CONCENTRATION FOR RANDOM EUCLIDEAN COMBINATORIAL OPTIMIZATION

MATTEO D'ACHILLE FRANCESCO MATTESINI DARIO TREVISAN

ABSTRACT. We prove concentration bounds for random Euclidean combinatorial optimization problems with p–costs. For bipartite matching and for the (mono- and bipartite) traveling salesperson problem in dimension $d \geq 3$, we obtain concentration at the natural energy scale $n^{1-p/d}$ for $1 \leq p < d^2/2$. Our method combines a Poincaré inequality with a robust geometric mechanism providing uniform bounds on the edges of optimizers. We also formulate a conjectural $p \to q$ transfer principle for the p–optimal matching which, if true, would extend the concentration range to all $p \geq 1$.

1. INTRODUCTION AND OVERVIEW

We study concentration properties for random Euclidean optimization problems built on i.i.d. uniform point clouds in the cube $[0,1]^d$ (with Euclidean distance). Working in $[0,1]^d$ keeps the presentation simple; the same arguments extend with minor changes to compact d–dimensional Riemannian manifolds (with or without boundary), provided one has the corresponding local volume growth and Poincaré inequality for the underlying reference measure.

Two prototypical examples are:

- Bipartite matching (random assignment). Two independent clouds $\mathbf{X} = (X_i)_{i=1}^n$ and $\mathbf{Y} = (Y_j)_{j=1}^n$ are matched by a permutation σ, minimizing the sum of p–powers of distances, where $p \geq 1$ is a parameter.
- Euclidean traveling salesperson problem (TSP). A Hamiltonian cycle is chosen on one cloud (monopartite TSP), or on two clouds with an alternation constraint (bipartite TSP), minimizing the sum of p–powers of edge lengths.

In all these models, the disorder is geometric (the random points) and the optimizer selects a low-energy combinatorial structure. The natural energy scale is $n^{1-p/d}$ and – possibly excluding dimension $d = 1$ and $d = 2$ in the bipartite case – our goal is to quantify sample-to-sample fluctuations at that scale.

1.1. Main results and proof strategy (informal).

Our results show that, in dimension $d \geq 3$ and for $1 \leq p < d^2/2$, the optimal costs for matching and TSP are *self-averaging* at the energy scale: with (polynomially) high probability,

$$| \mathcal{C} - \mathbb{E}[\mathcal{C}] | \ll n^{1-p/d}.$$

where $\mathcal{C}$ denotes the optimal cost. Precise statements are given in Theorems 2.1, 3.1, and 3.4. The argument consists of two main layers.

(A) *Analytic concentration tool: Poincaré inequality.* Let μ be the uniform probability measure on $([0,1]^d)^n$ (monopartite problems) or on $([0,1]^d)^{2n}$ (two-cloud problems). The tensorized Poincaré inequality gives, for every locally Lipschitz F,

$$\mathrm{Var}_\mu(F) \leq C_{\mathrm{P}}\, \mathbb{E}_\mu\left[|\nabla F|^2\right]. \tag{1.1}$$

Applying (1.1) to the optimal cost $\mathcal{C}$ reduces concentration to bounding $|\nabla\mathcal{C}|$. For matching and TSP, at points of differentiability one has gradient bounds of the form

$$|\nabla\mathcal{C}|^2 \lesssim \sum_e |e|^{2(p-1)} \leq \mathcal{C}\cdot\big(\sup_e |e|\big)^{p-2}, \tag{1.2}$$

where the sum/sup are over the edges selected by an optimizer.

(B) *Geometric input: excluding long edges by local stability.* The key difficulty is to control the maximal edge length $\sup_e |e|$. The mechanism is:

1. a *mesoscopic diffuseness* event, ensuring that every ball of radius $\gtrsim n^{-\alpha/d}$ contains many points (with very high probability);
2. a *local optimality* condition: a suitable 2–opt move cannot decrease the cost;
3. a deterministic lemma converting 2–opt into a *local edge-to-energy* inequality (that we prove in details for the bipartite matching, see Lemma 2.3 for matching).

Combining these ingredients yields a power-law bound $\sup_e |e| \lesssim n^{-p/(d(p+d))}$ with very high probability, which closes the Poincaré estimate. The exponent $p < d^2/2$ is exactly the condition that the resulting polynomial tail exponent is positive.

1.2. A conjectural $p \to q$ transfer principle.

Numerical simulations (see Section A) indicate that the current range $p < d^2/2$ is likely not intrinsic. A natural strengthening would be to control not only the maximal edge, but higher moments of the edge lengths under the p–optimizer.

Conjecture 1.1 ($p \to q$ transfer for the p–optimal matching)**.** *Fix $d \geq 3$ and $p \geq 1$. For every $q > p$ there exists $c = c(d,p,q) < \infty$ such that for every p–optimal matching σ_p,*

$$\mathbb{E}\Big[\sum_{i=1}^{n} |X_i - Y_{\sigma_p(i)}|^q\Big] \leq c\, n^{1-q/d}.$$

Conjecture 1.1 matches the heuristic that the p–optimizer still pairs points at the typical matching scale $n^{-1/d}$, hence has q–energy of order $n^{1-q/d}$ for every $q > p$. We notice that the case $q < p$ is trivially true by optimality of σ_p and an application of Hölder inequality with exponent p/q. If the conjecture were true (with $q = 2p-2$), it would remove the restriction $p < d^2/2$ in our concentration bounds. Similar conjectures can be stated for TSP and bipartite TSP cases.

1.3. Related literature and context.

Random Euclidean combinatorial optimization lies at the intersection of probability, geometric measure theory, statistical physics, and algorithms. The foundational results are the law of large numbers for Euclidean functionals, starting from the Beardwood–Halton–Hammersley theorem and its extensions to broad classes of subadditive Euclidean functionals; see also the monographs by Steele and Yukich for a unified probabilistic framework covering the minimum spanning tree (MST), matching, TSP, and related models [5, 20, 22].

On the probabilistic side, quantitative asymptotics and fluctuation theory for Euclidean matching and assignment problems originate in the work of Ajtai–Komlós–Tusnády and its subsequent refinements via multiscale, Fourier, and empirical-process techniques [1, 6]. Modern developments connect these ideas to geometric probability tools such as uniform occupancy, stabilization, and VC-type arguments [19, 11]

A complementary viewpoint arises from statistical physics and optimal transport. For matching-type problems, renormalization heuristics and elliptic PDE descriptions have led to sharp asymptotics and robust structural results in low dimensions, notably through the PDE/OT approach and subsequent work [3, 15]. In particular, Proposition 5.3 in [15] establishes concentration under general structural assumptions using a Poincaré inequality, highlighting the role of spectral-gap methods in this context.

