DYADIC CANCELLATION OF BELTRAMI COEFFICIENTS

PETER LIN

Abstract. We prove that Beltrami coefficients on the square that are concatenated with random orientations tend to cancel each other out. As an application, we give a new perspective on stochastic homogenization.

1. Introduction

Let $S \subset \mathbb{C}$ be an open square. A *Beltrami coefficient* on S is a measurable function $\mu : S \to \mathbb{D}$. Such a coefficient specifies an alternate complex structure on S. See Figure 1.

The rotation group of the square, $\mathbb{Z}_4$, acts on Beltrami coefficients $\mu : S \to \mathbb{D}$ in the following way: if $r_{\pi/2} \in \mathbb{Z}_4$ denotes rotation by $\pi/2$, then $(r_{\pi/2} \cdot \mu)(z) = -\mu(r_{\pi/2}^{-1} z)$. This corresponds to a rotation of the square with its alternate complex structure.

We define the *energy* of a Beltrami coefficient on S:

$$\mathcal{E}_S(\mu) = \mathcal{E}(\mu) = \int_S \frac{1+|\mu|^2}{1-|\mu|^2} dxdy - \text{Area}(S).$$

Remark 1. Recall that $\mathbb{D}$ is often identified with the space of ellipses in $\mathbb{R}^2$ centered at the origin, modulo homotheties with respect to the origin. If E_λ denotes the ellipse associated to $\lambda \in \mathbb{D}$, then $E_{-\lambda}$ is the $\pi/2$ rotate of E_λ. We have $\frac{1+|\lambda|^2}{1-|\lambda|^2} = K_\lambda + K_\lambda^{-1}$ where K_λ is the eccentricity of the ellipse E_λ associated to λ, so $\mathcal{E}(\mu)$ measures the mean eccentricity of the ellipse field μ. Clearly $\mathcal{E}(\mu) = 0$ if and only if $\mu = 0$ almost everywhere, i.e. each ellipse is a round circle.

$\mathcal{E}(\mu) + \text{Area}(S)$ is also equal to the Dirichlet energy of h_μ^{-1} where h_μ is any homeomorphic solution to the Beltrami equation, see (2) in Section 2.

The *optimal energy* is $\mathcal{E}^*(\mu) = \inf_{\mu' \sim_{\text{Teich}} \mu} \mathcal{E}(\mu')$, where the Teichmüller equivalence $\mu' \sim_{\text{Teich}} \mu$ means that μ and μ' induce the same conformal structure relative to the boundary of S. See Figure 2 for an example, and see Section 2 for precise definitions.

If $\mu_1, \ldots, \mu_4 : (0,1)^2 \to \mathbb{D}$ are Beltrami coefficients, we define the *dyadic concatenation* to be the Beltrami coefficient $\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right] : (0,2)^2 \to \mathbb{D}$ obtained by dividing $(0,2)^2$ into 4 subsquares and then placing copies of μ_i into the appropriate subsquares. Linearity of integration implies

$$\mathcal{E}^*\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right] \le \mathcal{E}\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right] = \sum_{i=1}^{4} \mathcal{E}(\mu_i). \tag{1}$$

It is well known that equality can hold, for instance if the μ_i are constant and equal to the same constant.

Date: March 30, 2023.

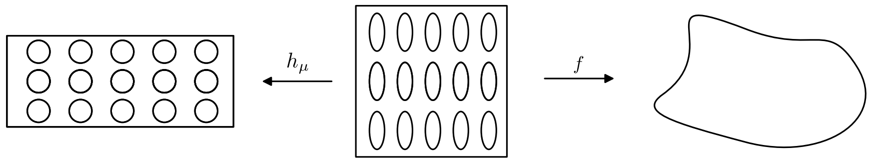

Figure 1. If $\mu : [0,1]^2 \to \mathbb{D}$ is a Beltrami coefficient, it can be interpreted as an ellipse field on $[0,1]^2$ (middle). This, in turn, can be interpreted as a deformation of the standard complex structure on $[0,1]^2$, in which the given ellipses are declared to be round. In particular, if μ has an integrating map h_μ, it defines a quadratic form $\mathrm{Dir}_\mu(f) := \mathrm{Dir}(f \circ h_\mu^{-1})$ on the space of functions $f : [0,1]^2 \to \mathbb{C}$. Equation (2) says that $\mathcal{E}(\mu) + 1 = \mathrm{Dir}(h_\mu^{-1}) = \mathrm{Dir}_\mu(\mathrm{Id})$. In the case depicted above, when μ is constant, h_μ is given by an affine stretch.

Our main result says that if any collection of nontrivial Beltrami coefficients are concatenated with randomly chosen orientations, then the inequality is likely to be strict, in a quantititive way. A *random Beltrami coefficient* is a random variable taking values in the space of Beltrami coefficients on $[0,1]^2$. See Section 2 for more precise definitions.

Theorem 1.1. *For $k < 1$ there exists a constant $c > 0$ such that the following is true.*

Let μ be a random Beltrami coefficient on $(0,1)^2$ with rotationally invariant distribution, and assume that $\|\mu\|_\infty < k$ almost surely.

Let $\mu_1, \ldots, \mu_4$ be i.i.d. samples of μ, and let $r_1, \ldots, r_4$ be i.i.d. samples from $\mathbb{Z}_4$. $\nu = \left[\begin{smallmatrix} r_1\mu_1 & r_2\mu_2 \\ r_3\mu_3 & r_4\mu_4 \end{smallmatrix}\right]$ be the concatenation of the μ_i. Then

$$\mathbb{E}\mathcal{E}^*\nu \le 4\mathbb{E}\mathcal{E}^*(\mu) - c \cdot T_1\left(\mathbb{E}\mathcal{E}^*(\mu)\right),$$

where $T_1(x) = \min(x^3, x)$.

Example 1. Suppose we take μ to be deterministically equal to a constant Beltrami coefficient $\mu \equiv k_0 < 1$, so that $\mu_i \equiv k_0$ for each i. If the rotations r_i are all aligned, then $\mathcal{E}^*(\nu) = \mathcal{E}(\nu) = 4\mathcal{E}^*(\mu)$. On the other hand, if the r_i are not aligned (see Figure 2), then $\mathcal{E}^*(\nu) < \mathcal{E}(\nu)$. The conclusion of Theorem 1.1 quantifies the strictness of the inequality.

Note that in this particular example, our theorem is not sharp. It can be shown directly that there exists $c > 0$ such that whenever μ is constant, $\mathbb{E}\mathcal{E}^*(\nu) \le 4\mathbb{E}\mathcal{E}^*(\mu) - c\mathcal{E}(\mu)$.

In fact, we obtain quantitative control even away from the uniformly elliptic or quasiconformal setting. The theorem below obviously implies the previous theorem.

In the theorem statement, *self compatibility* means that the concatenation $\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right]$ has a unique (up to postcomposition by conformal map) homeomorphic solution almost surely. A simple class of self compatible random Beltrami coefficients are those that satisfy $\|\mu\|_\infty < 1$ almost surely.

Theorem 1.2. *There exists $c_0 > 0$ such that the following is true. Let μ be a self compatible random Beltrami coefficient on $(0,1)^2$.*

Assume that there exists $\epsilon > 0$ such that $M = 1 + \mathbb{E}\mathcal{E}(\mu)^{1+\epsilon}$ is finite.

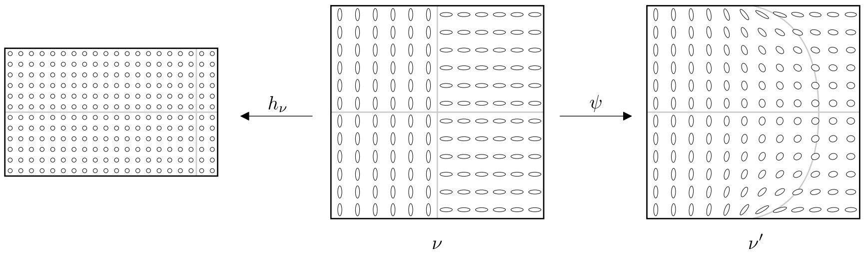

Figure 2. In the middle, we have the concatenation, ν, of 4 constant coefficient Beltrami coefficients (in fact, $|\mu| \equiv 1/2$ and the ellipses have eccentricity 3 here). The integrating map for this concatenation, h_ν is a piecewise linear map to a rectangle. Even though ν is optimal on each of the four subsquares, the coefficient as a whole is not optimal. Indeed, the cofficient ν' on the right satisfies $\mathcal{E}(\nu') < \mathcal{E}(\nu)$, and $\nu \sim_{\text{Teich}} \nu'$. In other words, there exists $g : [0,1]^2 \to [0,1]^2$ mapping the ellipse field ν' to ν, and $g = \text{Id}$ on $\partial[0,1]^2$.

Let $\mu_1, \ldots, \mu_4$ be i.i.d. samples of μ, and let $r_1, \ldots, r_4$ be i.i.d. samples from $\mathbb{Z}_4$. $\nu = \left[\begin{smallmatrix} r_1\mu_1 & r_2\mu_2 \\ r_3\mu_3 & r_4\mu_4 \end{smallmatrix}\right]$ be the concatenation of the μ_i. Then

$$\mathbb{E}\mathcal{E}^*(\nu) \leq 4\mathbb{E}\mathcal{E}^*(\mu) - c_0 5^{-2(1+\epsilon)/\epsilon} M^{-2/\epsilon} \cdot T_1\left(\mathbb{E}\mathcal{E}^*(\mu)\right),$$

where $T_1(x) = \min(x^3, x)$.

More generally, if $\psi : [0,\infty) \to [0,\infty)$ is monotone, nonnegative, convex, and vanishes at 0, then

$$\mathbb{E}\psi(\mathcal{E}^*(\nu)) \leq 4\mathbb{E}\psi(\mathcal{E}^*(\mu)) - c_0 5^{-2(1+\epsilon)/\epsilon} M^{-2/\epsilon} \cdot T_1\left(\mathbb{E}\psi\mathcal{E}^*(\mu)\right),$$

where now $M = 1 + \mathbb{E}\mathcal{E}(\mu)^\epsilon \psi(\mathcal{E}(\mu))$.

Proof of Theorem 1.1 from Theorem 1.2. If $\|\mu\|_\infty < k$, then $\mathcal{E}(\mu) \leq \frac{1+k^2}{1-k^2} - 1$. If this holds almost surely, then $\mathbb{E}\mathcal{E}(\mu)^2$ is bounded by a constant only depending on k. □

The theorem is an instance of the phenomenon that deformations of complex structures tend to cancel each other out, unless they are 'aligned'. We describe two other instances where this phenomenon appears.

