

How to trap a gradient flow

Sébastien Bubeck

Microsoft Research

SEBUBECK@MICROSOFT.COM

Dan Mikulincer*

Weizmann Institute

DAN.MIKULINCER@WEIZMANN.AC.IL

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We consider the problem of finding an ε -approximate stationary point of a smooth function on a compact domain of \mathbb{R}^d . In contrast with dimension-free approaches such as gradient descent, we focus here on the case where d is finite, and potentially small. This viewpoint was explored in 1993 by Vavasis, who proposed an algorithm which, for *any fixed finite dimension d* , improves upon the $O(1/\varepsilon^2)$ oracle complexity of gradient descent. For example for $d = 2$, Vavasis' approach obtains the complexity $O(1/\varepsilon)$. Moreover for $d = 2$ he also proved a lower bound of $\Omega(1/\sqrt{\varepsilon})$ for deterministic algorithms (we extend this result to randomized algorithms).

Our main contribution is an algorithm, which we call *gradient flow trapping* (GFT), and the analysis of its oracle complexity. In dimension $d = 2$, GFT closes the gap with Vavasis' lower bound (up to a logarithmic factor), as we show that it has complexity $O\left(\sqrt{\frac{\log(1/\varepsilon)}{\varepsilon}}\right)$. In dimension $d = 3$, we show a complexity of $O\left(\frac{\log(1/\varepsilon)}{\varepsilon}\right)$, improving upon Vavasis' $O(1/\varepsilon^{1.2})$. In higher dimensions, GFT has the remarkable property of being a *logarithmic parallel depth* strategy, in stark contrast with the polynomial depth of gradient descent or Vavasis' algorithm. In this higher dimensional regime, the total work of GFT improves quadratically upon the only other known polylogarithmic depth strategy for this problem, namely naive grid search.

1. Introduction

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a smooth function (i.e., the map $x \mapsto \nabla f(x)$ is 1-Lipschitz, and f is possibly non-convex). We aim to find an ε -approximate stationary point, i.e., a point $x \in \mathbb{R}^d$ such that $\|\nabla f(x)\|_2 \leq \varepsilon$. It is an elementary exercise to verify that for smooth and bounded functions, gradient descent finds such a point in $O(1/\varepsilon^2)$ steps, see e.g., [Nesterov \(2004\)](#). Moreover, it was recently shown in [Carmon et al. \(2019\)](#) that this result is *optimal*, in the sense that any procedure with only black-box access to f (e.g., to its value and gradient) must, *in the worst case*, make $\Omega(1/\varepsilon^2)$ queries before finding an ε -approximate stationary point. This situation is akin to the non-smooth convex case, where the same result (optimality of gradient descent at complexity $1/\varepsilon^2$) holds true for finding an ε -approximate optimal point (i.e., such that $f(x) - \min_{y \in \mathbb{R}^d} f(y) \leq \varepsilon$), [Nemirovski and Yudin \(1983\)](#); [Nesterov \(2004\)](#).

* This work was done while D. Mikulincer was an intern at Microsoft Research. Supported by an Azrieli foundation fellowship.

There is an important footnote to both of these results (convex and non-convex), namely that optimality only holds in *arbitrarily high dimension* (specifically the hard instance in both cases require $d = \Omega(1/\varepsilon^2)$). It is well-known that in the convex case this large dimension requirement is actually necessary, for the cutting plane type strategies (e.g., center of gravity) can find ε -approximate optimal points on compact domains in $O(d \log(1/\varepsilon))$ queries. It is natural to ask: **Is there some analogue to cutting planes for non-convex optimization?**¹ In dimension 1 it is easy to see that one can indeed do a binary search to find an approximate stationary point of a smooth non-convex function on an interval. The first non-trivial case is thus dimension 2, which is the focus of this paper (although we also obtain new results in high dimensions, and in particular our approach does achieve $O(\text{poly}(d) \log(1/\varepsilon))$ *parallel depth*, see below for details). We note that, besides being a fundamental optimization question, hyperparameter search in machine learning also gives an important motivation for studying low-dimensional non-convex optimization.

This problem, of finding an approximate stationary point of a smooth function on a compact domain of \mathbb{R}^2 , was studied in 1993 by Stephen A. Vavasis in (Vavasis, 1993). From an algorithmic perspective, his main observation is that in finite dimensional spaces one can speed up gradient descent by using a *warm start*. Specifically, observe that gradient descent only needs $O(\Delta/\varepsilon^2)$ queries when starting from a Δ -approximate optimal point. Leveraging smoothness (see e.g., Lemma 4 below), observe that the best point on a $\sqrt{\Delta}$ -net of the domain will be Δ -approximate optimal. Thus starting gradient descent from the best point on $\sqrt{\Delta}$ -net one obtains the complexity $O_d\left(\frac{\Delta}{\varepsilon^2} + \frac{1}{\Delta^{d/2}}\right)$ in \mathbb{R}^d . Optimizing over Δ , one obtains a $O_d\left(\left(\frac{1}{\varepsilon}\right)^{\frac{2d}{d+2}}\right)$ complexity. In particular for $d = 2$ this yields a $O(1/\varepsilon)$ query strategy. In addition to this algorithmic advance, Vavasis also proved a lower bound of $\Omega(1/\sqrt{\varepsilon})$ for deterministic algorithms. In this paper we close the gap up to a logarithmic term. Our main contribution is a new strategy loosely inspired by cutting planes, which we call *gradient flow trapping* (GFT), with complexity $O\left(\sqrt{\frac{\log(1/\varepsilon)}{\varepsilon}}\right)$. We also extend Vavasis lower bound to randomized algorithms, by connecting the problem with unpredictable walks in probability theory (Benjamini et al., 1998).

Although we focus on $d = 2$ for the description and analysis of GFT in this paper, one can in fact easily generalize to higher dimensions. Before stating our results there, we first make precise the notion of approximate stationary points, and we also introduce the *parallel query* model.

1.1. Approximate stationary point

We focus on the constraint set $[0, 1]^d$, although this is not necessary and we make this choice mainly for ease of exposition. Let us fix a differentiable function $f : [0, 1]^d \rightarrow \mathbb{R}$ such that $\forall x, y \in [0, 1]^d$, $\|\nabla f(x) - \nabla f(y)\|_2 \leq \|x - y\|_2$. Our goal is to find a point $x \in [0, 1]^d$ such that for any $\varepsilon' > \varepsilon$, there exists a neighborhood $N \subset [0, 1]^d$ of x such that for any $y \in N$,

$$f(x) \leq f(y) + \varepsilon' \cdot \|x - y\|_2.$$

1. We note that a different perspective on this question from the one developed in this paper was investigated in (Hinder, 2018), where the author asks whether one can adapt *actual cutting planes* to non-convex settings. In particular Hinder (2018) shows that one can improve upon gradient descent and obtain a complexity $O(\text{poly}(d)/\varepsilon^{4/3})$ with a cutting plane method, under a higher order smoothness assumption (namely third order instead of first order here).

We say that such an x is an ε -stationary point (its existence is guaranteed by the extreme value theorem). In particular if $x \in (0, 1)^d$ this means that $\|\nabla f(x)\|_2 \leq \varepsilon$. More generally, for $x = (x^1, \dots, x^d) \in [0, 1]^d$ (possibly on the boundary), let us define the *projected gradient* at x , $g(x) = (g_1(x), \dots, g_d(x))$ by:

$$g_i(x) = \begin{cases} \max\left(0, \frac{df}{dx^i}(x)\right) & \text{if } x^i = 0, \\ \frac{df}{dx^i}(x) & \text{if } x^i \in (0, 1), \\ \min\left(0, \frac{df}{dx^i}(x)\right) & \text{if } x^i = 1. \end{cases}$$

It is standard to show (see also [Vavasis \(1993\)](#)) that x is an ε -stationary point of f if and only if $\|g(x)\|_2 \leq \varepsilon$.

1.2. Parallel query model

In the classical black-box model, the algorithm can sequentially query an oracle at points $x \in [0, 1]^d$ and obtain the value² of the function $f(x)$. An extension of this model, first considered in ([Nemirovski, 1994](#)), is as follows: instead of submitting queries one by one sequentially, the algorithm can submit any number of queries in parallel. One can then count the *depth*, defined as the number of rounds of interaction with the oracle, and the *total work*, defined as the total number of queries.

It seems that the parallel complexity of finding stationary points has not been studied before. As far as we know, the only low-depth algorithm (say depth polylogarithmic in $1/\varepsilon$) is the naive grid search: simply query all the points on an ε -net of $[0, 1]^d$ (it is guaranteed that one point in such a net is an ε -stationary point). This strategy has depth 1, and total work $O(1/\varepsilon^d)$. As we explain next, the high-dimensional version of GFT has depth $O(\text{poly}(d) \log(1/\varepsilon))$, and its total work improves at least quadratically upon grid search.