From the algorithmic perspective, Euclidean TSP and matching have long served as benchmarks for local-improvement heuristics (e.g. 2-opt and k-opt moves) and average-case analysis in computer science and operations research, alongside worst-case PTAS results for geometric optimization [9, 17, 4, 18].

The present note contributes to the concentration theory of these models. While the use of a Poincaré (spectral gap) inequality is not new in this setting (see e.g. [15]), our contribution is to isolate a sharp and purely geometric mechanism converting local swap optimality (cyclical monotonicity for matching, admissible 2-edge swaps for TSP) and mesoscopic density into a quantitative L^∞ bound on edge lengths. This uniform edge control strengthens existing approaches and provides a direct route to subleading fluctuation bounds without relying on integrable representations of the optimizer.

1.4. **Extensions and open problems.** The mechanism developed in this note is not specific to matching or TSP, but relies only on two structural ingredients:

(i) a local swap (or 2–edge) optimality condition for the minimizer;
(ii) mesoscopic density control of the underlying random point cloud.

The same strategy applies to a broad class of Euclidean combinatorial problems whose minimizers satisfy a local exchange property. Examples include minimal spanning trees, k–MST, degree-constrained spanning subgraphs such as k–factors, and related matching-type structures. In all these cases, optimality can be expressed through a local exchange inequality, and the triangle-inequality argument yields a local edge-to-energy comparison analogous to Lemma 2.3. Combined with Poincaré's inequality, this leads to concentration estimates at the natural scaling, provided suitable annealed bounds are available.

Although we restrict attention to the cube with uniform sampling for simplicity, the argument extends to substantially more general settings. If the underlying measure has a density ρ that is bounded above and below on a bounded domain with Lipschitz boundary, the mesoscopic density estimates remain valid and the same proof applies. More generally, on compact Riemannian manifolds with bounded geometry, the argument carries over provided geodesics exist and satisfy standard triangle inequalities, volumes of geodesic balls scale like r^d at small scales, and a Poincaré (spectral gap) inequality holds for the product measure. Under these mild geometric assumptions, the local edge-to-energy argument and the Poincaré reduction are unchanged.

For simplicity, we also formulated concentration using the L^2 Poincaré inequality, bounding the variance by the squared L^2 norm of the gradient. This yields polynomial tail bounds via Chebyshev or Markov inequalities. Stronger tail estimates (e.g. higher polynomial moments or sub-Gaussian behavior in suitable regimes) could be obtained by combining the same geometric input with L^q versions of the Poincaré inequality or related functional inequalities. We do not pursue these refinements here, as our focus is on the geometric mechanism.

The main limitation of the present method is the restriction $p < d^2/2$, which arises from the current power-law control of the maximal edge length. Numerical evidence suggests that no qualitative change occurs at $p = d^2/2$, and that concentration at the natural scale $n^{1-p/d}$ should hold for all $p \geq 1$. Establishing concentration for arbitrary p remains, in our view, the central open problem in this direction.

Organization. Section 2 proves concentration for bipartite matching. Section 3 treats the monopartite TSP and then the bipartite TSP. Appendix A presents numerical simulations supporting Conjecture 1.1 and showing that the current range $p < d^2/2$ is likely not optimal.

Acknowledgements. M. D'A. acknowledges support from the ANR project LOUCCOUM (ANR-24-CE40-7809). M. D'A. and F. M. are grateful to the Department of Mathematics, University of Pisa for excellent working conditions in the occasion of the invitation which started this project (November 2021). F. M. acknowledges the support of the ERC Advanced Grant NEITALG, grant agreement No. 101198055. D.T. acknowledges the MUR Excellence Department Project awarded to the Department of Mathematics, University of Pisa, CUP I57G22000700001, the HPC Italian National Centre for HPC, Big Data and Quantum Computing - Proposal code CN1 CN00000013, CUP I53C22000690001, the PRIN 2022 Italian grant 2022WHZ5XH - "understanding the LEarning process of QUantum Neural networks (LeQun)", CUP J53D23003890006, the project G24-202 "Variational methods for geometric and optimal matching problems" funded by Università Italo Francese. Research also partly funded by PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - "FAIR - Future Artificial Intelligence Research" - Spoke 1 "Human-centered AI", funded by the European Commission under the NextGeneration EU programme. This research benefitted from the support of the FMJH Program Gaspard Monge for optimization and operations research and their interactions with data science.

2. Concentration for bipartite matching

Let $\mathbf{x} = (x_i)_{i=1}^n$ and $\mathbf{y} = (y_j)_{j=1}^n$ be points in $[0,1]^d$, and let $p \geq 1$. The (Euclidean) bipartite matching cost is

$$\mathcal{C}^p_{bM}(\mathbf{x}, \mathbf{y}) \coloneqq \min_{\sigma \in \mathcal{S}_n} \sum_{i=1}^n |x_i - y_{\sigma(i)}|^p. \tag{2.1}$$

We use the following notational conventions throughout. For a point set $\mathbf{z} = (z_i)_{i=1}^n$ and a measurable set $\Omega \subseteq [0,1]^d$ we write

$$N^{\mathbf{z}}_\Omega \coloneqq \#\{\, i \in \{1, \dots, n\} : z_i \in \Omega \,\}.$$

We denote by $B(x, r)$ the Euclidean ball in $\mathbb{R}^d$ of center x and radius r; when working in $[0,1]^d$ we tacitly intersect with $[0,1]^d$ (this only affects boundary balls and does not change the arguments).

