_2) = a'$. Then $\mathfrak{f}^{-1}(\{a'\}) = \{a_2\}$ and $\mathfrak{f}^{-1}(\{\bar{b}'\}) = \{a_1, \bar{b}\}$. By construction, $a_1 \neq a_2$. Moreover, \mathfrak{f} is a fold of $a = a_1a_2$ and \bar{b} , so a_1 and \bar{b} must be distinct directions. Thus, $a' \neq b'$, concluding the proof in the case where $\{d'_1, d'_2\}$ is an interior turn.

Case 2 ($\{d'_1, d'_2\}$ is an exterior turn): If $\{d'_1, d'_2\}$ is an exterior turn, then $\{d'_1, d'_2\} = \{D\mathfrak{f}(d_1), D\mathfrak{f}(d_2)\}$ for some turn $\{d_1, d_2\}$ taken by ℓ . We examine subcases:

Case 2a ($d_1 \neq a$ and $d_2 \neq a$): Since ℓ is tight, $d_1 \neq d_2$. Since \mathfrak{f} fixes all directions except for a (but for the addition of primes), the directions d_1 and d_2 must also be fixed under \mathfrak{f} (but for the addition of primes). Thus, $D\mathfrak{f}(d_1) \neq D\mathfrak{f}(d_2)$. Since $\{d'_1, d'_2\} = \{D\mathfrak{f}(d_1), D\mathfrak{f}(d_2)\}$, we have $d'_1 \neq d'_2$, as desired.

Case 2b ($d_1 = a$ or $d_2 = a$): Without loss of generality, we assume $d_1 = a$. Then $d'_1 = D\mathfrak{f}(d_1) = D\mathfrak{f}(a) = \bar{b}'$. To show $d'_1 \neq d'_2$, it suffices to show $d'_2 = D\mathfrak{f}(d_2) \neq \bar{b}'$.

Since $d_2 \in D\mathfrak{f}^{-1}(\{d'_2\})$, it suffices to show $d_2 \notin D\mathfrak{f}^{-1}(\{\bar{b}'\})$. Suppose for the sake of contradiction that $d_2 \in D\mathfrak{f}^{-1}(\{\bar{b}'\})$. Now $D\mathfrak{f}^{-1}(\{\bar{b}'\}) = \{a, \bar{b}\}$. Further, since ℓ is tight and $\{d_1, d_2\}$ is taken by ℓ and $d_1 = a$, we have $d_2 \neq a$. So, if we had $d_2 \in \{a, \bar{b}\}$, we would need $d_2 = b$. So suppose $d_2 = b$. Then ℓ would take the turn $\{d_1, d_2\} = \{a, \bar{b}\}$. But \mathfrak{f} is a fold of $\{a, \bar{b}\}$ and is permissible with respect to ℓ , a contradiction. Thus, $d_2 \neq \bar{b}$, as desired. \square

Lemma 4. *Let Γ be a high-valence graph, ℓ a tight comprehensive loop in Γ , and $\mathfrak{f} = (\mathfrak{f}_1, \mathfrak{f}_2, \dots, \mathfrak{f}_n)$ a sequence of permissible folds in Γ . Then for each $e \in E(\Gamma)$, we have that $\mathfrak{f}(e)$ is tight.*

Proof. Since ℓ is comprehensive, ℓ contains e . Applying \mathfrak{f} to ℓ consists of successively applying folds to ℓ . Since, by Lemma 3, a permissible fold with respect to a loop preserves its tightness, it follows that $\mathfrak{f}(\ell)$ is tight. Therefore, $\mathfrak{f}(e)$ is tight. \square

Definition 10 (Automaton $\mathcal{A}(\Gamma, \ell)$ generated by (Γ, ℓ)). Let Γ be a high-valence graph, ℓ a comprehensive loop in Γ , and \mathcal{F} the set of all sequences of permissible folds in Γ . Let

$$\mathcal{G} = \{(\Gamma_i, \ell_i) \mid \Gamma_i = \mathfrak{f}(\Gamma), \ell_i = \mathfrak{f}(\ell) \text{ for some } \mathfrak{f} \in \mathcal{F}\}$$

be the set of tuples obtained by applying all possible sequences of folds to Γ and ℓ . Let $\mathcal{E} \subseteq \mathcal{G}^2$ be the set such that $((\Gamma_i, \ell_i), (\Gamma'_i, \ell'_i)) \in \mathcal{E}$ precisely when there exists a permissible fold $\mathfrak{f} : \Gamma_i \rightarrow \Gamma'_i$ with respect to ℓ_i such that the set of turns taken by ℓ'_i is equal to the set of turns taken by $\mathfrak{f}(\ell_i)$, modulo primes.

We consider two elements (Γ_i, ℓ_i) and (Γ_j, ℓ_j) of \mathcal{G} equivalent if the following are satisfied:

- there exists an ornamentation-preserving graph isomorphism σ from Γ_i to Γ_j and

- the set of turns taken by ℓ_i can be identified via σ with the set of turns taken by ℓ_j .

Then *the automaton generated by* (Γ, ℓ) , denoted $\mathcal{A}(\Gamma, \ell)$, is the directed graph $(\mathcal{G}, \mathcal{E})$.

3.2. The Lonely Direction Property. We introduce in this section a property held by the states (and preserved within) the automata we construct.

Proposition 1. *Let Γ be a 5-edged rank-3 high-valence graph. Then Γ has exactly 3 vertices: two valence-3 vertices and one valence-4 vertex.*

Proof. Note that all vertices of Γ must have valence ≥ 2 , and that, since there are 5 edges, the sum of the valences of all the vertices must be 10.

Suppose Γ were a rank-3 graph with 5 edges and 1 vertex. This would be a rank-5 rose, contradicting that Γ has rank 3.

Now suppose Γ was a rank-3 graph with 5 edges and 2 vertices. One may partition the 5 edges into edges that connect the 2 vertices and edges that form loops on one of the vertices. The graph with 2 vertices and 5 edges connecting the vertices has rank 4. Since replacing an edge connecting the vertices with an edge forming a loop at a vertex does not change the rank, Γ also cannot have this form.

Now suppose Γ was a rank-3 graph with 5 edges and ≥ 4 vertices. Then, by the valence sum, there would be a vertex with valence < 3 , but we said Γ had to be of high valence.