1.3. Complexity bounds for GFT

In this paper we give a complete proof of the following near-optimal result in dimension 2:

Theorem 1 *Let $d = 2$. The gradient flow trapping algorithm (see [Section 4](#)) finds a 4ε -stationary point with less than $10^5 \sqrt{\frac{\log(1/\varepsilon)}{\varepsilon}}$ queries to the value of f .*

It turns out that there is nothing inherently two-dimensional about GFT. At a very high level, one can think of GFT as making hyperplane cuts, just like standard cutting planes methods in convex optimization. While in the convex case those hyperplane cuts are simply obtained by gradients, here we obtain them by querying a $\tilde{O}(\sqrt{\varepsilon})$ -net on a carefully selected small set of hyperplanes. Note also that the meaning of a “cut” is much more delicate than for traditional cutting planes methods (here we use those cuts to “trap” gradient flows). All of these ideas are more easily expressed in dimension 2, but generalizing them to higher dimensions presents no new difficulties (besides heavier notation). In [Section A](#) we prove the following result:

2. Technically we consider here the zeroth order oracle model. It is clear that one can obtain a first order oracle model from it, at the expense of a multiplicative dimension blow-up in the complexity. In the context of this paper an extra factor d is small, and thus we do not dwell on the distinction between zeroth order and first order.

Theorem 2 *The high-dimensional version of GFT finds an ε -stationary point in depth $O(d^2 \log(d/\varepsilon))$ and in total work $d^{O(d)} \cdot \left(\frac{\log(1/\varepsilon)}{\varepsilon}\right)^{\frac{d-1}{2}}$.*

In particular we see that the three-dimensional version of GFT has complexity $O\left(\frac{\log(1/\varepsilon)}{\varepsilon}\right)$. This improves upon the previous state of the art complexity $O(1/\varepsilon^{1.2})$ (Vavasis, 1993). However, on the contrary to the two-dimensional case, we believe that here GFT is suboptimal. As we discuss in Section B.3, in dimension 3 we conjecture the lower bound $\Omega(1/\varepsilon^{0.6})$.

In dimensions $d \geq 4$, the total work given by Theorem 2 is worse than the total work $O\left(\left(\frac{1}{\varepsilon}\right)^{\frac{2d}{d+2}}\right)$ of Vavasis' algorithm. On the other hand, the depth of Vavasis' algorithm is of the same order as its total work, in stark contrast with GFT which maintains a logarithmic depth even in higher dimensions. Among algorithms with polylogarithmic depth, the total work given in Theorem 2 is more than a quadratic improvement (in fixed dimension) over the previous state of the art (namely naive grid search).

1.4. Paper organization

The rest of the paper (besides the appendix and Section 5) is dedicated to motivating, describing and analyzing our gradient flow trapping strategy in dimension 2 (from now on we fix $d = 2$, unless specified otherwise). In Section 2 we make a basic “local to global” observation about gradient flow which forms the basis of our “trapping” strategy. Section 3 is an informal section on how one could potentially use this local to global phenomenon to design an algorithm, and we outline some of the difficulties one has to overcome. In Section 4 we formally describe our new strategy and analyze its complexity. Finally we conclude the paper in Section 5 by introducing several open problems related to higher dimensions. In Appendix A we consider higher dimensions and prove Theorem 2 and in Appendix B we extend Vavasis' $\Omega(1/\sqrt{\varepsilon})$ lower bound to randomized algorithms.

2. A local to global phenomenon for gradient flow

We begin with some definitions. For an axis-aligned rectangle $R = [a, b] \times [c, d]$ we denote its volume and diameter by

$$\text{diam}(R) := \sqrt{(b-a)^2 + (d-c)^2} \text{ and } \text{vol}(R) := (b-a)(d-c).$$

We further define the aspect ratio of R as $\frac{\max(b-a, d-c)}{\min(b-a, d-c)}$. The 4 edges of R are the subsets

$$\{a\} \times [c, d], \{b\} \times [c, d], [a, b] \times \{c\} \text{ and } [a, b] \times \{d\},$$

and the boundary of R , which we denote ∂R is the union of all edges.

If $E \subset [0, 1]^2$ is a segment and $\delta > 0$, we say that $N \subset E$ is a δ -net of E , if for any $x \in E$, there exists some $y \in N$ such that $\|x - y\|_2 \leq \delta$. We will always assume implicitly that if $N \subset E$ is a δ -net, then the endpoints of E are elements of N .

We denote $f_\delta^*(E)$ for the largest value one can obtain by minimizing f on a δ -net of E . Formally,

$$f_\delta^*(E) = \sup_N \inf_{x \in N} f(x),$$

where the supremum is taken over all δ -nets of E . We say that a pair (E, x) of segment/point in $[0, 1]^2$ (where E is *not* a subset of an edge of $[0, 1]^2$) satisfies the property P_c for some $c \geq 0$ if there exists $\delta > 0$ such that

$$f(x) < f_\delta^*(E) - \frac{\delta^2}{8} + c \cdot \text{dist}(x, E),$$

where

$$\text{dist}(x, E) := \inf_{y \in E} \|x - y\|_2.$$

When E is a subset of $\partial[0, 1]^2$ we *always* say that (E, x) satisfies P_c (for any $c \geq 0$ and any $x \in [0, 1]^2$).

For an axis-aligned rectangle R and $x \in R$, we say that (R, x) satisfies P_c if, for all four of the edges E of R , one has that (E, x) satisfies P_c . We refer to x as the *pivot* for R .

Our main observation is as follows:

Lemma 3 *Let R be a rectangle such that (R, x) satisfies P_c for some $x \in R$ and $c \geq 0$. Then R must contain a c -stationary point (in fact the gradient flow emanating from x must visit a c -stationary point before exiting R).*

This lemma will be our basic tool to develop cutting plane-like strategies for non-convex optimization. From “local” information (values on a net of the boundary of R) one deduces a “global” property (existence of approximate stationary point in R).

Proof Let us assume by contradiction that R does not contain a c -stationary point, and consider the unit-speed gradient flow $(x(t))_{t \geq 0}$ constrained to stay in $[0, 1]^2$. That is, $x(t)$ is the piecewise differentiable function defined by $x(0) = x$ and $\frac{d}{dt}x(t) = -\frac{g(x(t))}{\|g(x(t))\|_2}$, where g is the projected gradient defined in the previous section. Since there is no stationary point in R , it must be that the gradient flow exits R . Let us denote $T = \inf\{t \geq 0 : x(t) \notin R\}$, and E an edge of R such that $x(T) \in E$. Remark that E cannot be part of an edge of $[0, 1]^2$. Furthermore, for any $0 \leq t \leq T$, one has

$$f(x(t)) - f(x(0)) = \int_0^t g(x(s)) \cdot \frac{d}{ds}x(s) ds \leq -c \cdot t \leq -c \cdot \|x(t) - x(0)\|_2,$$

where the first inequality uses that R does not contain a c -stationary point. In particular, this implies $f(x(T)) - f(x) \leq -c \cdot \text{dist}(x, E)$, so that

$$\min_{y \in E} f(y) \leq f(x) - c \cdot \text{dist}(x, E).$$

Lemma 4 below shows that for any $\delta > 0$ one has $f_\delta^*(E) \leq \min_{y \in E} f(y) + \frac{\delta^2}{8}$, and thus together with the above display it shows that (E, x) does *not* satisfy P_c , which is a contradiction. ■

Lemma 4 *For any segment $E \subset [0, 1]^2$ and $\delta > 0$ one has:*

$$f_\delta^*(E) \leq \min_{y \in E} f(y) + \frac{\delta^2}{8}.$$

Proof Let $x \in E$ be such that $f(x) = \min_{z \in E} f(z)$. If x is an endpoint of E , then we are done since we require the endpoints of E to be in the δ -nets. Otherwise x is in the relative interior of E , and thus one has $\nabla f(x) \cdot (y - x) = 0$ for any $y \in E$. In particular by smoothness one has:

$$\begin{aligned} f(y) &= f(x) + \int_0^1 \nabla f(x + t(y - x)) \cdot (y - x) dt \\ &\leq f(x) + \int_0^1 t \cdot \|y - x\|_2^2 dt = f(x) + \frac{1}{2} \|y - x\|_2^2. \end{aligned}$$

Moreover for any δ -net of E there exists y such that $\|y - x\|_2 \leq \frac{\delta}{2}$, and thus $f(y) \leq f(x) + \delta^2/8$, which concludes the proof. \blacksquare

Our algorithmic approach to finding stationary points will be to somehow shrink the domain of consideration over time. At first it can be slightly unclear how the newly created boundaries interact with the definition of stationary points. To dispell any mystery, it might be useful to keep in mind the following observation, which states that if (R, x) satisfies P_c , then x cannot be on a boundary of R which was not part of the original boundary of $[0, 1]^2$.