Thurston proved that rational maps (with hyperbolic orbifold) are combinatorially rigid [DH84, Corollary 3.4] by showing that a certain map on a finite dimensional Teichmüller space is strictly contracting. Roughly speaking, the mechanism for this contraction can be interpreted in the following way: infinitesmal conformal deformations undergo cancellation under a type of concatenation operation related to the one we defined above.

Our theorem differs in that it is global, quantitative, and applies to an infinite-dimensional space. It can be used to construct conformal uniformizations of subdivision rules that do not arise from branched covers. Such subdivision rules are of interest in relation to Cannon's conjecture and the quasisymmetric uniformization problem. These applications will be covered in a separate article.

Cancellation of conformal deformations is also illustrated in the study of stochastic homogenization, or random walks on random environments. A prototypical result in the area is that Brownian motion with respect to ergodic conformal deformations is well approximated by an affine Brownian motion. See, for example, [ARST20, IM19]. In Section 4 we demonstrate how one such statement follows easily from our theorem.

We now turn to the proof of Theorem 1.2. It is a corollary of the following deterministic statement.

Theorem 1.3. *There exists $c_0 > 0$ such that the following is true. Let $\{\mu_1, \dots, \mu_5\}$ be a compatible collection of Beltrami coefficients.*

There exists rotations $r_1, \dots, r_4 \in \mathbb{Z}_4$ and an injective map $\iota : \{1, \dots, 4\} \to \{1, \dots, 5\}$ such that

$$\mathcal{E}^*\left[\begin{smallmatrix} r_1\mu_{\iota(1)} & r_2\mu_{\iota(2)} \\ r_3\mu_{\iota(3)} & r_4\mu_{\iota(4)} \end{smallmatrix}\right] \le \sum_{k=1}^{4} \mathcal{E}^*(\mu_{\iota(k)}) - c_0 T_{-1}(\max_i \mathcal{E}^*(\mu_i)),$$

where $T_{-1}(x) = \min(x^3, x^{-1})$.

We conclude the introduction by showing how Theorem 1.2 follows from this deterministic statement. Section 2 contains preliminary material about Beltrami coefficients and some elementary observations about the Dirichlet energy. Section 3 is devoted to the proof of Theorem 1.3. Finally, Section 4 demonstrates how the theorem applies to stochastic homogenization.

Proof of Theorem 1.2 from Theorem 1.3. Let $\{\mu_1, \dots, \mu_5\}$ be a fixed collection of self compatible Beltrami coefficients on $(0,1)^2$.

By Theorem 1.3, there exists an injective $\iota : \{1,2,3,4\} \hookrightarrow \{1,2,3,4,5\}$ and $r : \{1,2,3,4\} \to \mathbb{Z}_4$ such that

$$\mathcal{E}^*\begin{bmatrix} r(1)\mu_{\iota(1)} & r(2)\mu_{\iota(2)} \\ r(3)\mu_{\iota(3)} & r(4)\mu_{\iota(4)} \end{bmatrix} \le \sum_{k=1}^{4} \mathcal{E}(\mu^*_{\iota(k)}) - c_0 T_{-1}(\max_{i\le 5} \mathcal{E}^*(\mu_i)),$$

where $T_p(x) = \max(x^3, x^p)$. Thus if $\iota : \{1,2,3,4\} \hookrightarrow \{1,2,3,4,5\}$ is an uniformly random injective function and $r : \{1,2,3,4\} \to \mathbb{Z}_4$ is a uniformly random choice of rotation, we have

$$\mathbb{E}_{\iota,r}\left(\mathcal{E}^*\begin{bmatrix} r(1)\mu_{\iota(1)} & r(2)\mu_{\iota(2)} \\ r(3)\mu_{\iota(3)} & r(4)\mu_{\iota(4)} \end{bmatrix} - \sum_{k=1}^{4} \mathcal{E}^*(\mu_{\iota(k)})\right) \le -tc_0 \cdot T_{-1}(\max_i \mathcal{E}^*(\mu_i)),$$

where $t = \frac{4}{5\cdot 4\cdot 3\cdot 2\cdot 4^4}$. On the other hand, if we now let $\mu_1, \dots, \mu_5$ be five i.i.d. copies (and independent of the random functions ι and r) of some random Beltrami coefficient μ satisfying the hypotheses of the theorem, then $(r(1)\mu_{\iota(1)}, \dots, r(4)\mu_{\iota(4)})$ is equal in distribution to $(\mu_1, \dots, \mu_4)$. So taking the expectation of the previous inequality yields

$$\mathbb{E}_\mu \mathcal{E}^*\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right] = \mathbb{E}_{\mu,\iota,r}\mathcal{E}^*\begin{bmatrix} r(1)\mu_{\iota(1)} & r(2)\mu_{\iota(2)} \\ r(3)\mu_{\iota(3)} & r(4)\mu_{\iota(4)} \end{bmatrix} \le \mathbb{E}_\mu \sum_{k=1}^{4} \mathcal{E}^*(\mu_{\iota(k)}) - tc_0\mathbb{E}_\mu T_{-1}\left(\max_i \mathcal{E}^*(\mu_i)\right).$$

In the rest of the proof we will use the notation $a \gtrsim b$ to mean that $a \le Cb$ where the constant C does not depend on μ or ϵ, and write $a \asymp b$ to mean $a \gtrsim b$ and $b \gtrsim a$.

Set $X = \max_i \mathcal{E}^*(\mu_i)$ and note that $X \le \sum_i \mathcal{E}^*(\mu_i)$ so that $\mathbb{E}X^{1+\epsilon} \le 5^{1+\epsilon}M$, and $T_p(\max_i \mathcal{E}^*(\mu_i)) \asymp T_p(X)$ for $p \in \{-1, 1\}$.

The proof will be complete once we show that $\mathbb{E}_\mu T_{-1}(X) \gtrsim 5^{-2(1+\epsilon)/\epsilon} M^{-2/\epsilon} T_1(\mathbb{E}_\mu X)$.

Let $\Theta > 1$ be a constant to be chosen later. By considering the events $X < \Theta^{-1}$, $X \in [\Theta^{-1}, \Theta]$ and $X > \Theta$ separately, we get

$$\begin{aligned}\mathbb{E}_\mu X &\le \Theta^{-1} + \Theta \cdot \mathbb{P}(X \in [\Theta^{-1}, \Theta]) + \mathbb{E}_\mu X \mathbb{1}\{X \ge \Theta\} \\ &\le \Theta^{-1} + \Theta \cdot \mathbb{P}(X \in [\Theta^{-1}, \Theta]) + \Theta^{-\epsilon}\mathbb{E}_\mu X^{1+\epsilon}.\end{aligned}$$

Thus, we can find $\Theta \lesssim \max(5^{(1+\epsilon)/\epsilon} M^{1/\epsilon}, (\mathbb{E}_\mu X)^{-1})$ such that $\Theta \cdot \mathbb{P}(X \in [\Theta^{-1}, \Theta]) \gtrsim \mathbb{E}_\mu(X)$. For this Θ,

$$\begin{aligned}\mathbb{E}_\mu T_{-1}(X) &\ge \Theta^{-1}\mathbb{P}(X \in [\Theta^{-1}, \Theta]) \gtrsim \Theta^{-2}\mathbb{E}_\mu(X) \asymp \Theta^{-2}\mathbb{E}_\mu(\mathcal{E}(\mu)) \\ &\gtrsim 5^{2\epsilon/(1+\epsilon)} M^{-2/\epsilon} \min(\mathbb{E}_\mu(\mathcal{E}(\mu))^3, \mathbb{E}_\mu \mathcal{E}(\mu)).\end{aligned}$$

as desired. □

2. Preliminaries

In this paper, Ω will always be an open Jordan domain in the complex plane, or the complex plane itself, and S will always be an open rectangle in the plane. Let $C(\overline{\Omega})$ be the space of continuous $\mathbb{C}$-valued functions with continuous extension to the boundary. Let $C_0(\overline{\Omega})$ be the subspace of functions that vanish on the boundary. For $p \ge 1$, let $W^{1,p}(\Omega)$ be the Sobolev space of $\mathbb{C}$-valued functions whose real and imaginary parts have weak derivatives in $L^p(\Omega)$.

2.1. Beltrami Coefficients.

See [Ast09] for detailed background on Beltrami coefficients and the Beltrami equation, the following section is only intended to set up terminology.

A *Beltrami coefficient* on S is a measurable function $\mu : S \to \mathbb{D}$, considered up to a.e. equivalence. An *integrating homeomorphism* for μ is a $W^{1,1}(S, \mathbb{C})$ homeomorphism $h_\mu : S \to \Omega$ with the following properties: $\Omega \subset \mathbb{C}$ is a Jordan domain, and h_μ solves the *Beltrami equation* $\bar\partial h_\mu = \mu \partial h_\mu$, and h extends continuously to a homeomorphism $\overline{S} \to \overline{\Omega}$.

We say that μ is *integrable* if $\mathcal{E}(\mu) < \infty$ and there exists an integrating homeomorphism for μ. Geometrically, μ can be interpreted as an infinitesimal ellipse field, and an integrating homeorphism is simply a homeomorphism which maps each ellipse of μ to a round circle. See Figure 1.

If $h : S \to \Omega$ is an integrating homeomorphism and $\varphi : \Omega \to \tilde\Omega$ is a conformal map of Jordan domains, then $\varphi \circ h$ is also an integrating homeomorphism.

Therefore if μ is integrable, one can freely choose the codomain of the integrating homeomorphism.

We say that μ is *uniquely integrable* if any two integrating homeomorphisms $h, \tilde h$ are related in the preceding way: $h \circ \tilde h^{-1}$ is complex analytic on its domain of definition. We say that μ is *strongly uniquely integrable* if $\mu|_{S'}$ is uniquely integrable for all open rectangles $S' \subset S$.

We now give the precise definition of Teichmüller equivalence used in this paper.

Definition 2.1. *Let $\mu : S \to \mathbb{D}$ and $\mu' : S \to \mathbb{D}$ be integrable Beltrami coefficients. We write $\mu \sim_{\text{Teich}} \mu'$ if there exists integrating maps h_μ and $h_{\mu'}$ that are equal on ∂S.*

Teichmüller equivalence automatically implies the following stronger equivalence:

Lemma 2.1. *Suppose $\mu \sim_{\text{Teich}} \mu'$ and suppose μ is uniquely integrable. Let $\tilde{h}_\mu : S \to \Omega$ be an integrating map for μ. Then there exists an integrating map $\tilde{h}_{\mu'} : S \to \tilde{\Omega}$ for μ' such that $\tilde{h}_{\mu'} = \tilde{h}_\mu$ on ∂S.*

Proof. Since $\mu \sim_{\text{Teich}} \mu'$, there exists integrating maps $h_\mu : S \to \Omega$ and $h_{\mu'} : S \to \Omega$ that are equal on ∂S. Since μ is uniquely integrable, $\tilde{h}_\mu = \varphi \circ h_\mu$ for some conformal $\varphi : \Omega \to \tilde{\Omega}$. By Carathéodory's theorem, φ extends continuously to a homeomorphism of the closures, so $\tilde{h}_{\mu'} := \varphi \circ h_{\mu'}$ is the desired integrating map. □

To avoid technicalities regarding measurability and σ-algebras for probability measures, we will only consider random variables taking on a finite set of values.

