

Certifying Smoothness of Discrete Functions and Measuring Legitimacy of Images*

J. MICHAEL STEELE

*Program in Statistics and Operations Research, School of Engineering and Applied
Science, Princeton University, Princeton, New Jersey 08544*

Received April 3, 1989

Several problems that deal with smoothness of discrete functions and that are motivated by the task of deciding whether an $N \times N$ array of distances is a feasible laser radar (LADAR) image are investigated. The main observation is that there is a simple and explicit graph that is tightly related to the smoothness problem. The graph is used to show that a function can be verified to be Lipschitz on a random finite subset of \mathbb{R}^2 by testing the Lipschitz condition at only $O(n)$ pairs of points. © 1989 Academic Press, Inc.

1. INTRODUCTION

The problems studied here are motivated in part by the investigation of statistical models for laser radar images. The salient feature of such images is that they consist of an array of $N \times N$ pixels where each pixel is associated with the actual distance to a corresponding part of a physical scene being viewed. This direct link to distances contrasts powerfully with more traditional image data where the value associated with each pixel is typically the intensity (in some range of the EM spectrum) of either the reflection or the emission from the scene being viewed.

For many practical applications, laser radar (LADAR) images have important advantages over spectrally based images; and, while most of these advantages come directly from the distance interpretation of the pixel values, there are also technical considerations that make LADAR applications easier. For example, LADAR images do not force us to worry about the direction of the illumination source as we must in many applications of optical images. Still, even with its technical and concep-

* Invited Paper.

tual advantages, there are features of LADAR images that provide new challenges; and, in particular, one must come to grips with their genuinely three-dimensional nature.

In several approaches to the understanding of 3D images, like those that LADAR gives for a view of a physical scene, one needs an efficient way to judge whether a specified set of $N \times N$ pixel values is a physically feasible image. Certainly if one wishes to provide an effective prior distribution on the set of feasible LADAR images in order to permit a Bayesian analysis, then it is almost essential to have some reasonably simple criteria for a realization of $N \times N$ pixels to correspond to a real view.

In Section 4, we engage the LADAR problem as squarely as possible, but, in intermediate sections, we look at more stylized problems. These new problems may cast some light on the particular problems that need to be solved in the context of LADAR, but they are of more general applicability.

The most stripped down question we address is “How can we certify that a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is smooth?” The way we choose to frame a discrete (and possibly answerable) version of this question uses the verification of a Lipschitz condition at a finite set of points in order to gauge smoothness. Specifically, given a finite set $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\} \subset [0, 1]^d$ and $f: S \rightarrow \mathbb{R}$, we examine the information that is needed to certify truth of the statement $S(\alpha)$: for all \mathbf{x}_i and \mathbf{x}_j contained in S we have $|f(\mathbf{x}_i) - f(\mathbf{x}_j)| \leq \alpha |\mathbf{x}_i - \mathbf{x}_j|$.

One model of complexity that is reasonable in many contexts (though admittedly strained in the context of LADAR imaging) is to suppose there is an oracle such that for each i and j the oracle can verify the truth of the statement $S_{ij}(\alpha)$ that $|f(\mathbf{x}_i) - f(\mathbf{x}_j)| \leq \alpha |\mathbf{x}_i - \mathbf{x}_j|$. The natural question that presents itself in this model is “How many times must we query the oracle before we can certify the smoothness of f ?” Obviously, $S(\alpha)$ can be certified by $n(n - 1)/2$ queries of the form $S_{ij}(\alpha)$, but we aim to do better.

If $d = 1$ the problem is easy; just order the points $\mathbf{x}_1 \leq \mathbf{x}_2 \leq \dots \leq \mathbf{x}_n$ and ask the oracle the $n - 1$ questions $S_{i,i+1}(\alpha)$. If all of these are answered in the affirmative, we know f is smooth in the sense that $S(\alpha)$ is proved. In this case we can also be sure that no fewer queries can suffice since each \mathbf{x}_i must figure into some query.

After $d = 1$, the most natural setting to consider is that where $n = N^2$ points on the rectangular lattice of \mathbb{R}^2 . This and other lattice problems are solved in Section 2. The lattice problems are still easy, but they are worth considering because they do a good job of isolating the essential features that underlie the complexity of certifying smoothness. In particular, we are led to the useful notion of a *transit graph*, i.e., a spanning graph $G = (V, E)$ of $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ such that for each pair of vertices \mathbf{x}_i and \mathbf{x}_j

there is a path from \mathbf{x}_i to \mathbf{x}_j consisting of edges of E with slope of constant sign. We will show that the transit graph is an apt tool for studying Lipschitz smoothness. Somewhat surprisingly, it is also possible to give an explicit characterization of the transit graph of any V . One further aspect of the transit graph is that it is often sparse. In particular, we find the expected number of edges of the transit graph on n random points in $[0, 1]^2$ is $O(n)$ even though in the worst case it can have $\Omega(n^2)$ edges.

