

Output Sensitive Algorithms for Approximate Incidences and Their Applications*

Dror Aiger¹, Haim Kaplan², and Micha Sharir³

1 Google Inc., Mountain View, CA, USA
aigerd@google.com

2 Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel
haimk@tau.ac.il

3 Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel
michas@tau.ac.il

Abstract

An ε -approximate incidence between a point and some geometric object (line, circle, plane, sphere) occurs when the point and the object lie at distance at most ε from each other. Given a set of points and a set of objects, computing the approximate incidences between them is a major step in many database and web-based applications in computer vision and graphics, including robust model fitting, approximate point pattern matching, and estimating the fundamental matrix in epipolar (stereo) geometry.

In a typical approximate incidence problem of this sort, we are given a set P of m points in two or three dimensions, a set S of n objects (lines, circles, planes, spheres), and an error parameter $\varepsilon > 0$, and our goal is to report all pairs $(p, s) \in P \times S$ that lie at distance at most ε from one another. We present efficient output-sensitive approximation algorithms for quite a few cases, including points and lines or circles in the plane, and points and planes, spheres, lines, or circles in three dimensions. Several of these cases arise in the applications mentioned above. Our algorithms report all pairs at distance $\leq \varepsilon$, but may also report additional pairs, all of which are guaranteed to be at distance at most $\alpha\varepsilon$, for some problem-dependent constant $\alpha > 1$. Our algorithms are based on simple primal and dual grid decompositions and are easy to implement. We note that (a) the use of duality, which leads to significant improvements in the overhead cost of the algorithms, appears to be novel for this kind of problems; (b) the correct choice of duality in some of these problems is fairly intricate and requires some care; and (c) the correctness and performance analysis of the algorithms (especially in the more advanced versions) is fairly non-trivial. We analyze our algorithms and prove guaranteed upper bounds on their running time and on the “distortion” parameter α .

1998 ACM Subject Classification F.2.2 Geometrical problems and computations

Keywords and phrases Approximate incidences, near-neighbor reporting, duality, grid-based approximation

Digital Object Identifier 10.4230/LIPIcs.ESA.2017.5

* Work by Haim Kaplan has been supported by Grant 1161/2011 from the German-Israeli Science Foundation and by Grant 1841-14 from the Israel Science Foundation. Work by Micha Sharir has been supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation, by Grant 892/13 from the Israel Science Foundation, by the Blavatnik Research Fund in Computer Science at Tel Aviv University, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University. Both Kaplan and Sharir have also been supported by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11).



1 Introduction

Approximate incidences. Given a finite point set S_1 and finite set S_2 of geometric primitives (e.g., lines, planes, circles, or spheres in \mathbb{R}^2 or \mathbb{R}^3), and some $\varepsilon > 0$, we define the set of ε -incidences (also referred to as ε -approximate incidences, or just *approximate incidences*) between S_1 and S_2 to be

$$I_\varepsilon(S_1, S_2) = \{(s_1, s_2) \mid s_1 \in S_1, s_2 \in S_2, \text{dist}(s_1, s_2) \leq \varepsilon\},$$

where $\text{dist}(s_1, s_2) = \inf\{\text{dist}(s_1, y) \mid y \in s_2\}$ is the Euclidean distance between s_1 and s_2 . We are interested in efficient algorithms for computing $I_\varepsilon(S_1, S_2)$, ideally in time linear in $|S_1| + |S_2| + |I_\varepsilon(S_1, S_2)|$. Most classical work in discrete and computational geometry is focused on exact incidences ($\varepsilon = 0$). When S_2 is a set of lines in the plane and $\varepsilon = 0$, detecting whether $I_0(S_1, S_2)$ is empty or not is the well studied *Hopcroft's problem* (see, e.g., [8]). In contrast, the notion of approximate incidences, as we define here, probably received less theoretical attention, but has many important applications which we review below. We consider the problem of reporting all pairs in $I_\varepsilon(S_1, S_2)$. Our algorithms, though, can also estimate $|I_\varepsilon(S_1, S_2)|$, rather than report its members, and do it faster when $|I_\varepsilon(S_1, S_2)|$ is small.