Thus, Γ must have precisely 3 vertices. Since each vertex must have valence ≥ 3 , the valences of the vertices must be 3, 3, and 4. □

Definition 11 (Lonely Direction Property). Let Γ be a 5-edged rank-3 graph, ℓ a comprehensive loop in Γ taking all but two turns of Γ , and v the valence-4 vertex of Γ . If there exists a particular direction d at v that is contained in the only two turns not taken by ℓ , then we say (Γ, ℓ) satisfies the *Lonely Direction Property with distinguished direction d* .

Proposition 2. *Let Γ be a 5-edged rank-3 graph, and ℓ a comprehensive loop in Γ such that (Γ, ℓ) satisfies the Lonely Direction Property. Then, for all 5-edged graphs Γ_i in $\mathcal{A}(\Gamma, \ell)$, letting ℓ_i denote image of ℓ in Γ_i , we have that (Γ_i, ℓ_i) satisfies the Lonely Direction Property.*


Proof. Let v be the valence-4 vertex of Γ and u, w the valence-3 vertices. Note that all permissible folds in Γ are at v , since the only two turns not taken by ℓ are at v . Suppose that $\mathfrak{f}: \Gamma \rightarrow \Gamma'$ is a permissible fold of the turn $\{a, b\}$ at v with b longer and that either a or b is the distinguished direction. As in Definition 7, subdivide b into $b_1 b_2$ so that $\mathfrak{f}(a) = \mathfrak{f}(b_1)$. Call $\mathfrak{f}(a)$ by a' (so that also $\mathfrak{f}(b_1) = a'$) and $\mathfrak{f}(b)$ by b' . Then $D\mathfrak{f}(a) = D\mathfrak{f}(b) = a'$. Let $\mathfrak{f}(\ell) := \ell'$ and let v', u' , and w' denote the respective \mathfrak{f} -images of v, u , and w .

The valence of v' will be either 3 or 4. Specifically, the valence of v' is 4 if a is a loop and is 3 otherwise. We consider separately:

Case 1: v' has valence 3 and **Case 2:** v' has valence 4.

Case 1 (v' has valence 3): We aim to show that the valence-4 vertex of Γ' contains all turns not taken by ℓ' , and each of these turns contains a common distinguished direction.

Suppose the directions at v are $\{a, b, c, d\}$, i.e. c and d are the 2 directions not involved in the fold. Then letting $c' := D\mathfrak{f}(c)$ and $d' := D\mathfrak{f}(d)$, the directions at v' are $\{a', c', d'\}$. Then ℓ' takes all turns at v' , as follows. By Lemma 1, ℓ' takes $\{a', c'\}$ and $\{a', d'\}$ as follows. Either ℓ took $\{b, c\}$ and $\{b, d\}$ if a was the distinguished direction or $\{a, c\}$ and $\{a, d\}$ if b was the distinguished direction. Moreover, ℓ' takes $\{c', d'\}$, because ℓ took $\{c, d\}$.



Suppose, without loss of generality, that the vertex u of Γ contains the direction \bar{a} . Note that v cannot contain \bar{a} , or a would be a loop and v' would have valence 4. Let e and f denote the other directions at u . Let the images of u , \bar{a} , e , and f under \mathfrak{f} be u' , \bar{a}' , e' , and f' respectively. The directions at u' are then $\{\bar{a}', e', f', b'\}$ (see image above), where b' is the image of the portion of b not identified with \bar{a} , i.e. b_2 .

By Lemmas 1 and 2:

- Since ℓ takes all turns at u , we have ℓ' takes all turns among $\{\bar{a}', e', f'\}$ and
- ℓ' takes the turn $\{\bar{a}', b'\}$ and
- ℓ' takes no other turns at u' involving b' .


That is, all turns at u' that ℓ' does not take involve b' .

Also by Lemma 1, since ℓ takes all turns at w and the directions at w are not involved in the fold \mathfrak{f} , we have that ℓ' takes all turns at w' .

This completes the proof of Case 1 with u' as the valence-4 vertex of Γ' and b' as the distinguished direction.

Case 2 (v' has valence 3): We aim to show that the valence-4 vertex of Γ' contains all of the turns not taken, and each of these turns contains a common distinguished direction.

In this case, the edge a must be a loop at v . Note that the \mathfrak{f} -image of a will also be a loop. We let c denote the remaining direction at v and c' denote $\mathfrak{f}(c)$. At v , ℓ either takes $\{b, \bar{a}\}$, $\{b, c\}$, and $\{\bar{a}, c\}$, if a is the distinguished direction, or $\{a, \bar{a}\}$, $\{a, c\}$, and $\{\bar{a}, c\}$, if b is the distinguished direction. Thus, by Lemma 1, ℓ' takes $\{a', \bar{a}'\}$, $\{a', c'\}$, and $\{\bar{a}', c'\}$. That is, ℓ' takes all turns among the directions $\{a', \bar{a}', c'\}$. Moreover, by Lemma 2, ℓ' takes the turn $\{\bar{a}', b'\}$ and no other turns at v' involving b' . That is, all turns that ℓ' does not pass over at v' involve b' . Further, since u and w were not involved in any folds under \mathfrak{f} and ℓ takes all turns at u and w , by Lemma 1, all turns at u' and at w' are taken by ℓ' . This completes Case 2 with v' as the valence-4 vertex of Γ' and b' as the distinguished direction.



Inductively, all 5-edged graphs Γ_i in $\mathcal{A}(\Gamma, \ell)$ must also have the same structure as Γ' . In particular, Γ_i must have a distinguished direction b at the valence-4 vertex such that the image of ℓ takes all turns in Γ_i except for precisely two turns involving b .

□

3.3. Rank-3 lonely direction automaton algorithms & construction. The rank-3 lonely direction automaton \mathcal{A}_3 is constructed as follows (the code can be found at [Gag20b]).

To more generally construct $\mathcal{A}(\Gamma, \ell)$ for a high-valence graph Γ and comprehensive loop ℓ in Γ :

- (1) For each nondegenerate turn $\{d_{i,1}, d_{i,2}\}$ in Γ such that ℓ does not take $\{d_{i,1}, d_{i,2}\}$, let \mathcal{F}_i denote the set containing the different-length fold $\{d_{i,1}, d_{i,2}\}$ with $d_{i,1}$ longer and the different-length fold $\{d_{i,1}, d_{i,2}\}$ with $d_{i,2}$ longer. Note that, letting $\mathcal{F} = \cup_i \mathcal{F}_i$, we have that \mathcal{F} is the set of all different-length permissible folds of Γ with respect to ℓ .
- (2) Let $S_0 := (\Gamma, \ell)$.
- (3) Given S_k , use (1) with inputs Γ_k and ℓ_k to determine all permissible folds $\mathcal{F}^k = (\mathbf{f}_{k,1}, \mathbf{f}_{k,2}, \dots, \mathbf{f}_{k,n_k})$ on Γ_k with respect to ℓ_k .
- (4) Perform each fold $\mathbf{f}_{k,j}$ on Γ_k and ℓ_k to obtain $(\mathbf{f}_{k,j}(\Gamma_k), \mathbf{f}_{k,j}(\ell_k))$. Let

$$E_{k,j} = ((\Gamma_k, \ell_k), (\mathbf{f}_{k,j}(\Gamma_k), \mathbf{f}_{k,j}(\ell_k))) \text{ and } S_m = (\mathbf{f}_{k,j}(\Gamma_k), \mathbf{f}_{k,j}(\ell_k)),$$

where m is one more than the maximum of the indices k used thus far.