Lemma 5 *Let R be a rectangle such that (R, x) satisfies P_c for some $x \in R$ and $c \geq 0$. Then $x \notin \partial R \setminus \partial[0, 1]^2$.*

Proof Let E be an edge of R which is not a subset of $\partial[0, 1]^2$. Then by definition of P_c , and by invoking Lemma 4, one has:

$$f(x) < f_\delta^*(E) - \frac{\delta^2}{8} + c \cdot \text{dist}(x, E) \leq \min_{y \in E} f(y) + c \cdot \text{dist}(x, E).$$

In particular if $x \in E$ then $\text{dist}(x, E) = 0$, and thus $f(x) < \min_{y \in E} f(y)$ which is a contradiction. \blacksquare

3. From Lemma 3 to an algorithm

Lemma 3 naturally leads to the following algorithmic idea (for sake of simplicity in this discussion we replace squares by circles): given some current candidate point x in some well-conditioned domain (e.g., such that the domain contains and is contained in balls centered at x and of comparable sizes), query a $\sqrt{\varepsilon}$ -net on the circle $C = \{y : \|y - x\|_2 = 1\}$, and denote y for the best point found on this net. If one finds a significant enough improvement, say $f(y) < f(x) - \frac{3}{4}\varepsilon$, then this is great news, as it means that one obtained a **per query** improvement of $\Theta(\varepsilon^{-3/2})$ (to be compared with gradient descent which only yields an improvement of $\Theta(\varepsilon^{-2})$). On the other hand if no such improvement is found, then the gradient flow from x must visit an ε -stationary point inside C .³ In other words one can now hope to restrict the domain of consideration to a region inside C , which is a constant fraction smaller than the original domain. Figure 1 illustrates the two possibilities.

3. In “essence” (C, x) satisfies P_ε , this is only slightly informal since we defined P_ε for rectangles and C is a circle. In particular we chose the improvement $\frac{3}{4}\varepsilon$ instead of the larger $\frac{7}{8}\varepsilon$ (which is enough to obtain P_c) to account for an extra term due to polygonal approximation of the circle. We encourage the reader to ignore this irrelevant technicality.

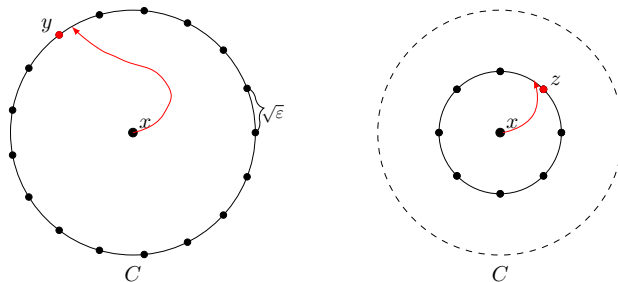


Figure 1: The red curve illustrates the gradient flow emanating from x . On the left, the flow does not visit an ε -stationary point and y has a significantly smaller function value than x . Otherwise, as in the right, we shrink the domain.

Optimistically this strategy would give a $\tilde{O}(B/\varepsilon^{3/2})$ rate for B -bounded smooth functions (since at any given scale one could make at most $O(B/\varepsilon^{3/2})$ improvement steps). In particular together with the warm start this would tentatively yield a $\tilde{O}(1/\varepsilon^{3/4})$ rate, thus already improving the state-of-the-art $O(1/\varepsilon)$ by Vavasis.

There is however a difficulty in the induction part of the argument. Indeed, what we know after a shrinking step is that the current point x satisfies $f(x) \leq f(y) + \varepsilon$ for any $y \in C$. Now we would like to query a net on $\{y : \|y - x\|_2 = 1/2\}$. Say that after such querying we find that we can't shrink, namely we found some point z with $f(z) < f(x) - \frac{\varepsilon}{2} + \frac{\delta^2}{8}$, and in particular $f(z) < f(y) + \frac{1}{2}\varepsilon + \frac{\delta^2}{8}$ for any $y \in C$. Could the gradient flow from z escape the original circle C without visiting an ε -stationary point? Unfortunately the answer is yes. Indeed (because of the discretization error $\delta^2/8$) one cannot rule out that there would be a point $y \in C$ with $f(y) < f(z) - \frac{\varepsilon}{2}$, and since C is only at distance $1/2$ from z , such a point could be attained from z with a gradient flow without ε -stationary points. Of course one could say that instead of satisfying P_ε we now only satisfy $P_{\varepsilon+\delta^2/4}$, and try to control the increase of the approximation guarantee, but such an approach would not improve upon the $1/\varepsilon^2$ of gradient descent (simply because we could additively worsen the approximation guarantee too many times).

The core part of the above argument will remain in our full algorithm (querying a $\sqrt{\varepsilon}$ -net to shrink the domain). However it is made more difficult by the discretization error as we just saw. We also note that this discretization issue does not appear in discrete spaces, which is one reason why discrete spaces are much easier than continuous spaces for local optimization problems.

Technically we observe that the whole issue of discretization comes from the fact that when we update the center, we move closer to the boundary, which we “pay” in the term $\text{dist}(x, E)$ in P_ε , and we cannot “afford” it because of the discretization error term that we suffer when we update. Thus this issue would disappear if in our induction hypothesis we had P_0 for the boundary. Our strategy will work in two steps: first we give a querying strategy for a domain with P_0 that ensures that one can **always** shrink with P_ε guaranteed for the boundary, and secondly we give a method to essentially turn a P_ε boundary into P_0 .

4. Gradient flow trapping

We say that a pair (R, x) is a *domain* if R is an axis-aligned rectangle with aspect ratio bounded by 3, and $x \in R$. The gradient flow trapping (GFT) algorithm is decomposed into two subroutines:

1. The first algorithm, which we call the *parallel trap*, takes as input a domain (R, x) satisfying P_0 . It returns a domain (\tilde{R}, \tilde{x}) satisfying P_ε and such that $\text{vol}(\tilde{R}) \leq 0.95 \text{vol}(R)$. The cost of this step is at most $2\sqrt{\frac{\text{diam}(R)}{\varepsilon}}$ queries.
2. The second algorithm, which we call *edge fixing*, takes as input a domain (R, x) satisfying $P_{\varepsilon'}$ (for some $\varepsilon' \in [\varepsilon, 2\varepsilon]$) and such that for $k \in \{0, 1, 2, 3\}$ edges E of R one also has P_0 for (E, x) . It returns a domain (\tilde{R}, \tilde{x}) such that either (i) it satisfies $P_{\varepsilon'}$ and for $k+1$ edges it also satisfies P_0 , or (ii) it satisfies $P_{\left(1 + \frac{1}{500 \log(1/\varepsilon)}\right)\varepsilon'}$ and furthermore $\text{vol}(\tilde{R}) \leq 0.95 \text{vol}(R)$.

The cost of this step is at most $90\sqrt{\frac{\text{diam}(R) \log(1/\varepsilon)}{\varepsilon}}$ queries.

Equipped with these subroutines, GFT proceeds as follows. Initialize $(R_0, x_0) = ([0, 1]^2, (0.5, 0.5))$, $\varepsilon_0 = \varepsilon$, and $k_0 = 4$. For $t \geq 0$:

- if $k_t = 4$, call *parallel trap* on (R_t, x_t) , and update $k_{t+1} = 0$, $(R_{t+1}, x_{t+1}) = (\tilde{R}_t, \tilde{x}_t)$, and $\varepsilon_{t+1} = \varepsilon$.
- Otherwise call *edge fixing*, and update $(R_{t+1}, x_{t+1}) = (\tilde{R}_t, \tilde{x}_t)$. If $R_{t+1} = R_t$ then set $k_{t+1} = k_t + 1$ and $\varepsilon_{t+1} = \varepsilon_t$, and otherwise set $k_{t+1} = 0$ and $\varepsilon_{t+1} = \left(1 + \frac{1}{500 \log(1/\varepsilon)}\right) \varepsilon_t$.

We terminate once the diameter of R_t is smaller than 2ε .

Next we give the complexity analysis of GFT assuming the claimed properties of the subroutines *parallel trap* and *edge fixing* in 1. and 2. above. We then proceed to describe in details the subroutines, and prove that they satisfy the claimed properties.