This problem can be viewed as a range searching problem. Specifically, we treat each member s_2 of S_2 as the range $s_2(\varepsilon) = \{p \in \mathbb{R}^d \mid \text{dist}(p, s_2) \leq \varepsilon\}$, $d = 2, 3$, which is the Minkowski sum of s_2 with a disk (ball in \mathbb{R}^3) of radius ε (centered at the origin); thus points become disks, lines become slabs (in \mathbb{R}^2) or cylinders (in \mathbb{R}^3), circles become annuli (in \mathbb{R}^2) or tori (in \mathbb{R}^3), and so on. The goal now is to report all pairs $(s_1, s_2) \in S_1 \times S_2$ such that $s_1 \in s_2(\varepsilon)$. As mentioned, the known algorithms for such tasks have a rather large overhead. For example, when S_1 is a set of m points and S_2 is a set of n lines in the plane, i.e., the ranges $s_2(\varepsilon)$ are fixed-width slabs, the best known algorithms for solving the problem have an overhead close to $m^{2/3}n^{2/3}$, and there are matching lower bounds in certain models of computation. The overhead is larger when the objects in S_2 are of more complex shapes (e.g., arbitrary circles) or when we move to three (or higher) dimensions; see [1]. In addition, these algorithms, while interesting and sophisticated from a theoretical point of view, are a nightmare to implement in practice.

Instead, with the goal of obtaining algorithms that are really simple to implement (and therefore with good performance in practice), and that run in time linear in the input and output sizes, we adopt the approach of using *approximation schemes*, in which we still report all the pairs (s_1, s_2) that satisfy $\text{dist}(s_1, s_2) \leq \varepsilon$, but are willing to report additional pairs, provided that all pairs that we report satisfy $\text{dist}(s_1, s_2) \leq \alpha\varepsilon$, for some constant problem-dependent parameter $\alpha > 1$. To be more precise, assuming that the test whether $\text{dist}(s_1, s_2) \leq \varepsilon$ is cheap, we can filter the reported pairs by such a test, and actually report only the pairs that pass it. The actual number of pairs that we have to inspect will typically be larger than $|I_\varepsilon(S_1, S_2)|$, but it will always be at most $|I_{\alpha\varepsilon}(S_1, S_2)|$ (and in practice considerably less than that), and the hope is that the number of inspected pairs will not be much larger than those that we actually report. (We expect it to be larger by only a constant factor, which depends on α and on the geometry of the setup under consideration.)

Our results. We present simple and efficient output-sensitive algorithms (in the above sense) for approximate-incidence reporting problems between points and various simple geometric shapes, in two and three dimensions.

To calibrate the merits of our solutions, we first note that these approximate incidence reporting problems can also be solved by naive grid-based algorithms, as follows. Consider,

for example, the problem of reporting approximate incidences between a set S_1 of m points and a set S_2 of n lines in the plane. We assume that all the incidences that we seek occur in the unit disk (ball in \mathbb{R}^3). We partition the unit disk by a uniform grid, each of whose cells is a square of side length ε . We store each point in S_1 in a bucket corresponding to the grid cell that contains it, and, for each line $\ell \in S_2$, we report all the pairs involving ℓ and the points in the grid cells that ℓ crosses, and in their neighboring cells. The running time is $O(m + n/\varepsilon + k)$, where k is the number of reported approximate incidences. Clearly, all pairs $(p, \ell) \in S_1 \times S_2$ with $\text{dist}(p, \ell) \leq \varepsilon$ are reported, and each reported pair (p, ℓ) satisfies $\text{dist}(p, \ell) \leq 2\sqrt{2}\varepsilon$, as is easily checked. If n is much larger than m , we can use duality (where some care is needed to preserve point-line distances), to map the points to lines and the lines to points, and thereby reduce the complexity to $O(n + m + \min\{m, n\}/\varepsilon + k)$. This method can also be applied in three dimensions, and yields the same time bounds as in the preceding primal-only approach (duality is much trickier in these situations), namely, $O(m + n/\varepsilon + k)$, when S_2 consists of one-dimensional objects (e.g., lines or circles), but the running time deteriorates to $O(m + n/\varepsilon^2 + k)$ when S_2 consists of surfaces (e.g., planes or spheres). In these latter cases (involving planes or *congruent* spheres) duality can be applied, to improve the time bound to $O(n + m + \min\{m, n\}/\varepsilon^2 + k)$.