Definition 2.2 (Random Beltrami Coefficient)**.** *A random Beltrami coefficient on S is a random variable taking values in some finite set of uniquely integrable Beltrami coefficients on S.*

2.2. **Compatibility of Beltrami coefficients.** The dyadic concatenation of uniquely integrable Beltrami coefficients is not necessarily uniquely integrable or even integrable. This is related to the conformal welding problem. This presents a small technical obstacle to iterating the dyadic concatenation operation (and the reader is encouraged to ignore the notion of compatibility and skip this section on first read).

However, the only way that this can happen is if the 'interface' of the concatenation is badly behaved.

Therefore, this issue can be avoided by making the minimal assumptions on the local regularity of the boundary behaviour, which will automatically be preserved under dyadic concatenation and Teichmüller equivalence. We clarify this below.

A quadruplet $\mu_1, \ldots, \mu_4$ of Beltrami coefficients is said to be *compatible* if $\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right]$ is strongly uniquely integrable. It is said to be *strongly compatible* if for every choice of rotations $r : \{1, 2, 3, 4\} \to \mathbb{Z}_4$ and permutation $\iota : 1, \ldots, 4$, the quadruplet $(r\mu_{\iota i}, \ldots, r\mu_{\iota_i})$ is compatible. A collection of Beltrami coefficients is said to be self compatible if every quadruplet from that collection is strongly compatible.

A random Beltrami coefficient μ is said to be self compatible if it is almost surely true that when $\mu_1, \ldots, \mu_4$ are i.i.d. samples of μ, then $(\mu_1, \ldots, \mu_4)$ is strongly uniquely integrable.

As an example, it is immediate from the measurable Riemann mapping theorem that any collection of coefficients μ satisfying $\|\mu\|_\infty < 1$ is a compatible collection, and any random Beltrami coefficient satisfying $\|\mu\|_\infty < 1$ almost surely is automatically self compatible.

Since this already covers many interesting applications, the reader is encouraged to skip the rest of this section on first read.

To iterate Theorem 1.2, we need that self compatibility is closed under dyadic concatenation.

Proposition 2.1. *Suppose μ is a random Beltrami coefficient that is self compatible. If $\mu_1, \ldots, \mu_4$ are i.i.d. with distribution μ, then the random Beltrami coefficient $\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right]$ is also self compatible.*

Proof. Let $\nu_1, \ldots, \nu_4$ be i.i.d. copies of $\left[\begin{smallmatrix} \mu_1 & \mu_2 \\ \mu_3 & \mu_4 \end{smallmatrix}\right]$. Then the concatenated coefficient $\nu := \left[\begin{smallmatrix} \nu_1 & \nu_2 \\ \nu_3 & \nu_4 \end{smallmatrix}\right]$ may be thought of as the concatenation of 16 i.i.d. coefficients with distribution μ. This 16-fold concatenation can be covered by a finite number of overlapping open 2×2 squares.

By the self compatibility assumption on μ and the union bound, it is almost surely true that the restriction of ν to each of these open squares is strongly uniquely integrable. If this is the case, then Lemma 2.2 below implies that ν itself is uniquely integrable, so we are done. $\square$

Above, we used the following locality of strong unique integrability.

Lemma 2.2 (Locality of strong unique integrability)**.** *Let $\mu : S \to \mathbb{D}$ be a Beltrami coefficient and let $\mathcal{U}$ be an open covering of S by rectangles. Suppose for each $U \in \mathcal{U}$, $\mu|_U$ is strongly uniquely integrable. Then μ is integrable and in fact strongly uniquely integrable.*

Proof. Let S' be a subrectangle of S and let $\mathcal{U}' = \{U \cap S' : U \in \mathcal{U}\}$ be the restriction of $\mathcal{U}$ to S'. For each $U \in \mathcal{U}$, fix an integrating map $h_U : U \to \Omega_U$ for $\mu|_U$. The intersection between each $U, V \in \mathcal{U}'$ is either empty or a rectangle, so if all the $\mu|_U$ are strongly uniquely integrable, then the collection of maps $\{h|_U\}$ forms a complex chart for S'. By the uniformization theorem, there exists a global homeomorphism $h : S' \to X$ where X is $\mathbb{D}$ or $\mathbb{C}$. Extremal length considerations rule out this latter possibility, so $X = \mathbb{D}$.

This h is compatible with all the charts, that is, for each U, the map $h \circ h_U^{-1}$ is a conformal. Therefore h is an integrating homeomorphism for μ on S'.

Now let $h, \tilde{h}$ be any integrating homeomorphisms for $\mu|_{S'}$. By strong unique integrability of the $\mu|_U$, these homeomorphisms are compatible with the charts $(h_U)_{U \in \mathcal{U}'}$, and in hence with each other, on each U.

That is, $h \circ \tilde{h}^{-1}$ is conformal on each U, and hence is conformal on S'. $\square$

In the proof of Theorem 1.3, we start by replacing each Beltrami coefficient with its optimal representative. For this to be valid, we need to show that this operation does not affect compatibility.

In the next two lemmas, suppose $\mu_1, \ldots, \mu_4$ and $\tilde{\mu}_1, \ldots, \tilde{\mu}_4$ are uniquely integrable Beltrami coefficients on $(0,1)^2$. Let $\mu, \tilde{\mu} : (0,2)^2 \to \mathbb{D}$ be the respective dyadic concatenations.

Lemma 2.3. *Let $h_\mu : [0,2]^2 \to \Omega$ be an integrating map for μ. Then there exists an integrating map $h_{\tilde{\mu}}$ for $\tilde{\mu}$ which is equal to h_μ on $\cup_{i=1}^4 \partial S_i$.*

Proof. Let $h_\mu|_1, \ldots, h_\mu|_4$ be the restriction of h_μ to each of the four subsquares of $(0,2)^2$. Then each $h_\mu|_i$ is an integrating map for μ_i. By Lemma 2.1, there exists $h_{\tilde{\mu}}|_i \sim_\partial h_\mu|_i$ which integrate $\tilde{\mu}_i$. These maps glue together to a homeomorphism $h_{\tilde{\mu}} : [0,2]^2 \to \Omega$, which integrates $\tilde{\mu}$ and has the desired property. $\square$

Lemma 2.4. *Suppose $\mu_i \sim_{\mathrm{Teich}} \tilde{\mu}_i$ for each i.*

Then $\mu \sim_{\mathrm{Teich}} \tilde{\mu}$, and μ is uniquely integrable iff $\tilde{\mu}$ is uniquely integrable.

Proof. The first conclusion follows from the preceding Lemma 2.3. Now suppose that $\tilde{\mu}$ is uniquely integrable.

Let h_μ and h'_μ be integrating maps for μ. By postcomposition with a conformal map, we can assume that they have the same codomain Ω and that they are equal on 3 of the vertices of $[0,2]^2$. By the preceding lemma, we get integrating maps $h_{\tilde{\mu}}$ and $h'_{\tilde{\mu}}$ for $\tilde{\mu}$ which are equal to h_μ and h'_μ respectively on $\cup_{i=1}^4 \partial S_i$. In particular, $h_{\tilde{\mu}}$ is equal to $h'_{\tilde{\mu}}$ on three of the vertices of $[0,2]^2$. Since $\tilde{\mu}$ is uniquely integrable, this implies that $h_{\tilde{\mu}} = h'_{\tilde{\mu}}$, which implies that $h_\mu = h'_\mu$ on $\cup_{i=1}^4 \partial S_i$. Since the μ_i are uniquely integrable, this implies that the restriction of each h_μ to each subsquare S_i is equal to the restriction of h'_μ to S_i. It follows that $h_\mu = h'_\mu$. □

2.3. **Dirichlet Energy.** For $f, g \in C(\overline{\Omega}) \cap W^{1,2}(\Omega)$ we have the *Dirichlet inner product*:

$$\mathrm{Dir}_\Omega(f,g) = \frac{1}{2}\int_\Omega (\nabla \mathrm{Re} f)\cdot(\nabla \mathrm{Re} g) + (\nabla \mathrm{Im} f)\cdot(\nabla \mathrm{Im} g)\; dxdy.$$

By abuse of notation we write $\mathrm{Dir}_\Omega(f) := \mathrm{Dir}_\Omega(f,f)$ for the *Dirichlet energy* of f. If $\Omega' \subset \Omega$ is a subdomain we write $\mathrm{Dir}_{\Omega'}(f,g)$ to mean $\mathrm{Dir}(f|_{\Omega'}, g|_{\Omega'})$.

We define $\mathrm{Dir}^*_\Omega(f) = \inf_{\tilde{f}\sim f} \mathrm{Dir}_\Omega(f)$ where the infimum is over $\tilde{f} \in C(\overline{\Omega}) \cap W^{1,2}(\Omega)$ such that $\tilde{f} = f$ on $\partial\Omega$.

If the Beltrami coefficient $\mu : S \to \mathbb{D}$ is uniquely integrable, then it endows S with the structure of a Riemann surface (S,μ), and all the definitions above relating to the Dirichlet inner product and harmonic functions can be extended to such surfaces.

Concretely, suppose $h_\mu : S \to \Omega$ is any integrating homeomorphism for μ. Let $W^{1,2}(S,\mu)$ be the space of functions $f : S \to \mathbb{C}$ such that $f \circ h_\mu^{-1} \in W^{1,2}(\Omega)$. For $f, g \in W^{1,2}(S,\mu)$, define the Dirichlet inner product

$$\mathrm{Dir}_{S,\mu}(f,g) := \mathrm{Dir}_\Omega(f\circ h_\mu^{-1}, g \circ h_\mu^{-1}),$$

Since the Dirichlet energy is invariant under conformal precomposition, these definitions do not depend on the choice of h_μ.

If $S' \subset S$ is an open Jordan domain, we write (abusing notation) $\mathrm{Dir}_{S',\mu}(f,g) := \mathrm{Dir}_{S',\mu|_{S'}}(f|_{S'}, g|_{S'})$.