- (5) Let $S = \cup_k S_k$ and $E = \cup_{k,j} E_{k,j}$.

It is not difficult to see that the directed graph $(S, E)_{(\Gamma, \ell)}$ is $\mathcal{A}(\Gamma, \ell)$.

Definition 12 (The rank-3 lonely direction automaton \mathcal{A}_3). We define \mathcal{A}_3 is the maximal strongly connected components of $\cup \mathcal{A}(\Gamma, \ell)$ as (Γ, ℓ) varies over all pairs (Γ, ℓ) such that Γ is a high-valence graph possessing the Lonely Direction Property with respect to the comprehensive loop ℓ .

Remark 3. \mathcal{A}_3 is obtained by taking $\cup (S, E)_{(\Gamma, \ell)}$ for suitable (Γ, ℓ) as above and then taking the maximal strongly connected components.

3.4. The rank-3 lonely direction automaton. Consider \mathcal{G} to be the set of graphs in the two representations below, but with the inner embedded (colored) graphs replaced with vertices. Each graph $\Gamma_i \in \mathcal{G}$ has associated with it a comprehensive loop ℓ_i such that ℓ_i


- takes each turn at each valence-3 vertex of Γ_i and
- takes a turn $\{d_1, d_2\}$ at the valence-4 vertex of Γ_i if and only if $\{d_1, d_2\}$ is an edge on the (colored) graph depicted in the valence-4 vertex of Γ_i in the figures.

Then, (Γ_i, Γ_j) is an edge in the figures if and only if the following are satisfied:


- There exists a permissible fold $\mathbf{f}_{i,j}$ such that $(\mathbf{f}_{i,j}(\Gamma_i), \mathbf{f}_{i,j}(\ell_i))$ is equivalent to (Γ_j, ℓ_j) up to a permutation of the edge labels. In this case, we write $(\mathbf{f}_{i,j}(\Gamma_i), \mathbf{f}_{i,j}(\ell_i)) \sim (\Gamma_j, \ell_j)$.
- There exists a sequence of permissible folds $\mathcal{F}_{j,i}$ such that $(\mathbf{f}_{j,i}(\Gamma_i), \mathbf{f}_{j,i}(\ell_i)) \sim (\Gamma_j, \ell_j)$.

That is, starting with all high-valence, 5-edged graphs Γ and comprehensive loops ℓ such that (Γ, ℓ) satisfies the Lonely Direction Property, the figures below give the maximal strongly connected components of $\cup \mathcal{A}(\Gamma, \ell) / \sim$, computed as in §3.3 using the program [Gag20a].

The following 3 graphs give single-graph maximal strongly connected components:



The primary maximal strongly connected component, which we will denote by PA_3 , is (the interested reader can find the complete automaton written out at <https://mast.queensu.ca/~cpfaff/Automaton>):



3.6. The turns taken by F .

Definition 13 (\mathcal{T}_f). Let $f : \Gamma \rightarrow \Gamma$ be an edge map. Then define \mathcal{T}_f to be the set of all turns taken by $f^k(E)$ before tightening, as E ranges over $E(\Gamma)$ and as k ranges over $\mathbb{Z}_{>0}$. If v is a vertex of Γ , then $\mathcal{T}_f|_v$ is the restriction of \mathcal{T}_f to turns at v .

We now verify the validity for the algorithm used in [Gag20b] to compute \mathcal{T}_F .

Lemma 5. Let Γ be a high-valence graph, $g : \Gamma \rightarrow \Gamma$ an edge map, $v \in V(\Gamma)$, and R_0 the set of interior turns of g in Γ .

Given R_0, R_1, \dots, R_i , define R_{i+1} as follows:

Define $N_i := \{\{dg(t_1), dg(t_2)\} \mid \{t_1, t_2\} \in R_i\}$ and

$R_{i+1} := \{\{t_1, t_2\} \mid \{t_1, t_2\} \in N_i, \{t_1, t_2\} \notin R_j \text{ for } j < i + 1\}$.

Let l be the first integer such that $R_l = \emptyset$ and set $R = \cup_{i=0}^{l-1} R_i$. Then $R = \mathcal{T}_g$.

Proof. We first show $\mathcal{T}_g \subseteq R$. Let $\{d_1, d_2\} \in \mathcal{T}_g$. Then, by definition, there exists an $E \in E(\Gamma)$ and $k \in \mathbb{Z}_{>0}$ such that $g^k(E)$ takes $\{d_1, d_2\}$ before tightening. Thus, it suffices to show for each $E \in E(\Gamma)$ and $k \in \mathbb{Z}_{>0}$, R contains all turns taken by $g^k(E)$ before tightening. To prove this, we fix an arbitrary $E \in E(\Gamma)$ and proceed by induction on k .

Base Case: We want to show that all turns taken by $g(E)$ before tightening are contained in R . By definition, all turns taken by $g(E)$ before tightening are contained in $R_0 \subseteq R$.

Inductive Step: Fix $j \in \mathbb{Z}_{>0}$. Suppose all turns taken by $g^j(E)$ before tightening are contained in R .

Inductive Step: We show that all turns taken by $g^{j+1}(E)$ before tightening are contained in R . Let $\{t_{1,j+1}, t_{2,j+1}\}$ be a turn taken by $g^{j+1}(E)$ before tightening. By Lemma 1, $\{t_{1,j+1}, t_{2,j+1}\}$ is either an interior turn or an exterior turn of g in $g^j(E)$.

