Improved Smoothed Analysis of 2-Opt for the Euclidean TSP

Bodo Manthey ✉ 🏠 🆔
Faculty of Electrical Engineering, Mathematics, and Computer Science, University of Twente, Enschede, The Netherlands

Jesse van Rhijn ✉ 🏠 🆔
Faculty of Electrical Engineering, Mathematics, and Computer Science, University of Twente, Enschede, The Netherlands

─── **Abstract** ───

The 2-opt heuristic is a simple local search heuristic for the Travelling Salesperson Problem (TSP). Although it usually performs well in practice, its worst-case running time is poor. Attempts to reconcile this difference have used smoothed analysis, in which adversarial instances are perturbed probabilistically. We are interested in the classical model of smoothed analysis for the Euclidean TSP, in which the perturbations are Gaussian. This model was previously used by Manthey & Veenstra, who obtained smoothed complexity bounds polynomial in n, the dimension d, and the perturbation strength σ^{-1}. However, their analysis only works for $d \geq 4$. The only previous analysis for $d \leq 3$ was performed by Englert, Röglin & Vöcking, who used a different perturbation model which can be translated to Gaussian perturbations. Their model yields bounds polynomial in n and σ^{-d}, and super-exponential in d. As the fact that no direct analysis exists for Gaussian perturbations that yields polynomial bounds for all d is somewhat unsatisfactory, we perform this missing analysis. Along the way, we improve all existing smoothed complexity bounds for Euclidean 2-opt with Gaussian perturbations.

1 Introduction

The Travelling Salesperson problem is a standard combinatorial optimization problem, which has attracted considerable interest from academic, educational and industrial directions. It can be stated rather compactly: given a Hamiltonian graph $G = (V, E)$ and edge weights $w : E \to \mathbb{R}$, find a minimum weight Hamiltonian cycle (tour) on G.

Despite this apparent simplicity, the TSP is NP-hard [6]. A particularly interesting variant of the TSP is the Euclidean TSP, in which the n vertices of the graph are identified with a point cloud in $\mathbb{R}^d$, and the edge weights are the Euclidean distances between these points. Even this restricted variant is NP-hard [10].

As a consequence of this hardness, practitioners often turn to heuristics. One often-used heuristic is 2-opt [1]. This heuristic takes as its input a tour T, and finds two sets of two edges each, $\{e_1, e_2\} \subseteq T$ and $\{f_1, f_2\} \not\subseteq T$, such that exchanging $\{e_1, e_2\}$ for $\{f_1, f_2\}$ yields

34th International Symposium on Algorithms and Computation (ISAAC 2023).
Editors: Satoru Iwata and Naonori Kakimura; Article No. 52; pp. 52:1–52:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

again a tour T', and the total weight of T' is strictly less than the total weight of T. This procedure is repeated with the new tour, and stops once no such edges exist. The resulting tour is said to be locally optimal.

Englert, Röglin and Vöcking constructed Euclidean TSP instances on which 2-opt can take exponentially many steps to find a locally optimal tour [4]. Despite this pessimistic result, 2-opt performs remarkably well in practice, usually requiring time sub-quadratic in n and obtaining tours which are only a few percent worse than the optimum [1, chapter 8].

To explain this discrepancy, the tools of probabilistic analysis have been employed [9, 2, 5, 3, 4]. In particular, smoothed analysis, a hybrid framework between worst-case and average-case analysis, has been successfully used in the analysis of 2-opt [5, 4, 9]. In the original version of this framework, the instances one considers are initially adversarial, and then perturbed by Gaussians. The resulting smoothed time complexity is then generally a function of the instance size n and the standard deviation of the Gaussian perturbations, σ.

Englert et al. obtained smoothed time complexity bounds for 2-opt on Euclidean instances by considering a more general model, in which the points are chosen in the unit hypercube according to arbitrary probability densities. The only restrictions to these densities are that (i) they are independent, and (ii) they are all bounded from above by ϕ. Their results can be transferred to Gaussian perturbations roughly by setting $\phi = \sigma^{-d}$, which yields a smoothed complexity that is $O(\text{poly}(n, \sigma^{-d}))$.

As the exponential dependence on d is somewhat unsatisfactory, Manthey & Veenstra [9] performed a simpler smoothed analysis yielding bounds polynomial in n, $1/\sigma$, and d. The analysis they performed is however limited to $d \geq 4$. While polynomial bounds for all d can be obtained by simply taking the result of Englert et al. for $d \in \{2, 3\}$, no smoothed analysis that directly uses Gaussian perturbations exists for these cases. We set out to perform this missing analysis, improving the smoothed complexity bounds for all $d \geq 2$ along the way.

Our analysis combines ideas from both Englert et al. and Manthey & Veenstra. From the former, we borrow the idea of conditioning on the outcomes of some of the distances between points in an arbitrary 2-change. We can then analyze the 2-change by examining the angles between certain edges in the 2-change, which are themselves random variables. From the latter, we borrow the Gaussian perturbation model (originally introduced by Spielman & Teng for the Simplex Method [11]).

We also note that in addition to improving the results of Manthey & Veenstra, our approach is significantly simpler than the analysis of Englert et al. The crux of the simplification is a carefully constructed random experiment to model a single 2-change, which allows us to bypass the need for the involved convolution integrals used by Englert et al.

We will begin by introducing some definitions and earlier results, before providing basic probability theoretical results (Section 2) that we will make heavy use of throughout the paper. We then proceed by analyzing a single 2-change in a similar manner as Englert et al., simplifying some of their analysis in the process (Section 3). Next, we prove a first smoothed complexity bound by examining so-called linked pairs of 2-changes (Section 4), an idea used by both Englert et al. and Manthey & Veenstra. Finally, we improve on this bound for $d \geq 3$ (Section 5), yielding the best known bounds for all dimensions.

2   Preliminaries

2.1   Travelling Salesperson Problem

Let $\mathcal{Y} \subseteq [-1, 1]^d$ be a point set of size n. The Euclidean Travelling Salesperson Problem (TSP) asks for a tour that visits each point $y \in \mathcal{Y}$ exactly once, such that the total length of the tour is minimized. The length of a tour in this variant of the TSP is the sum of the

Euclidean distances between consecutive points in the tour. Formally, if the points in $\mathcal{Y}$ are visited in the order $T = (y_{\pi(i)})_{i=0}^{n-1}$ defined by a permutation π of $[n]$, then the length of the tour T is

$$L(T) = \sum_{i=0}^{n-1} \|y_{\pi(i)} - y_{\pi(i+1)}\|,$$

where the indices are taken modulo n, and $\|\cdot\|$ denotes the standard Euclidean norm in $\mathbb{R}^d$. Since the Euclidean TSP is undirected, the tour T' in which the vertices are visited in the reverse order has the same length as T. We consider these tours to be identical.

2.2 Smoothed Analysis

Smoothed analysis is a framework for the analysis of algorithms, which was introduced in 2004 by Spielman & Teng [11]. The method is particularly suitable to algorithms with a fragile worst case input [7]. Since its introduction, the method has been applied to a wide variety of algorithms [8, 12].