A computation involving the chain rule shows that, formally, if h_μ is any integrating homeomorphism for μ, then the mean dilatation of μ is equal to the Dirichlet energy of h_μ^{-1}:

$$\mathcal{E}_S(\mu) = \mathrm{Dir}_\Omega(h_\mu^{-1}) - \mathrm{Area}(S) = \mathrm{Dir}_{S,\mu}(\mathrm{Id}_S) - \mathrm{Area}(S), \tag{2}$$

where $\mathrm{Area}(S)$ is the Lebesgue measure of S. In fact, the identity holds under the minimal regularity assumptions on μ and h_μ, see [HKO05, Theorem 2.1] (see also [AIMO05]),

From this we immediately conclude that $\mathcal{E}^*(\mu) \geq \mathrm{Dir}^*_{S,\mu}(\mathrm{Id}) - \mathrm{Area}(S)$, because the minimization problem on the right hand side is over a larger class of functions (possibly non-homeomorphisms). However, the Radó-Kneser-Choquet theorem [Dur04, Chapter 3] guarantees that if S is a convex Jordan region, and if $f : \Omega \to S$ is extends to a homeomorphism of the boundaries, then the Poisson extension of f is a smooth homeomorphism $f : \overline{\Omega} \to \overline{S}$ with nonvanishing Jacobian on the interior. Hence we get the reverse inequality:

$$\mathcal{E}^*(\mu) = \mathrm{Dir}^*_{S,\mu}(\mathrm{Id}) - \mathrm{Area}(S). \tag{3}$$

In summary, the connection between L^1 dilatation optimizers and Dirichlet energy optimizers implies:

Lemma 2.5 ([HKO05, Theorem 1.1])**.** *Let $\mu : S \to \mathbb{D}$ be a strongly uniquely integrable Beltrami coefficient with $\mathcal{E}(\mu) < \infty$.*

There exists $\mu^* : S \to \mathbb{D}$ with $\mu \sim_{\mathrm{Teich}} \mu^*$ such that $\mathcal{E}^*(\mu) = \mathcal{E}(\mu^*)$, and μ^* is strongly uniquely integrable.

Proof. The only part left to explain is that μ^* is uniquely integrable. By the inverse function theorem, μ^* is smooth. Smooth Beltrami coefficients are locally quasiconformal, so they are (strongly) uniquely integrable by the measurable Riemann mapping theorem. □

Remark 2. Instead of using Beltrami coefficients to specify complex structures on S, we may as well define a complex structure on S to be an equivalence class of homeomorphisms $h : X \to S$ where X is a simply connected hyperbolic Riemann surface, and h_1, h_2 are considered equivalent iff $h_2 \circ h_1^{-1}$ is a conformal homeomorphism. Then the energy of a complex structure h defined simply as $\mathrm{Dir}_X(h,h) - \mathrm{Area}(S)$.

This approach avoids issues such as the one considered above, and other differentiability issues, which are immaterial to the content and proofs of our main Theorems 1.2 and 1.3. The concatenation operator $[\,\because\,]$ can be interpreted as isometric welding of the Riemann surfaces X_i induced by the homeomorphisms $h : \partial X_i \to \partial S$.

In the proof of Theorem 1.3 we show that the energy improves by finding useful variations η of Id_S with respect to $\mathrm{Dir}_{S,\mu}$. The following variational characterization of the optimal energy clarifies exactly what is needed from our variations.

Lemma 2.6. *Let $\mu : S \to \mathbb{D}$ be a uniquely integrable Beltrami coefficient. For $f \in C(\overline{S}) \cap W^{1,2}(S,\mu)$,*

$$\mathrm{Dir}^*_{S,\mu}(f) \le \mathrm{Dir}_{S,\mu}(f) - \sup_{\eta \in C(\overline{S}) \cap W_0^{1,2}(S,\mu)} \frac{\mathrm{Dir}_{S,\mu}(f,\eta)^2}{\mathrm{Dir}_{S,\mu}(\eta)} \tag{4}$$

and

$$\mathcal{E}^*(\mu) \le \mathcal{E}(\mu) - \sup_{\eta \in C(\overline{S}) \cap W_0^{1,2}(S,\mu)} \frac{\mathrm{Dir}_{S,\mu}(\mathrm{Id}_S,\eta)^2}{\mathrm{Dir}_{S,\mu}(\eta)}. \tag{5}$$

Proof. Bilinearity of the Dirichlet inner product implies that for all $t \in \mathbb{R}$, we have

$$\mathrm{Dir}(f + t\eta) = \mathrm{Dir}(f) + t^2 \mathrm{Dir}(\eta) + 2t\mathrm{Dir}(f,\eta).$$

Optimizing the right hand side over $t \in \mathbb{R}$ shows that

$$\inf_{t \in \mathbb{R}} \mathrm{Dir}(f + t\eta) = \mathrm{Dir}(f) - \frac{\mathrm{Dir}(f,\eta)^2}{\mathrm{Dir}(\eta)} \tag{6}$$

and that the infimum is attained. This immediately proves the first statement.

The second statement follows from taking $f = \mathrm{Id}_S$, subtracting $\mathrm{Area}(S)$ from both sides, and applying (2) and (3), □

We will need the following improved version of the Cauchy-Schwarz inequality.

Corollary 2.1 (Cauchy-Schwarz)**.** *If $\mu : S \to \mathbb{D}$ is a Beltrami coefficient and $\eta \in W_0^{1,2}(S,\mu)$, then*

$$\mathrm{Dir}_{S,\mu}(\eta, \mathrm{Id}) \le \mathrm{Dir}_{S,\mu}(\eta)^{1/2} \mathcal{E}(\mu)^{1/2}.$$

Proof. Rearrange (5), using the fact that $\mathcal{E}^*(\mu) \ge 0$. □

2.4. **Localizing variations.** A variation $\eta \in C(\overline{S}) \cap W^{1,2}(S,\mu)$ which does not vanish on ∂S can be localized so that it does vanish on ∂S by multiplying it by a bump function. The chain rule controls how the Dirichlet energy may increase under this localization:

Lemma 2.7. *Suppose $\eta \in C(\overline{S}) \cap W^{1,2}(S,\mu)$ and let $\rho : \overline{S} \to [0,1]$ be a C^2 function. Then $\mathrm{Dir}_{S,\mu}(\eta \cdot \rho) \leq \mathrm{Dir}_{S,\mu}(\eta) + C_\rho^2 \mathrm{Dir}_{S,\mu}(\mathrm{Id}) \cdot \|\eta\|_\infty^2$, where $C_\rho = \|\nabla \rho\|_\infty$.*

Proof. By the product rule, $\mathrm{Dir}_{S,\mu}(\eta \cdot \rho) \leq \mathrm{Dir}_{S,\mu}(\eta)\|\rho\|_\infty^2 + \mathrm{Dir}_{S,\mu}(\rho)\|\eta\|_\infty^2$. The chain rule then implies $\mathrm{Dir}_{S,\mu}(\rho) = \mathrm{Dir}(\rho \circ h_\mu^{-1}) \leq \|\nabla\rho\|_\infty^2 \mathrm{Dir}(h_\mu^{-1})$, and the result follows. □

3. Proof of Theorem 1.3

Let $\{\mu_1, \ldots, \mu_5\}$ be self compatible Beltrami coefficients on the unit square. We may assume that they have finite energy, otherwise the conclusion of the theorem is vacuous.

By Lemma 2.4, we can assume that μ_i are already optimal, that is $\mathcal{E}^*(\mu_i) = \mathcal{E}(\mu_i)$. Such optimal representatives exist by Lemma 2.5.

The idea of the proof is as follows. By concatenating the coefficients $\mu_1, \ldots, \mu_5$ in a symmetric way, we get a symmetric coefficient on $\mathbb{C}$. In this symmetric setting, we can construct a variation that reduces the energy. Using a partition of unity, and the pigeonhole principle, we can localize this variation so that it is supported on one of the small 2×2 subsquares. Thus we get the desired improvement on some 2×2 subsquare.

Before presenting the details in the argument above, we fix some notation for the rest of this section. Assume that the μ_i are indexed so that $\mathcal{E}(\mu_1) \geq \cdots \geq \mathcal{E}(\mu_5)$. Fix $N \geq 15$ odd and large. We write $a \lesssim b$ to mean that $a \leq Cb$ where C is a constant that does not depend on $N, \mu_1, \ldots, \mu_5$. Similarly for $a \gtrsim b$. We write $a \asymp b$ to mean $a \gtrsim b$ and $a \lesssim b$.

3.1. **Constructing a symmetric Beltrami coefficient.** It is crucial that the symmetric concatenation of the coefficients is done in such a way that each 2×2 subsquare (with a few exceptions) contains 4 distinct coefficients.

We specify this concatenation pattern by defining a labelling of the square grid in $\mathbb{C}$.

Let $\mathbb{Z}_{1/2} = \mathbb{Z} + 1/2 = \{\ldots, -3/2, -1/2, 1/2, 3/2, \ldots\}$, so that $\mathbb{Z}_{1/2}^2$ may be identified with the collection of (centers of) unit length squares $\{[n, n+1] \times [m, m+1] : n, m \in \mathbb{Z}\}$. Recall that $\mathbb{Z}_4$ is the rotation group generated by $r_{\pi/2}$. For notational convenience, we identify $\mathbb{Z}_{1/2}^2$ with a subset of $\mathbb{C}$ in the standard way.