We return to the problem of laser radar images in Section 4 and try to lay out the basic decision criteria for physical feasibility. Naturally, aspects of smoothness are among the criteria, but the characteristic feature of physically realistic LADAR images is strong limitation on the discontinuities that are present.

2. LATTICE CONSIDERATIONS

The nearest neighbor graph $G = (V, E)$ of the $N \times N$ lattice has $n = N^2$ vertices, and, since G has $N - 1$ horizontal edges in each of N rows and a similar number of vertical edges, we see that the total number of edges $|E|$ is exactly $2N(N - 1)$. Now, suppose we know $S_{ij}(\alpha)$ for each $(i, j) \in E$. By joining any pair of points \mathbf{x}_a and \mathbf{x}_b of V with a path consisting of horizontal path P_H from \mathbf{x}_a to \mathbf{x}_c and a vertical path P_V from \mathbf{x}_c to \mathbf{x}_b , we can derive essentially the same bound for lattice points in \mathbb{R}^2 that we found for general points in $d = 1$; specifically we find by telescoping and a little geometry that

$$\begin{aligned} |f(\mathbf{x}_a) - f(\mathbf{x}_b)| &\leq \alpha \sum_{e \in P_V} |e| + \alpha \sum_{e \in P_H} |e| \\ &\leq \alpha |\mathbf{x}_a - \mathbf{x}_c| + \alpha |\mathbf{x}_b - \mathbf{x}_c| \\ &\leq 2^{1/2} \alpha |\mathbf{x}_a - \mathbf{x}_b|. \end{aligned}$$

We can amplify the logic of the last step and also see how one might improve the factor of $2^{1/2}$ by considering the hexagonal lattice. Consider a set of $n = N(N - 1)/2$ points placed at the vertices of that part of a regular hexagonal lattice of side 1 that is contained in an equilateral triangle of side $N - 1$ (see Fig. 1).

Given any \mathbf{x}_a and \mathbf{x}_b in S there is again a shortest path that consists of two straight segments that connect \mathbf{x}_a and \mathbf{x}_b . If \mathbf{x}_c is the lattice point at the turn, and, if $S_{ij}(\alpha)$ is valid for all nearest neighbor pairs (i, j) of S in the hexagonal lattice, then again by telescoping we have

$$|f(\mathbf{x}_a) - f(\mathbf{x}_b)| \leq \alpha |\mathbf{x}_a - \mathbf{x}_c| + \alpha |\mathbf{x}_c - \mathbf{x}_b|.$$




FIG. 1. Hexagonal lattice configuration.

This time we need a more explicit calculation to bound the sum. By the law of sines we have

$$|\mathbf{x}_a - \mathbf{x}_b|/\sin \theta_1 = |\mathbf{x}_a - \mathbf{x}_c|/\sin \theta_2 = |\mathbf{x}_c - \mathbf{x}_b|/\sin \theta_3$$

and, of course $\theta_1 = 2\pi/3$, so since $\sin \theta_1 = \sqrt{3}/2$, we have

$$|\mathbf{x}_a - \mathbf{x}_c| + |\mathbf{x}_b - \mathbf{x}_c| \leq 2|\mathbf{x}_a - \mathbf{x}_b|(\sin \theta_2 + \sin \theta_3)/\sqrt{3}.$$

Since $0 \leq \theta_2, \theta_3 \leq \pi/3$, and $\theta_2 + \theta_3 = \pi/3$, we maximize $\sin \theta_2 + \sin \theta_3$ by taking $\theta_2 = \theta_3 = \pi/6$, and since $\sin \pi/6 = \frac{1}{2}$, we find

$$|\mathbf{x}_a - \mathbf{x}_c| + |\mathbf{x}_b - \mathbf{x}_c| \leq (2/\sqrt{3})|\mathbf{x}_a - \mathbf{x}_b|.$$

We therefore find that for the hexagonal lattice $S(\beta)$ is verified with $\beta = (2/\sqrt{3})\alpha$.