Heuristically, one imagines that an adversary chooses an input to the algorithm. The input is then perturbed in a probabilistic fashion. The hope is that any particularly pathological instances that the adversary might choose are destroyed by the random perturbation. One then computes a bound on the expected number of steps that the algorithm performs, where the expectation is taken with respect to the perturbation.

For our model of a smoothed TSP instance, we allow the adversary to choose a point set $\mathcal{Y} \subseteq [-1, 1]^d$ of size n. We then perturb each point $y_i \in \mathcal{Y}$ with an independent d-dimensional Gaussian random variable g_i, $i \in [n]$, with mean 0 and standard deviation σ. This yields a new point set, $\mathcal{X} = \{y_i + g_i \mid y_i \in \mathcal{Y}\}$. We will bound the expected number of steps taken by the 2-opt heuristic on the TSP instance defined by $\mathcal{X}$, with the expectation taken over this Gaussian perturbation. We will refer to this quantity as the smoothed complexity of 2-opt.

For the purposes of our analysis, we always assume that $\sigma \leq 1$. This is a mild restriction, as the bound for $\sigma = 1$ also applies to all larger values of σ, and small perturbations are particularly interesting in smoothed analysis.

For a general outline of the strategy, consider a 2-change where the edges $\{a, z_1\}$ and $\{b, z_2\}$ are replaced by $\{a, z_2\}$ and $\{b, z_1\}$. The change in tour length of this 2-change is

$$\Delta = \|a - z_1\| + \|b - z_2\| - \|a - z_2\| - \|b - z_1\|.$$

Since the locations of the points $\{a, b, z_1, z_2\}$ are random variables, so is Δ. We seek to bound the probability that there exists a 2-change whose improvement is exceedingly small, enabling us to use a potential argument.

Let $\Delta_{\min}$ denote the improvement of the least-improving 2-change in the instance. If $\mathbb{P}(\Delta_{\min} \leq \epsilon)$ is suitably small for small ϵ, then each iteration is likely to decrease the tour length by a large amount. As long as the initial tour has bounded length, this then provides a limit to the number of iterations that the heuristic can perform, since the tour length is bounded from below by 0.

2.3 Basic Results

We state some general results that we will need at points throughout the paper.

The next lemma provides a simple framework that we can use to prove smoothed complexity bounds for 2-opt.

Let $\Delta_{\min}$ denote the smallest improvement of any 2-change, and let $\Delta_{\min}^{\text{link}}$ denote the smallest improvement of any pair of linked 2-changes (see Section 4 for a definition of linked pairs).

▶ **Lemma 1** ([9, Lemma 2.2]). *Suppose that the longest tour has a length of at most L with probability at least $1 - 1/n!$. Let $\alpha > 1$ be a constant. If for all $\epsilon > 0$ it holds that $\mathbb{P}(\Delta_{\min} \in (0, \epsilon]) = O(P\epsilon^{\alpha})$, then the smoothed complexity of 2-opt is bounded from above by $O(P^{1/\alpha}L)$. The same holds if we replace $\Delta_{\min}$ by $\Delta_{\min}^{\text{link}}$, provided that $P^{1/\alpha}L = \Omega(n^2)$.*

2.4 Probability Theory

We provide some basic probability theoretical results. Throughout the paper, given a random variable X, we denote its probability density by f_X and its cumulative distribution function by F_X. If we furthermore condition on some event Y, we write $f_{X|Y}$ for the conditional density of X given Y.

2.4.1 Chi Distributions

Suppose we are given two points $y_1, y_2 \in \mathcal{Y}$ and perturb both points with independent Gaussian random variables g_1 and g_2, resulting in $x_i = y_i + g_i$, $i \in [2]$. Then the distance $\|x_1 - x_2\|$ between the two perturbed points is distributed according to a noncentral d-dimensional chi distribution with noncentrality parameter $s = \|y_1 - y_2\|$, which we denote χ_d^s. We call χ_d^0 a central d-dimensional χ distribution.

2.4.2 General Results

In the following, we use the notion of stochastic dominance. Let X and Y be two real-valued random variables. We say that X stochastically dominates Y if for all x, it holds that $\mathbb{P}(X \geq x) \geq \mathbb{P}(Y \geq x)$, and this inequality is strict for some x. We may equivalently say that the density of X stochastically dominates the density of Y.

To use Lemma 1, we need to limit the probability that any TSP tour in our smoothed instance is too long. This was previously done by Manthey & Veenstra; we state their result in Lemma 2.

▶ **Lemma 2** ([9, Lemma 2.3]). *Let $c \geq 2$ be a sufficiently large constant, and let $D = c \cdot (1 + \sigma\sqrt{n \log n})$. Then $\mathbb{P}(\mathcal{X} \nsubseteq [-D, D]^d) \leq 1/n!$.*

The next lemma is a reformulation of another result by Manthey & Veenstra [9]. The lemma is very useful in conjunction with Lemma 4, as we will have cause to condition on the outcome of drawing noncentral d-dimensional chi random variables.

▶ **Lemma 3** ([9, Lemma 2.8]). *The noncentral d-dimensional chi distribution with parameter $\mu > 0$ and standard deviation σ stochastically dominates the central d-dimensional chi distribution with the same standard deviation.*

The following lemma from Manthey & Veenstra is slightly generalized compared to its original statement. We do not provide a proof, since the original proof remains valid when simply replacing the original assumption with ours.

▶ **Lemma 4** ([9, Lemma 2.7]). *Assume $c \in \mathbb{R}_{\geq 0}$ is a fixed constant and $d \in \mathbb{N}$ is fixed and arbitrary with $d > c$. Let χ_d denote the d-dimensional chi distribution with variance σ^2. Then*

$$\int_0^{\infty} \chi_d(x)x^{-c}\,\mathrm{d}x = \Theta\left(\frac{1}{d^{c/2}\sigma^c}\right).$$
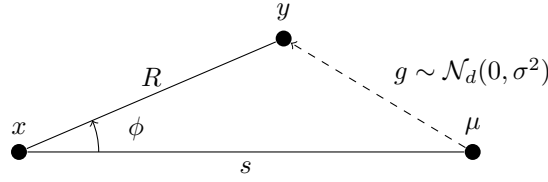
2.4.3 Limiting the Adversary

In our analysis we will closely study the angles between edges in the smoothed TSP instance. These angles can be initially specified to our detriment by the adversary. However, the power of the adversary is limited by the strength of the Gaussian perturbations. We quantify the power of the adversary in Theorem 5. See Figure 1 for a sketch accompanying the theorem.