4.1. Complexity analysis of GFT

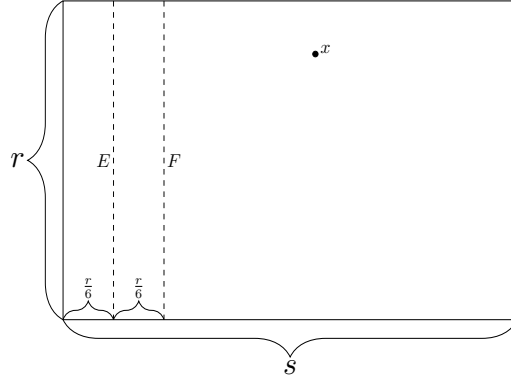
The following three lemmas give a proof of Theorem 1.

Lemma 6 *GFT stops after at most $200 \log(1/\varepsilon)$ steps.*

Proof First note that at least one out of five steps of GFT reduces the volume of the domain by 0.95 (since one can do at most four steps in a row of edge fixing without volume decrease). Thus on average the volume decrease per step is at least 0.99, i.e., $\text{vol}(R_T) \leq 0.99^T$. In particular since R_T has aspect ratio smaller than 3, it is easy to verify $\text{diam}(R_T) \leq 2\sqrt{\text{vol}(R_T)} \leq 2 \times 0.99^{T/2}$. Thus for any $T \geq \log_{100/99}(1/\varepsilon^2)$, one must have $\text{diam}(R_T) \leq 2\varepsilon$. Thus we see that GFT performs at most $\log_{100/99}(1/\varepsilon^2) \leq 200 \log(1/\varepsilon)$ steps. ■

Lemma 7 *When GFT stops, its pivot is a 4ε -stationary point.*

Proof First note that $\varepsilon_T \leq \left(1 + \frac{1}{500 \log(1/\varepsilon)}\right)^T \varepsilon$, thus after $T \leq 200 \log(1/\varepsilon)$ steps we know that (R_T, x_T) satisfies at least $P_{2\varepsilon}$. In particular by Lemma 3, R_T must contain a 2ε -stationary point, and since the diameter is less than 2ε , it must be (by smoothness) that x_T is a 4ε -stationary point. ■


 Figure 2: The *parallel trap*

Lemma 8 *GFT makes at most $10^5 \sqrt{\frac{\log(1/\varepsilon)}{\varepsilon}}$ queries before it stops.*

Proof As we saw in the proof of Lemma 6, one has $\text{diam}(R_t) \leq 2 \times 0.99^{t/2}$. Furthermore the t^{th} step requires at most $90 \sqrt{\frac{\text{diam}(R_t) \log(1/\varepsilon)}{\varepsilon}}$ queries. Thus the total number of queries is bounded by:

$$90 \sqrt{\frac{2 \log(1/\varepsilon)}{\varepsilon}} \sum_{t=0}^{\infty} 0.99^{t/4} \leq 10^5 \sqrt{\frac{\log(1/\varepsilon)}{\varepsilon}}.$$

■

4.2. A parallel trap

Let (R, x) be a domain. We define two segments E and F in R as follows. Assume that R is a translation of $[0, s] \times [0, r]$. For sake of notation assume that in fact $R = [0, s] \times [0, r]$ with $s \in [r, 3r]$ and $x^1 \geq r/2$, where $x = (x^1, x^2)$ (in practice one always ensures this situation with a simple change of variables). Now we define $E = \{r/6\} \times [0, r]$ and $F = \{r/3\} \times [0, r]$ (See Figure 2). The parallel trap algorithm queries a $\sqrt{r\varepsilon}$ -net on both E and F (which cost at most $2 \frac{r}{\sqrt{r\varepsilon}} = 2\sqrt{\frac{r}{\varepsilon}}$). Denote \bar{x} to be the best point (in terms of f value) found on the union of those nets. That is, denoting $N \subset F \cup E$ for the queried $\sqrt{r\varepsilon}$ -net, then

$$\bar{x} = \arg \min_{y \in N} f(y).$$

One has the following possibilities (see Figure 3 for an illustration):

- If $f(x) \leq f(\bar{x})$ then we set $\tilde{x} = x$ and $\tilde{R} = [r/3, s] \times [0, r]$.
- Otherwise we set $\tilde{x} = \bar{x}$. If $\bar{x} \in F$ we set $\tilde{R} = [r/6, s] \times [0, r]$, and if $\bar{x} \in E$ we set $\tilde{R} = [0, r/3] \times [0, r]$.

The above construction is justified by the following lemma (a trivial consequence of the definitions), and it proves in particular the properties of *parallel trap* described in 1. at the beginning of Section 4.

Lemma 9 *The rectangle \tilde{R} has aspect ratio smaller than 3, and it satisfies $\text{vol}(\tilde{R}) \leq 0.95 \text{vol}(R)$. Furthermore if (R, x) satisfies P_0 , then (\tilde{R}, \tilde{x}) satisfies P_ε .*

Proof The first sentence is trivial to verify. For the second sentence, first note that for any edge E of R one has P_0 for (E, \tilde{x}) since by assumption one has P_0 for (E, x) and furthermore $f(\tilde{x}) \leq f(x)$. Next observe that \tilde{R} has at most one new edge \tilde{E} with respect to R , and this edge is at distance at least $r/6$ from \tilde{x} , thus in particular one has $\varepsilon \cdot \text{dist}(\tilde{x}, \tilde{E}) - \delta^2/8 > 0$ for $\delta = \sqrt{r\varepsilon}$. Furthermore by definition $f(\tilde{x}) \leq f_\delta^*(\tilde{E})$, and thus $f(\tilde{x}) < f_\delta^*(\tilde{E}) - \frac{\delta^2}{8} + \varepsilon \cdot \text{dist}(\tilde{x}, \tilde{E})$, or in other words (\tilde{E}, \tilde{x}) satisfies P_ε . \blacksquare

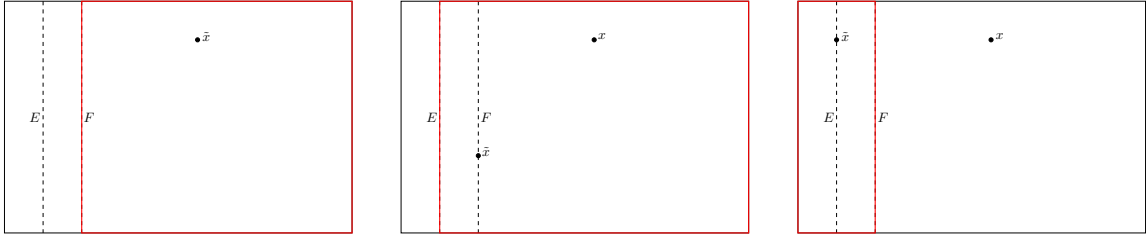


Figure 3: The three possible cases for (\tilde{R}, \tilde{x}) . \tilde{R} is marked in red.

4.3. Edge fixing

Let (R, x) be a domain satisfying $P_{\varepsilon'}$ for some $\varepsilon' \in [\varepsilon, 2\varepsilon]$, and with some edges possibly also satisfying P_0 . Denote \mathcal{E} for the closest edge to x that does not satisfy P_0 , and let $r = \text{dist}(x, \mathcal{E})$. We will consider three⁴ candidate smaller rectangles, R_1 , R_2 and R_3 , as well as three candidate pivots (in addition to x) $x_1 \in \partial R_1$, $x_2 \in \partial R_2$ and $x_3 \in \partial R_3$. The rectangles are defined by $R_i = R \cap \{y : \|x_{i-1} - y\|_\infty \leq \frac{r}{3}\}$, where we set $x_0 = x$. The possible output (\tilde{R}, \tilde{x}) of *edge fixing* will be either (R_i, x_{i-1}) for some $i \in \{1, 2, 3\}$, or (R, x_3) (see Figure 4 for a demonstration of how to construct the rectangles).

To guarantee the properties described in 2. at the beginning of Section 4 we will prove the following: if the output is (R_i, x_{i-1}) for some i then all edges will satisfy $P_{\left(1 + \frac{1}{500 \log(1/\varepsilon)}\right)\varepsilon'}$ (Lemma 12 below) and the domain has shrunk (Lemma 10 below), and if the output is (R, x_3) then one more edge satisfies P_0 compared to (R, x) while all edges still satisfy at least $P_{\varepsilon'}$ (Lemma 11 below).