In the random case $\mathbf{x} = \mathbf{X}$, $\mathbf{y} = \mathbf{Y}$, we assume that $x_i = X_i$ and $y_j = Y_j$ are realizations of i.i.d. random variables uniformly distributed on $[0,1]^d$. The matching cost becomes a random variable $\mathcal{C}^p_{bM}(\mathbf{X}, \mathbf{Y})$, and its annealed asymptotics are by now standard [1, 21, 3, 10, 12]: as $n \to \infty$,

$$\mathbb{E}[\mathcal{C}^p_{bM}(\mathbf{X}, \mathbf{Y})] \sim n \begin{cases} n^{-\frac{p}{2}} & \text{if } d = 1, \\ \left(\frac{\log n}{n}\right)^{\frac{p}{2}} & \text{if } d = 2, \\ n^{-\frac{p}{d}} & \text{if } d \geq 3, \end{cases} \tag{2.2}$$

where $\sim$ denotes asymptotic upper and lower bounds – existence of a limit after rescaling remains open only in $d = 2$, $p \neq 2$. The exceptional scaling in dimensions $d = 1$ and $d = 2$ is due to the presence of anomalously long edges in the optimal matching, which are not ruled out by the local stability mechanism described above. For this reason, we focus on the case $d \geq 3$.

Theorem 2.1 (Concentration for bipartite matching)**.** *Let $d \geq 3$ and $1 \leq p < d^2/2$. Then, there exist constants $\theta = \theta(p,d) > 0$ and $C = C(p,d) < \infty$ such that for all $n \geq 1$ and $\lambda > 0$,*

$$\mathbb{P}\Big(|\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y}) - \mathbb{E}[\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y})]| \geq \lambda n^{1-\frac{p}{d}}\Big) \leq C n^{-\theta}\lambda^{-2}.$$

2.1. **2-opt inequality and local edge-to energy control.** Before we provide the proof of Theorem 2.1, we focus on the key steps, described informally in the previous section, that allow to obtain a uniform control on edge costs, starting from a 2–opt inequality. In the case of bipartite matching, the optimality condition we use is the following two-point swap inequality (a special instance of the cyclical monotonicity in optimal transport theory).

Lemma 2.2 (2–opt inequality)**.** *Let σ be an optimizer in* (2.1)*. Then for all $i, j \in \{1, \dots, n\}$,*

$$|x_i - y_{\sigma(i)}|^p + |x_j - y_{\sigma(j)}|^p \leq |x_i - y_{\sigma(j)}|^p + |x_j - y_{\sigma(i)}|^p. \tag{2.3}$$

Proof. Otherwise swapping $\sigma(i)$ and $\sigma(j)$ strictly decreases the total cost. □

We now make the 2–opt condition quantitative by turning it into a uniform bound of a single edge cost by local energy. In the semi-discrete matching (where one cloud of points is replaced with the uniform density), this type of local-energy control goes back at least to [2, Lemma 4.4] and [13, Lemma 4.1]; the lemma below adapts this viewpoint to the present two-cloud setting.

Lemma 2.3. *For any $p > 1$, there exists $\varepsilon = \varepsilon(p) > 0$ and $C = C(p) < \infty$ such that the following holds. Let $\mathbf{x} = (x_i)_{i=1}^n, \mathbf{y} = (y_j)_{j=1}^n \subseteq [0,1]^d$ be two sets of points, and let σ be an optimizer for the p-cost. For any $i \in \{1, \dots, n\}$, set*

$$B_i = B\left(\frac{x_i + y_{\sigma(i)}}{2}, \varepsilon\, |x_i - y_{\sigma(i)}|\right).$$

Then,

$$N^{\mathbf{x}}_{B_i}\, |x_i - y_{\sigma(i)}|^p \;\leq\; C \sum_{x_j \in B_i} |x_j - y_{\sigma(j)}|^p. \tag{2.4}$$

Proof of Lemma 2.3. Write $z_i = (x_i + y_{\sigma(i)})/2$ and assume that $0 < \varepsilon < 1/2$. Let $x_j \in B_i$ (if $B_i \cap \mathbf{x} = \varnothing$ then (2.4) is trivial). It suffices to show that (2.3) and the triangle inequality imply

$$|x_i - y_{\sigma(i)}|^p \leq C\, |x_j - y_{\sigma(j)}|^p. \tag{2.5}$$

Indeed, summing (2.5) over $x_j \in B_i$ yields (2.4).

By the triangle inequality,

$$|x_i - y_{\sigma(j)}| \leq |x_i - z_i| + |z_i - x_j| + |x_j - y_{\sigma(j)}| \leq (1+\varepsilon)|x_i - z_i| + |x_j - y_{\sigma(j)}|, \tag{2.6}$$

$$|x_j - y_{\sigma(i)}| \leq |x_j - z_i| + |z_i - y_{\sigma(i)}| \leq (1+\varepsilon)|x_i - z_i|. \tag{2.7}$$

Using the inequality $(a+b)^p \leq (1+\eta)a^p + C_{\eta,p}\, b^p$ with $\eta > 0$ small, together with (2.3), (2.6) and (2.7), we obtain

$$|x_i - y_{\sigma(i)}|^p + |x_j - y_{\sigma(j)}|^p \leq (2+\eta)(1+\varepsilon)^p |x_i - z_i|^p + C|x_j - y_{\sigma(j)}|^p.$$

Since $|x_i - z_i| = \frac{1}{2}|x_i - y_{\sigma(i)}|$, choosing $\varepsilon = \varepsilon(p)$ (and then $\eta = \eta(p)$) so that $1 - (2+\eta)(1+\varepsilon)^p 2^{-p} > 0$, we can rearrange to get (2.5). □

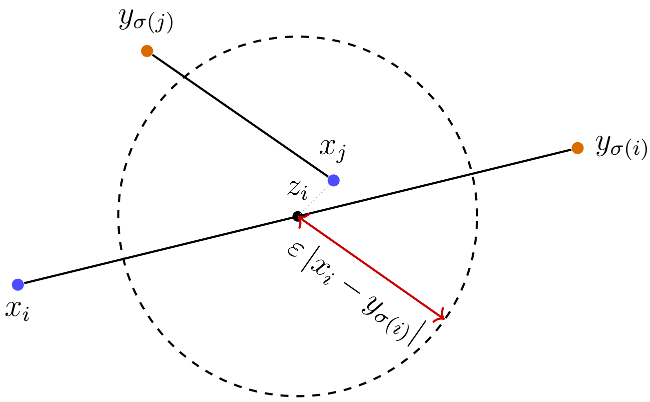

FIGURE 1. Local geometry used in Lemma 2.3.

2.2. **Proof of Theorem 2.1.** As already remarked, the case $p < d$ can be dealt with a simpler argument and is essentially covered in [15, Proposition 5.3]. For this reason, we focus only on the case $p \geq d$ here.