▶ **Theorem 5.** *Let L be some line in $\mathbb{R}^d$, and let $x \in L$. Let y be a point drawn from a d-dimensional Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and variance σ^2. Let ϕ denote the angle between L and $x - y$, and let $R = \|x - y\|$ and $s = \|x - \mu\|$. Let $f_{\phi|R=r}$ denote the density of ϕ, conditioned on a specific outcome $r > 0$ for R. Then for all $d \geq 2$,*

$$\sup_{\phi \in [0,\pi]} f_{\phi|R=r}(\phi) = O\left(\sqrt{d} + \frac{\sqrt{rs}}{\sigma}\right).$$

Moreover, for $d \geq 3$,

$$\sup_{\phi \in (0,\pi)} \frac{f_{\phi|R=r}(\phi)}{\sin \phi} = O\left(\sqrt{d} + \frac{rs}{\sigma^2\sqrt{d}}\right).$$

Theorem 5 yields the following corollary, which provides information on the angle between two Gaussian random points in $\mathbb{R}^d$ with respect to some third point.

▶ **Corollary 6.** *Let $x \in \mathbb{R}^d$. Let y and z be drawn from d-dimensional Gaussian distributions with arbitrary means and the same variance σ^2. Let ϕ denote the angle between $y - x$ and $z - x$, and let $R = \|x - y\|$ and $S = \|x - z\|$. Let $f_{\phi|R=r,S=s}$ denote the density of ϕ conditioned on some outcome $r > 0$ for R and $s > 0$ for S. Then for all $d \geq 2$,*

$$\sup_{\phi \in [0,\pi]} f_{\phi|R=r,S=s}(\phi) = O\left(\sqrt{d} + \frac{\sqrt{\min\{r\bar{r}, s\bar{s}\}}}{\sigma}\right),$$

where $\bar{r} = \|x - \mathbb{E}(y)\|$ and $\bar{s} = \|x - \mathbb{E}(z)\|$. Moreover, for $d \geq 3$,

$$\sup_{\phi \in (0,\pi)} \frac{f_{\phi|R=r,S=s}(\phi)}{\sin \phi} = O\left(\sqrt{d} + \frac{\min\{r\bar{r}, s\bar{s}\}}{\sigma^2\sqrt{d}}\right).$$

3 Analysis of Single 2-Changes

To improve upon the previous analyses, it pays to examine where the analysis of Euclidean 2-opt with Gaussian perturbations [9] fails for $d \in \{2, 3\}$. The problem is that in the course of the proof, Manthey & Veenstra compute

$$\int_0^\infty \frac{1}{x^2} \chi_{d-1}(x) \, dx,$$

where χ_d denotes the d-dimensional chi distribution. This integral is finite only when $d \geq 4$.
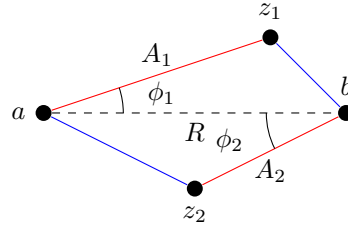
Figure 2 Labels of points and angles involved in a single 2-change.

This problem does not appear in the results obtained by Englert et al. [4]. They consider a more general model of smoothed analysis wherein the adversary specifies a probability density for each point in the TSP instance independently. Since the only information available on the probability densities is their upper bound, they consider a simplified model of a 2-change to keep the analysis tractable. The analysis is then translated to their generic model, which incurs a factor which is super-exponential in d.

Even when one considers d to be a constant as Englert et al. do, the genericity of their model still comes at a cost when translated to a smoothed analysis with Gaussian perturbations, eventually yielding a bound which is polynomial in σ^{-d}.

Specifying the perturbations as Gaussian enables us to analyze the true random experiment modeling a 2-change more closely, as we know the distributions of the distances between points in the smoothed instance. Combined with Theorem 5, which provides information on the angles between edges in the instance, we can carry out an analysis that improves on both Englert et al.'s as well as Manthey & Veenstra's analysis.

We first set up our model of a 2-change perturbed by Gaussian perturbations. To obtain a bound for this case, we first formulate a different analysis of single 2-changes. Consider a 2-change involving the points $\{a, b, z_1, z_2\} \subseteq [-D, D]^d$, where the edges $\{a, z_1\}$ and $\{b, z_2\}$ are replaced by $\{b, z_1\}$ and $\{a, z_2\}$. The improvement to the tour length due to this 2-change is

$$\Delta = \|a - z_1\| - \|b - z_1\| + \|b - z_2\| - \|a - z_2\|.$$

To analyze Δ, we first define $A_1 := \|a - z_1\|$, $A_2 := \|b - z_2\|$, and $R := \|a - b\|$. Moreover, we identify the angle ϕ_1 as the angle between $a - z_1$ and $a - b$, and restrict it to $[0, \pi]$. The corresponding angle ϕ_2 is defined similarly. The restriction of these angles to $[0, \pi]$ is without loss of generality; one may readily observe from Figure 2 that flipping the sign of either ϕ_1 or ϕ_2 does not change the value of Δ.

While Figure 2 may give the impression that we are restricting the analysis to the $d = 2$ case, the analysis is valid for any $d \geq 2$. The two triangles $\triangle az_1 b$ and $\triangle az_2 b$ will lie in two separate planes in general. The distances involved must thus be understood as d-dimensional Euclidean distances.

With these definitions, we have $\Delta = \eta_1 + \eta_2$, where for $i \in [2]$

$$\eta_i = A_i - \sqrt{A_i^2 + R^2 - 2A_i R \cos \phi_i},$$

which follows from the Law of Cosines.

Suppose we condition on the events $A_1 = a_1$, $A_2 = a_2$, and $R = r$, for some $a_1, a_2, r > 0$. Under these events, η_1 and η_2 are independent random variables. Moreover, Δ is completely fixed by revealing the angles ϕ_1 and ϕ_2. Since we condition on $A_i = a_i$, we can then bound the density of ϕ_i using Corollary 6.

We can use this independence to obtain bounds for $\mathbb{P}(\Delta \in (0, \epsilon))$ for some small $\epsilon > 0$ under these events, for various orderings of a_1, a_2 and r. These bounds are given in Lemma 10.

We begin by obtaining a bound to the density of η_i, $i \in [2]$, using the fact that all randomness in η_i is contained in the angle ϕ_i under the conditioning that $A_i = a_i$ and $R = r$. We denote by $f_{\phi_i|R=r,A_i=a_i}$ the density of the angle ϕ_i, conditioned on $R = r$ and $A_i = a_i$.

▶ **Lemma 7.** *Let $i \in [2]$. The density of $\eta_i = \|a - z_i\| - \|b - z_i\|$, conditioned on $A_i = a_i$ and $R = r$, is bounded from above by*

$$\frac{a_i + r}{a_i r} \cdot \frac{f_{\phi_i|R=r,A_i=a_i}(\phi_i(\eta))}{|\sin \phi_i(\eta)|},$$

where $\phi_i(\eta) = \arccos\left(\frac{a_i^2 + r^2 - (a_i - \eta)^2}{2a_i r}\right)$.