Lemma 10 *For any $i \in \{1, 2, 3\}$ one has $\text{vol}(R_i) \leq \frac{2}{3} \text{vol}(R)$. Furthermore if the aspect ratio of R is smaller than 3, then so is the aspect ratio of R_i .*

Proof Let us denote $\ell_1(R)$ for the length of R in the axis of \mathcal{E} (the edge whose distance to x defines r), and $\ell_2(R)$ for the length in the orthogonal direction (and similarly define $\ell_1(R_i)$ and $\ell_2(R_i)$). Since $R_i \subset R$ one has $\ell_1(R_i) \leq \ell_1(R)$. Furthermore $\ell_2(R_i) \leq \frac{2}{3}r$ and $\ell_2(R) \geq r$, so that $\ell_2(R_i) \leq \frac{2}{3}\ell_2(R)$. This implies that $\text{vol}(R_i) \leq \frac{2}{3} \text{vol}(R)$.

For the second statement observe that $\ell_1(R) \geq \frac{\ell_2(R)}{3} \geq \frac{r}{3}$ (the first inequality is by assumption on the aspect ratio of R , the second inequality is by definition of r). Given this estimate, the

4. We need three candidates to ensure that the domain will shrink.

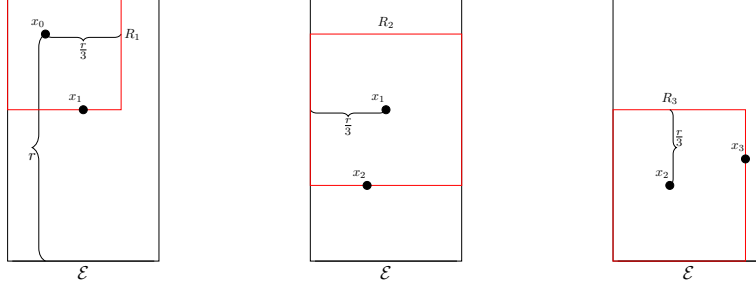


Figure 4: *Edge fixing*: the rectangles R_1 , R_2 and R_3 are marked in red, from left to right.

construction of R_i implies that $\frac{1}{3}r \leq \ell_2(R_i), \ell_1(R_i) \leq \frac{2}{3}r$, which concludes the fact that R_i has aspect ratio smaller than 2. \blacksquare

Queries and choice of output. The edge fixing algorithm queries a $\sqrt{\frac{\varepsilon' r}{500 \log(1/\varepsilon)}}$ -net on ∂R_i for all $i \in \{1, 2, 3\}$ (thus a total of $4\sqrt{\frac{500r \log(1/\varepsilon)}{\varepsilon'}} \leq 90\sqrt{\frac{r \log(1/\varepsilon)}{\varepsilon}}$ queries), and we define x_i to be the best point found on each respective net.

If for all $i \in \{1, 2, 3\}$ one has

$$f(x_i) \leq f(x_{i-1}) - \frac{\varepsilon' r}{3}, \quad (1)$$

then we set $(\tilde{R}, \tilde{x}) = (R, x_3)$. Otherwise denote $i^* \in \{1, 2, 3\}$ for the smallest number which violates (1), and set $(\tilde{R}, \tilde{x}) = (R_{i^*}, x_{i^*-1})$.

Lemma 11 *If $(\tilde{R}, \tilde{x}) = (R, x_3)$ then (\mathcal{E}, x_3) satisfies P_0 . Furthermore for any edge E of R , if (E, x) satisfies P_0 (respectively $P_{\varepsilon'}$) then so does (E, x_3) .*

Proof Since $(\tilde{R}, \tilde{x}) = (R, x_3)$ it means that $f(x_3) \leq f(x_0) - \varepsilon' r$. In particular since (\mathcal{E}, x_0) satisfies $P_{\varepsilon'}$ one has $f(x_0) < f_\delta^*(\mathcal{E}) - \frac{\delta^2}{8} + \varepsilon' r$, and thus now one has $f(x_3) < f_\delta^*(\mathcal{E}) - \frac{\delta^2}{8}$ which means that (\mathcal{E}, x_3) satisfies P_0 .

Let us now turn to some other edge E of R . Certainly if (E, x_0) satisfies P_0 then so does (E, x_3) since $f(x_3) \leq f(x_0)$. But, in fact, even $P_{\varepsilon'}$ is preserved since by the triangle inequality (and $\|x_3 - x_0\|_2 \leq r$) one has

$$f(x_3) - \varepsilon' \cdot \text{dist}(x_3, E) \leq f(x_3) + \varepsilon' r - \varepsilon' \cdot \text{dist}(x_0, E) \leq f(x_0) - \varepsilon' \cdot \text{dist}(x_0, E).$$

\blacksquare

Lemma 12 *If $(\tilde{R}, \tilde{x}) = (R_i, x_{i-1})$ for some $i \in \{1, 2, 3\}$, then (\tilde{R}, \tilde{x}) satisfy $P_{\left(1 + \frac{1}{500 \log(1/\varepsilon)}\right)\varepsilon'}$.*

Proof By construction, if $(\tilde{R}, \tilde{x}) = (R_i, x_{i-1})$, then for any edge E of R_i one has $f(x_{i-1}) < f_\delta^*(E) + \frac{\varepsilon' r}{3}$. Furthermore one has $\frac{\varepsilon' r}{3} = -\frac{\varepsilon' r}{8 \times 500 \log(1/\varepsilon)} + \left(1 + \frac{3}{8 \times 500 \log(1/\varepsilon)}\right) \frac{\varepsilon' r}{3}$, and thus one has

$P_{\left(1 + \frac{3}{8 \times 500 \log(1/\varepsilon)}\right)^{\varepsilon'}}$ for (E, x_{i-1}) whenever $\text{dist}(x_{i-1}, E) = \frac{r}{3}$. Indeed, since $\delta = \sqrt{\frac{\varepsilon' r}{500 \log(1/\varepsilon)}}$,

$$f(x_{i-1}) < f_{\delta}^*(E) - \frac{\delta^2}{8} + \left(1 + \frac{3}{8 \times 500 \log(1/\varepsilon)}\right) \varepsilon' \cdot \text{dist}(x_{i-1}, E).$$

If $\text{dist}(x_{i-1}, E) < \frac{r}{3}$ then by the triangle inequality, $\text{dist}(x_0, E) < r$, and moreover E is also an edge with respect to R . Thus from the definition of r , (E, x_0) satisfies P_0 . Also by our choice of x_{i-1} , we know that $f(x_{i-1}) \leq f(x_0)$. Hence (E, x_{i-1}) satisfies P_0 as well. \blacksquare

5. Discussion

In this paper we introduced a near-optimal algorithm for finding ε -stationary points in dimension 2. Finding a near-optimal algorithm in dimensions $d \geq 3$ remains open. Specific challenges include:

1. Finding a strategy in dimension 3 which improves upon GFT's $\tilde{O}(1/\varepsilon)$ complexity.
2. The heuristic extension of the lower bound in Appendix B.3 suggests $\Omega\left(\frac{1}{\varepsilon^{\frac{1}{d+2}}}\right)$ as a complexity lower bound for any dimension d (note in particular that the exponent tends to 1 as d tends to infinity). On the other hand, Carmon et al. (2019) proved that for $d = \Omega(1/\varepsilon^2)$, one has the complexity lower bound $\Omega(1/\varepsilon^2)$. How do we reconcile these two results? Specifically we raise the following question: Is there an algorithm with complexity C_d/ε for some constant C_d which depends only on d ? (Note that C_d as small as $O(\sqrt{d})$ would remain consistent with Carmon et al. (2019).) Alternatively we might ask whether the Carmon et al. (2019) lower bound holds for much smaller dimensions, e.g. when $d = \Theta(\log(1/\varepsilon))$, are we in the $1/\varepsilon$ regime as suggested by the heuristic, or are we already in the high-dimensional $1/\varepsilon^2$ of Carmon et al. (2019)?
3. Especially intriguing is the limit of low-depth algorithms, say as defined by having depth smaller than $\text{poly}(d \log(1/\varepsilon))$. Currently this class of algorithms suffers from the curse of dimensionality, as GFT's total work degrades significantly when the dimension increases (recall from Theorem 2 that it is $\tilde{O}\left(\frac{1}{\varepsilon^{\frac{d-1}{2}}}\right)$). Is this necessary? A much weaker question is to simply show a separation between low-depth and high-depth algorithms. Namely can one show a lower bound $\Omega(1/\varepsilon^c)$ with $c > 2$ for low-depth algorithms? We note that lower bounds on depth have been investigated in the convex setting, see (Nemirovski, 1994), (Bubeck et al., 2019).
4. A technically challenging problem is to adapt the construction in [Section 3, Vavasis (1993)] to non-monotone paths in higher dimensions. In particular, to formalize the heuristic argument from Appendix B.3, such construction should presumably avoid creating saddle points.

Many more questions remain open on how to exploit the low-dimensional geometry of smooth gradient fields, and the above four questions are only a subset of the fundamental questions that we would like to answer. Other interesting questions include closing the logarithmic gap in dimension 2, or understanding better the role of randomness for this problem (note that GFT is deterministic, but other type of strategies include randomness, such as Hinder's non-convex cutting plane (Hinder, 2018)).