If $\{t_{1,j+1}, t_{2,j+1}\}$ is an interior turn, then $\{t_{1,j+1}, t_{2,j+1}\} \in R_0 \subseteq R$. So we assume $\{t_{1,j+1}, t_{2,j+1}\}$ is an exterior turn. Then, there exists a turn $\{t_{1,j}, t_{2,j}\}$ taken by $g^j(E)$ before tightening such that $\{t_{1,j+1}, t_{2,j+1}\} = \{Dg(t_{1,j}), Dg(t_{2,j})\}$.

By the inductive hypothesis, $\{t_{1,j}, t_{2,j}\} \in R$. Thus, there exists some $\alpha < l$ such that $\{t_{1,j}, t_{2,j}\} \in R_\alpha$. Then, by the definition of R_α , $\{t_{1,j+1}, t_{2,j+1}\} = \{Dg(t_{1,j}), Dg(t_{2,j})\} \in R_\beta$ for some $\beta \leq \alpha + 1$. Then, since $\alpha < l$, we have $\alpha + 1 \leq l$. This implies $\beta \leq l$. But since $R_l = \emptyset$ and $\{t_{1,j+1}, t_{2,j+1}\} \in R_\beta$, we have $\beta \neq l$. Thus, $\beta < l$, and $\{t_{1,j+1}, t_{2,j+1}\} \in R_\beta \subseteq R$. Therefore, $\{t_{1,j+1}, t_{2,j+1}\} \in R$.

This proves the claim and thus that $\mathcal{T}_g \subseteq R$.

To prove $R \subseteq \mathcal{T}_g$, we let $\{d_1, d_2\} \in R$ and show $\{d_1, d_2\} \in \mathcal{T}_g$. Either $\{d_1, d_2\} \in R_0$ or $\{d_1, d_2\} \in R_k$ for some $k > 0$. We prove these cases separately:

Case 1 ($\{d_1, d_2\} \in R_0$): If $\{d_1, d_2\} \in R_0$, then, by definition, there exists some $E \in E(\Gamma)$ such that $\{d_1, d_2\}$ is taken by $g(E)$ before tightening. Therefore, $\{d_1, d_2\} \in \mathcal{T}_g$. This proves the claim in Case 1.

Case 2 ($\{d_1, d_2\} \in R_k$ for some $k > 0$): By construction, for all $i > 0$ and turns $\{a, b\} \in R_i$, there exists a turn $\{c, d\} \in R_{i-1}$ such that $\{a, b\} = \{Dg(c), Dg(d)\}$. Then, $\{d_1, d_2\}$ can be obtained by applying k times Dg to a turn in R_0 , i.e. there must exist some sequence of turns $\{d_{0,1}, d_{0,2}\}, \{d_{1,1}, d_{1,2}\}, \dots, \{d_{k,1}, d_{k,2}\}$ such that:

- $\{d_{k,1}, d_{k,2}\} = \{d_1, d_2\}$,
- $\{d_{i,1}, d_{i,2}\} \in R_i$,
- $\{d_{i+1,1}, d_{i+1,2}\} = \{Dg(d_{i,1}), Dg(d_{i,2})\}$ for $i = 1, 2, \dots, k - 1$, and

- $\{d_{0,1}, d_{0,2}\} \in R_0$.

Thus, $\{d_1, d_2\} = \{Dg^k(d_{0,1}), Dg^k(d_{0,2})\}$. Since $\{d_{0,1}, d_{0,2}\} \in R_0$, there exists some $E \in E(\Gamma)$ satisfying that $\{d_{0,1}, d_{0,2}\}$ is taken by $g(E)$ before tightening. Therefore, $\{d_1, d_2\}$ is a turn taken by $g^{k+1}(E)$ before tightening.

Thus, $\{d_1, d_2\} \in \mathcal{T}_g$, which proves the claim in Case 2.

This proves that $R \subseteq \mathcal{T}_g$. Since $\mathcal{T}_g \subseteq R$ and $\mathcal{T}_g \supseteq R$, we must have $\mathcal{T}_g = R$, which completes the proof. □

Lemma 6. *Suppose $f = h^n$ for some $n \in \mathbb{N}$. Then $\mathcal{T}_f = \mathcal{T}_h$*

Proof. First we show $\mathcal{T}_f \subseteq \mathcal{T}_h$. Let $\{d_1, d_2\} \in \mathcal{T}_f$. We aim to show $\{d_1, d_2\} \in \mathcal{T}_h$. Since $\{d_1, d_2\} \in \mathcal{T}_f$, there exists some $k \geq 1$ and $e \in E^+(\Gamma)$ such that $f^k(e)$ takes $\{d_1, d_2\}$. Let $nj > k$. Since f maps each $\alpha \in E^+(\Gamma)$ over each $\beta \in E^+(\Gamma)$ at least once, f^i maps each $\alpha \in E^+(\Gamma)$ over each $\beta \in E^+(\Gamma)$ at least once for every $i \geq 1$. In particular, $f^{nj-k}(e)$ contains e . Thus, $h^j(e) = f^{nj}(e) = (f^k \circ f^{nj-k})(e)$ contains $(f^k)(e)$. Thus $h^j(e)$ takes $\{d_1, d_2\}$. Thus, $\{d_1, d_2\} \in \mathcal{T}_h$, and so $\mathcal{T}_f \subseteq \mathcal{T}_h$.

Now we show $\mathcal{T}_h \subseteq \mathcal{T}_f$. Let $\{d_1, d_2\} \in \mathcal{T}_h$. We aim to show $\{d_1, d_2\} \in \mathcal{T}_f$. Since $\{d_1, d_2\} \in \mathcal{T}_h$, there exists $k \geq 1$ and $e \in E^+(\Gamma)$ such that $h^k(e)$ takes $\{d_1, d_2\}$. Since $h^k = f^{nk}$, then $f^{nk}(e)$ also takes $\{d_1, d_2\}$. Thus, $\{d_1, d_2\} \in \mathcal{T}_f$, and so $\mathcal{T}_h \subseteq \mathcal{T}_f$.