Proof. Let the conditional density of η_i be $f_{\eta_i|R=r,A_i=a_i}$. Since ϕ_i is restricted to $[0, \pi]$ by assumption, there exists a bijection between η_i and ϕ_i. To be precise, we have

$$\phi_i(\eta_i) = \arccos\left(\frac{a_i^2 + r^2 - (a_i - \eta_i)^2}{2a_i r}\right).$$

By standard transformation rules of probability densities, it holds that

$$f_{\eta_i|R=r,A=a_i}(\eta) = \left|\frac{\mathrm{d}\phi_i(\eta)}{\mathrm{d}\eta}\right| f_{\phi_i|R=r,A_i=a_i}(\phi_i(\eta)).$$

The derivative is easily evaluated:

$$\frac{\mathrm{d}\phi_i(\eta)}{\mathrm{d}\eta} = \frac{-1}{\sqrt{1 - \left(\frac{a_i^2 + r^2 - (a_i - \eta)^2}{2a_i r}\right)}} \cdot \frac{a_i - \eta}{a_i r} = \frac{-1}{\sin \phi(\eta)} \cdot \frac{a_i - \eta}{a_i r}.$$

Finally, we have $a_i - \eta \leq a_i + r$, which follows from the triangle inequality. This concludes the proof. ◀

By Corollary 6, we have an upper bound for $f_{\phi_i|R=r,A_i=a_i}$. Unfortunately, simply inserting this upper bound is not enough for us to bound $f_{\eta_i|A_i=a_i,R=r}$, since the density as obtained from Lemma 7 diverges for $\phi = 0$ and $\phi = \pi$. There is however a way to cure this divergence.

We now consider a full 2-change (cf. Figure 2). To analyze the improvement Δ caused by this 2-change, we construct a random experiment, conditioned on the outcomes $A_1 = a_1$, $A_2 = a_2$, and $R = r$. We write this random experiment in Algorithm 1, since we will need to execute different experiments depending on the ordering of the values of a_1, a_2 and r. The parameters b_1 and b_2 of this algorithm will take values in $\{a_1, a_2, r\}$, depending on this ordering.

The function `RandomExpt` outlined in Algorithm 1 branches on the outcome of the variable $Z_i = \sqrt{b_i} \sin \phi_i$, $i \in [2]$, where b_i is some distance; we will choose b_i among $\{r, a_i\}$ in subsequent lemmas.

Note that `RandomExpt` returns a tuple (i, ϕ), where $i \in [2]$. We call the angle returned by `RandomExpt` the *good angle*. Moreover, we label the event $i = 1$ as E_1, and $i = 2$ by E_2. The crux of the analysis is now to analyze η_1 if E_1 occurs, and η_2 if E_2 occurs, as under E_i the density of η_i is bounded from above.

■ **Algorithm 1** The algorithm we use to model a random 2-change with fixed $A_1 = a_1$, $A_2 = a_2$, and $R = r$.

Data: Distances $b_1, b_2 > 0$.
Function RandomExpt(b_1, b_2):

     Draw $\phi_1 \sim f_{\phi|R=r, A_1=a_1}$
     Draw $\phi_2 \sim f_{\phi|R=r, A_2=a_2}$
     if $\sqrt{b_1} \sin \phi_1 > \sqrt{b_2} \sin \phi_2$ **then**
        return $(1, \phi_1)$
     else
        return $(2, \phi_2)$
     end

▶ **Lemma 8.** *Let $(i, \phi) = $* RandomExpt$(b_1, b_2)$ *for some $b_1, b_2 > 0$. Let $j = 3 - i$. The density of ϕ, conditioned on $R = r$, $A_1 = a_1$, $A_2 = a_2$, is then bounded from above by*

$$\frac{2 M_{\phi_1} M_{\phi_2}}{\mathbb{P}(E_i)} \cdot \arcsin\left(\min\left\{ 1, \sqrt{\frac{b_i}{b_j}} \sin \phi \right\} \right),$$

where $M_{\phi_i} = \max_{0 \leq \phi \leq \pi} f_{\phi_i | R=r, A_i=a_i}(\phi)$.

Proof. We omit the conditioning on $A_1 = a_1$, $A_2 = a_2$ and $R = r$ in the following, for the sake of clarity. We prove only the case $i = 1$, thus conditioning on E_1, as the proof for $i = 2$ proceeds essentially identically.

Let $X_i = \sqrt{b_i} \sin \phi_i$, $i \in [2]$. The event E_1 is then equivalent to $X_1 > X_2$. Let Z in turn denote the random variable given by X_1 conditioned on E_1. The cumulative distribution function of Z is equal to

$$F_Z(x) = \mathbb{P}(X_1 \leq x \mid X_1 > X_2) = \frac{\mathbb{P}(X_1 \leq x \wedge X_1 > X_2)}{\mathbb{P}(E_1)}.$$

By the independence of X_1 and X_2, this is equal to

$$F_Z(x) = \frac{1}{\mathbb{P}(E_1)} \cdot \int_0^x f_{X_1}(y) \int_0^y f_{X_2}(z) \, \mathrm{d}z \, \mathrm{d}y.$$

Computing the density of Z is then simply a matter of differentiation. Since $\mathbb{P}(E_1)$ does not depend on x, we obtain

$$f_Z(x) = \frac{1}{\mathbb{P}(E_1)} \cdot f_{X_1}(x) \int_0^x f_{X_2}(z) \, \mathrm{d}z.$$

We next require the density of $X_i = \sqrt{b_i} \sin \phi_i$. Observe that

$$\mathbb{P}(X_i \leq x) = \mathbb{P}\left(\phi_i \leq \arcsin(x/\sqrt{b_i}) \right) + \mathbb{P}\left(\phi_i \geq \pi - \arcsin(x/\sqrt{b_i}) \right). \tag{1}$$

Differentiating this expression to x, we find for $x < \sqrt{b_i}$

$$\begin{aligned}
f_{X_i}(x) &= \frac{\mathrm{d}}{\mathrm{d}x} \left(\mathbb{P}\left(\phi_i \leq \arcsin(x/\sqrt{b_i}) \right) + 1 - \mathbb{P}\left(\phi_i \geq \pi - \arcsin(x/\sqrt{b_i}) \right) \right) \\
&= \frac{\mathrm{d}}{\mathrm{d}x} \left(\arcsin\left(\frac{x}{\sqrt{b_i}} \right) \right) \cdot \left[f_{\phi_i}\left(\arcsin\left(\frac{x}{\sqrt{b_i}} \right) \right) + f_{\phi_i}\left(\pi - \arcsin\left(\frac{x}{\sqrt{b_i}} \right) \right) \right] \\
&= \frac{1}{\sqrt{b_i - x^2}} \cdot \left[f_{\phi_i}\left(\arcsin\left(\frac{x}{\sqrt{b_i}} \right) \right) + f_{\phi_i}\left(\pi - \arcsin\left(\frac{x}{\sqrt{b_i}} \right) \right) \right],
\end{aligned}$$

and 0 for $x \geq \sqrt{b_i}$. Letting $M_{\phi_i} = \max_{0 \leq \phi \leq \pi} f_{\phi_i | R=r, A_i=a_i}(\phi)$, which exists by Corollary 6, we obtain

$$
f_{X_i}(x) \leq 2M_{\phi_i} \cdot \begin{cases} \frac{1}{\sqrt{b_i - x^2}}, & \text{if } x < \sqrt{b_i}, \\ 0, & \text{otherwise.} \end{cases}
$$