While superficially these simple solutions might look ideal, as they are linear in m , n , and k , their dependence on ε is too naive and weak, and when m and n are large and ε small (as is typically the case in practice), the algorithms are rather slow in practice.

In this paper we address this issue, and develop a series of “primal-dual” grid-based algorithms for several approximate incidence reporting problems, that are faster than this naive scheme for suitable ranges of the parameters m , n , and ε (which cover most of the practical instances of these problems). Specifically, we present the following results. In all of them, S_1 is a set of m points, contained in the unit ball in two or three dimensions.

- (a) In the plane, for a set S_2 of n lines, all k approximate incidences can be reported in time $O(m + n + \sqrt{mn}/\sqrt{\varepsilon} + k)$. (The dependency of the complexity on ε is improved by a factor of $\sqrt{\varepsilon}$ compared to the naive scheme when n and m are comparable.)
- (b) In three dimensions, for a set S_2 of n planes, all k approximate incidences can be reported in time $O(m + n + \sqrt{mn}/\varepsilon + k)$. (The dependency of the complexity on ε is improved by a factor of ε compared to the naive scheme, when n and m are comparable.)
- (c) In the plane, for a set S_2 of n congruent circles, all k approximate incidences can be reported in time $O(m + n + \sqrt{mn}/\sqrt{\varepsilon} + k)$.
- (d) In the plane, for a set S_2 of n arbitrary circles, all k approximate incidences can be reported in time $O(m + n + m^{1/3}n^{2/3}/\varepsilon^{2/3} + k)$.
- (e) In three dimensions, for a set S_2 of n congruent spheres, all k approximate incidences can be reported in time $O((m + n)/\varepsilon + k)$.
- (f) In three dimensions, for a set S_2 of n lines, all k approximate incidences can be reported in time $O(m + n + m^{1/3}n^{2/3}/\varepsilon^{2/3} + k)$.
- (g) In three dimensions, for a set S_2 of n congruent circles, all k approximate incidences can be reported in time $O((m + n)/\varepsilon^{1/2} + m^{1/3}n^{2/3}/\varepsilon^{7/6} + k)$.

In Section 4, we use the algorithms in (e) and (g), to obtain an efficient algorithm to find nearly congruent triangles which is the first step in solving the approximate point pattern matching problem in \mathbb{R}^3 .

A comparison with the naive solutions sketched above clearly shows the superiority of our technique. For example, for lines or congruent circles in the plane, assuming that $n \leq m$, our algorithms (in (a) and (c), respectively) are asymptotically faster than the naive method when $\sqrt{mn}/\varepsilon \leq n/\varepsilon$, that is, when $\varepsilon \leq n/m$, an assumption that holds in most practical applications.

To recap, we show that, by allowing to report some additional approximate incidences between pairs that are at most $\alpha\varepsilon$ apart, one can obtain substantially better bounds than the naive ones. Our methods are based on grids and on duality – they construct much coarser primal grids, and pass each subproblem, consisting of the points in a grid cell and of the objects that pass through or near that cell, to a secondary dual stage, in which another coarse grid is constructed in a suitably defined dual space. The output pairs are obtained from the cells of these secondary grids, and the gain is in the overhead, as each primal or dual object crosses much fewer grid cells than in the naive solutions. Although this primal-dual paradigm is fairly standard, its power in the approximate incidences context, as considered here, has not been demonstrated before (to the best of our knowledge). The analysis (and the particular duality one has to use) for some of the three-dimensional variants is fairly challenging, but the algorithms all remain simple to describe and to implement.