Since we have shown containment in both directions, $\mathcal{T}_f = \mathcal{T}_h$. □

Using the program [Gag20b], one obtains:

$$\mathcal{T}_F = \{\{d, e\}, \{\bar{d}, \bar{e}\}, \{b, d\}, \{\bar{b}, c\}, \{\bar{c}, \bar{e}\}, \{\bar{c}, \bar{d}\}, \{\bar{a}, c\}, \{\bar{e}, a\}, \{b, e\}, \{\bar{a}, \bar{b}\}\}$$

4. TRAIN TRACK MAPS & (PRINCIPAL) FULLY IRREDUCIBLE OUTER AUTOMORPHISMS

4.1. Train track maps & fully irreducible outer automorphisms.

Definition 14 (Train track maps, transition matrix, Perron-Frobenius matrix). Let $g : \Gamma \rightarrow \Gamma$ an edge map. We call g a *train track map* if for each $k \in \mathbb{Z}_{>0}$ and $e \in E(\Gamma)$, we have that $g^k(e)$ does not contain $e'e'$ for any $e' \in E(\Gamma)$. We call the train track map g *expanding* if for each edge $e \in E(\Gamma)$ we have that $|g^n(e)| \rightarrow \infty$ as $n \rightarrow \infty$, where for a path γ we use $|\gamma|$ to denote the number of edges γ traverses (with multiplicity). Note that, apart from our not requiring a “marking,” these definitions coincide with those in [BH92] when g is instead viewed as a homotopy equivalence of graphs. We remind the reader that we often drop discussion of the marking when it does not affect our arguments.

The *transition matrix* of a train track map $g : \Gamma \rightarrow \Gamma$ is the square $|E^+(\Gamma)| \times |E^+(\Gamma)|$ matrix $[a_{ij}]$ such that a_{ij} , for each i and j , is the number of times $g(e_i)$ contains either e_j or \bar{e}_j . A transition matrix $A = [a_{ij}]$ is *Perron-Frobenius (PF)* if there exists an N such that, for each $k \geq N$, we have that A^k is strictly positive.

Since we do not explicitly use the definition of a Nielsen path, periodic Nielsen path (PNP), or indivisible Nielsen path (iNP), we refer the reader to [BH92] for the definitions.

Definition 15 (Irreducible, fully irreducible). We call a train track map *irreducible* if it has no proper invariant subgraph with a noncontractible component. It is useful to note that

an edge map whose transition matrix is Perron-Frobenius must be irreducible. An outer automorphism $\varphi \in \text{Out}(F_r)$ is *fully irreducible* if no positive power preserves the conjugacy class of a proper free factor of F_r . Bestvina and Handel [BH92] proved that every fully irreducible outer automorphism admits expanding irreducible train track representatives.

4.2. Whitehead graphs. Local Whitehead graphs, local stable Whitehead graphs, and ideal Whitehead graphs were introduced in [HM11] and give information about how images of edges pass through vertices. We give definitions here only in the circumstance of no PNPs, as this will always be the case for us. We further assume that we have replaced any edge map g by a transparent power (in the sense of Definition 6).

Let $g : \Gamma \rightarrow \Gamma$ be a train track map. The *local Whitehead graph* $LW(g; v)$ at a $v \in V(\Gamma)$ has a vertex for each direction at v and an edge connecting the vertices corresponding to a pair of directions $\{d_1, d_2\}$ at v precisely when the turn $\{d_1, d_2\}$ is g -taken. Given a fixed vertex v , we obtain the *local stable Whitehead graph* $SW(g; v)$ from $LW(g; v)$ by restricting to the fixed direction vertices and the edges between them. In this situation of no PNPs, if g represents a fully irreducible outer automorphism φ , then the *ideal Whitehead graph* $IW(\varphi)$ for φ is isomorphic to the disjoint union $\bigsqcup SW(g; v)$ taken over all fixed vertices of g . Justification of the ideal Whitehead graph being an outer automorphism and conjugacy class invariant can be found in [HM11, Pfa12].

Remark 4. For a transparent train track map $g : \Gamma \rightarrow \Gamma$, the set of edges in $LW(g; v)$ is equal to $\mathcal{T}_g|_v$.

4.3. Full irreducibility criterion. We will use the following criterion for proving that a train track map represents a fully irreducible outer automorphism.

Proposition 3. [Pfa13, Proposition 4.1] (Full Irreducibility Criterion (FIC)) *Let $g : \Gamma \rightarrow \Gamma$ be a PNP-free, irreducible train track representative of $\varphi \in \text{Out}(F_r)$. Suppose that the transition matrix for g is Perron-Frobenius and that all the local Whitehead graphs are connected. Then φ is a fully irreducible outer automorphism.*

4.4. Principal fully irreducible outer automorphisms. Principal fully irreducible outer automorphisms (sometimes just called principal outer automorphisms) were introduced in [AKKP19] as analogues of pseudo-Anosov surface homeomorphisms with only 3-pronged singularities.

Definition 16 (principal fully irreducible outer automorphism). As in [AKKP19], we call a fully irreducible outer automorphism $\varphi \in \text{Out}(F_r)$ *principal* if $IW(\varphi)$ is the disjoint union of $2r - 3$ triangles.

By the main theorem of [MP16], principal outer automorphisms are lone axis outer automorphisms. Relevant for us will be that each of their train track representatives has a unique Stallings fold decomposition, the Stallings fold decompositions of any two of these train track representatives are cyclic permutations of each other, and none of these train track representatives has a PNP.

5. MAIN PROOF

The following lemma will tell us that loops in \mathcal{A}_3 yield train track maps.

Lemma 7. *Let Γ be a high-valence graph, ℓ a tight comprehensive loop on Γ , and $\mathcal{F} = (f_1, f_2, \dots, f_n)$ a sequence of permissible folds in Γ such that the following are satisfied by $f =: f_n \circ \dots \circ f_1$:*

1. *There exists an ornamentation-preserving (but for the addition of primes) isomorphism between $f(\Gamma)$ and Γ and*
2. *the isomorphism induces an identification of the set of turns taken by ℓ with the set of turns taken by $f(\ell)$.*

Then, f is a train track map.

Proof. Let $k \in \mathbb{Z}_{>0}$ and $e \in E(\Gamma)$ be fixed and arbitrary. It suffices to show that $f^k(e)$ is tight. By (1) and (2), f is also a sequence of permissible folds on $f(\Gamma)$. Inductively, f is a sequence of permissible folds on $f^j(\Gamma)$ for any j . Therefore, f^j is a permissible sequence of folds on $f(\Gamma)$ for any j . In particular, f^k is a permissible sequence of folds on Γ . Then, by Lemma 4, $f^k(e)$ is tight, as required. □

Proposition 4. *The map F of §3.5 represents a principal fully irreducible outer automorphism.*

Proof. In light of Proposition 3 (and the fact that a Perron-Frobenius transition matrix implies irreducibility for a transparent map), for full irreducibility, it suffices to show that F is an expanding, Perron-Frobenius train track map with no PNPs, and that the local Whitehead graphs connected. Recall that F has been taken to be transparent.