Using this density, together with the identity $\int_0^x (\sqrt{b} - y^2)^{-1/2} \, \mathrm{d}y = \arcsin(x/\sqrt{b})$ for $x < \sqrt{b}$, we obtain

$$
f_Z(x) \leq \frac{2M_{\phi_1} M_{\phi_2}}{\mathbb{P}(E_1)} \cdot \frac{\arcsin\left(\min\left\{1, \frac{x}{\sqrt{b_2}}\right\}\right)}{\sqrt{b_1 - x^2}}
$$

if $x < \sqrt{b_1}$, and $f_Z(x) = 0$ otherwise. It remains to convert Z back to ϕ, where ϕ is the good angle. Since we have conditioned on E_1, we know that $Z = \sqrt{b_1} \sin \phi$. Using similar considerations as used in Equation (1), we have

$$
f_Z(x) = \frac{1}{\sqrt{b_1 - x^2}} f_\phi(\arcsin(x/\sqrt{b_1})) + \frac{1}{\sqrt{b_1 - x^2}} f_\phi(\pi - \arcsin(x/\sqrt{b_1})).
$$

Since this expression holds for all $x \in (0, \sqrt{b_1})$, and since probability densities are non-negative, it follows that

$$
f_\phi(\phi) \leq \frac{2M_{\phi_1} M_{\phi_2}}{\mathbb{P}(E_1)} \cdot \arcsin\left(\min\left\{1, \sqrt{\frac{b_1}{b_2}} \sin \phi\right\}\right),
$$

for all $\phi \in (0, \pi)$. ◀

For the next part, we apply Lemma 8 to Lemma 7 to bound the density of η_i, given that E_i occurs.

▶ **Lemma 9.** *Let $i \in [2]$ and $j = 3 - i$. Let $f_{\eta_i | E_i}$ denote the density of η_i, conditioned on E_i as well as the outcomes $R = r$, $A_1 = a_1$, and $A_2 = a_2$. Then*

$$
f_{\eta_i | E_i}(\eta) \leq \frac{1}{\mathbb{P}(E_i)} \cdot \frac{2\pi M_{\phi_1} M_{\phi_2}}{\min\{a_1, r\} \min\{a_2, r\}},
$$

where $M_{\phi_i} = \max_{0 \leq \phi \leq \pi} f_{\phi_i | R=r, A_i=a_i}(\phi)$.

Proof. We prove only the case $i = 1$. From Lemma 7, we know that

$$
f_{\eta_i | E_i}(\eta) \leq \frac{a_i + r}{a_i r} \cdot \frac{f_{\phi_i | E_i, A_1=a_1, A_2=a_2}(\phi)}{\sin \phi}.
$$

Let $(i, \phi) = \texttt{RandomExpt}(b_1, b_2)$, for some $b_1, b_2 > 0$. We will choose values for b_1 and b_2 depending on the ordering of a_1, a_2 and r. Note that we may do this, since we know the choices of a_1, a_2 and r before executing $\texttt{RandomExpt}$.

Since we condition on E_1, we know that $i = 1$, and hence that ϕ_1 is the good angle. By Lemma 8, we can obtain a bound for $f_{\phi | E_i, A_1=a_1, A_2=a_2, R=r}$. We thus find

$$
f_{\eta_1 | E_1}(\eta) \leq \frac{2M_{\phi_1} M_{\phi_2}}{\mathbb{P}(E_1)} \cdot \frac{a_1 + r}{a_1 r} \cdot \frac{\arcsin\left(\min\left\{1, \sqrt{\frac{b_1}{b_2}} \sin \phi\right\}\right)}{\sin \phi}.
$$

First, suppose $\sin \phi \geq \sqrt{b_2/b_1}$. Then the arcsine evaluates to $\pi/2$, and so the above is bounded from above by

$$\frac{\pi}{2} \sqrt{\frac{b_1}{b_2}}.$$

Second, suppose $\sin \phi < \sqrt{b_2/b_1}$. Since $\arcsin(x) \leq \pi x/2$ for $x \in (0,1)$, this case yields the same bound, and we obtain

$$f_{\eta_1|E_1}(\eta) \leq \frac{\pi M_{\phi_1} M_{\phi_2}}{\mathbb{P}(E_1)} \cdot \frac{a_1 + r}{a_1 r} \cdot \sqrt{\frac{b_1}{b_2}}$$

We now examine the four cases in the lemma statement.

Case 1: $a_1, a_2 \leq r$.

We let $b_1 = a_1$ and $b_2 = a_2$. Then we have

$$\frac{a_1 + r}{a_1 r} \cdot \sqrt{\frac{a_1}{a_2}} = \frac{a_1 + r}{r\sqrt{a_1 a_2}} \leq \frac{2r}{r\sqrt{a_1 a_2}} = \frac{2}{\sqrt{a_1 a_2}}.$$

Case 2: $a_1, a_2 \geq r$.

We let $b_1 = b_2 = r$, and obtain

$$\frac{a_1 + r}{a_1 r} \leq \frac{2a_1}{a_1 r} = \frac{2}{r}.$$

Case 3: $a_1 \geq r \geq a_2$.

We let $b_1 = r$ and $b_2 = a_2$, which yields

$$\frac{a_1 + r}{a_1 r} \cdot \sqrt{\frac{r}{a_2}} = \frac{a_1 + r}{\sqrt{a_2 r} a_1} \leq \frac{2}{\sqrt{a_2 r}}.$$

Case 4: $a_2 \geq r \geq a_1$.

We let $b_1 = a_1$ and $b_2 = r$, to find

$$\frac{a_1 + r}{a_1 r} \sqrt{\frac{a_1}{r}} \leq \frac{2r\sqrt{a_1}}{a_1 r \sqrt{r}} = \frac{2}{\sqrt{a_1 r}}.$$

This final case concludes the proof. ◀

The bound on the density of η_i from Lemma 9 puts us in the position to prove a bound on the probability that $\Delta \in (0, \epsilon]$.

▶ **Lemma 10.** *Let Δ denote the improvement of a 2-change. Then*

$$\mathbb{P}(\Delta \in (0, \epsilon] \mid A_1 = a_1, A_2 = a_2, R = r) \leq \frac{\pi M_{\phi_1} M_{\phi_2} \epsilon}{\min\{a_1, r\} \min\{a_2, r\}},$$

where $M_{\phi_i} = \max_{0 \leq \phi \leq \pi} f_{\phi_i|R=r, A_i=a_i}(\phi)$.

Proof. We condition first on E_1, and then let an adversary choose an outcome for η_2, say, $\eta_2 = t$. Then we have $\Delta \in (0, \epsilon]$ iff $\eta_1 \in (-t, -t + \epsilon]$, which is an interval of size ϵ.