This value of β is a slight improvement over the value $\beta = \sqrt{2}\alpha$ that we obtained in the rectangular case; but this improvement comes at a price, and in terms of the number of queries required we do slightly worse with the hexagonal lattice than with the rectangular lattice. For large n , the number of tested edges per vertex in the rectangular lattice is approximately 2 whereas the number of tested edges per vertex in the hexagonal lattice is approximately 3. These values would be exact if it were not for the influence of boundary effects.

The main lead we take from the lattice point investigation concerns the use of special paths, but before we pursue this lead, there are two observations that should be made concerning $d > 2$. By the same counting we applied in $d = 2$, one also finds for the set S of N^d rectangular lattice points in \mathbb{R}^d that there are $d(N - 1)N^{d-1}$ edges in the nearest neighbor graph and that therefore the verification of $d(N - 1)N^{d-1}$ queries of the form $S_{ij}(\alpha)$ is enough to verify $S(d^{1/2}\alpha)$. We will not pursue the consideration of $d > 2$ any further here, but one might expect the investigation of

the higher-dimensional problems to be particularly interesting because of the rich increase in the possible lattices.

3. GENERAL POSITION

The possibility of verifying $S(\beta)$ with only $O(n)$ queries in two-dimensional lattices came from using telescoping sums along special paths. In this respect the lattice computation suggests a generalization that offers promise in a number of different problems. If a sequence $M \mathbf{x}_j = (x_j, y_j)$, $1 \leq j \leq k$, of points in \mathbb{R}^2 that satisfies either $x_j \leq x_{j+1}$ and $y_j \leq y_{j+1}$ for all j , or $x_j \leq x_{j+1}$ and $y_j \geq y_{j+1}$ for all $1 \leq j \leq k$ then we will call M a *monotone* sequence. In other words, a path M connecting a sequence of points in \mathbb{R}^2 is monotone if all of the slopes of the edges of M have the same sign.

We will call $G = (V, E)$ a *transit* graph of $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ provided that for each pair of points \mathbf{x}_i and \mathbf{x}_j of V there is a monotone path M_{ij} consisting of edges of E . One can easily check that the arguments we applied for lattices apply directly to any set of points spanned by a transit graph. This is recorded more formally in the following proposition.

PROPOSITION 1. *If $G = (V, E)$ is a transit graph for $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ and $|f(\mathbf{x}) - f(\mathbf{y})| \leq \alpha$ for all $(\mathbf{x}, \mathbf{y}) \in E$, then $|f(\mathbf{x}) - f(\mathbf{y})| \leq 2^{1/2}\alpha$ for all \mathbf{x}, \mathbf{y} contained in V .*

One fact that makes the notion of a transit graph useful is that we can easily characterize the smallest spanning transit graph of a finite set of points in \mathbb{R}^2 . Given any two points \mathbf{x} and \mathbf{x}' in \mathbb{R}^2 , we can determine a *box* $B(\mathbf{x}, \mathbf{x}')$ in \mathbb{R}^2 by letting $B(\mathbf{x}, \mathbf{x}')$ denote the smallest closed rectangle containing \mathbf{x} and \mathbf{x}' . For example, if $\mathbf{x} = (x, y)$ and $\mathbf{x}' = (x', y')$ with $x < x'$ and $y < y'$, then we have $B(\mathbf{x}, \mathbf{x}') = [x, x'] \times [y, y']$. It turns out that boxes lead to an explicit characterization of the minimal spanning transit graph.

LEMMA 1. *Let $G = (V, E)$ be the graph with vertex set $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ and edge set E such that $(\mathbf{x}, \mathbf{x}') \in E$ if and only if $B(\mathbf{x}, \mathbf{x}') \cap V = \{\mathbf{x}, \mathbf{x}'\}$. The graph G is then a spanning transit graph for V . Moreover, if G' is any other spanning transit graph, then $G \subset G'$.*

Proof. To see that the G defined by the lemma is a transit graph, we use induction on n . If \mathbf{x} and \mathbf{x}' are given, we need to find a monotone path in G from \mathbf{x} to \mathbf{x}' . By symmetry, there is no loss of generality in assuming that $\mathbf{x} = (x, y)$ and $\mathbf{x}' = (x', y')$ satisfy $x < x'$ and $y < y'$. If $(\mathbf{x}, \mathbf{x}')$ is not already an edge of E there is a point $\mathbf{x}'' \in V$ such that $\mathbf{x}'' \in B(\mathbf{x}, \mathbf{x}')$ and \mathbf{x}'' distinct from \mathbf{x} and \mathbf{x}' . By the induction hypothesis applied to $G - \{\mathbf{x}'\}$ there is a monotone M_1 from \mathbf{x} to \mathbf{x}'' and by the induction hypothesis applied to $G - \{\mathbf{x}\}$ there is a monopath M_2 from \mathbf{x}'' to \mathbf{x}' . We can take $M =$