Proof. **Step 1.** We argue that

$$\mathbb{P}\Big(|\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y}) - \mathbb{E}[\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y})]| \geq \lambda n^{1-\frac{p}{d}}\Big) \lesssim_{d,p} \lambda^{-2} n^{\frac{p}{d}-1}\,\mathbb{E}\big[\sup_i |X_i - Y_{\sigma(i)}|^{2(p-2)}\big]^{\frac{1}{2}}, \tag{2.8}$$

where σ is an optimizer. Indeed, for $p \geq 1$ the map $(\mathbf{x},\mathbf{y}) \mapsto \mathcal{C}^p_{bM}(\mathbf{x},\mathbf{y})$ is Lipschitz and at points of differentiability (see e.g. the proof of [15, Proposition 5.3]),

$$\nabla_{x_i}\mathcal{C}^p_{bM} = p|x_i - Y_{\sigma(i)}|^{p-2}(x_i - y_{\sigma(i)}), \qquad \nabla_{y_j}\mathcal{C}^p_{bM} = p|x_{\sigma^{-1}(j)} - y_j|^{p-2}(x_{\sigma^{-1}(j)} - y_j).$$

Thus,

$$|\nabla \mathcal{C}^p_{bM}(\mathbf{x},\mathbf{y})|^2 \lesssim \sum_{i=1}^n |x_i - y_{\sigma(i)}|^{2(p-1)} \leq \mathcal{C}^p_{bM}(\mathbf{x},\mathbf{y}) \sup_i |x_i - y_{\sigma(i)}|^{p-2}. \tag{2.9}$$

The Cauchy-Schwarz inequality gives

$$(\mathcal{C}^p_{bM}(\mathbf{x},\mathbf{y}))^2 \leq n\mathcal{C}^{2p}_{bM}(\mathbf{x},\mathbf{y})$$

so that

$$\begin{aligned}\mathbb{E}\left[\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y}) \sup_i |x_i - y_{\sigma(i)}|^{p-2}\right] &\leq n^{\frac{1}{2}}\mathbb{E}\left[\mathcal{C}^{2p}_{bM}(\mathbf{X},\mathbf{Y}))^2\right]^{\frac{1}{2}}\mathbb{E}\big[\sup_i |X_i - Y_{\sigma(i)}|^{2(p-2)}\big]^{\frac{1}{2}} \\ &\lesssim_{d,p} n^{\frac{1}{2}+\left(1-\frac{2p}{d}\right)\frac{1}{2}}\mathbb{E}\big[\sup_i |X_i - Y_{\sigma(i)}|^{2(p-2)}\big]^{\frac{1}{2}},\end{aligned}$$

by (2.2) with exponent $2p$. Combining Markov's inequality and the Poincaré inequality yields (2.8).

Step 2. We use Lemma 2.3 to control the maximal matching edge.

Step 2.1. (Good events) Fix $\alpha \in (0,1)$ and set

$$A := \Big\{\forall x \in [0,1]^d,\ \forall r > n^{-\alpha/d} :\ N^{\mathbf{X}}_{B(x,r)} \geq \frac{nr^d}{2}\Big\}.$$

Then for every $\beta > 0$, as $n \to \infty$, it holds see for instance [16, eq. (2.52)] for the Poisson case, the proof being analogous in the case of i.i.d. distributed points)

$$\mathbb{P}(A^c) \lesssim_{d,\alpha,\beta} n^{-\beta}. \tag{2.10}$$

Fix $\alpha' \in (0,1)$ and set

$$B := \{\mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y}) \leq n^{1-\alpha' p/d}\}.$$

Then for every $\beta > 0$, as $n \to \infty$,

$$\mathbb{P}(B^c) \lesssim_{d,p,\alpha',\beta} n^{-\beta}. \tag{2.11}$$

Indeed, by Markov and (2.2), for any $q > 1$,

$$\mathbb{P}(\mathcal{C}^p_{bM} > n^{1-\alpha' p/d}) \leq \frac{\mathbb{E}[(\mathcal{C}^p_{bM})^q]}{n^{(1-\alpha' p/d)q}} \leq \frac{n^{q-1}\mathbb{E}[\mathcal{C}^{pq}_{bM}]}{n^{(1-\alpha' p/d)q}} \lesssim_{d,p,q} n^{(\alpha'-1)\frac{pq}{d}},$$

and choosing q large gives (2.11).

Step 2.2. (L^∞ bound on the matching edges) On $A \cap B$, Lemma 2.3 yields a power-law bound on the maximal edge. Indeed, fix i and set $r := \varepsilon|X_i - Y_{\sigma(i)}|$. If $r \leq n^{-\alpha/d}$ we are done. Otherwise, on A,

$$N^{\mathbf{X}}_{B_i} \geq \frac{nr^d}{2} = \frac{1}{2} n\varepsilon^d |X_i - Y_{\sigma(i)}|^d.$$

Applying (2.4) and bounding the local cost by the global one gives

$$n\varepsilon^d |X_i - Y_{\sigma(i)}|^{p+d} \ \lesssim \sum_{X_j \in B_i} |X_j - Y_{\sigma(j)}|^p \ \leq \ \mathcal{C}^p_{bM}(\mathbf{X},\mathbf{Y}) \ \leq \ n^{1-\alpha' p/d}.$$

Choosing α, α' so that

$$\alpha(p+d) = \alpha' p \tag{2.12}$$

yields, on $A \cap B$,

$$\sup_i |X_i - Y_{\sigma(i)}| \lesssim n^{-\alpha/d}. \tag{2.13}$$

Since $\alpha', \alpha \in (0,1)$, condition (2.12) allows any $0 < \alpha < p/(p+d)$.

Step 3. Using (2.13) in (2.8) and the fact that (2.10) and (2.11) hold with arbitrarily large β (and $[0,1]^d$ has bounded diameter), we obtain

$$\mathbb{E}[\sup_i |X_i - Y_{\sigma(i)}|^{2(p-2)}]^{\frac{1}{2}} \lesssim_{d,p,\alpha} n^{-\frac{\alpha(p-2)}{d}}.$$

Hence,

$$\mathbb{P}\Big(|\mathcal{C}^p_{bM} - \mathbb{E}[\mathcal{C}^p_{bM}]| \geq \lambda n^{1-\frac{p}{d}}\Big) \ \lesssim_{d,p,\alpha} \ \lambda^{-2} n^{\frac{p}{d}-1-\frac{\alpha(p-2)}{d}}. \tag{2.14}$$

Choosing α arbitrarily close to $p/(p+d)$, we see that the right hand side is infinitesimal, with a rate $n^{-\theta}$ for a suitable $\theta > 0$, provided that

$$p\left(1 - \frac{p-2}{p+d}\right) < d,$$

that yields exactly the condition $p < d^2/2$. $\square$

3. Concentration for the traveling salesperson problem

In this section we argue that the tools leading to the concentration result are flexible enough to deal with other combinatorial optimization problems, by showing how they apply in the case of the traveling salesperson problem (both for the monopartite and bipartite versions).