Lemma 3.1 (Alternating pattern)**.** *For $N \geq 15$ odd there exists a function $\mathfrak{p} : \mathbb{Z}_{1/2}^2 \to \{1,2,3,4,5\} \times \mathbb{Z}_4$ with the following properties. Below, $\mathfrak{i}$ and $\mathfrak{r}$ are the components of $\mathfrak{p}$.*

1. *$\mathfrak{p}$ is doubly periodic of period $2N$, that is $\mathfrak{p}(z) = \mathfrak{p}(z + 2N) = \mathfrak{p}(z + 2Ni)$ for all z.*
2. *$\mathfrak{p}$ is rotationally covariant, that is $\mathfrak{p}(iz) = (\mathfrak{i}(z), r_{\pi/2} \cdot \mathfrak{r}(z))$ for all z.*
3. *There exists an exceptional set $E \subset [-N,N]^2 \cap \mathbb{Z}_{1/2}^2$ with $|E| = 16$ such that for $p \in ([-N,N]^2 \cap \mathbb{Z}_{1/2}^2) \backslash E$, the value of $\mathfrak{i}$ at p is distinct from the value at each of its 8 neighbours.*

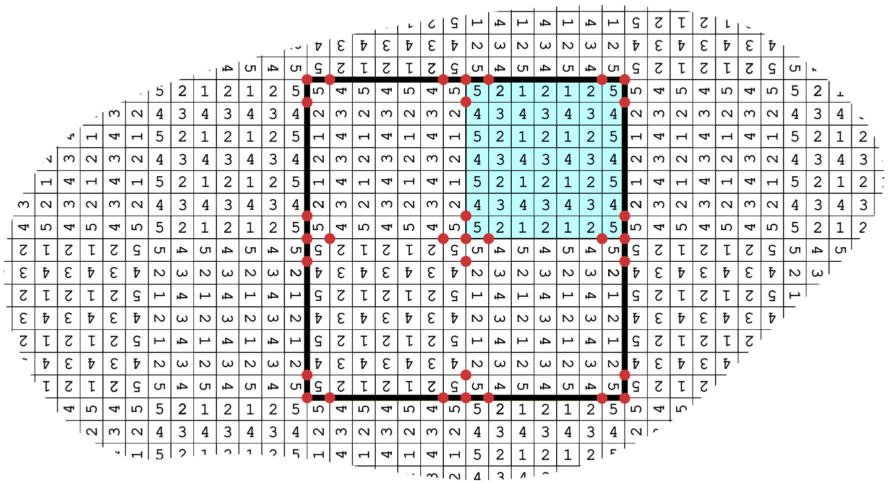

Figure 3. The labelling of the lattice $\mathbb{Z}^2_{1/2}$ constructed in Lemma 3.1, for $N = 7$. The large bold square is $[-N, N]^2$. The labelling is periodic of period $2N$, and also covariant with respect to $\pi/2$ rotations around the origin. The symmetries of the labelling imply that it is determined by the values on $[0, N]^2$, which is highlighted in blue. In $[-N, N]^2$, with a few exceptions, each label is numerically distinct from its 8 neighbours. In particular, with a few exceptions (marked with a red disk), the 4 labels around any vertex are numerically distinct. In the proof of Theorem 1.3, the Beltrami coefficients $\mu_1, \dots, \mu_5$ are concatenated according to this pattern to form a Beltrami coefficient on $\mathbb{C}$.

(4) *For* $a \in \{1, 2, 3, 4\}$*, we have* $\frac{|\{z\in[-N,N]^2\cap\mathbb{Z}^2_{1/2}:\mathfrak{i}(z)=a\}|}{|[-N,N]\cap\mathbb{Z}^2_{1/2}|^2} \geq \frac{(N-1)^2}{(2N+1)^2} \geq \frac{1}{5}$.

Proof. It suffices to define $\mathfrak{p}$ on the quadrant $[0, N]^2\cap\mathbb{Z}^2_{1/2}$ because then the rotational symmetry and periodicity determines the values of $\mathfrak{p}$ on all other points.

In the definitions below, we restrict to points in this quadrant. In this quadrant, we take $\mathfrak{r}$ to be the identity element in $\mathbb{Z}_4$, and we define $\mathfrak{i}$ as follows. See Figure 3.

- Define $\mathfrak{i}(1/2, 1/2) = \mathfrak{i}(1/2, N-1/2) = \mathfrak{i}(N-1/2, 1/2) = \mathfrak{i}(N-1/2, N-1/2) = 5$. This defines $\mathfrak{i}$ on the four corners of the quadrant.
- For $k \in \{1/2, N - 1/2\}$ and $j \notin \{1/2, N - 1/2\}$, define $\mathfrak{i}(j + 1/2, k) = 1$ if $j \equiv 0 \bmod 2$, otherwise $\mathfrak{i}(j + 1/2, k) = 2$. This defines $\mathfrak{i}$ on the top and bottom edges of the quadrant.
- If $j \in \{1/2, N - 1/2\}$, define $\mathfrak{i}(j, k) = 5$ if $k \equiv 0 \bmod 2$, otherwise $\mathfrak{i}(j, k) = 4$. This defines $\mathfrak{i}$ on the left and right edges of the quadrant.
- On the remaining interior points of the quadrant, define (all equivalences are mod 2)

$$\mathfrak{p}(x, y) = \begin{cases} 1 & \text{if } x - 1/2 \equiv 0 \text{ and } y - 1/2 \equiv 0 \\ 2 & \text{if } x - 1/2 \equiv 1 \text{ and } y - 1/2 \equiv 0 \\ 3 & \text{if } x - 1/2 \equiv 1 \text{ and } y - 1/2 \equiv 1 \\ 4 & \text{if } x - 1/2 \equiv 0 \text{ and } y - 1/2 \equiv 0. \end{cases} \tag{7}$$

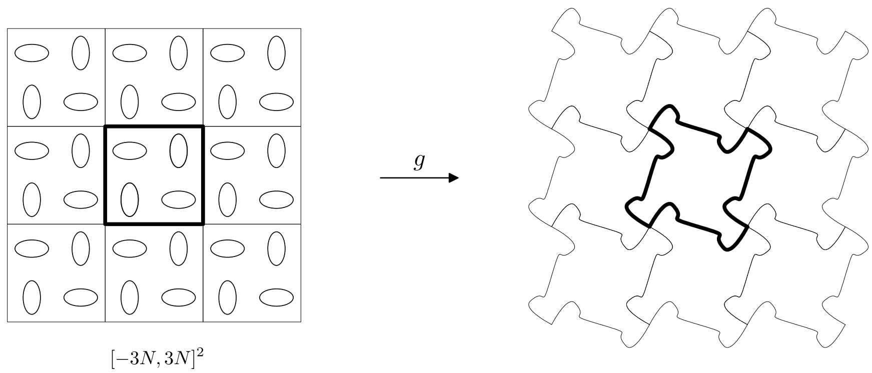

Figure 4. The map $g : \mathbb{C} \to \mathbb{C}$ is the integrating homeomorphism for the coefficients μ_i that are concatenated in the pattern shown in Figure 3. g is normalized so that it fixes the lattice $(2N\mathbb{Z})^2$. The symmetries of the concatenation pattern imply that g has the corresponding rotational and translation symmetries. The image above depicts the restriction of g to $[-3N, 3N]^2$. In particular, if $S = [-N, N]^2$ is the middle small square, then $\mathrm{Area}(S) = \mathrm{Area}(g(S))$. In Lemma 3.3, the diameter of $g(S)$ is bounded in terms of the energy of the coeffients μ_i.

Let E' be the four corner points of $[0, N]^2 \cap \mathbb{Z}^2_{1/2}$, together with its orbit under the group generated by the translations and rotations $\{z \mapsto z+2N, z \mapsto z+2Ni, z \mapsto iz\}$. A case analysis verifies that for $z \notin E'$, the value $\mathfrak{p}(z)$ is distinct from the value of $\mathfrak{p}$ on each of its 8 neighbours. Since $E := E' \cap [-N, N]^2$ has 16 elements, this proves the third item of the conclusion.

The last item in the conclusion follows immediately from (7), together with the assumption $N \geq 15$. □

Let $\mathfrak{p} = (\mathfrak{i}, \mathfrak{r})$ be the labelling of $\mathbb{Z}^2_{1/2}$ constructed by Lemma 3.1 above. Construct a Beltrami coefficient ν on $\mathbb{C}$ by concatenating the Beltrami coefficients $\mu_1, \ldots, \mu_5$ in the way prescribed by $\mathfrak{p}$, where the value of $\mathfrak{i}$ corresponds to the the choice of Beltrami coefficient $\mu_1, \ldots, \mu_5$, and $\mathfrak{r}$ describes a rotation applied to the Beltrami coefficient. See Figure 3. More formally, if S_p is the unit square centered at $p \in \mathbb{Z}^2_{1/2} \cap S$, then $\nu|_{S_p}(z) = \mathfrak{r}(p) \cdot \mu_{\mathfrak{i}(p)}(z - p + 1/2 + i/2)$.

The symmetries of the labelling $\mathfrak{p}$ ensure that ν is rotationally covariant and doubly periodic of period $2N$. Therefore, there is a unique integrating homeomorphism $g : \mathbb{C} \to \mathbb{C}$ that solves the Beltrami equation, with the following additional properties: $g(0) = 0$, $g(\cdot + w) = g(\cdot)$ for $w \in (2N\mathbb{Z})^2$, and $g \circ r = g$ whenever r is a $\pi/2$ rotation around some point in $(N\mathbb{Z})^2$. See Figure 4 for an illustration of the mapping g.

3.2. **A variation for the symmetric coefficient.** Let $\eta = \mathrm{Id}_{\mathbb{C}} - g$ and let $S = (-N, N)^2$. By definition, $\mathrm{Dir}_{S,\nu}(\mathrm{Id} - \eta, \mathrm{Id} - \eta) = \mathrm{Dir}_{S,\nu}(g, g) = \mathrm{Dir}_{g(S)}(\mathrm{Id}, \mathrm{Id})$. This is equal to $\mathrm{Area}(g(S))$, which is in turn equal to $\mathrm{Area}(S)$ because g is $2N$-periodic and g fixes $(2N\mathbb{Z})^2$. Thus, from (6), we get

$$\frac{\mathrm{Dir}_{S,\nu}(\mathrm{Id}, \eta)^2}{\mathrm{Dir}_{S,\nu}(\eta, \eta)} \geq \mathrm{Dir}_{S,\nu}(\mathrm{Id}, \mathrm{Id}) - \mathrm{Area}(g(S)) = \mathcal{E}_S(\nu). \tag{8}$$

We will need the following lower bound on the numerator $\mathrm{Dir}_{S,\nu}(\mathrm{Id}, \eta)$.

Lemma 3.2. *With notation as above,*

$$\mathrm{Dir}_{S,\nu}(\mathrm{Id},\eta) \gtrsim N^2\mathcal{E}(\mu_1).$$

Proof. We have, by the AM-GM inequality,

$$\begin{aligned}\mathrm{Dir}_{S,\nu}(\mathrm{Id},\eta) &= \mathrm{Dir}_{S,\nu}(\mathrm{Id},\mathrm{Id}-g) = \mathrm{Dir}_{S,\nu}(\mathrm{Id},\mathrm{Id}) - \mathrm{Dir}_{S,\nu}(\mathrm{Id},g)\\ &\geq \mathrm{Dir}_{S,\nu}(\mathrm{Id}) - \frac{1}{2}\left(\mathrm{Dir}_{S,\nu}(\mathrm{Id}) + \mathrm{Dir}_{S,\nu}(g)\right) = \frac{1}{2}\left(\mathrm{Dir}_{S,\nu}(\mathrm{Id}) - \mathrm{Area}(S)\right)\\ &= \frac{1}{2}\sum_{U\subset S}\mathrm{Dir}_{U,\nu}(\mathrm{Id}) - \mathrm{Area}(U) = \frac{1}{2}\sum_{U\subset S}\mathcal{E}_U(\nu).\end{aligned}$$

where the sum is over the $4N^2$ unit subsquares U of S. Property 4 of the pattern in Lemma 3.1 ensures that this last term is $\gtrsim N^2\mathcal{E}(\mu_1)$. □

Now choose a rotation and translation invariant partition of unity of the plane, $\{\rho_v\}_{v\in\mathbb{Z}^2}$, with each ρ_v supported on the 2×2 square U_v centered at v. Translation invariance means that $\rho_v(\cdot) = \rho_w(\cdot - v + w)$ for $v, w \in \mathbb{Z}^2$, and rotation invariance means that $\rho_v \circ r_v = \rho_v$ when r_v is a rotation by $\pi/2$ around v.