Acknowledgments

We thank Ronen Eldan, Yin Tat Lee, Yuanzhi Li and Mark Selke for many helpful discussions on this problem.

References

- Scott Aaronson. Lower bounds for local search by quantum arguments. *SIAM Journal on Computing*, 35(4):804–824, 2006.
- Itai Benjamini, Robin Pemantle, and Yuval Peres. Unpredictable paths and percolation. *The Annals of Probability*, 26(3):1198–1211, 1998.
- Sebastien Bubeck, Qijia Jiang, Yin-Tat Lee, Yuanzhi Li, and Aaron Sidford. Complexity of highly parallel non-smooth convex optimization. In *Advances in Neural Information Processing Systems 32*, pages 13900–13909. 2019.
- Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. *Mathematical Programming*, 2019.
- Olle Häggström and Elchanan Mossel. Nearest-neighbor walks with low predictability profile and percolation in $2 + \varepsilon$ dimensions. *The Annals of Probability*, 26(3):1212–1231, 1998.
- Oliver Hinder. Cutting plane methods can be extended into nonconvex optimization. In *Conference On Learning Theory, COLT 2018*, pages 1451–1454, 2018.
- A. Nemirovski and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley Interscience, 1983.
- Arkadi Nemirovski. On parallel complexity of nonsmooth convex optimization. *Journal of Complexity*, 10(4):451 – 463, 1994.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*. Kluwer Academic Publishers, 2004.
- Xiaoming Sun and Andrew Chi-Chih Yao. On the quantum query complexity of local search in two and three dimensions. *Algorithmica*, 55(3):576–600, 2009.
- Stephen A. Vavasis. Black-box complexity of local minimization. *SIAM Journal on Optimization*, 3(1):60–80, 1993.
- Shengyu Zhang. New upper and lower bounds for randomized and quantum local search. In *STOC’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 634–643. ACM, New York, 2006.

Appendix A. Generalization to higher dimensions

As explained in the introduction, there is no reason to restrict GFT to $[0, 1]^2$ and, in fact, the algorithm may be readily adapted to higher-dimensional spaces, such as $[0, 1]^d$, for some $d > 2$. We now detail the necessary changes and derive the complexity announced in Theorem 2.

First, if F is an affine hyperplane, and $x \in [0, 1]^d$, we define P_c for (F, x) in the obvious way (i.e., same definition except that we consider a δ -net of F). Similarly for (R, x) , when R is an axis-aligned hyperrectangle.

Gradient flow trapping in higher dimensions replaces every line by a hyperplane, and every rectangle by a hyperrectangle. In particular at each step GFT maintains a domain (R, x) , where R is a hyperrectangle with aspect ratio bounded by 3, and $x \in R$. The two subroutines are adapted as follows:

1. *Parallel trap* works exactly in the same way, except that the two lines E and F are replaced by two corresponding affine hyperplanes. In particular the query cost of this step is now $O\left(\left(\frac{\text{diam}(R)}{\varepsilon}\right)^{\frac{d-1}{2}}\right)$, and the volume shrinks by at least 0.95.
2. In *edge fixing*, we now have three hyperrectangles R_i , and we need to query nets on their $2d$ faces. Thus the total cost of this step is $O\left(d\left(\frac{\text{diam}(R)\log(1/\varepsilon)}{\varepsilon}\right)^{\frac{d-1}{2}}\right)$. Moreover, suppose that domain does not shrink at the end of this step and the output is a domain (R, \tilde{x}) for some other $\tilde{x} \in R$. In this case we know that R has some face F , such that (F, x) did not satisfy P_0 , but (F, \tilde{x}) does satisfy P_0 . It follows that we can run *edge fixing*, at most $2d$ times before the domain shrinks.

We can now analyze the complexity of the high-dimensional version of GFT:

Proof [Of Theorem 2] First observe that, if R is a hyperrectangle in $[0, 1]^d$ with aspect ratio bounded by 3, then we have the following inequality,

$$\text{diam}(R) \leq 3\sqrt{d} \cdot \text{vol}(R)^{\frac{1}{d}}.$$

By repeating the same calculations done in Lemma 6 and the observation about *parallel trap* and *edge fixing* made above, we see that the domain shrinks at least once in every $2d + 1$ steps, so that at step T ,

$$\text{vol}(R_T) \leq 0.95^{\frac{T}{2d+1}},$$

and

$$\text{diam}(R_T) \leq 3\sqrt{d} \cdot 0.95^{\frac{T}{(2d+1)d}}.$$

Since the algorithm stops when $\text{diam}(R_T) \leq 2\varepsilon$, we get

$$T = O\left(d^2 \log\left(\frac{d}{\varepsilon}\right)\right).$$

The total work done by the algorithm is evident now by considering the number of queries at each step. ■

Appendix B. Lower bound for randomized algorithms

In this section, we show that any randomized algorithm must make at least $\tilde{\Omega}\left(\frac{1}{\sqrt{\varepsilon}}\right)$ queries in order to find an ε -stationary point. This extends the lower bound in (Vavasis, 1993), which applied only to deterministic algorithms. In particular, it shows that, up to logarithmic factors, adding randomness cannot improve the algorithm described in the previous section.

For an algorithm \mathcal{A} , a function $f : [0, 1]^2 \rightarrow \mathbb{R}$ and $\varepsilon > 0$ we denote by $\mathcal{Q}(\mathcal{A}, f, \varepsilon)$ the number of queries made by \mathcal{A} , in order to find an ε -stationary point of f . Our goal is to bound from below

$$\mathcal{Q}_{\text{rand}}(\varepsilon) := \inf_{\mathcal{A} \text{ random}} \sup_f \mathbb{E}_{\mathcal{A}} [\mathcal{Q}(\mathcal{A}, f, \varepsilon)],$$

where the infimum is taken over all random algorithms and the supremum is taken over all smooth functions, f . The expectation is with respect to the randomness of \mathcal{A} . By Yao's minimax principle we have the equality

$$\mathcal{Q}_{\text{rand}}(\varepsilon) = \sup_{\mathcal{D}} \inf_{\mathcal{A} \text{ deterministic}} \mathbb{E}_{f \sim \mathcal{D}} [\mathcal{Q}(\mathcal{A}, f, \varepsilon)].$$

Here, \mathcal{A} is a deterministic algorithm and \mathcal{D} is a distribution over smooth functions. The rest of this section is devoted to proving the following theorem:

Theorem 13 *Let $h : \mathbb{N} \rightarrow \mathbb{R}$ be a decreasing function such that*

$$\sum_{k=1}^{\infty} \frac{h(k)}{k} < \infty,$$

and set

$$S_h(n) := \sum_{k=1}^n \frac{1}{k \cdot h(k)}. \quad (2)$$

Then,

$$\mathcal{Q}_{\text{rand}}(\varepsilon) = \Omega\left(\frac{1}{\sqrt{\varepsilon} \cdot S_h\left(\left\lceil \frac{1}{\sqrt{\varepsilon}} \right\rceil\right)}\right).$$

Remark that one may take $h(k) := \frac{1}{\log(k)^2 + 1}$ in the theorem. In this case $S_h(k) = O(\log^3(k))$, and $\mathcal{Q}_{\text{rand}}(\varepsilon) = \Omega\left(\frac{1}{\log^3(1/\varepsilon)\sqrt{\varepsilon}}\right)$, which is the announced lower bound.

One of the main tools utilized in our proof is the construction introduced in (Vavasis, 1993). We now present the relevant details.

B.1. A reduction to monotone path functions

Let $G_n = (V_n, E_n)$ stand for the $n + 1 \times n + 1$ grid graph. That is,

$$V_n = \{0, \dots, n\} \times \{0, \dots, n\} \text{ and } E_n = \{(v, u) \in V_n \times V_n : \|v - u\|_1 = 1\}.$$

We say that a sequence of vertices, (v_0, \dots, v_n) is a *monotone path* in G_n if $v_0 = (0, 0)$ and for every $0 < i \leq n$, $v_i - v_{i-1}$ either equals $(0, 1)$ or $(1, 0)$. In other words, the path starts at the origin and continues each step by either going right or up. If (v_0, \dots, v_n) is a monotone path, we associate to it a *monotone path function* $P : V_n \rightarrow \mathbb{R}$ by

$$P(v) = \begin{cases} -\|v\|_1 & \text{if } v \in \{v_0, \dots, v_n\} \\ \|v\|_1 & \text{otherwise} \end{cases}.$$

By a slight abuse of notation, we will sometimes refer to the path function and the path itself as the same entity. If $i = 0, \dots, n$ we write P_i for $P^{-1}(-i)$ and $P[i]$ for the prefix (P_0, P_1, \dots, P_i) . If $v \in V_n$ is such that $P(v) > 0$, we say that v does not lie on the path.