3.1. **Monopartite case.** Let $\mathbf{x} = (x_i)_{i=1}^n$ be points in $[0,1]^d$. We say that a cycle τ is Hamilonian if it visits each vertex. For $p \geq 1$ and a Hamiltonian cycle τ on $\{1, \dots, n\}$, set

$$L_p(\tau; \mathbf{x}) \coloneqq \sum_{k=1}^n |x_{\tau(k)} - x_{\tau(k+1)}|^p, \qquad \tau(n+1) \coloneqq \tau(1),$$

and define the optimal TSP p-cost

$$\mathcal{C}_{TS}^p(\mathbf{x}) \coloneqq \min_\tau L_p(\tau; \mathbf{x}). \tag{3.1}$$

Assuming again that $x_i = X_i$ and $y_j = Y_j$ are realizations of i.i.d. random variables uniformly distributed on $[0,1]^d$, upper bounds for $T_{n,p}$ are also standard [5, 20, 22]: as $n \to \infty$,

$$\mathbb{E}[\mathcal{C}_{TS}^p(\mathbf{X})] \sim n^{1-p/d}. \tag{3.2}$$

Upper bounds in particular can be obtained via space-filling curves heuristics. Existence of a limit after rescaling appears to be open for $p > d \geq 2$ (although it is claimed in [22] that the argument for $p = d$ extends to larger exponents). Differently from the bipartite problem, there are no exceptional scaling in dimensions $d = 1$ and $d = 2$.

Theorem 3.1 (Concentration for Euclidean TSP)**.** *Let $d \geq 2$ and $1 \leq p < d^2/2$. Then, there exist constants $\theta = \theta(p,d) > 0$ and $C = C(p,d) < \infty$ such that for all $n \geq 1$ and $\lambda > 0$,*

$$\mathbb{P}\Big(|\mathcal{C}_{TS}^p(\mathbf{X}) - \mathbb{E}[\mathcal{C}_{TS}^p(\mathbf{X})]| \geq \lambda n^{1-\frac{p}{d}}\Big) \leq C n^{-\theta} \lambda^{-2}.$$

The TSP analogue of (2.3) is the standard 2-opt move: remove two non-adjacent edges and reconnect to obtain another Hamiltonian cycle (equivalently, reverse a segment of the tour). Global optimality implies that no improving 2-opt move exists.

Lemma 3.2 (2-opt inequality for an optimal tour)**.** *Let τ be an optimiser in* (3.1)*. Consider two edges of the tour,*

$$(x_{\tau(a)}, x_{\tau(a+1)}) \quad \text{and} \quad (x_{\tau(b)}, x_{\tau(b+1)}), \qquad a \neq b.$$

Then the 2-opt reconnection does not decrease the p–cost, i.e.

$$|x_{\tau(a)} - x_{\tau(a+1)}|^p + |x_{\tau(b)} - x_{\tau(b+1)}|^p \;\leq\; |x_{\tau(a)} - x_{\tau(b)}|^p + |x_{\tau(a+1)} - x_{\tau(b+1)}|^p. \tag{3.3}$$

Proof. Perform the 2-opt move replacing the two removed edges by one of the two admissible reconnections. The resulting cycle has cost at least that of τ by optimality, yielding (3.3). ☐

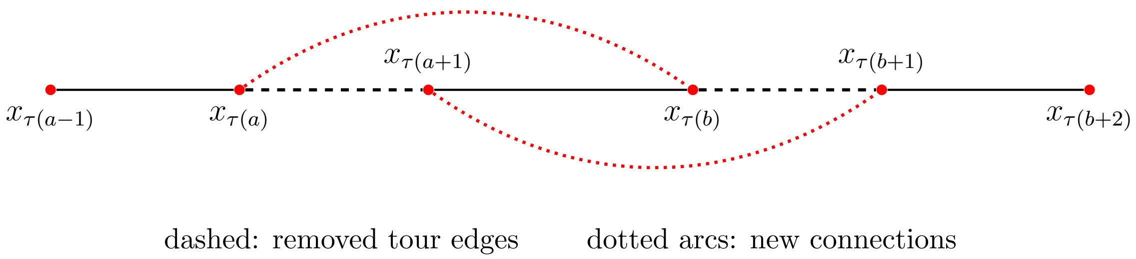

FIGURE 2. TSP: a standard 2–opt move cutting $(x_{\tau(a)}, x_{\tau(a+1)})$ and $(x_{\tau(b)}, x_{\tau(b+1)})$ and reconnecting $(x_{\tau(a)}, x_{\tau(b)})$, $(x_{\tau(a+1)}, x_{\tau(b+1)})$.

The geometric ingredient is completely analogous to the matching case. The only additional observation is that, given a long tour edge $e_i = (x_{\tau(a)}, x_{\tau(a+1)})$ and a vertex

$x_{\tau(b)}$ in a small ball centered in the midpoint of the edge, one can always perform an alternating 2–opt move.