Motivation and applications. Approximate incidence reporting and counting problems arise in several basic practical applications, in computer vision, pattern recognition, and related areas. Three major applications of this sort are robust *model fitting*, *approximate point pattern matching* under rigid motions, and estimating the fundamental matrix in (stereo) *epipolar geometry*. All three problems share a common paradigm, which we first explain for model fitting. In this problem, we are given a set P of n points, say in \mathbb{R}^3 (typically, these are so-called *interest points*, extracted from some image or 3D sensors), and we want to fit objects (called *models*) from some given family, such as lines, circles, planes, or spheres, so that each model passes near (i.e., is approximately incident to) many points of P ; the quality of the model is measured in terms of the number of approximately incident points. The standard approach is to construct (usually, by repeated random sampling) a sufficiently rich collection of candidate models. (For example, for line models, one can simply sample pairs of points of P , and for each pair construct the line passing through its points.) One then counts, for each candidate line, the number of approximately incident points (for some specified error parameter $\varepsilon > 0$), and reports the models that have sufficiently many such points.

Similar reductions arise in the other problems. In approximate point pattern matching, we are given two sets A, B of points, and want to find rigid motions that map sufficiently large subsets of A to sets whose (unidirectional) Hausdorff distance to B is at most ε . Here too we construct candidate rigid motions, and test the quality of each of them. For example, in the plane, we sample pairs of points from A , and find, for each sampled pair, the pairs of points of B that are nearly at the same distance. For each such pair of pairs we construct a rigid motion that maps the first pair to near the other pair, and then test the quality of each of these motions, namely, the number of points of A that lie, after the motion, near points of B . The first step can be reduced to approximate incidence counting involving circles (whose radii correspond to the distances between the sampled points of A , and which are centered at the points of B) and the points of B . In three dimensions, we need to sample triples of points of A , and for each triple a, b, c , we need to find those triples of B that span triangles that are nearly congruent to Δabc (because to determine a rigid motion in \mathbb{R}^3 we need to specify how it maps three (noncollinear) source points to three respective image points). This step is described in detail in Section 4.

In epipolar geometry, we have two stereo images A, B of the same scene, and we want to estimate the fundamental matrix F that best matches A to B , where a point $p \in A$ is (exactly) matched to a point $q \in B$ if $p^T F q = 0$. We construct a sample of candidate matrices, by repeatedly sampling $O(1)$ interest points from both images, and test the quality

of each matrix. To do so for a candidate matrix F , we left-multiply each point $p \in A$ by F , interpret the resulting vectors $p^T F$, for $p \in A$, as lines, and count the approximate incidences of each line with the points of B . If sufficiently many lines have sufficiently high counts, we regard F as a good fit and output it.

To recap, in each of these applications, and in other applications of a similar nature, we generate a random sample of candidate models, motions, or matrices, and need to test the quality of each candidate. Approximate incidence reporting and counting arises either in the generation step, or in the quality testing step, or in both. Improving the efficiency of these steps is therefore a crucial ingredient of successful solutions for these problems. The standard approach, used “all over” in computer vision in practice, is the RANSAC technique [6, 9], which checks in brute force each model against each point. Replacing it by efficient methods for approximate incidence counting, which is our focus here, can drastically improve the running time of these applications.

To support the claim that this is indeed the case in practice, we have conducted preliminary experiments (not reported here) with some of our algorithms, tested them on real and random data, and compared them with other existing methods. Roughly, they demonstrate that our approach is significantly faster than the other approaches. Our experiments also support our feeling that the cost of reporting more pairs than really needed (pairs that might be at most $\alpha\varepsilon$ apart, rather than just ε), is negligible compared to the cost of the other steps (in themselves much more efficient than the competing techniques). We leave the project of conducting a thorough experimental study for future work, and focus this paper on developing the algorithms and establishing their worst-case guarantees.