Then $\eta_v := \rho_v \cdot \eta$ vanishes on the boundary of U_v.

Property 3) of the labelling $\mathfrak{p} = (\mathfrak{i},\mathfrak{r})$ in Lemma 3.1 ensures that for all but a small exceptional set of $v \in \mathbb{Z}^2$, the four subsquares of U_v have different $\mathfrak{i}$ labels.

In particular, if v is not exceptional, $\nu|_{U_v} \equiv \left[\begin{smallmatrix} r_1\mu_{\iota(1)} & r_2\mu_{\iota(2)} \\ r_3\mu_{\iota(3)} & r_4\mu_{\iota(4)}\end{smallmatrix}\right]$ for some choice of rotations $r : \{1,2,3,4\} \to \mathbb{Z}_4$ and some injection $\pi : \{1,2,3,4\} \to \{1,2,3,4,5\}$. Here $\mu \equiv \nu$ means that the Beltrami coefficients μ and ν are equal up to translation of the domain.

To apply the variational formula (Lemma 2.6), we need to give an upper bound for $\mathrm{Dir}_{U_v,\nu}(\eta_v)$, then we need to find a non-exceptional v for which $\mathrm{Dir}_{U_v,\nu}(\eta_v,\mathrm{Id})^2$ is sufficiently large.

3.3. **Upper bound for** $\mathrm{Dir}_{U_v,\nu}(\eta_v)$**.** We use the product rule (Lemma 2.7), giving $\mathrm{Dir}_{U_v,\nu}(\eta_v) \lesssim \mathrm{Dir}_{U_v,\nu}(\eta) + \|\eta|_{U_v}\|_\infty^2\mathrm{Dir}_{U_v,\nu}(\mathrm{Id}_S)$. The first term is bounded by $\mathrm{Dir}_{U_v,\nu}(\eta) \lesssim \mathrm{Dir}_{U_v,\nu}(g) + \mathrm{Dir}_{U_v,\nu}(\mathrm{Id}) \lesssim \mathrm{Dir}_{U_v,\nu}(\mathrm{Id})$. For the second term, we need the following bound on $\|\eta|_{U_v}\|_\infty$.

Lemma 3.3. *Let μ_i, η, U_v and N be as above. Then*

$$\|\eta|_{U_v}\|_\infty \lesssim N(\mathcal{E}(\mu_1)+1)^{1/2}.$$

Proof. Let $S' = (-3N, 3N)^2$ so that $S = (-N, N)^2$ is a centered subsquare of S', and $S'\backslash S$ is the union of $\lesssim N$ unit squares. See Figure 4.

Since $U_v \subset S$, it suffices to bound $\|\eta|_S\|_\infty$. Also recall that $\eta = g - \mathrm{Id}_{\mathbb{C}}$. The idea is that $\mathrm{diam}(g(S))$ can be bounded in terms of $\mathcal{E}_{S'\backslash S}(\nu)$ due to the well known relation between conformal modulus and relative distances. The desired bound on $\|\eta|_S\|_\infty$ will then follow from the triangle inequality.

Since g is an integrating homeomorphism for ν, $\frac{1+|\nu|^2}{1-|\nu|^2} = \frac{\|Dg\|_{\mathrm{HS}}^2}{\det Dg}$ a.e. Here $\|\cdot\|_{\mathrm{HS}}$ is the Hilbert-Schmidt norm, which is greater than the operator norm. By the Cauchy-Schwarz inequality we get

$$\begin{aligned}\iint_{S'\backslash S}\|Dg\|_{\mathrm{HS}}dxdy &\leq \left(\iint_{S'\backslash S}\frac{1+|\nu|^2}{1-|\nu|^2}dxdy\right)^{1/2}\cdot\left(\iint_{S'\backslash S}\det Dg\ dxdy\right)^{1/2}\\ &\lesssim (N^2\mathcal{E}(\mu_1)+N^2)^{1/2}\cdot \mathrm{Area}(g(S'\backslash S))^{1/2} \lesssim N^2(\mathcal{E}(\mu_1)+1)^{1/2}.\end{aligned}$$

In the last inequality we used the fact that $\text{Area}(g(S'\backslash S)) = \text{Area}(S'\backslash S) = 32N^2$. Foliating $S'\backslash S$ by concentric boundaries of squares and applying Fubini's theorem, we find that there is at least one loop γ in the foliation such that $\int_\gamma \|Dg\|_{\text{HS}} ds \lesssim N^{-1}N^2(\mathcal{E}(\mu_1)+1)^{1/2}$. We then get $\text{diam} g(S) \le \frac{1}{2}\text{Length}(g(\gamma)) \lesssim N(\mathcal{E}(\mu_1)+1)^{1/2}$.

Let p be a corner of S so that $g(p) = p$. Then for $z \in S$, we have $|\eta(z)| = |g(z) - z| = |g(z-p+p) - g(p) - z + p| \le \text{diam}(g(S)) + \text{diam}(S) \lesssim N(\mathcal{E}(\mu)+1)^{1/2}$, as desired. □

Combining the above, we get

$$\text{Dir}_{U_v,\nu}(\eta_v) \lesssim N^2(\mathcal{E}(\mu_1)+1)^2. \tag{9}$$

3.4. **Finding nonexceptional v such that $\text{Dir}_{U_v,\nu}(\eta_v, \text{Id})$ is large.** Let V be the set of points $v \in \mathbb{Z}^2$ for which $U_v \cap S$ is nonempty, and let $E \subset V$ be the exceptional points; these are the v for which the labelling $\mathfrak{i}$ is not injective on the 4 subsquares of U_v. By construction (Item 3 of Lemma 3.1), $|E| \lesssim 1$.

Our goal is to find $v \in V - E$ for which $\text{Dir}_{U_v,\nu}(\eta_v, \text{Id})$ is large. Since $\sum_v \eta_v = \eta$, this follows from a pigeonhole principle argument, which we detail below. Here we need to pick N (the side length of the concatenation pattern) large, to overcome the contribution from the exceptional set E.

From Lemma 3.2, we have

$$N^2\mathcal{E}(\mu_1) \lesssim \text{Dir}_{S,\nu}(\text{Id}, \eta) = \sum_{v\in V\backslash E} \text{Dir}_{U_v,\nu}(\text{Id}, \eta_v) - \sum_{v\in V\backslash E} \text{Dir}_{U_v\backslash S,\nu}(\text{Id}, \eta_v) + \sum_{v\in E} \text{Dir}_{U_v\cap S,\nu}(\text{Id}, \eta_v). \tag{10}$$

We first need to deal terms that are of the form $\text{Dir}_V(\cdot)$ where V is not a 2×2 square – they are in the second and third sums above.

For each nonzero term in the second sum, $U_v\backslash S$ is the union of two unit squares, and the symmetries of ν and ρ imply that there is a matching term corresponding to the union of two squares $U_{v'}\backslash S$ on the opposite side of S, such that $\text{Dir}_{U_v\backslash S,\nu}(\text{Id}, \eta_v) + \text{Dir}_{U_{v'}\backslash S,\nu}(\text{Id}, \eta_{v'}) = \text{Dir}_{U_v,\nu}(\text{Id}, \eta_v)$. Pairing up those terms, we get

$$\sum_{v\in V\backslash E} \text{Dir}_{U_v\backslash S,\nu}(\text{Id}, \eta_v) = \frac{1}{2} \sum_{\substack{v\in V\backslash E \\ U_v\backslash S\neq\emptyset}} \text{Dir}_{U_v,\nu}(\text{Id}, \eta_v). \tag{11}$$

A similar idea works for the terms in the third sum. The four terms in $\sum_{v\in E} \text{Dir}_{U_v\cap S,\nu}(\text{Id}, \eta_v)$ corresponding to the four corners of $(-N,N)^2$ can be collected together, and their sum is equal to $\text{Dir}_{U_v,\nu}(\text{Id}, \eta_v)$ where v is one of the four corners. The remaining terms of $\sum_{v\in E}$ lying on the edges of $(-N,N)^2$ can be paired together as in the proof of (11).

This expresses $\sum_{v\in E} \text{Dir}_{U_v\cap S,\nu}(\text{Id}, \eta_v)$ as a sum $\sum_v \text{Dir}_{U_v,\nu}(\text{Id}, \eta_v)$, still consisting of $\lesssim 1$ terms. Now that η_v vanishes on ∂U_v, we can use the improved Cauchy-Schwarz inequality (Lemma 2.1) on each term, giving

$$\sum_{v\in E} \text{Dir}_{U_v\cap S}(\text{Id}, \eta_v) \lesssim \text{Dir}_{U_v,\nu}(\eta_v)^{1/2}\mathcal{E}_{U_v}(\nu)^{1/2} \lesssim N(\mathcal{E}(\mu_1)+1)\mathcal{E}(\mu_1)^{1/2}. \tag{12}$$

Substituting (11) and (12) into (10) yields

$$N^2\mathcal{E}(\mu_1) \lesssim \sum_{v\in V\setminus E} \mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v) - \frac{1}{2}\sum_{\substack{v\in V\setminus E\\ U_v\setminus S\neq\emptyset}} \mathrm{Dir}_{U_v\setminus S}(\mathrm{Id},\eta_v) + N(\mathcal{E}(\mu_1)+1)\mathcal{E}(\mu_1)^{1/2}.$$

Dividing through by N^2, we see that it is possible to choose $N \lesssim (\mathcal{E}(\mu_1)+1)\mathcal{E}(\mu_1)^{-1/2}$ such that

$$\mathcal{E}(\mu_1) \lesssim \frac{1}{N^2}\sum_{v\in V\setminus E} \mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v) - \frac{1}{2N^2}\sum_{\substack{v\in V\setminus E\\ U_v\setminus S\neq\emptyset}} \mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v).$$

Viewing this a single sum containing $\asymp N^2$ terms, we conclude that there is at least one $v \in V\setminus E$ for which $|\mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v)| \gtrsim \mathcal{E}(\mu_1)$.