Lemma 3.3 (Local edge-to-energy control for TSP)**.** *For any $p > 1$, there exist $\varepsilon = \varepsilon(p) > 0$ and $C = C(p) < \infty$ such that the following holds.*

Let τ be an optimiser in (3.1)*, and let*

$$e = (x_{\tau(a)}, x_{\tau(a+1)})$$

be a tour edge. Let $z_e = \frac{1}{2}(x_{\tau(a)} + x_{\tau(a+1)})$ and define

$$B_e := B\big(z_e, \varepsilon\,|x_{\tau(a+1)} - x_{\tau(a)}|\big)\,.$$

Then

$$N^{\mathbf{x}}_{B_e}\,|x_{\tau(a+1)} - x_{\tau(a)}|^p \le C \sum_{b:\ x_{\tau(b)} \in B_e} |x_{\tau(b+1)} - x_{\tau(b)}|^p. \tag{3.4}$$

Sketch of proof. Fix $x_{\tau(b)} \in B_e$ and apply the 2–opt inequality (3.3) to the two disjoint edges $e = (x_{\tau(a)}, x_{\tau(a+1)})$ and $(x_{\tau(b)}, x_{\tau(b+1)})$:

$$|x_{\tau(a+1)} - x_{\tau(a)}|^p + |x_{\tau(b+1)} - x_{\tau(b)}|^p \le |x_{\tau(a)} - x_{\tau(b)}|^p + |x_{\tau(a+1)} - x_{\tau(b+1)}|^p.$$

Since $x_{\tau(b)} \in B_e$ and z_e is the midpoint of e, triangle inequalities imply

$$|x_{\tau(a+1)} - x_{\tau(b+1)}| \le \Big(\tfrac{1}{2} + \varepsilon\Big)|x_{\tau(a+1)} - x_{\tau(a)}| + |x_{\tau(b+1)} - x_{\tau(b)}|.$$

Using the elementary inequality $(u+v)^p \le (1+\eta)u^p + C_{\eta,p}v^p$ and choosing $\varepsilon = \varepsilon(p)$ small enough, one obtains

$$|x_{\tau(a+1)} - x_{\tau(a)}|^p \le C\,|x_{\tau(b+1)} - x_{\tau(b)}|^p.$$

Summing this estimate over all vertices $x_{\tau(b)} \in B_e$ yields

$$N^{\mathbf{x}}_{B_e}\,|x_{\tau(a+1)} - x_{\tau(a)}|^p \le C \sum_{b:\ x_{\tau(b)} \in B_e} |x_{\tau(s(b)+1)} - x_{\tau(s(b))}|^p.$$

Since each tour edge has at most two endpoints, the right-hand side is bounded (up to a factor 2) by

$$\sum_{b:\ x_{\tau(b)} \in B_e} |x_{\tau(b+1)} - x_{\tau(b)}|^p,$$

which gives (3.4). □

Then, the proof of Theorem 3.1 follows along the same lines of the matching problem.

3.2. **Bipartite case.** Let $\mathbf{x} = (x_i)_{i=1}^n$ and $\mathbf{y} = (y_j)_{j=1}^n$ be two sets of points in $[0,1]^d$. A *bipartite Hamiltonian cycle* is a cycle alternating between x- and y-vertices. Equivalently, it can be encoded by a permutation σ of $\{1, \dots, n\}$ and the cycle

$$x_1 \to y_{\sigma(1)} \to x_2 \to y_{\sigma(2)} \to \cdots \to x_n \to y_{\sigma(n)} \to x_1.$$

For $p \ge 1$, define its p–cost by

$$L_p(\sigma; \mathbf{x}, \mathbf{y}) := \sum_{i=1}^n \Big(|x_i - y_{\sigma(i)}|^p + |x_{i+1} - y_{\sigma(i)}|^p\Big), \qquad x_{n+1} := x_1,$$

and set

$$\mathcal{C}^p_{bTS}(\mathbf{x}, \mathbf{y}) := \min_\sigma L_p(\sigma; \mathbf{x}, \mathbf{y}). \tag{3.5}$$

Assuming again that $x_i = X_i$ and $y_j = Y_j$ are realizations of i.i.d. random variables uniformly distributed on $[0,1]^d$, the annealed asymptotics are of the same order as in the bipartite matching problem [8, 7, 14]:

$$\mathbb{E}[\mathcal{C}^p_{bTS}(\mathbf{X},\mathbf{Y})] \sim n \begin{cases} n^{-\frac{p}{2}} & \text{if } d=1, \\ \left(\frac{\log n}{n}\right)^{\frac{p}{2}} & \text{if } d=2, \\ n^{-\frac{p}{d}} & \text{if } d\geq 3, \end{cases} \tag{3.6}$$

and existence of a limit for the rescaled expected cost is known in the case $d=1$, $d=p=2$ or $d\geq 3$ and $p<d$.

Theorem 3.4 (Concentration for bipartite TSP)**.** *Let $d \geq 3$ and $1 \leq p < d^2/2$. Then there exist constants $\theta = \theta(p,d) > 0$ and $C = C(p,d) < \infty$ such that for all $n \geq 1$ and $\lambda > 0$,*

$$\mathbb{P}\Big(|\mathcal{C}^p_{bTS}(\mathbf{X},\mathbf{Y}) - \mathbb{E}[\mathcal{C}^p_{bTS}(\mathbf{X},\mathbf{Y})]| \geq \lambda n^{1-\frac{p}{d}}\Big) \leq C n^{-\theta}\lambda^{-2}. \tag{3.7}$$

As in the monopartite case, the proof reduces to uniform control of the maximal edge length. Given two alternating edges, there are exactly two ways to reconnect their endpoints. One may disconnect the cycle into two components, but the other necessarily produces a single alternating Hamiltonian cycle (after reversing the intermediate segment, as in the standard 2–opt construction). Hence at least one admissible alternating 2–opt move is always available.

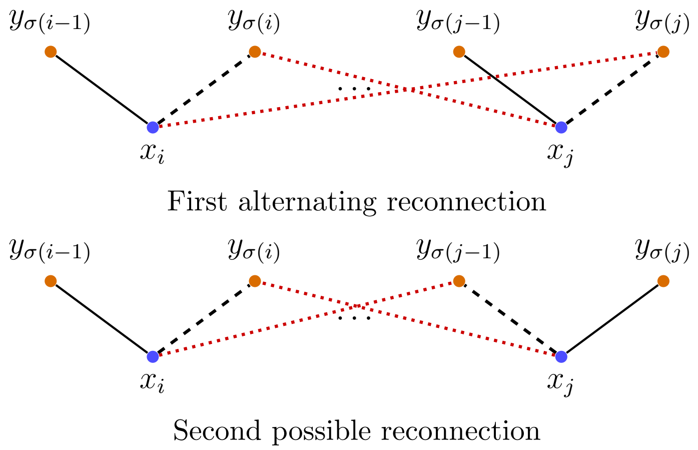

FIGURE 3. Two possible reconnections. One of the two reconnections always yields a single alternating cycle.

Lemma 3.5 (Alternating swap inequality)**.** *Let σ be an optimizer in* (3.5)*. Consider two alternating edges*

$$(x_i, y_{\sigma(i)}) \quad \textit{and} \quad (x_j, y_{\sigma(j)}), \qquad i \neq j.$$

Then the admissible alternating reconnection does not decrease the cost.