Related work. Model fitting and point pattern matching have been the focus of many studies, both theoretical and practical; see for example [2, 3, 4, 5, 7, 10, 11, 12, 14].

We first note that many of the common approaches used in practice (e.g., RANSAC for model fitting [6, 9]), reporting or counting approximate incidences between models and points is done using brute force, examining every pair of a model and a point. Some heuristic improvements have also been proposed (see, e.g., [5] and the references therein). A similar brute-force technique is commonly used for approximate point pattern matching too (e.g., in the *Alignment* method [12] and its many variants).

The use of (exact) geometric incidences in algorithms for *exact* point pattern matching is well established; see, e.g., Brass [4] for details. Similar connections have also been used for the more practical problem of *approximate* point pattern matching. Gavrilov et al. [10] gave efficient algorithms for approximate pattern matching in two and three dimensions (where the entire sets A and B are to be matched), that use algorithms for reporting approximate incidences. One of the main results in [10] is that in the plane, all pairs of points at distance in $[(1 - \varepsilon)r, (1 + \varepsilon)r]$ can be reported in $O(n\sqrt{r/\varepsilon})$ time, using a grid-based search. (In a way, part of the study in this paper formalizes, extends, and improves this method.)

Aiger et al. [3] proposed a method for point pattern matching in \mathbb{R}^3 , called 4PCS (4-Points Congruent Sets), which iterates over all coplanar pairs of quadruples of points, one from A and one from B , that can be matched via an affine transformation, and then tests the quality of each pair, focusing on pairs where the transformation is rigid. This algorithm does not use approximate incidences, and assumes the existence of coplanar tuples.

In a more recent work, Aiger and Kedem [2] describe another algorithm for computing approximate incidences of points and circles, following a similar approach by Fonseca and Mount [7] for points and lines, which is better than the one of [10] for $n = \Omega(1/\varepsilon^{3/2})$, and use this for approximate point pattern matching. This algorithm has been used in Mellado

et al. [14], to reduce the running time of the 4PCS algorithm in [3] to be asymptotically linear in n and in the output size.

The method of [2, 7] provides an alternative approach to approximate incidence reporting, for the cases of points and lines or congruent circles (the analysis in [2] is rather sketchy, though). This technique runs in $O(m + n + \log(1/\varepsilon)/\varepsilon^2 + k)$ time. For the case of lines in the plane, the scheme exploits the fact that we can approximate (up to an error of $O(\varepsilon)$) all lines in the plane that cross the unit disk, by $O(1/\varepsilon^2)$ representative lines, such that if a point in the unit disk is close to a representative line ℓ , then it is also close to all the lines in the input that ℓ represents (and vice versa). Assuming, for example, that m is constant, this alternative scheme is better than our new algorithm (for these restricted scenarios) when $\sqrt{n}/\sqrt{\varepsilon} \geq 1/\varepsilon^2$, that is, when $n \geq 1/\varepsilon^3$ (we ignore the factor $\log(1/\varepsilon)$ in this calculation). (This technique seems to be extendible to three dimensions, and to surfaces, but the formal details have not yet been worked out, as far as we know.)

Paper organization. The full version of the paper presents seven algorithms for various instances of approximate incidence reporting, as listed in (a)–(g) above. Although the high-level structure of the algorithms is fairly uniform, the specific details are rather different, and each case requires careful analysis to ensure its correctness and efficiency. Working out the details, including the appropriate form of duality (which, in some cases, is rather intricate and requires extra care), the choice of the various parameters, and the analysis that makes everything work, turned out to be fairly demanding and nontrivial. Due to lack of space, this version contains full details of only the first algorithm (for points and lines in the plane), and of the last one (finding all nearly congruent triangles in \mathbb{R}^3), and then describes, briefly and informally, the main features of the rest.