Combining with (9), we get $v \in V\setminus E$ such that

$$\frac{\mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v)^2}{\mathrm{Dir}_{U_v,\nu}(\eta_v)} \gtrsim \frac{\mathcal{E}(\mu_1)^3}{(\mathcal{E}(\mu_1)+1)^4}.$$

By considering the cases $\mathcal{E}(\mu_1) \geq 1$ and $\mathcal{E}(\mu_1) < 1$ separately, we get $\frac{\mathrm{Dir}_{U_v,\nu}(\mathrm{Id},\eta_v)^2}{\mathrm{Dir}_{U_v,\nu}(\eta_v)} \gtrsim \min(\mathcal{E}(\mu_1)^3, \mathcal{E}(\mu_1)^{-1})$. This completes the proof of Theorem 1.3.

4. Application to homogenization

In this section we use the improvement inequality Theorem 1.2 to prove the following homogenization result, which is closely related to results found in [IM19] and [ARST20].

Roughly speaking, it says that randomly oriented Beltrami coefficients on a square lattice are, from a large scale perspective, indistinguishable from the standard complex structure.

Let $\mathcal{D}_m$ be the collection of dyadic squares of side length 2^{-m} in $[0,1]^2$, so that $|\mathcal{D}_m| = 4^m$.

Theorem 4.1. *Let $\lambda : (0,1)^2 \to \mathbb{D}$ be a self compatible random Beltrami coefficient on the unit square with rotationally invariant distribution. Assume that there exists $\epsilon > 0$ such that $\mathbb{E}\mathcal{E}(\lambda)^{1+\epsilon} < \infty$.*

Let μ_n be the Beltrami coefficient on $(0,1)^2$ obtained by placing an i.i.d. scaled copies of λ on each of the squares in $\mathcal{D}_n$.

Let $h_n : [0,1]^2 \to [0,1]^2$ be the unique integrating homeomorphism of μ_n which fixes the vertices $\{(0,0),(1,0),(0,1)\}$.

Then $h_n \to \mathrm{Id}_{[0,1]^2}$ in probability, with respect to the uniform norm on $C([0,1]^2)$. That is,

$$\textit{For all } \epsilon > 0, \quad \lim_{n\to\infty} \mathbb{P}\left(\sup_{z\in[0,1]^2} |h_n(z) - z| > \epsilon\right) = 0.$$

Our integrability condition $\mathbb{E}\mathcal{E}(\lambda)^{1+\epsilon} < \infty$ is much less stringent than the uniform ellipticity assumption $\lambda \leq c < 1$ used in [ARST20]. Furthermore, the theorem stated in [IM19] assumes that the random Beltrami coefficient λ is almost surely constant on $(0,1)^2$, although they only assume existence of first moments rather than $(1+\epsilon)$ moments.

The proof of Theorem 4.1 will occupy the rest of this section.

Jensen's inequality immediately implies that $\mathbb{E}\mathcal{E}(\mu_{n+1})^{1+\epsilon} \leq \mathbb{E}\mathcal{E}(\mu_n)^{1+\epsilon}$ for all n. So, letting $E_n = \mathbb{E}\mathcal{E}^*(\mu_n)$ and iterating Theorem 1.2, we get $E_{n+1} \leq E_n - c\min(E_n^3, E_n)$, where $c > 0$ is a constant that only depends on the distribution of λ.

Hence, $\lim_{n\to\infty} E_n = \lim_{n\to\infty} \mathbb{E}\mathcal{E}^*(\mu_n) = 0$. Thus, by Markov's inequality,

$$\limsup_{n\to\infty} \mathbb{P}(\mathcal{E}^*_{[0,1]^2}(\mu_n) \geq 3) = 0. \tag{13}$$

If the λ are uniformly elliptic as in [ARST20], then the sequence of random variables h_n take values in a compact space, namely the space of K-quasiconformal mappings with the uniform norm. Without uniform ellipticity, this is no longer true in general, but the energy bound (13) can be used probabilistically at multiple scales to give equicontinuity bounds and hence obtain subsequential limits.

Let $\mathfrak{X}$ be the space of homeomorphisms $[0,1]^2 \to [0,1]^2$ fixing the vertices $\{(0,0),(1,0),(0,1)\}$. Let $d^\infty(f,g) = \sup_{z\in[0,1]^2} |f(z) - g(z)|$ be the sup norm. Consider the metric $d_{\mathfrak{X}}(f,g) = d^\infty(f,g) + d^\infty(f^{-1}, g^{-1})$ on $\mathfrak{X}$.

Recall that a sequence of random variables Y_n taking values in a metric space $(\mathfrak{X}, d_{\mathfrak{X}})$ is said to be *tight* if for all $\epsilon > 0$ there exists a precompact $K \subset \mathfrak{X}$ such that $\sup_n \mathbb{P}(Y_n \notin K) < \epsilon$.

Lemma 4.1. *The sequence h_n is tight with respect to $d_{\mathfrak{X}}$.*

Proof. Write $d_{\mathfrak{X}} = d^1 + d^{-1}$ as the sum of two metrics. Fix $\epsilon > 0$. Recall that h_n is an integrating homeomorphism for the Beltrami coefficient μ_n, and $\mathbb{E}\mathcal{E}(\mu_n) = \mathbb{E}\mathcal{E}(\lambda) < \infty$ by assumption. By (2), $\mathbb{E}\mathrm{Dir}(h_n^{-1})$ is bounded uniformly in n too, so Markov's inequality implies that there exists $M_2 > 0$ such that $\mathbb{P}(\mathrm{Dir}(h_n^{-1}) > M_2) < \epsilon/2$. The length-area method (see, e.g, the proof of the Courant-Lebesgue lemma [Cou37, Lemma 5]), and the Arzela-Ascoli theorem implies that the set $\mathcal{B}_{M_2} := \{h \in \mathfrak{X} : \mathrm{Dir}(h^{-1}) \leq M_2\}$ is precompact with respect to the metric d^{-1}.

Now we consider the metric d^1. Fix an integer $m \geq 0$ and a dyadic square $S \in \mathcal{D}_m$ at level m, of side length 2^{-m}. The three marked vertices of $\partial[0,1]^2$ separate $\partial[0,1]^2$ into three components, I_1, I_2 and I_3. We will inductively construct a sequence of disjoint sets $B_1, \ldots, B_M$, each of which separates Q from a fixed arc I_i, in the sense that any path from Q to I_i must cross each of the B_j. See Figure 5 for an example of this procedure. The B_i will either be topological annuli or topological rectangles. Let $m_1 = m$, and let B_1 be the union of all dyadic squares at level m that are *adjacent*, but not equal to, Q. Subsets of $\mathbb{C}$ are considered adjacent if their closures intersect in a nonempty set of zero interior.

Once B_k is chosen, let m_{k+1} be the largest integer such that $\mathrm{diam}_{L^\infty} B_k \leq 2^{-m_{k+1}}$, where diam_{L^∞} is the diameter with respect to the Chebyshev metric on $[0,1]^2$, so that there is an axes-aligned square (not necessarily dyadic) of length $2^{-m_{k+1}}$ containing B_k.

It is possible to cover B_k by at most 4 dyadic squares of side length $2^{-m_{k+1}}$, call the minimal such covering $\mathcal{Q}(B_k)$. Let B_{k+1} be the union of all dyadic squares in $[0,1]^2$ of level m_{k+1} that are adjacent to, but not equal to, the squares in $\mathcal{Q}(B_k)$. There are at most $4\cdot(4-1) = 12$ such squares.

This process generates a sequence of sets B_k, which we continue running while B_k touches at most two edges of $\partial[0,1]^2$, which ensures that they separate Q from at least one of the three arcs I_j.

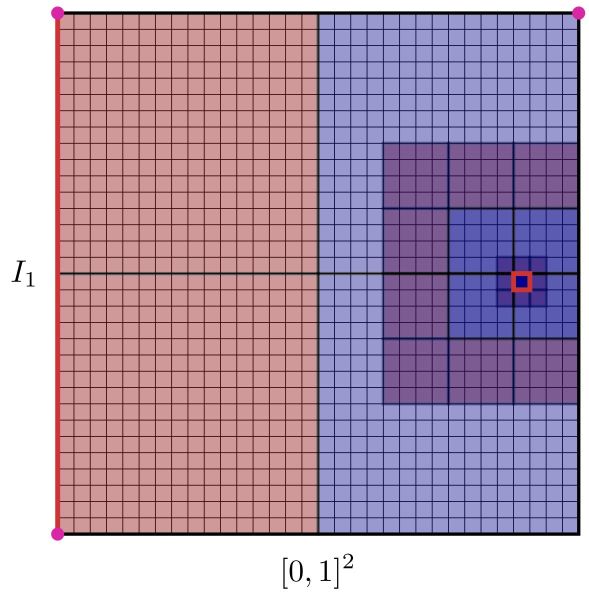

Figure 5. The three marked points on $\partial[0,1]^2$ divide $\partial[0,1]^2$ into three arcs. Let Q be the small dark blue square on the middle right of side length 2^{-5}. The image depicts the 'exploration process' $B_1, \mathcal{Q}(B_1), B_2, \mathcal{Q}(B_2), B_3$, which consists of $8, 4, 8, 2, 2$ squares at scales $5, 5, 3, 3, 1, 1$ respectively. Each of the red sets B_1, B_2, B_3 separate Q from the arc I_1. The moduli of crossing path families in the images $h_n(B_k)$ can be bounded by putting together bounds on the energy for each individual dyadic square, (13). These moduli estimates translate to bounds on the diameter of $h_n(S)$.

Since $\mathrm{diam}(B_{k+1}) \leq 12\mathrm{diam}(B_k)$, we can continue for at least $\gtrsim m$ steps, because as long as $\mathrm{diam}(B_k) \leq \frac{1}{2}$ then B_k touches at most two edges of $\partial[0,1]^2$.

Each B_k is either a topological annulus or a topological rectangle. If B_k is an annulus, let Γ_k be the path family joining the inner boundary of $h_n(B_k)$ to the outer boundary of $h_n(B_k)$. Otherwise, B_k is a rectangle with two sides lying on $\partial[0,1]^2$. We let Γ_k be the family of paths in B_k joining the other two sides.

The point of this definition is that every path from Q to I_j contains a path in each Γ_k.

We now estimate the extremal length [Ahl10, Chapter 4] of the path families $\mathrm{EL}(\Gamma_k)$. Here it is convenient that for each dyadic square Q, we have estimates on the energy $\mathcal{E}^*(\mu_n|_Q)$ rather than the extremal length of the path families joining opposite sides of Q.

Reciprocal energy is a lower bound for the extremal length:

$$\mathrm{EL}(\Gamma_k) \gtrsim (2^{-m_k})^2 \mathrm{Dir}^*(h_n^{-1}|_{h_n(B_k)})^{-1} = (2^{-m_k})^2 (\mathrm{Area}(B_k) + \mathcal{E}^*_{B_k}(\mu_n))^{-1}. \tag{14}$$

To see this, let g_n be the Poisson extension of $h_n^{-1}|_{h_n(B_k)}$, then take $\rho = |\nabla g_n|$ on $h_n(B_k)$ as a test function in the variational definition of extremal length. If $\gamma \in \Gamma_k$ then $\int_\gamma \rho ds$ is the length of $g_n(\gamma)$, which is at least 2^{-m_k}, and $\int_{h_n(B_k)} \rho^2 dxdy = \mathrm{Dir}^*(h_n^{-1}|_{h_n(B_k)})$. (This is closely related to the proof of Lemma 3.3).