We denote the set of all monotone path functions on G_n by F_n . It is clear that if $P \in F_n$ then P_n is the only local minimum of P and hence the global minimum.

Informally, the main construction in (Vavasis, 1993) shows that for every $P \in F_n$ there is a corresponding smooth function $\hat{P} : [0, 1]^2 \rightarrow \mathbb{R}$, which 'traces' the path in P and preserves its structure. In particular, finding an ε -stationary point of \hat{P} is not easier than finding the minimum of P .

To formally state the result we fix $\varepsilon > 0$ and assume for simplicity that $\frac{1}{\sqrt{\varepsilon}}$ is an integer. We henceforth denote $n(\varepsilon) := \frac{1}{\sqrt{\varepsilon}}$ and identify $V_{n(\varepsilon)}$ with $[0, 1]^2$ in the following way: if $(i, j) = v \in V_{n(\varepsilon)}$ we write $\text{square}(v)$ for the square:

$$\text{square}(v) = \left[\frac{i}{n(\varepsilon) + 1}, \frac{i + 1}{n(\varepsilon) + 1} \right] \times \left[\frac{j}{n(\varepsilon) + 1}, \frac{j + 1}{n(\varepsilon) + 1} \right].$$

If $\varphi : [0, 1]^2 \rightarrow \mathbb{R}$, then $\text{supp}(\varphi)$ denotes the closure of the set $\{x \in [0, 1]^2 : \varphi(x) \neq 0\}$.

Lemma 14 (Section 3, (Vavasis, 1993)) *Let $P \in F_{n(\varepsilon)}$. Then there exists a function $\hat{P} : [0, 1]^2 \rightarrow \mathbb{R}$ with the following properties:*

1. \hat{P} is smooth.
2. $\hat{P} = f_P + \ell$, where ℓ is a linear function, which does not depend on P , and

$$\text{supp}(f_P) \subset \bigcup_{i=0}^n \text{square}(P_i).$$

3. If $x \in [0, 1]^2$ is an ε -stationary point of \hat{P} then $x \in \text{square}(P_n)$.
4. if $P' \in F_{n(\varepsilon)}$ is another function and for some $i = 0, \dots, n$, $(P'_{i-1}, P'_i, P'_{i+1}) = (P_{i-1}, P_i, P_{i+1})$.
Then

$$\hat{P}'|_{\text{square}(P_i)} = \hat{P}|_{\text{square}(P_i)}$$

We now make precise of the fact that finding the minimum of P is as hard as finding an ε -stationary point of \hat{P} . For this we define $\mathcal{G}(\mathcal{A}, P)$, the number of queries made by algorithm \mathcal{A} , in order to find the minimal value of the function P .

Lemma 15 *For any algorithm \mathcal{A} , which finds an ε -stationary point of smooth functions on $[0, 1]^2$, there exists an algorithm $\tilde{\mathcal{A}}$ such that*

$$\mathcal{Q}(\mathcal{A}, \hat{P}, \varepsilon) \geq \frac{1}{5} \mathcal{G}(\tilde{\mathcal{A}}, P),$$

for any $P \in F_{n(\varepsilon)}$.

Proof Given an algorithm \mathcal{A} we explain how to construct $\tilde{\mathcal{A}}$. Fix $P \in F_{n(\varepsilon)}$. If \mathcal{A} queries a point $x \in \text{square}(v) \subset [0, 1]^2$. Then $\tilde{\mathcal{A}}$ queries v and all of its neighbors. When \mathcal{A} terminates it has found an ε -stationary point. By Lemma 14, this point must lie in $\text{square}(P_n)$. By querying P_n and its neighbors, $\tilde{\mathcal{A}}$ will determine that P_n is a local minimum and hence the minimum of P .

Since each vertex has at most 4 neighbors, it will now suffice to show that $\tilde{\mathcal{A}}$ can remain consistent with \mathcal{A} . We thus need to show that after querying the neighbors of v , $\tilde{\mathcal{A}}$ may deduce the value of $\hat{P}(x)$.

As we are only interested in the number of queries made by $\tilde{\mathcal{A}}$, it is fine to assume that $\tilde{\mathcal{A}}$ has access to the construction used in Lemma 14. Now, suppose that $P(v) > 0$ and v does not lie on the path. In this case, by Lemma 14, $\hat{P}(x) = \ell(x)$, which does not depend on P itself and $\ell(x)$ is known. Otherwise $v = P_i$ for some $i = 0, \dots, n$. So, after querying the neighbors of v , $\tilde{\mathcal{A}}$ also knows P_{i-1} and P_{i+1} . The lemma then tells us that $\hat{P}|_{\text{square}(v)}$ is uniquely determined and, in particular, the value of $\hat{P}(x)$ is known. \blacksquare

B.2. A lower bound for monotone path functions

Denote $\mathcal{D}_p(n)$ to be the set of all distributions supported on F_n . By Lemma 15,

$$\mathcal{Q}_{\text{rand}}(\varepsilon) \geq \sup_{\mathcal{D} \in \mathcal{D}_p(n(\varepsilon))} \inf_{\mathcal{A} \text{ deterministic}} \mathbb{E}_{P \sim \mathcal{D}} \left[\mathcal{Q}(\mathcal{A}, \hat{P}, \varepsilon) \right] \geq \frac{1}{5} \sup_{\mathcal{D} \in \mathcal{D}_p(n(\varepsilon))} \inf_{\mathcal{A} \text{ deterministic}} \mathbb{E}_{P \sim \mathcal{D}} \left[\mathcal{G}(\tilde{\mathcal{A}}, P) \right].$$

In (Sun and Yao, 2009), the authors present a family of random paths $(X_\delta)_{\delta > 0} \subset \mathcal{D}_p(n)$. Using these random paths it is shown that for every $\delta > 0$,

$$\mathcal{G}_{\text{rand}}(n) := \sup_{\mathcal{D} \in \mathcal{D}_p(n)} \inf_{\mathcal{A} \text{ deterministic}} \mathbb{E}_{P \sim \mathcal{D}} [\mathcal{G}(\mathcal{A}, P)] = \Omega(n^{1-\delta}).$$

This immediately implies,

$$\mathcal{Q}_{\text{rand}}(\varepsilon) = \Omega \left(\left(\frac{1}{\sqrt{\varepsilon}} \right)^{1-\delta} \right).$$

Their proof uses results from combinatorial number theory in order to construct a random path which, roughly speaking, has unpredictable increments. This distribution is then used in conjunction with a method developed by Aaronson (Aaronson, 2006) in order to produce a lower bound.

We now present a simplified proof of the result, which also slightly improves the bound. We simply observe that known results concerning unpredictable random walks, can be combined with Aaronson's method. Theorem 13 will then be a consequence of the following theorem:

Theorem 16 *Let the notations of Theorem 13 prevail. Then*

$$\mathcal{G}_{\text{rand}}(n) = \Omega\left(\frac{n}{S_h(n)}\right).$$

The theorem of Aaronson, reformulated using our notations (see also (Sun and Yao, 2009, Lemma 2)), is given below.

Theorem 17 (Theorem 5, (Aaronson, 2006)) *Let $w : F_n \times F_n \rightarrow \mathbb{R}^+$ be a weight function with the following properties:*

- $w(P, P') = w(P', P)$.
- $w(P, P') = 0$, whenever $P_n = P'_n$.

Define

$$T(w, P) := \sum_{Q \in F_n} w(P, Q),$$

and for $v \in V_n$

$$T(w, P, v) := \sum_{Q \in F_n: Q(v) \neq P(v)} w(P, Q).$$

Then

$$\mathcal{G}_{\text{rand}}(n) = \Omega\left(\min_{\substack{P, P', v \\ P(v) \neq P'(v), w(P, P') > 0}} \max\left(\frac{T(w, P)}{T(w, P, v)}, \frac{T(w, P')}{T(w, P', v)}\right)\right).$$

For $P \in F_n$, one should think about w as inducing a probability measure according to $w(P, \cdot)$. If Q is sampled according to this measure, then the quantity $\frac{T(w, P, v)}{T(w, P)}$ is the probability that $P(v) \neq Q(v)$. That is, either $v \in P$ or $v \in Q$, but not both. The theorem then says that if this probability is small, for at least one path in each pair (P, P') such that $P_n \neq P'_n$, then any randomized algorithm must make as many queries as the reciprocal of the probability.