Thus the same triangle-inequality argument yields:

Lemma 3.6 (Local edge-to-energy control for bipartite TSP)**.** *For any $p > 1$, there exist $\varepsilon = \varepsilon(p) > 0$ and $C = C(p) < \infty$ such that for every optimizer σ and every edge $e = (x_i, y_{\sigma(i)})$ with midpoint z_e and $B_e = B(z_e, \varepsilon|x_i - y_{\sigma(i)}|)$,*

$$N^{\mathbf{x}}_{B_e}\, |x_i - y_{\sigma(i)}|^p \leq C \sum_{j:\ x_j \in B_e} |x_j - y_{\sigma(j)}|^p.$$

The proof is identical to the matching case. Once Lemma 3.6 is available, the argument for Theorem 3.4 follows verbatim as in the matching and monopartite TSP cases.

Appendix A. Numerical simulations

In this section we provide numerical evidence supporting two claims:

(i) the $p \to q$ transfer conjecture appears to hold empirically in both the matching and TSP settings;
(ii) the threshold $p < d^2/2$ arises from the analytical technique rather than from an intrinsic change of behavior of the models.

All costs are normalized by the expected scaling $r(d,q)(n) = n^{1-q/d}$ (we work in $d \geq 3$), so that convergence to a positive constant corresponds to the conjectured behavior. The code used is freely available at the repository `https://github.com/DarioTrevisan/REMPF`.

A.1. **Numerical evidence for the $p \to q$ conjecture.** We test whether, for a fixed optimizer exponent p, the q–cost of the p–optimal configuration scales as $n^{1-q/d}$ for $q > p$. Monopartite TSP. Figure 4a shows the normalized costs for $d = 3$, $p = 1$ and $q = 1, 2, 3$. The curves are stable and appear to converge to positive constants, supporting the scaling prediction.

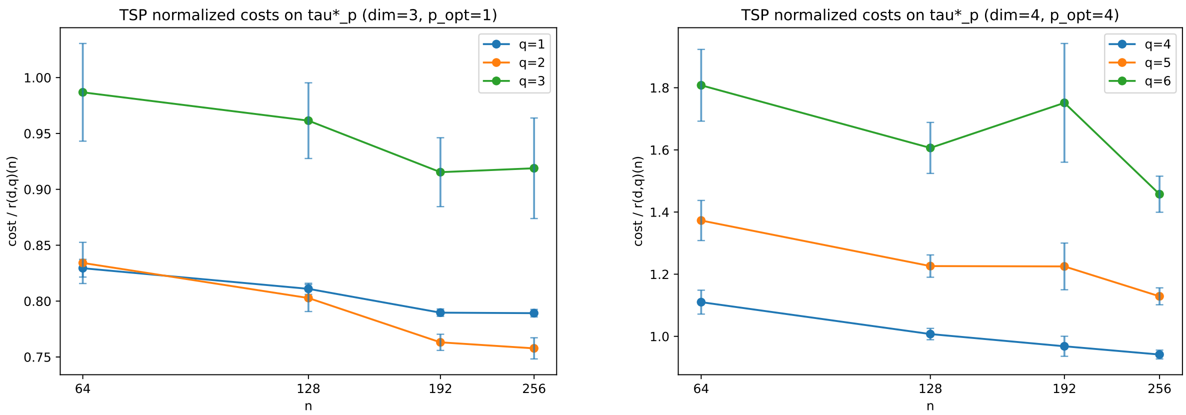

(A) $d = 3$, $p = 1$, $q = 1, 2, 3$.

(B) $d = 4$, $p = 4$, $q = 4, 5, 6$.

Figure 4. Monopartite TSP: normalized q–costs on τ_p^*. In both regimes, the curves remain approximately flat after normalization by $n^{1-q/d}$.

Bipartite matching. We perform the same experiment for bipartite matching.

Figure 5a corresponds to $d = 3$, $p = 1$, $q = 1, 2, 3$. The normalized costs stabilize as n grows. Figure 5b shows the case $d = 4$, $p = 4$, $q = 4, 5, 6$. The behavior is again compatible with the conjectured $n^{1-q/d}$ scaling.

Overall, in all tested regimes the normalized q–costs appear bounded and stable, lending strong numerical support to the $p \to q$ conjecture.

A.2. **The threshold $p < d^2/2$ is technical.** We now examine the critical case $d = 4$, $p = 8$, for which $p = d^2/2$. The analytical method developed above does not yield concentration at this threshold.

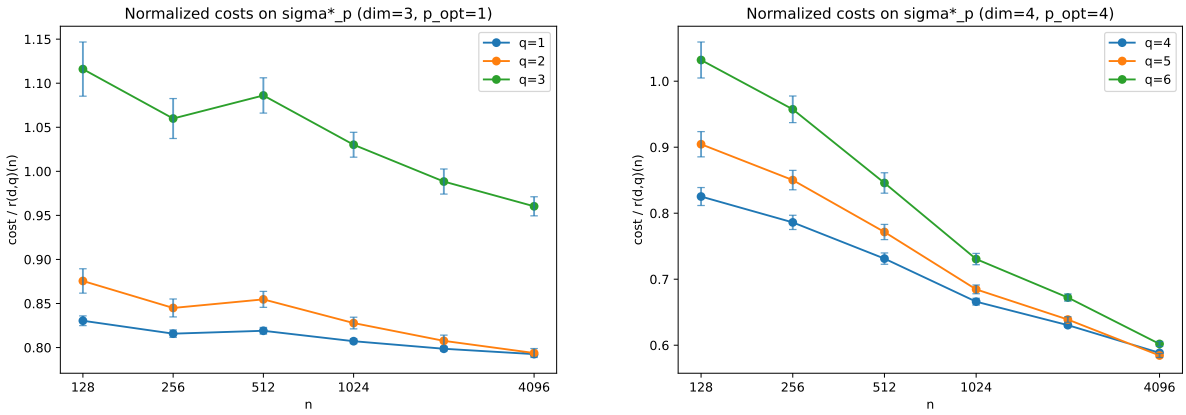

(A) $d = 3$, $p = 1$, $q = 1, 2, 3$. (B) $d = 4$, $p = 4$, $q = 4, 5, 6$.