In addition, energy, unlike modulus, is trivially additive (see (1)) under concatenation. Thus (14) implies

$$\mathrm{EL}(\Gamma_k) \geq C_0 \left(\sum_{Q \subset B_k} (2^{m_k})^2 (\mathrm{Area}(Q) + \mathcal{E}^*_Q(\mu_n)) \right)^{-1} \tag{15}$$

where the sum is over the (at most 12) dyadic squares Q of level m_k constituting B_k.

If Q is a dyadic square of side length 2^{-m}, then up to translation and scaling, $\mu_n|_Q \overset{d}{=} \mu_{n-m}$ for $n \geq m$. In particular, $(2^m)^2 \mathcal{E}^*(\mu_n|_Q) \overset{d}{=} \mathcal{E}^*(\mu_{n-m})$. So by (13),

$$\text{For } Q \text{ dyadic of level } m, \quad \liminf_{n\to\infty} \mathbb{P}((2^m)^2(\mathrm{Area}(Q) + \mathcal{E}^*(\mu_n|_Q)) < 4) = 1.$$

Taking the union bound over the (at most 12) squares Q in the sum (15) gives

$$\liminf_{n\to\infty} \mathbb{P}(\mathrm{EL}(\Gamma_k) > C_0/48) = 1.$$

Large deviations estimates for binomial random variables (e.g. Chernoff's inequality [Ver18, Theorem 2.3.1]) can be used to show that with for sufficiently small $c_0 > 0$, we have with probability at least $1 - 5^{-m}$, that $\mathrm{EL}(\Gamma_k) > C_0/48$ for at least $c_0 m$ of the scales k. On this event, the series rule ([Ahl10, Theorem 4.2]) implies that the extremal length of the path family joining $h_n(S)$ to $h_n(I_j) = I_j$ is bounded below by $C_0 c_0 m/48$.

On the other hand, $\log \mathrm{diam}(h_n S)^{-1} \gtrsim \mathrm{EL}(h_n(S), I_j)$, which can be seen by considering the associated extremal problem (see e.g. [Ahl10, Theorem 4.7]).

Combining the above yields

$$\text{For } \delta_0 \text{ sufficiently close to 1}, \ \limsup_{n\to\infty} \mathbb{P}\left(\mathrm{diam}(h_n S) > \delta_0^m\right) < 5^{-m}.$$

Taking the union bound over the 4^m dyadic squares S of level m in $[0,1]^2$ gives

$$\limsup_{n\to\infty} \mathbb{P}\left(\sup_{S \in \mathcal{D}_m} \mathrm{diam}(h_n S) \geq \delta_0^m \right) \lesssim (4/5)^m.$$

Every subset $K \subset [0,1]^2$ of diameter less than 2^{-m} can be covered by a union of at most 4 squares in $\mathcal{D}_m$, so

$$\limsup_{n\to\infty} \mathbb{P}\left(\sup_{K \subset [0,1]^2,\ \mathrm{diam}(K) \leq 2^{-m}} \mathrm{diam}(h_n K) \geq 4\delta_0^m \right) \lesssim (4/5)^m.$$

Choose M_1 large enough that the sum over $m \geq M_1$ of the right hand side is smaller than $\epsilon/2$. The Arzela-Ascoli theorem implies that the set

$$\mathcal{A}_{M_1} := \left\{ h \in \mathfrak{X} : \forall m \geq M_1, \quad \sup_{K \subset [0,1]^2,\ \mathrm{diam}(K) \leq 4\delta_0^m} \mathrm{diam}(hK) < 2^{-m} \right\}$$

is precompact with respect to d^1.

By the union bound, $\mathbb{P}(h \notin (\mathcal{A}_{M_1} \cap \mathcal{B}_{M_2})) < \epsilon$. Furthermore, $\mathcal{A}_{M_1} \cap \mathcal{B}_{M_2}$ is precompact with respect to $d_{\mathfrak{X}}$, because if $d^\infty(f_n, f) \to 0$ and $d^\infty(f_n^{-1}, g^{-1}) \to 0$, then $f = g$.

Since ϵ was arbitrary, we are done. $\square$

The space of homeomorphisms $(\mathfrak{X}, d_{\mathfrak{X}})$ is complete and separable, and we just proved tightness of the random variables h_n. So by Prokhorov's theorem, there is a subsequence h_{n_j} which converges weakly to a random homeomorphism $h : [0,1]^2 \to [0,1]^2$.

To pass from information about the energy of h_{n_j} to the energy of the limit h, we need:

Lemma 4.2. *For each square $S \subset [0,1]^2$, the functional $h \mapsto \mathrm{Dir}^*_{h(S)}(h|_S^{-1})$ is lower semicontinuous on $(\mathfrak{X}, d_{\mathfrak{X}})$.*

Proof. For X, Y metric spaces, let $\mathrm{Hom}(X, Y)$ be the space of homeomorphisms with the uniform norm $d(f, g) = \sup_{x \in X} d_Y(f(x), g(x))$.

Fix 3 ordered points on ∂S in counterclockwise order. Define the operator $\mathcal{C}_S : \mathfrak{X} \to \mathrm{Hom}(\overline{S}, \overline{\mathbb{D}})$ via the following procedure. For $h \in \mathfrak{X}$, let (Ω, p_1, p_2, p_3) be the image of S with its three marked boundary points under h. Let ψ be the conformal map taking $(\Omega_S, p_1, p_2, p_3)$ to $(\mathbb{D}, 1, i, -1)$. Then define $\mathcal{C}_S[h] = \psi \circ h|_S$.

By conformal invariance of Dirichlet energy, $\mathrm{Dir}^*_{h(S)}(h|_S^{-1}) = \mathrm{Dir}^*_{\mathbb{D}}(\mathcal{C}_S[h]^{-1})$. So the desired semicontinuity will follow once we verify that $h \mapsto \mathrm{Dir}^*_{\mathbb{D}}(\mathcal{C}_S[h]^{-1})$ is composed of (semi)continuous maps.

The operator $\mathcal{C}_S$ is continuous by continuity properties of the Riemann mapping (see [Pom92, Theorem 2.11]). The inversion operator $h \mapsto h^{-1}$ is also continuous $\mathrm{Hom}(\overline{S}, \overline{\mathbb{D}}) \to \mathrm{Hom}(\overline{\mathbb{D}}, \overline{S})$. Finally, the optimal energy operator $u \mapsto \mathrm{Dir}^*_{\mathbb{D}}(u) = \mathrm{Dir}_{\mathbb{D}}(P_{\mathbb{D}}[u])$ is lower semicontinuous. This is a consequence of Fatou's lemma and the fact that if harmonic functions converge uniformly then their derivatives converge uniformly too. ☐

Let $k \geq 0$ be an integer and suppose $S \in \mathcal{D}_k$ is a dyadic square at level k. The semicontinuity above allows us to conclude that

$$\mathbb{E}\mathrm{Dir}^*_{h(S)}(h|_S^{-1}) \leq \liminf_{k \to \infty} \mathbb{E}\mathrm{Dir}^*_{h_{n_j}(S)}(h_{n_j}|_S^{-1}) = \mathrm{Area}(S).$$

Hence $\mathrm{Dir}^*_{h(S)}(h|_S^{-1}) = \mathrm{Area}(S)$ almost surely, which means that the $h|_S$ has the boundary values of a conformal map:

Lemma 4.3. *Let Ω be a Jordan domain and suppose $g : \Omega \to S$ is harmonic, with $\mathrm{Dir}_\Omega(g) = \mathrm{Area}(S)$. Then g is conformal.*

Proof. By (2), $\mathcal{E}_S(\mu_{g^{-1}}) = 0$, so the Beltrami coefficient of g^{-1} vanishes, so g^{-1} is conformal. ☐

Taking the union bound over the squares $S \in \mathcal{D}_k$, we can modify h on each $S \in \mathcal{D}_k$ so that it is conformal on S. Call this modified homeomorphism $\tilde{h}$, so that $\tilde{h} = h$ on $\bigcup_{S \in \mathcal{D}_k} \partial S$. By (the proof of) Morera's theorem, $\tilde{h}$ is conformal on $[0,1]^2$. It also fixes three points on $\partial[0,1]^2$, so in fact $\tilde{h} = \mathrm{Id}_{[0,1]^2}$.

Since k was arbitrary and h is continuous, we get $h = \mathrm{Id}_{[0,1]^2}$. The subsequential limit was arbitrary, so $h_n \to \mathrm{Id}_{[0,1]^2}$ weakly. Since the limit is deterministic, this implies that $h_n \to \mathrm{Id}_{[0,1]^2}$ in probability. This concludes the proof of Theorem 4.1.

References

[Ahl10] Lars Ahlfors. *Conformal invariants : topics in geometric function theory.* AMS Chelsea Pub, Providence, R.I, 2010.

[AIMO05] K. Astala, T. Iwaniec, G. J. Martin, and J. Onninen. Extremal Mappings of Finite Distortion. *Proceedings of the London Mathematical Society*, 91(3):655–702, 11 2005.

[ARST20] Kari Astala, Steffen Rohde, Eero Saksman, and Terence Tao. Homogenization of iterated singular integrals with applications to random quasiconformal maps, 2020.

[Ast09] Kari Astala. *Elliptic partial differential equations and quasiconformal mappings in the plane*. Princeton University Press, Princeton, NJ, 2009.

[Cou37] Richard Courant. Plateau's problem and dirichlet's principle. *Annals of Mathematics*, 38(3):679–724, 1937.

[DH84] Adrien Douady and J. H Hubbard. *Etude dynamique des polynomes complexes*. Orsay : Universite de Paris-Sud, Dept. de mathematique, 1984.

[Dur04] Peter Duren. *Harmonic Mappings in the Plane*. Cambridge Tracts in Mathematics. Cambridge University Press, 2004.

[HKO05] Stanislav Hencl, Pekka Koskela, and Jani Onninen. A note on extremal mappings of finite distortion. *Mathematical Research Letters*, 12:231–237, 03 2005.

[IM19] Oleg Ivrii and Vladimir Markovic. Homogenization of random quasiconformal mappings and random delauney triangulations, 2019.

[Pom92] Christian Pommerenke. *Boundary behaviour of conformal maps*. Springer-Verlag, Berlin New York, 1992.

[Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.