We now formalize this notion; For a random path $X \in \mathcal{D}_p(n)$, define the following weight function:

$$w_X(P, P') = \begin{cases} 0 & \text{if } P_n = P'_n \\ \mathbb{P}(X = P) \cdot \sum_{i=0}^{n-1} \mathbb{P}(X = P' | X[i] = P[i]) & \text{otherwise} \end{cases}.$$

Here $w_X(P, P')$ is proportional to the probability that $X = P'$, conditional on agreeing with P on the first i steps, where i is uniformly chosen between 0 and $n - 1$. Note that, for any i ,

$$\begin{aligned} \mathbb{P}(X = P) \cdot \mathbb{P}(X = P' | X[i] = P[i]) \\ &= \mathbb{P}(X[i] = P[i]) \cdot \mathbb{P}(X = P | X[i] = P[i]) \cdot \mathbb{P}(X = P' | X[i] = P[i]) \\ &= \mathbb{P}(X = P') \cdot \mathbb{P}(X = P | X[i] = P'[i]). \end{aligned}$$

Hence, $w_X(P, P') = w_X(P', P)$. We will use the following theorem from Häggström and Mossel (1998), which generalizes the main result of (Benjamini et al., 1998).

Theorem 18 (Theorem 1.4, (Häggröm and Mossel, 1998)) *Let h be as in Theorem 13, Then there exists a random path $X^h \in \mathcal{D}_p(n)$ and a constant $c_h > 0$, such that for all $m \geq k$, and for every $(v_0, v_1, \dots, v_{m-k})$, sequence of vertices,*

$$\sup_{\|u\|_1=m} \mathbb{P} \left(X_m^h = u | X_0^h = v_0, \dots, X_{m-k}^h = v_{m-k} \right) \leq \frac{c_h}{k f(k)}. \quad (3)$$

For X^h as in the theorem abbreviate $w_h := w_{X^h}$ and recall $S_h(n) := \sum_{k=1}^n \frac{1}{k \cdot h(k)}$. We now prove the main quantitative estimates which apply to w_h .

Lemma 19 *For any $P \in F_n$,*

$$\sum_{Q \in F_n} w_h(P, Q) \geq \mathbb{P}(X^h = P) \cdot (n - c_h S_h(n)).$$

Proof We write

$$\begin{aligned} \sum_{Q \in F_n} w_h(P, Q) &= \mathbb{P}(X^h = P) \sum_{Q: Q_n \neq P_n} \sum_{i=0}^{n-1} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \\ &= \mathbb{P}(X^h = P) \sum_{i=0}^{n-1} \left(1 - \sum_{Q: Q_n = P_n} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \right) \\ &= \mathbb{P}(X^h = P) \left(n - \sum_{i=0}^{n-1} \sum_{Q: Q_n = P_n} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \right). \end{aligned}$$

Using (3), we get

$$\sum_{Q: Q_n = P_n} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \leq \mathbb{P}(X_n^h = P_n | X^h[i] = P[i]) \leq \frac{c_h}{(n-i) \cdot h(n-i)},$$

and

$$\sum_{i=0}^{n-1} \sum_{Q: Q_n = P_n} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \leq \sum_{k=1}^n \frac{c_h}{k \cdot h(k)} = c_h S_h(n). \quad \blacksquare$$

Lemma 20 *Let $P \in F_n$ and $v \in V_n$ such that $\|v\|_1 = \ell$ and $P_\ell \neq v$. Then,*

$$\sum_{\substack{Q \in F_n \\ Q_\ell = v}} w_h(P, Q) \leq 2 \mathbb{P}(X^h = P) c_h S_h(n).$$

Proof

$$\begin{aligned} \sum_{\substack{Q \in F_n \\ Q_\ell = v}} w_h(P, Q) &= \mathbb{P}(X^h = P) \sum_{i=0}^{\ell-1} \sum_{\substack{Q: Q_n \neq P_n \\ Q_\ell = v}} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \\ &\leq \mathbb{P}(X^h = P) \sum_{i=0}^{\ell-1} \sum_{Q: Q_\ell = v} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right). \end{aligned}$$

Observe that if $Q_\ell = v$, then $Q_{\ell+1}$ must equal $v + (0, 1)$ or $v + (1, 0)$. In particular, for $i < \ell$, (3) shows

$$\begin{aligned} & \sum_{Q: Q_\ell = v} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) \\ & \leq \mathbb{P} \left(X_{\ell+1}^h = v + (0, 1) \text{ or } X_{\ell+1}^h = v + (1, 0) | X^h[i] = P[i] \right) \\ & \leq \mathbb{P} \left(X_{\ell+1}^h = v + (0, 1) | X^h[i] = P[i] \right) + \mathbb{P} \left(X_{\ell+1}^h = v + (1, 0) | X^h[i] = P[i] \right) \\ & \leq \frac{2c_h}{(\ell + 1 - i) \cdot h(\ell + 1 - i)}. \end{aligned}$$

So,

$$\begin{aligned} \sum_{i=0}^{\ell-1} \sum_{Q: Q_\ell = v} \mathbb{P} \left(X^h = Q | X^h[i] = P[i] \right) & \leq \sum_{i=0}^{\ell-1} \frac{2c_h}{(\ell + 1 - i) \cdot h(\ell + 1 - i)} \\ & \leq 2c_h S_h(n). \end{aligned}$$

■

We are now in a position to prove Theorem 16.

Proof [of Theorem 16] Let $P \in F_n$ and let $v \in V_n$, with

$$\|v\|_1 = \ell \text{ and } P_\ell \neq v.$$

Note that $P(v) = \ell$. So, if $Q \in F_n$ is such that $Q(v) \neq P(v)$, then necessarily $Q_\ell = v$. We now set $P' \in F_n$, with $P(v) \neq P'(v)$. In this case, the previous two lemmas show

$$\begin{aligned} & \max \left(\frac{T(w_h, P)}{T(w_h, P, v)}, \frac{T(w_h, P')}{T(w_h, P', v)} \right) \\ & \geq \frac{T(w_h, P)}{T(w_h, P, v)} = \frac{\sum_{Q \in F_n} w_h(P, Q)}{\sum_{\substack{Q \in F_n \\ Q(v) \neq P(v)}} w_h(P, Q)} = \frac{\sum_{Q \in F_n} w_h(P, Q)}{\sum_{\substack{Q \in F_n \\ Q_\ell = v}} w_h(P, Q)} \geq \frac{n - c_h S_h(n)}{2c_h S_h(n)}. \end{aligned}$$

Since we are trying to establish a lower bound, we might as well assume that $S_h(n) = o(n)$. So, for n large enough

$$\frac{n - c_h S_h(n)}{2c_h S_h(n)} \geq \frac{n}{4c_h S_h(n)}.$$

Plugging this estimate into Theorem 17 yields the desired result

■

B.3. Heuristic extension to higher dimensions

In this section we propose a heuristic approach to extend the lower bound to higher dimensions. In the 2 dimensional case, the proof method of Section B consisted of two steps: first reduces the problem to the discrete setting of monotone paths in $[n]^2$, and then analyze the query complexity

of finding the minimal point for such path functions. Thus, to extend the result we should consider path functions on the d -dimensional grid, as well as a way to build smooth functions on $[0, 1]^d$ from those paths.

The lower bound for finding minimal points of path functions in high-dimensional grids was obtained in (Zhang, 2006), where it was shown that, in the worst case, any randomized algorithm must make $\Omega\left(n^{\frac{d}{2}}\right)$ queries in order to find the end point of a path defined over $[n]^d$. Thus, if we can find a discretization scheme, analogous to Lemma 14, in higher dimensions, we could obtain a lower bound for finding ε -stationary points. What are the constraints on such a discretization?

First note that necessarily the construction of (Zhang, 2006) must be based on paths of lengths $\Omega\left(n^{\frac{d}{2}}\right)$, for otherwise one could simply trace the path to find its endpoint. In particular, since each cube has edge length $\frac{1}{n}$, an analogous construction to Lemma 14 will reach value smaller than $-\varepsilon \cdot n^{\frac{d}{2}-1}$ at the stationary point (i.e., the endpoint of the path). On the other hand, in at least one of the neighboring cubes (which are at distance less than $1/n$ from the stationary point), the background linear function should prevail, meaning that the function should reach a positive value. Since around the stationary point the function is quadratic, we get the constraint:

$$-\varepsilon \cdot n^{\frac{d}{2}-1} + \left(\frac{1}{n}\right)^2 > 0 \Leftrightarrow n < \left(\frac{1}{\varepsilon}\right)^{\frac{2}{d+2}}.$$

In particular the lower bound $\Omega\left(n^{\frac{d}{2}}\right)$ now suggests that for finding stationary point one has the complexity lower bound $\left(\frac{1}{\varepsilon}\right)^{\frac{d}{d+2}}$.