We first prove F is an expanding irreducible train track map with Perron-Frobenius transition matrix. There does not exist a direction d in Γ such that $\{d, d\}$ is an element of \mathcal{T}_F . That is, for each $e \in E(\Gamma)$ and $k \in \mathbb{Z}_{>0}$, we have that $F^k(e)$ does not contain $e'\bar{e}'$ for any $e' \in E(\Gamma)$. Thus, F is a train track map. Since F is a train track map, so that there is no cancellation in any edge F^k -images, and $|F(\alpha)| > 1$ for each $\alpha \in E(\Gamma)$, it follows that $|F^k(\alpha)| \rightarrow \infty$ as $k \rightarrow \infty$ for each $\alpha \in E(\Gamma)$. By inspection, F maps each $\alpha \in E^+(\Gamma)$ over each $\beta \in E^+(\Gamma)$ at least once. Thus, each entry in the transition matrix, A , of F is strictly positive. It follows that A^k is strictly positive for each $k \in \mathbb{Z}_{>0}$. Thus F is Perron-Frobenius.

One can show F has no PNPs using the sage train track package of Coulbois ([Cou14]) or the methodology of [Pfa13], [HM11, Example 3.4], [Pfa15b], [KP18, Lemma 4.7].

We now show that for each $v \in V(\Gamma)$, we have that $LW(v; F)$ is connected. Let v_1 be the vertex of Γ with directions $\{\bar{d}, \bar{e}, \bar{c}, a\}$, and v_2 the vertex of Γ with directions $\{d, e, b\}$, and v_3 be the vertex of Γ with directions $\{c, \bar{a}, \bar{b}\}$, and D_{v_i} the set of directions at v_i . Restricting \mathcal{T}_F to each vertex, one obtains:

$$\begin{aligned} \mathcal{T}_F|_{v_1} &= \{\{\bar{d}, \bar{e}\}, \{\bar{c}, \bar{e}\}, \{\bar{c}, \bar{d}\}, \{\bar{e}, a\}\} \\ \mathcal{T}_F|_{v_2} &= \{\{d, e\}, \{b, d\}, \{b, e\}\} \\ \mathcal{T}_F|_{v_3} &= \{\{\bar{b}, c\}, \{\bar{a}, c\}, \{\bar{a}, \bar{b}\}\}. \end{aligned}$$

Thus, $LW(v_1; F) = (D_{v_1}, \mathcal{T}_F|_{v_1})$, $LW(v_2; F) = (D_{v_2}, \mathcal{T}_F|_{v_2})$, and $LW(v_3; F) = (D_{v_3}, \mathcal{T}_F|_{v_3})$, which are connected graphs, as desired. And F is a fully irreducible outer automorphism. Further, since the only nonperiodic direction is a , this implies that the ideal Whitehead graph is a union of 3 triangles, and thus φ is a principal fully irreducibly outer automorphism. □

Proposition 5. *There exists a power F^p of F satisfying the following. Suppose L is a loop in the automata containing the loop L' giving the fold sequence for F^p . Let g denote the train track map obtained by the sequence of folds the directed edges in L represent. Then g represents a principal fully irreducible outer automorphism.*

Proof. Note that we have already taken F to be transparent.

We first need that g contains no PNPs. This argument is in the proof of [AKKP19, Lemma 5.3] only mildly disguised. The first observation is that the condition in [AKKP19, Lemma 5.3] that a periodic fold line satisfies $\mathcal{L} \in B^k(A_\varphi, R, \varepsilon)$ is simply to ensure that a Stallings fold composition for g contains the Stallings fold decomposition for F^p , which holds true in our case. The condition of having a “witness loop” is replaced by the comprehensive loop (and its images) determining our automaton. The rest of the proof is then exactly the same.

By the proof of [AKKP19, Lemma 5.3], we further have that F^p and g have the same local Whitehead graphs and stable Whitehead graphs and that the transition matrix for g is Perron-Frobenius. Thus, since the local Whitehead graphs for F , hence F^p , were connected, we have that those for g are also. We have established that g satisfies all conditions in the Full Irreducibility Criterion, hence represents a fully irreducible outer automorphism.

Further, since g has no periodic Nielsen paths, its ideal Whitehead graph is now the union of the stable Whitehead graphs of F^p , hence of F , which we have already established is the ideal Whitehead graph for a principal fully irreducible outer automorphism. \square

Proposition 6. *Suppose φ is a principal fully irreducible outer automorphism in $\text{Out}(F_3)$ and g is a train track representative of φ . Then some power of g 's Stallings fold decomposition is fold-conjugate to one determining a directed loop in the Lonely Direction Automaton \mathcal{A}_3 .*

Proof. Suppose $g: \Gamma \rightarrow \Gamma$ is a train track representative of a principal fully irreducible outer automorphism φ . Since φ is a principal fully irreducible outer automorphism, it and its powers are lone axis outer automorphisms. By [MP16, Corollary 3.8], there exists another representative $g': \Gamma' \rightarrow \Gamma'$ of φ and a power $h = g'^p$ of g' so that h fixes all vertices of Γ' and all but one direction of Γ' . Further, h has precisely one illegal turn and each fold in its Stallings fold decomposition is a different-length fold (or h would identify vertices, hence could not fix all vertices). Since φ^p is a lone axis outer automorphism and h represents φ^p , we know h has no PNPs, implying that $IW(\varphi)$ will be the disjoint union of its stable Whitehead graphs. The stable Whitehead graphs are the restrictions of the local Whitehead graphs to the periodic direction vertices. This tells us:

- (1) Γ' is 5-edged, with all but a single vertex being of valence-3, and the final vertex (which we call w) being of valence 4, and
- (2) $LW(h, v)$ is a triangle for each valence-3 vertex v , and
- (3) $LW(h, w)$ contains a triangle, between the vertices representing the 3 fixed directions at w . Call the vertex representing the 4th direction α .


Now, notice that the final different-lengthed fold yields a direction coming from the interior of an edge, so that it can only ever be involved in a single taken turn (see Lemma 2). By the process of elimination, this tells us that α , which necessarily represents this direction, is attached to the rest of $LW(h, w)$ by precisely a single edge.