$M_1 \cup M_2$ as a monotone path from \mathbf{x} to \mathbf{x}' . If $B(\mathbf{x}, \mathbf{x}')$ does not contain such an \mathbf{x}'' , then by definition of G we have $(\mathbf{x}, \mathbf{x}') \in E$, and we can simply take $M = \{(\mathbf{x}, \mathbf{x}')\}$ as the desired path. ■

It would be pleasing if the transit graph $G = (V, E)$ of a set $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ could be guaranteed to be relatively small, say with $|E| = O(n)$ or even $|E| = o(n^2)$. While we can show that in a sense this is typically the case, we can also show that $|E| = \Omega(n^2)$ is possible for worst case configurations.

To obtain an example that shows the transit graph can be large, consider the set of $2n$ points given by $V = \{(i, -i), (n + i, 2n - i), 1 \leq i \leq n\}$. For each $1 \leq i \leq n$ and $1 \leq j \leq n$ the boxes $B((i, -i), (n + j, 2n - j))$ contain no points of V other than their defining vertices, so the transit graph of G contains at least the n^2 edges required by the boxes for the pairs $(i, -i)$ and $(n + j, 2n - j)$ for $1 \leq i \leq n$ and $1 \leq j \leq n$. In fact, by considering the boxes $B((i, -i), (i + 1, -i - 1))$ and $B(n + j, 2n - j), (n + j + 1, 2n - j - 1)$ we see that the minimal spanning transit graph of V has exactly $n^2 + 2(n - 1)$ edges.

A more optimistic view of the world is provided by considering points chosen at random from the unit square.

PROPOSITION 2. *If $\mathbf{x}_i, 1 \leq i \leq n$, are independent uniformly distributed random variables with values in $[0, 1]^2$, then the expected number of edges of the minimal spanning transit graph $G = (V, E)$ of $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ satisfies $E(|E|) = O(n)$.*

Proof. If B_{12} denotes the rectangle determined by $\mathbf{x}_1 = (x_1, y_1)$ and $\mathbf{x}_2 = (x_2, y_2)$ then B_{12} has the random area $|x_1 - x_2| |y_1 - y_2|$. For B_{12} to be empty the other $n - 2$ points of V must fall outside of B_{12} , so by conditioning we have

$$P(B_{12} \text{ empty}) = \int_{[0,1]^4} (1 - |x_1 - x_2| |y_1 - y_2|)^{n-2} dx_1 dx_2 dy_1 dy_2.$$

By applying $1 - x \leq e^{-x}$ for $x \geq 0$ and changing variables, we further find

$$P(B_{12} \text{ empty}) \leq \int_{[0,1]^4} \exp\{-|x_1 - x_2| |y_1 - y_2|(n - 2)\} dx_1 dx_2 dy_1 dy_2 = O(n^{-1}).$$

Finally, since there are $n(n - 1)/2$ pairs of points that have associated boxes we have

$$E(|E|) = \binom{n}{2} P(B_{12} \text{ empty}) = O(n). \quad \blacksquare$$

We will not pursue further calculations of this type, but it should be clear from the preceding propositions that the transit graph is quite tractable. Also, by Proposition 1, it is clear that possession of the transit graph is a useful step in Lipschitz verification.

4. CRITERIA FOR LADAR REALISM

We now return to the issue of judging the physical feasibility of LADAR images. This pursuit differs in spirit from the preceding analysis in several respects, and in particular it no longer seems possible to give results that can be expressed in tidy propositions. In fact, almost anything one says about criteria for a LADAR image must rest substantially on judgments that one makes about the physical world and current technology.

The applications that are currently envisioned for LADAR involve carrying the LADAR device in a helicopter. The elevation of the helicopter and its distance from the scene of interest are typically determined so that the line of sight makes about a 15° angle with the horizontal. By examining typical scenes at the resolution that can be assumed, it is possible to make some assumptions about the sizes of the discontinuities one can expect in the distances represented in the pixels of the $N \times N$ image array in A .

To avoid the extra layer of complexity associated with the discretized image, first consider a continuous version of the image. That is, we view the image as if it were given by a distance function $f: [0, 1]^2 \rightarrow \mathbb{R}$. The class of functions that one is led to consider in this representation is a little unusual by the standards of traditional analysis. In particular, by considering the simple image sketched in Figs. 2a and 2b, one finds that it is essential to consider functions f with discontinuities, and, in order for the representing function to have discontinuities that are physically realistic, they must be of a highly confined nature. In fact the discontinuities of f cannot amount to much more than the two-dimensional analogs of simple jump discontinuities.