FIGURE 5. Bipartite matching: normalized q–costs on σ_p^*. The curves remain stable under normalization by $n^{1-q/d}$.

Matching and monopartite TSP at the critical exponent. Figure 6a shows normalized costs for $d = 4$, $p = 8$, $q = 8, 9, 10$. Despite being exactly at the theoretical threshold, the normalized curves remain stable and decrease smoothly with n. Figure 6b displays the corresponding matching experiment. Again, no qualitative change in behavior is observed at $p = d^2/2$.

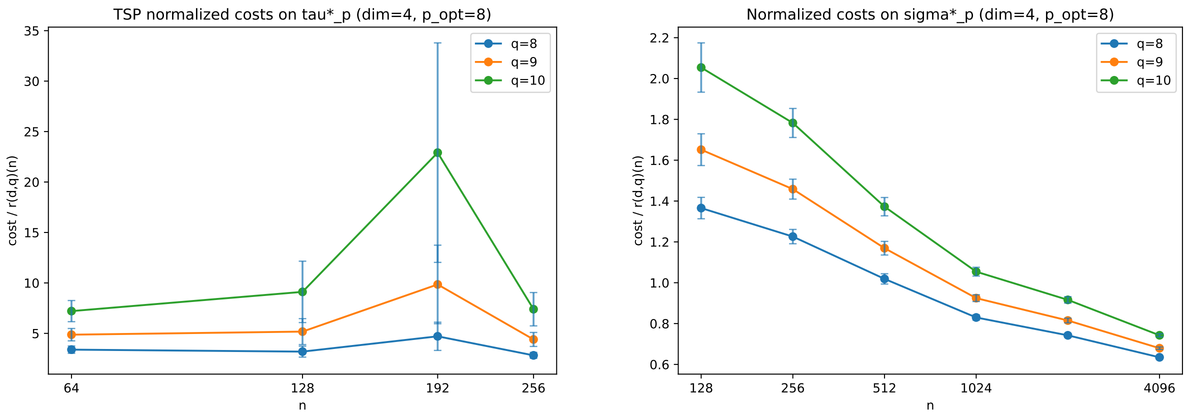

(A) Monopartite TSP, $d = 4$, $p = 8$. (B) Bipartite matching, $d = 4$, $p = 8$.

FIGURE 6. Critical regime $p = d^2/2$ ($d = 4$, $p = 8$). No qualitative change in behavior is observed at the theoretical threshold.

These experiments strongly suggest that the restriction $p < d^2/2$ is an artifact of the current analytical technique (specifically, of the edge-uniform bound combined with Poincaré), rather than a genuine phase transition in the models.

European Research Council
Established by the European Commission

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This project has received funding from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement No. 101198055, project acronym NEITALG).

References

[1] M. Ajtai, J. Komlós, and G. Tusnády. On optimal matchings. *Combinatorica*, 4:259–264, 1984.

[2] L. Ambrosio, F. Glaudo, and D. Trevisan. On the optimal map in the 2-dimensional random matching problem. *Discrete Contin. Dyn. Syst.*, 39(12):7291–7308, 2019.

[3] L. Ambrosio, F. Stra, and D. Trevisan. A PDE approach to a 2-dimensional matching problem. *Probability Theory and Related Fields*, 173:433–477, 2019. Published online 19 Feb 2018.

[4] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.

[5] J. Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. *Proceedings of the Cambridge Philosophical Society*, 55(4):299–327, 1959.

[6] S. Bobkov and M. Ledoux. A simple Fourier analytic proof of the AKT optimal matching theorem, 2019.

[7] R. Capelli, S. Caracciolo, A. Di Gioacchino, and E. M. Malatesta. Exact value for the average optimal cost of the bipartite traveling salesman and two-factor problems in two dimensions. *Physical Review E*, 98(3):030101, 2018.

[8] S. Caracciolo, A. Di Gioacchino, M. Gherardi, and E. M. Malatesta. Solution for a bipartite euclidean traveling-salesman problem in one dimension. *Physical Review E*, 97(5):052109, 2018.

[9] G. A. Croes. A method for solving traveling-salesman problems. *Operations Research*, 6(6):791–812, 1958.

[10] E. del Barrio, E. Giné, and C. Matrán. Central limit theorems for the Wasserstein distance between the empirical and the true distributions. *The Annals of Probability*, 27(2):1009–1071, 1999.

[11] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation*. Springer Series in Statistics. Springer, 2001.

[12] N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Related Fields*, 162(3-4):707–738, 2015.

[13] M. Goldman and F. Otto. A variational proof of partial regularity for optimal transportation maps. *Ann. Sci. Ec. Norm. Super. (4)*, 53(5):1209–1233, 2020.

[14] M. Goldman and D. Trevisan. Convergence of asymptotic costs for random euclidean matching problems. *Probability and Mathematical Physics*, 2020.

[15] M. Goldman and D. Trevisan. Optimal transport methods for combinatorial optimization over two random point sets. *Probability Theory and Related Fields*, 188(3):1315–1384, Apr. 2024.

[16] M. Huesmann, F. Mattesini, and F. Otto. There is no stationary p-cyclically monotone Poisson matching in 2d. *Electronic Journal of Probability*, 29:1–20, 2024.

[17] S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21(2):498–516, 1973.

[18] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple PTAS for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.

[19] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, 2003.

[20] J. M. Steele. *Probability Theory and Combinatorial Optimization*. CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, 1997.

[21] M. Talagrand. Matching random samples in many dimensions. *The Annals of Applied Probability*, 2(4):846–856, 1992.

[22] J. E. Yukich. *Probability Theory of Classical Euclidean Optimization Problems*, volume 1675 of *Lecture Notes in Mathematics*. Springer, 1998.

(MATTEO D'ACHILLE) INSTITUT ÉLIE CARTAN DE LORRAINE, CNRS, UNIVERSITÉ DE LORRAINE, F-57070, METZ, FRANCE

Email address: MATTEO.D-ACHILLE@UNIV-LORRAINE.FR

(FRANCESCO MATTESINI) DEPARTMENT OF MATHEMATICS, CIT SCHOOL, TECHNICAL UNIVERSITY OF MUNICH
Email address: FRANCESCO.MATTESINI@TUM.DE

(DARIO TREVISAN) UNIVERSITÀ DI PISA, ITALY
Email address: DARIO.TREVISAN@UNIPI.IT