Since the probability that η_1 falls into an interval of size ϵ is at most $\epsilon \cdot \max_\eta f_{\eta_1|E_1}(\eta)$, all we need to conclude the proof for E_1 is a bound on $f_{\eta_1|E_1}(\eta)$. This is provided by Lemma 9.

We then repeat the same argument for E_2. The result is obtained by applying the Law of Total Probability. ◀

4 Linked Pairs of 2-Changes

To obtain bounds on the smoothed complexity of 2-opt, we consider so-called linked pairs of 2-changes, introduced previously by Englert et al [4]. A pair of 2-changes is said to be linked if some edge removed from the tour by one 2-change is added to the tour by the other 2-change.

Such linked pairs have been considered in several previous works [4, 9]. In each case, the distinction has been made between several types of linked pairs. In our analysis, only two of these types are relevant, and so we will describe only these types for the sake of brevity.

We consider 2-changes which share exactly one edge, and subdivide them into pairs of type 0 and of type 1. A generic 2-change removes the edges $\{z_1, z_2\}$ and $\{z_3, z_6\}$ while adding $\{z_1, z_6\}$ and $\{z_2, z_3\}$. The other 2-change removes $\{z_3, z_4\}$ and $\{z_5, z_6\}$ while adding $\{z_3, z_6\}$ and $\{z_4, z_5\}$. Note that $\{z_3, z_6\}$ occurs in both 2-changes.

- If $|\{z_1, \ldots, z_6\}| = 6$, then we say the linked pair is of type 0.
- If $|\{z_1, \ldots, z_6\}| = 5$, then we say the linked pair is of type 1.

Type 1 can itself be subdivided into two types, 1a and 1b. We will detail this distinction in Section 4.2.

Before moving on to analyzing linked pairs, we state a useful lemma that justifies limiting the discussion to just linked pairs of types 0 and 1.

▶ **Lemma 11** ([4, Lemma 9]). *In every sequence of t consecutive 2-changes the number of disjoint pairs of 2-changes of type 0 or type 1 is at least $\Omega(t) - O(n^2)$.*

4.1 Type 0

We begin with type 0, as this is by far the simplest linked pair. For clarity, see Figure 3 (left) for an illustration of a type 0 linked pair. It should be noted that, while Figure 3 shows a specific configuration of vertices in two dimensions, the results of this section hold generally; the analysis does not depend on any point having a particular orientation with respect to its neighbors. The same holds for the results in Section 4.2.

The improvement of a type 0 linked pair is completely specified by a small number of random variables. We require five distances between vertices, $R_1 = \|z_1 - z_3\|$, $A_1 = \|z_3 - z_6\|$, $A_2 = \|z_1 - z_2\|$, $R_2 = \|z_4 - z_6\|$ $A_3 = \|z_4 - z_5\|$. Additionally, we need the following angles:
1. ϕ_1 between $z_2 - z_1$ and $z_3 - z_1$,
2. ϕ_2 between $z_1 - z_3$ and $z_6 - z_3$,
3. ϕ_1' between $z_3 - z_6$ and $z_4 - z_6$,
4. ϕ_3 between $z_6 - z_4$ and $z_5 - z_4$.

Note that, if we condition on $A_1 = a_1$, the events $\Delta_1 \in (0, \epsilon]$ and $\Delta_2 \in (0, \epsilon]$ are independent. We can then apply Lemma 10, together with several applications of Lemma 4.

▶ **Lemma 12.** *Let $\Delta_{\min}^{\text{link}}$ denote the minimum improvement of any type 0 pair of linked 2-changes, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta_{\min}^{\text{link}} \in (0, \epsilon]) = O\left(\frac{dD^2 n^6 \epsilon^2}{\sigma^4}\right).$$

4.2 Type 1

As mentioned previously, type 1 linked pairs can be subdivided into two distinct subtypes. Subtype 1a shares exactly one edge between the two 2-changes, while subtype 1b shares two edges.
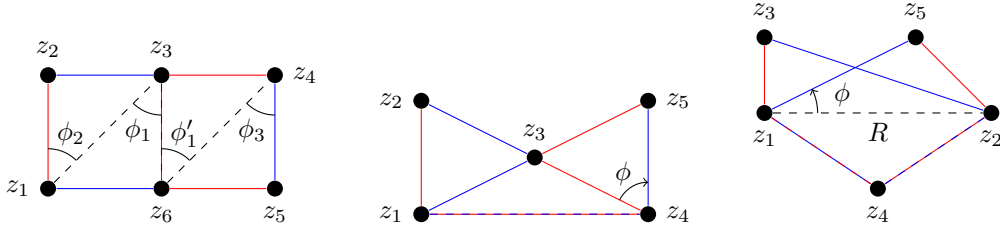
Figure 3 Labels of points involved in the three types of pairs of linked 2-changes. Left: type 0. Center: type 1a. Right: type 1b.

4.2.1 Type 1a

We first consider type 1a. See Figure 3 (center) for a graphical representation of the type, as well as the labels of the points and edges involved.

Let the 2-change replacing $\{z_1, z_2\}$ and $\{z_3, z_4\}$ by $\{z_2, z_3\}$ and $\{z_1, z_4\}$ be called S_1, and the 2-change replacing $\{z_1, z_4\}$ and $\{z_3, z_5\}$ by $\{z_1, z_3\}$ and $\{z_4, z_5\}$ be called S_2.

We proceed by conditioning on $A_2 = \|z_3 - z_4\| = a_2$ and $A_3 = \|z_4 - z_5\| = a_3$. Using Lemma 10, we can then compute the probability that $\Delta_1 \in (0, \epsilon]$. Moreover, the location of z_5 is then still random. Hence, the random variable $\eta = \|z_3 - z_5\| - \|z_4 - z_5\|$ can be analyzed independently from Δ_1.

For the density of η, we have the following lemma from Englert et al [4].

▶ **Lemma 13** ([4, Lemma 15, modified]). *Let $i \in [2]$, and assume that $\mathcal{X} \subseteq [-D, D]^d$. For $a_2, a_3 \in (0, 2\sqrt{d}D]$ and $\eta \in (-a_2, \min\{a_2, 2a_3 - a_2\})$,*

$$f_{\eta|A_2=a_2, A_3=a_3}(\eta) \leq M_\phi \cdot \begin{cases} \sqrt{\frac{2}{a_2^2 - \eta^2}}, & \text{if } a_3 \geq a_2, \\ \sqrt{\frac{2}{(a_2 + \eta)(2a_3 - a_2 - \eta)}}, & \text{if } a_3 < a_2, \end{cases}$$

where $M_\phi = \max_{0 \leq \phi \leq \pi} f_{\phi|A_2=a_2, A_3=a_3}(\phi)$. For $\eta \notin (-r, \min\{a_2, 2a_3 - a_2\})$, the density vanishes.