One can construct, as follows, a comprehensive loop in Γ' whose image will not be folded by any fold in the Stallings fold decomposition of h , even if the Stallings fold decomposition

is performed multiple times. Consider any edge e of Γ' . Because h is expanding Perron-Frobenius, for some k , we'll have that $h^k(e)$ will contain at least 2 copies of some edge e' of Γ' . Consider any turn $\{\bar{e}_1, e_2\}$ taken by $h(e')$. $h^{k+1}(e)$ will contain at least 2 copies of e_1e_2 . Let σ denote the subpath of $h^{k+1}(e)$ starting directly after e_1 in a first copy of e_1e_2 and then ending directly after e_1 in a second copy of e_1e_2 . Note that σ is a loop. Also, since h is a train track map and σ is in the h^{k+1} -image of an edge, the Stallings folds cannot fold any turns in an image of σ , so will always be permissible. If α is not comprehensive, since the h -image of each edge contains each other edge, we can replace it by a h -image that is comprehensive, we can also do so to ensure that the turns taken by σ are precisely those appearing in the $LW(h, v)$, as v varies over the vertices in $V(\Gamma')$. Then Γ' satisfies the Lonely Direction Property and the Stallings fold sequence for h gives a loop in $\mathcal{A}(\Gamma')$ with σ as the comprehensive loop.

Since g and h are train track representatives (of powers) of the same lone axis fully irreducible, they share an axis, hence (up to taking a power) have fold-conjugate Stallings fold decompositions and the proposition follows. □

Main Theorem A. The graphs carrying train track representatives of principal fully irreducible outer automorphisms in $Out(F_3)$ are precisely:



That is, each of the graphs of CV_3 where either all vertices are valence-3 or all vertices are valence-3 except that one is valence-4 carries a train track representative a of principal fully irreducible outer automorphism in $Out(F_3)$.

Proof. First notice that these are precisely the graphs in the primary maximal strongly connected component of \mathcal{A}_3 (which one may recall we denote by $P\mathcal{A}_3$), together with those obtained by partially completing the folds represented by the edges of the component (we will call these the *edge-determined partial fold graphs*). The graphs and edges in the other maximal strongly connected components appear also in this component.

Next, notice that, given a directed loop \mathcal{L} in the automaton, the Stallings fold sequences obtained by starting the loop at different nodes in \mathcal{L} yield $Out(F_r)$ -conjugate outer automorphisms. Since the property of being a principal fully irreducible outer automorphism is a conjugacy class invariant, the property of an outer automorphism obtained from \mathcal{L} being a principal fully irreducible outer automorphism is invariant with respect to the node in \mathcal{L} one starts traversing the loop from. Hence, to realize a graph G appearing in the first row of the theorem statement, we can take the loop L' of Proposition 5, append to it a loop containing the desired graph G in our automaton (existing because L' and G are in the same

maximal strongly connected component) and, again using Proposition 5, obtain a train track representative on G of a principal fully irreducible outer automorphism.

We now prove that the edge-determined partial fold graphs (depicted in the theorem statement's 2nd row) carry train track representatives of principal fully irreducible elements of $Out(F_3)$. For an arbitrary such graph Γ' represented by a directed edge \mathcal{E} of \mathcal{PA}_3 , suppose Γ' is obtained by partially folding the edge e' over the edge e in a graph Γ . Further, suppose $g: \Gamma \rightarrow \Gamma$ is an irreducible transparent train track representative of a principal fully irreducible outer automorphism whose Stallings fold decomposition is given by a directed loop \mathcal{L} in the automaton starting with a fold of e' over e . Such a loop exists by appending to the loop L' of Proposition 5 a loop traversing \mathcal{E} , then following the arguments in the previous paragraph. We claim, as follows, that there is a train track representative $g': \Gamma' \rightarrow \Gamma'$ of a principal fully irreducible element of $Out(F_3)$ obtained by conjugating g by this partial fold.

Since g is irreducible and transparent, $length(g(e)) \geq |E(\Gamma)| > 2$. Supposing first that the direction of e is fixed, we can write $g(e) = ee_1 \cdots e_k$ for some $k \geq 2$. Subdivide e into subedges so that $e = a_0 \cdots a_k$, and $g(a_0) = e$, and $g(a_j) = e_j$ for $j = 1, \dots, k$. Similarly subdivide e' as $a'_0 \cdots a'_{k+1}$ with $g(a'_0) = e$ and $g(a'_j) = e_j$ for $j = 1, \dots, k$. Fold a_0 and a'_0 to obtain an edge E . Call the fold f and let $g': \Gamma' \rightarrow \Gamma'$ be the quotient of g under f . Let g_1 be such that $g = g_1 \circ f$ (and so that then also $g' = f \circ g_1$). If e' is instead fixed, replace e with e' in $g(e)$ and $g(e')$ and then define g' the same. We show that g' is a train track map.

By construction, g' takes vertices to vertices. We now show each g'^k is locally injective on edge interiors. Notice that $g^{k+1} = g_1 \circ g'^k \circ f$. Suppose for the sake of contradiction that for some edge E' in Γ' , we have $g'^k(E')$ is not tight. We show that in this circumstance g would not be a train track map. We consider separately the cases where $E' = E$ and $E' \neq E$. In the latter case, let $a \in E(\Gamma)$ be such that $E' = f(a)$. Since g is a train track map, $g^{k+1}(a) = g_1 \circ g'^k \circ f(a) = g_1 \circ g'^k(E')$ must be tight. But $g'^k(E')$ was not tight and one cannot map a nontight path to a tight path (in this case via g_1), implying that $g^{k+1}(a)$ could not be tight, a contradiction. Now suppose $E' = E$. Then $f(e) = Ea_1 \cdots a_k$ and $g^{k+1}(e) = g_1 \circ g'^k \circ f(e) = g_1 \circ g'^k(Ea_1 \cdots a_k)$ must be tight. But $g'^k(E' = E)$ was not tight, and so $g'^k(Ea_1 \cdots a_k)$ cannot be tight, and thus $g^{k+1}(e)$ cannot be tight, again contradicting that g is a train track map.

Since the outer automorphisms represented by g and g' are $Out(F_3)$ -conjugate, g' must also represent a principal fully irreducible outer automorphism.

Proposition 6 proves the reverse containment. □

6. OUTER SPACE INTERPRETATION

Culler–Vogtmann Outer space was first defined in [CV86]. We refer the reader to [FM11, Bes14, Vog15] for background on Outer space and give only abbreviated discussion here. For $r \geq 2$ we denote the (volume-one normalized) Outer space for F_r by CV_r . Points of CV_r are equivalence classes of volume-1 marked metric graphs $h: R_r \rightarrow \Gamma$ where R_r is the r -rose, where Γ is a finite volume-1 metric graph with betti number $b_1(\Gamma) = r$ and with all vertices of degree ≥ 3 , and where h is a homotopy equivalence called a *marking*. There is an asymmetric metric d_{CV} on CV_r .