By considering Fig. 2, it also seems appropriate to assume that the set of discontinuities of f yields a set of curves in $[0, 1]^2$ such that as one traverses one of these curves of discontinuity C_i obliquely, one sees only either an increase in f (a *plus* jump) or a decrease in f (a *minus* jump). Moreover, because of the relationship between the physical scene and the orientation of the image, there are constraints on the types of jumps that can occur along specified paths in $[0, 1]^2$. For example, in a world where we view features like overhanging eaves as too rare to worry about, one




FIG. 2. Distance function geometry. (a) Scene. (b) Discontinuities of f .

sees either plus jumps or minus jumps as we move from right to left, and only plus jumps as one moves from bottom to top.

The restriction to pure jump discontinuities of these confined types seems to be about as far as it seems appropriate to go in terms of the topological constraints one should place on realistic LADAR images. There are obviously additional metric constraints one should impose, but at that point one has to dig more deeply into the technology of the moment, and we will not pursue the abstract modeling of LADAR images beyond this first cut. In fact, the steps that remain in the engineering problem are probably best addressed via the Edisonian shotgun approach of trying the many reasonable variants that present themselves. Still, one should have confidence that the remaining problems will receive attention for the good reason that the technology of LADAR offers such promise that its refinement and development is almost inevitable.

From the point of view of the complexity of approximately solved problems, there is an interesting distinction between line of applied research suggested by LADAR and the work in remote sensing that preceded LADAR. The distinction is that, in LADAR, in order to identify the items of interest in the remotely sensed scene one has a compelling need for geometry, while in older areas of remote sensing, geometry has not been nearly so essential. There are two reasons that underlie this shift. First, in older technologies one simply did not have enough geometric detail to make serious use of geometry in the identification process. Second, there is the lucky fact that in older technologies one often found characteristic signatures for items of interest in the remotely sensed scene and useful identification was possible even without geometric detail.

The intention of this section has been to describe the flavor of the questions that are of interest in the analysis of LADAR images. The subject is new, yet vast, so the discussion has been narrowed to the discussion of the topological constraints although we readily admit that a

technologically useful answer to the criteria for realism needs to engage many metric issues as well.

Before we leave the question of criteria for realism of LADAR images, it is useful to record two benefits of a full answer. In the first place, by knowing the domain of realistic LADAR images we are in a far better position to specify a family of probability distributions on the set of LADAR images. Naturally, such a step is necessary if one is to use Bayesian methods in the analysis of LADAR images. A more immediate application for a criterion of realistic images is that it permits real time feedback to the LADAR device. This can in some cases provide the mechanism for improving the LADAR device itself, and at a minimum it provides a command and control tool that can suggest when a LADAR device should be taken out of use.

5. CONCLUSION

We entered this exploration by considering the distinctive features of LADAR images, and we have left by the same door. In the middle we addressed the complexity of certifying the smoothness of a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, and it was there that the most satisfying steps were taken.

Both of these problem areas fall within the scope of the theory of complexity of approximately solved problems as delineated by Traub (1985). Also, we should note that even though image understanding is ultimately connected to the aims of LADAR, the simpler step contemplated here of checking for legitimacy must be satisfactorily resolved before one can address broader issues of image understanding such as those contemplated by Lee (1985).

The certification of smoothness seems like a natural and interesting problem. It also seems not to have been treated before, though naturally, there is work with a related flavor. One example is given by the work on the realizability of metric spaces in graphs (e.g., Altenhöfer, 1988), and a second is the work on the approximate isometric embedding of finite subsets of \mathbb{R}^n into \mathbb{R}^k , where k is much smaller than n (e.g., Frankl and Maehara, 1988). We hope that Sections 2 and 3 illustrate that there are interesting issues in the complexity of smoothness, and at least that they should leave little doubt that many issues remain essentially unexplored.

REFERENCES

- ALTENHÖFER, I. (1988), On optimal realizations of finite metric spaces by graphs, *Discrete Comput. Geom.* **3**, 103–122.

- FRANKL, P., AND MAEHARA, H. (1988). The Johnson–Lindenstrauss lemma and the sphericity of some graphs, *J. Combin. Theory Ser. B* **44**, 355–362.
- LEE, D. (1985). Optimal algorithms for image understanding: Current status and future plans, *J. Complexity* **1**, 138–146.
- TRAUB, J. F. (1985). Complexity of approximately solved problems, *J. Complexity* **1**, 3–10.