Note that the factor M_ϕ was not present in the original statement of Lemma 13. This is because the original statement concerned a simplified random experiment, wherein the points z_5 and z_3 are chosen uniformly from a hyperball centered on z_4. As such, ϕ is assumed to be distributed uniformly[1]. Since we do not analyze a simplified random experiment, we cannot make this assumption. However, examining the original proof of Lemma 13, this can be resolved by simply inserting the upper bound of the density of ϕ, conditioned on $A_2 = a_2$ and $A_3 = a_3$. This bound is provided to us by Corollary 6.

▶ **Lemma 14.** *Let Δ_2 be the improvement yielded by S_2, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta_2 \in (0, \epsilon] \mid A_2 = a_2) = O\left(\left(\frac{d^{1/4}\sqrt{D}}{\sigma} + \sqrt{\frac{d}{a_2}}\right) \cdot \sqrt{\epsilon}\right).$$

---

[1] This assumption is only valid for $d = 2$. To see this, observe that by conditioning on $A_i = a_i$, the point z_i is distributed uniformly on the $(d-1)$-sphere with radius a_i. For $d > 2$, the density of ϕ is thus concentrated near $\phi = \pi/2$. An upper bound for this density can be obtained by setting $s = 0$ in Theorem 5, yielding $O(\sqrt{d})$. As Englert et al. assume d to be constant, this has no effect on their eventual result.

Using Lemmas 4 and 14, we can easily prove the following statement about type 1a pairs of 2-changes.

▶ **Lemma 15.** *Let $\Delta_{\min}^{\text{link}}$ denote the minimum improvement of any type 1a pair of 2-changes, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta_{\min}^{\text{link}} \in (0, \epsilon]) = O\left(\frac{n^5 d^{3/4} D^{3/2}}{\sigma^3} \epsilon^{3/2}\right).$$

4.2.2 Type 1b

The final type of linked pair we consider is type 1b. See Figure 3 (right) for a graphical representation.

Let S_1 denote the 2-change replacing $\{z_1, z_3\}$ and $\{z_2, z_4\}$ with $\{z_2, z_3\}$ and $\{z_1, z_4\}$, and let S_2 denote the 2-change replacing $\{z_2, z_5\}$ and $\{z_1, z_4\}$ with $\{z_1, z_5\}$ and $\{z_2, z_5\}$. From Figure 3, it is evident that we can treat Δ_1 and $\eta = \|z_2 - z_5\| - \|z_1 - z_5\|$ as independent variables, as long as we condition on $R = r$.

▶ **Lemma 16.** *Let $\Delta_{\min}^{\text{link}}$ denote the minimum improvement of any type 1b pair of 2-changes, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta_{\min}^{\text{link}} \in (0, \epsilon]) = O\left(\frac{n^5 d^{3/4} D^{3/2}}{\sigma^3} \epsilon^{3/2}\right).$$

Lemmas 12, 15, and 16 enable us to prove an upper bound to the smoothed complexity of 2-opt in the present probabilistic model.

▶ **Theorem 17.** *The expected number of iterations performed by 2-opt for smoothed Euclidean instances of TSP in $d \geq 2$ dimensions is bounded from above by $O\left(dD^2 n^{4 + \frac{1}{3}}/\sigma^2\right)$.*

Proof. We assume for this proof that the entire instance is contained within $[-D, D]^d$, with $D = \Theta(1 + \sigma\sqrt{n \log n})$. This occurs with probability at least $1 - 1/n!$. Thus, with probability at least $1 - 1/n!$, the longest tour in the instance has length at most $2\sqrt{d}Dn$. The assumption that the entire instance lies within this hypercube enables us to use Lemmas 12, 15, and 16, which were proved under this assumption.

Let E denote the event that, among all type 0 and type 1 linked pairs of 2-changes, the pair with the smallest improvement is of type 0, and let E^c denote the event that this pair is of type 1a or type 1b. Let the random variable T denote the number of iterations taken by 2-opt to reach a local optimum.

We first compute $\mathbb{E}(T \mid E)$. We apply Lemma 1 with $\alpha = 2$ and $\beta = 2$, which is feasible due to Lemma 12. We then obtain immediately that $\mathbb{E}(T \mid E) = O(dD^2 n^4/\sigma^2)$.

Next, we compute $\mathbb{E}(T \mid E^c)$. In this case, we apply Lemma 1 with $\alpha = 3/2$ and $\beta = 1$ (cf. Lemmas 15 and 16). This yields $\mathbb{E}(T \mid E^c) = O(dD^2 n^{4 + \frac{1}{3}}/\sigma^2)$.

Combining the bounds for E and E^c yields the result. ◀

5 Improving the Analysis for $d \geq 3$

The bottleneck in Theorem 17 stems from Lemmas 15 and 16, which bound the probability that any linked pair of type 1a or type 1b improves the tour by at most ϵ. The probability given by these lemmas is proportional to $\epsilon^{3/2}$, which yields an extra factor of $n^{1/3}$ compared to type 0 linked pairs.

For $d \geq 3$, we can improve this to ϵ^2, yielding improved smoothed complexity bounds. The key to this improvement is to use the second part of Corollary 6 to bound the density of η_i as in Lemma 7. This immediately yields the following result on $\eta_i = \|a - z_i\| - \|b - z_i\|$.

▶ **Lemma 18.** *Let $i \in [2]$, and assume that $\mathcal{X} \subseteq [-D, D]^d$. The density of η_i in $d \geq 3$ dimensions, conditioned on $A_i = a_i$ and $R = r$, is bounded from above by*

$$O\left(\frac{a_i + r}{a_i r} \cdot \left(\sqrt{d} + \frac{D \min\{r, a_i\}}{\sigma^2}\right)\right).$$

Proof. We call the desired density $f_{\eta_i|A=a_i,R=r}$. From Lemma 7, we know that

$$f_{\eta_i|A_i=a_i,R=r}(\eta) \leq \frac{a_i + r}{a_i r} \cdot \frac{f_{\phi_i|A_i=a_i,R=r}(\phi_i(\eta))}{|\sin \phi_i(\eta)|}.$$

Since $d \geq 3$, we can use the second part of Corollary 6 to obtain the desired bound, making use of the assumption that all points fall within $[-D, D]^d$. ◀

Lemma 18 enables us to find an improved version of Lemma 10.

▶ **Lemma 19.** *Let Δ denote the improvement of a 2-change in $d \geq 3$ dimensions. Let $i \in [2]$, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta \in (0, \epsilon] \mid A_i = a_i, R = r) = O\left(\left(\frac{\sqrt{d}}{\min\{a_i, r\}} + \frac{D}{\sigma^2}\right) \cdot \epsilon\right).$$

The following lemma now yields the probability that any linked pair of 2-changes improves the tour by at most ϵ. We omit the proof, since it follow easily from Lemma 19 along the same lines as the lemmas in Section 4.

▶ **Lemma 20.** *Let $\Delta_{\min}^{\text{link}}$ denote the minimum improvement of any linked pair of 2-changes of type 0 or type 1 for $d \geq 3$, and assume that $\mathcal{X} \subseteq [-D, D]^d$. Then*

$$\mathbb{P}(\Delta_{\min}^{\text{link}} \in (0, \epsilon]) = O\left(\frac{D^2 n^6 \epsilon^2}{\sigma^4}\right).$$

We then obtain our result for $d \geq 3$.