Given a Stallings fold decomposition of a train track map g , one can define a “periodic fold line” in CV_r . In [AKKP19] it is proved that such periodic fold lines of train track

maps are geodesics in the sense that, given 3 points $\gamma(t_1), \gamma(t_2), \gamma(t_3)$ on the geodesic γ with $t_1 < t_2 < t_3$, we have that $d_{CV}(\gamma(t_1), \gamma(t_2)) + d_{CV}(\gamma(t_2), \gamma(t_3)) \leq d_{CV}(\gamma(t_1), \gamma(t_3))$. One may call such a geodesic γ_g an *axis* for g and for φ , where φ is the outer automorphism represented by g . This is all explained carefully in [AKKP19].

CV_r has a simplicial complex structure with some faces missing. It has an open simplex for each marked graph (obtained by varying the lengths on the edges of that graph). The faces of a simplex are obtained by collapsing the edges of a forest (so that their lengths become zero).

The following is an easy Corollary of Main Theorem A.

Main Theorem B. The open simplices in CV_3 with principal axes passing through them are precisely those whose underlying graph is listed in Main Theorem A, i.e all those in the top 2 dimensions.

A few words on the proof: By the proofs of Proposition 6 and Main Theorem A, we have that the Stallings fold decomposition of any principal fully irreducible element of $Out(F_3)$ will consist only of folds between graphs represented by nodes in PA_3 , together with possibly part of a fold represented by an edge in PA_3 . At the start and end of each of the full folds, the axis will pass through the simplex of a 1st row graph. During the fold (or at the start of end of a partial fold), the axis will pass through the open simplices of 2nd row graphs. Varying the marking on a simplex simply $Out(F_3)$ -conjugates the outer automorphism, so does not change whether it is principle fully irreducible outer automorphism.

REFERENCES

- [AKKP19] Y. Algom-Kfir, I. Kapovich, and C. Pfaff. Stable strata of geodesics in outer space. *International Mathematics Research Notices*, 2019(14):4549–4578, 2019.
- [Bes14] M. Bestvina. Geometry of outer space. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 173–206. Amer. Math. Soc., Providence, RI, 2014.
- [BH92] M. Bestvina and M. Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, 1992.
- [CL15] T. Coulbois and M. Lustig. Index realization for automorphisms of free groups. *Illinois Journal of Mathematics*, 59(4):1111–1128, 2015.
- [CM15] D. Calegari and J. Maher. Statistics and compression of scl. *Ergodic Theory and Dynamical Systems*, 35(1):64–110, 2015.
- [Cou14] T. Coulbois. Free group automorphisms and train-track representative in python/sage. <https://github.com/coulbois/sage-train-track>, 2012–2014.
- [CP10] M. Clay and A. Pettet. Twisting out fully irreducible automorphisms. *Geometric And Functional Analysis*, 20(3):657–689, 2010.
- [CV86] M. Culler and K. Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones mathematicae*, 84(1):91–119, 1986.
- [FM11] S. Francaviglia and A. Martino. Metric properties of outer space. *Publicacions Matemàtiques*, 55(2):433–473, 2011.
- [Gag20a] D. Gagnier. Automaton maker. https://github.com/dsgagnier/automata_maker, 2020.
- [Gag20b] D. Gagnier. Edge map generator. https://github.com/dsgagnier/loop_investigation, 2020.
- [GM17] V. Gadre and J. Maher. The stratum of random mapping classes. *Ergodic Theory and Dynamical Systems*, pages 1–17, 2017.
- [HM11] M. Handel and L. Mosher. *Axes in outer space*. Number 1004. Amer Mathematical Society, 2011.

- [Hor16] C. Horbez. The horoboundary of outer space, and growth under random automorphisms. In *Annales Scientifiques de l'École Normale Supérieure*, volume 49, pages 1075–1123, 2016.
- [KM98] V. A. Kaimanovich and H. Masur. The Poisson boundary of Teichmüller space. *J. Funct. Anal.*, 156(2):301–332, 1998.
- [KMPT22a] I. Kapovich, J. Maher, C. Pfaff, and S. J. Taylor. Random trees in the boundary of outer space. *Geometry & Topology*, 26(1):127–162, 2022.
- [KMPT22b] I. Kapovich, J. Maher, C. Pfaff, and S.J. Taylor. Random outer automorphisms of free groups: Attracting trees and their singularity structures. *Transactions of the American Mathematical Society*, 375(01):525–557, 2022.
- [KP18] I. Kapovich and C. Pfaff. Counting conjugacy classes of fully irreducibles: double exponential growth. *arXiv preprint arXiv:1801.07471*, 2018.
- [Mah11] J. Maher. Random walks on the mapping class group. *Duke Mathematical Journal*, 156(3):429–468, 2011.
- [Mas82] H. Masur. Interval exchange transformations and measured foliations. *Ann. of Math*, 115(1):169–200, 1982.
- [MP16] L. Mosher and C. Pfaff. Lone axes in outer space. *Algebr. Geom. Topol.*, 16(6):3385–3418, 2016.
- [MS17] J. Maher and A. Sisto. Random subgroups of acylindrically hyperbolic groups and hyperbolic embeddings. *International Mathematics Research Notices*, 2017.
- [Pfa12] C. Pfaff. *Constructing and Classifying Fully Irreducible Outer Automorphisms of Free Groups*. PhD thesis, Rutgers University, 2012. PhD Thesis,<http://www.math.rutgers.edu/~cpfaff/Thesis.pdf>.
- [Pfa13] C. Pfaff. Ideal Whitehead graphs in $\text{Out}(F_r)$ II: the complete graph in each rank. *Journal of Homotopy and Related Structures*, 10(2):275–301, 2013.
- [Pfa15a] C. Pfaff. Ideal Whitehead graphs in $\text{Out}(F_r)$ III: Achieved graphs in rank 3. *Journal of Topology and Analysis*, 0(0):1–36, 2015.
- [Pfa15b] C. Pfaff. $\text{Out}(F_3)$ index realization. *Math. Proc. Cambridge Philos. Soc.*, 159(3):445–458, 2015.
- [Riv08] I. Rivin. Walks on groups, counting reducible matrices, polynomials, and surface and free group automorphisms. *Duke Math. J.*, 142(2):353–379, 2008.
- [Sis18] A. Sisto. Contracting elements and random walks. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2018(742):79–114, 2018.
- [Sta83] J.R. Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71(3):551–565, 1983.
- [Vog15] K. Vogtmann. On the geometry of outer space. *Bull. Amer. Math. Soc. (N.S.)*, 52(1):27–46, 2015.