▶ **Theorem 21.** *The expected number of iterations performed by 2-opt for smoothed Euclidean instances of TSP in $d \geq 3$ dimensions is bounded from above by $O\left(\sqrt{d} D^2 n^4 / \sigma^2\right)$.*

6   Discussion

For convenience, we provide comparisons of the previous smoothed complexity bounds with our bound from Theorem 17 in Tables 1 and 2. These bounds are provided both for small values of σ and for large values, meaning $\sigma = \Omega(1/\sqrt{n \log n})$ and $\sigma = O(1/\sqrt{n \log n})$.

Observe from Tables 1 and 2 that the bound for $d = 2$ has a worse dependence on n compared to the bound for $d \geq 3$. The technical reasons for this difference can be understood from Section 5. A more intuitive explanation for the difference is that our analysis benefits from large angles between edges in the smoothed TSP instance. In $d = 2$, the density of these angles is maximal when they are small, while for $d \geq 3$ it is maximal when the angles are large. In effect, this means that the adversary has less power to specify these angles to our detriment when $d \geq 3$.

	Englert, Röglin & Vöcking [4]	Manthey & Veenstra [9]	This paper
$d = 2$	$O\left(n^{4+\frac{1}{3}}/\sigma^{5+\frac{1}{3}} \cdot \log \frac{n}{\sigma}\right)$	–	$O\left(n^{4+\frac{1}{3}}/\sigma^2\right)$
$d = 3$	$O\left(n^{4+\frac{1}{3}}/\sigma^8 \cdot \log \frac{n}{\sigma}\right)$	–	$O\left(n^4/\sigma^2\right)$
$d \geq 4$	$O\left(c_d \cdot n^{4+\frac{1}{3}}/\sigma^{8d/3}\right)$	$O\left(\sqrt{d}n^4/\sigma^4\right)$	$O\left(\sqrt{d}n^4/\sigma^2\right)$

	Englert, Röglin & Vöcking [4]	Manthey & Veenstra [9]	This paper
$d = 2$	$O\left(n^7 \log^{3+\frac{2}{3}} n\right)$	–	$O\left(n^{5+\frac{1}{3}} \log n\right)$
$d = 3$	$O\left(n^{8+\frac{1}{3}} \log^5 n\right)$	–	$O\left(n^5 \log n\right)$
$d \geq 4$	$O\left(c_d \cdot n^{4+\frac{1+4d}{3}} \log^{1+\frac{4d}{3}} n\right)$	$O\left(\sqrt{d}n^6 \log^2 n\right)$	$O\left(\sqrt{d}n^5 \log n\right)$

From these tables, the greatest improvement is made for $d = 3$, where we improve by $n^{3+\frac{1}{3}} \log^4 n$ in the large σ case, and by $\sqrt[3]{n} \log(n/\sigma)/\sigma^6$ for small σ. For $d = 2$, the improvement is more modest at $n^{1+\frac{2}{3}} \log^{2+\frac{2}{3}} n$ for large σ and $\log(n/\sigma)/\sigma^{3+\frac{1}{3}}$ for small σ. For $d \geq 4$, we improve by $n \log n$ for large σ, and by σ^{-2} for small σ.

Note that we improve upon previous bounds mainly in the dependence on the perturbation strength. In an intuitive sense, this is most substantial for instances that are weakly perturbed from the adversarial instance, or in other words, that are close to worst case. In addition, the small-σ case is considered more interesting for a smoothed analysis, since small σ model the intuition of smoothed analysis of a small perturbation, while large σ reduce the analysis basically to an average-case analysis In order to improve the explicit dependence on n, which is the same as for Manthey & Veenstra [9], we believe new techniques are necessary.

As a final comment, we note that the techniques we employed in Sections 3 and 5 can also be used to improve and significantly simplify the analysis of the one-step model used by Englert et al [4]. For $d \geq 3$, the improvement amounts to a factor of $n^{1/3}\phi^{1/6} \log(n\phi)$, while for $d = 2$, the improvement is just $\log(n\phi)$, where ϕ denotes the upper bound of the density functions used in the one-step model.

References

1   Emile Aarts and Jan Karel Lenstra, editors. *Local Search in Combinatorial Optimization.* Princeton University Press, 2003. `doi:10.2307/j.ctv346t9c`.

2   Barun Chandra, Howard Karloff, and Craig Tovey. New Results on the Old k-opt Algorithm for the Traveling Salesman Problem. *SIAM Journal on Computing*, 28(6):1998–2029, January 1999. `doi:10.1137/S0097539793251244`.

3   Christian Engels and Bodo Manthey. Average-case approximation ratio of the 2-opt algorithm for the TSP. *Operations Research Letters*, 37(2):83–84, March 2009. `doi:10.1016/j.orl.2008.12.002`.

4   Matthias Englert, Heiko Röglin, and Berthold Vöcking. Worst Case and Probabilistic Analysis of the 2-Opt Algorithm for the TSP. *Algorithmica*, 68(1):190–264, January 2014. Corrected version: `arXiv:2302.06889`. `doi:10.1007/s00453-013-9801-4`.

5   Matthias Englert, Heiko Röglin, and Berthold Vöcking. Smoothed Analysis of the 2-Opt Algorithm for the General TSP. *ACM Transactions on Algorithms*, 13(1):10:1–10:15, September 2016. `doi:10.1145/2972953`.

6   Bernhard Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*. Algorithms and Combinatorics. Springer-Verlag, Berlin Heidelberg, 2000. `doi:10.1007/978-3-662-21708-5`.

7   Bodo Manthey. Smoothed Analysis of Local Search. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 285–308. Cambridge University Press, Cambridge, 2021. `doi:10.1017/9781108637435.018`.

8   Bodo Manthey and Heiko Röglin. Smoothed Analysis: Analysis of Algorithms Beyond Worst Case. *it – Information Technology*, 53(6):280–286, December 2011. `doi:10.1524/itit.2011.0654`.

9   Bodo Manthey and Rianne Veenstra. Smoothed Analysis of the 2-Opt Heuristic for the TSP: Polynomial Bounds for Gaussian Noise. In Leizhen Cai, Siu-Wing Cheng, and Tak-Wah Lam, editors, *Algorithms and Computation*, Lecture Notes in Computer Science, pages 579–589, Berlin, Heidelberg, 2013. Springer. Full, improved version: Full, improved version: `arXiv:2308.00306`. `doi:10.1007/978-3-642-45030-3_54`.

10   Christos H. Papadimitriou. The Euclidean travelling salesman problem is NP-complete. *Theoretical Computer Science*, 4(3):237–244, June 1977. `doi:10.1016/0304-3975(77)90012-3`.

11   Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, May 2004. `doi:10.1145/990308.990310`.

12   Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, October 2009. `doi:10.1145/1562764.1562785`.