2 Approximate incidences in point-line configurations


We consider the approximate incidences problem between a set P of m points in the unit disk B in \mathbb{R}^2 , and a set L of n lines that cross B , with a given accuracy parameter $0 < \varepsilon \leq 1/2$.

We approximate the distance $\text{dist}(p, \ell)$ by the vertical distance between $p \in P$ and $\ell \in L$, which we denote by $\text{dist}_v(p, \ell)$. For this approximation to be good, the angle between ℓ and the x -direction should not be too large. To ensure this, we partition L into two subfamilies, one consisting of the lines with positive slopes, and one of the lines with negative slopes. We fix one subfamily, rotate the plane by 45° , and get the desired property.


Without loss of generality, we replace the unit disk B by the unit square $S = [0, 1]^2$, and apply the following two-stage partitioning procedure. First we partition S into $1/\delta_1^2$ pairwise openly disjoint smaller squares, each of side length δ_1 , where δ_1 is a parameter whose exact value will be set later. See Figure 1.

Enumerate these squares as $S_1, S_2, \dots, S_{1/\delta_1^2}$. For $i = 1, \dots, 1/\delta_1^2$, let P_i denote the set of all points of P that lie either in S_i or in one of the two squares that are directly above and below S_i (if they exist), and let L_i be the set of all the lines of L that cross S_i . Put $m_i := |P_i|$ and $n_i := |L_i|$. We have $\sum_i m_i \leq 3m$ and $\sum_i n_i \leq 2n/\delta_1$, because each line of L crosses at most $2/\delta_1$ squares S_i .

We now apply a duality transformation to each small square S_i separately. For notational simplicity, and without loss of generality, we may assume that $S_i = [-\delta_1/2, \delta_1/2]^2$. (Technically, this means that we shift the cells by $\delta_1/2$ in both coordinate directions, so that the grid vertices now represent the centers of the cells.) We map each point $p = (\xi, \eta)$ in P_i to the line $p^* : y = \xi x - \eta$, and each line $\ell : y = cx + d$ in L_i to the point



■ **Figure 1** The partition of S into subsquares, and the subproblem associated with the middle highlighted subsquare.



■ **Figure 2** The reference triangle Δabc aligned with Δpqo . The shaded region is K . The circle is the a cross section of $T_{p,q}$.

$\ell^* = (c, -d)$. This duality preserves the vertical distance dist_v between a point p and a line ℓ ; that is, $\text{dist}_v(p, \ell) = \text{dist}_v(\ell^*, p^*)$. Note that the slope condition ensures that $\text{dist}(p, \ell) \leq \text{dist}_v(p, \ell) \leq \sqrt{2}\text{dist}(p, \ell)$.

Let $\ell : y = cx + d$ be a line in L_i , that is, ℓ crosses S_i . By the slope condition we have $-1 \leq c \leq 1$ and $-\delta_1 \leq d \leq \delta_1$, so the dual point ℓ^* lies in the rectangle $R := [-1, 1] \times [-\delta_1, \delta_1]$. Each point $p = (\xi, \eta) \in P_i$ satisfies $-\delta_1/2 \leq \xi \leq \delta_1/2$ and $-3\delta_1/2 \leq \eta \leq 3\delta_1/2$ so the coefficients of the dual line $p^* : y = \xi x - \eta$ satisfy these inequalities.

We now partition R into $1/\delta_2^2$ small rectangles, each of width $2\delta_2$ and height $2\delta_1\delta_2$, where δ_2 is another parameter that we will shortly specify. Each dual line p^* crosses at most $2/\delta_2$ small rectangles. To facilitate the following analysis, we choose δ_1, δ_2 so that they satisfy $\delta_1\delta_2 = \varepsilon$; we still have one degree of freedom in choosing them, which we will exploit later.

► **Lemma 1.** *For each small rectangle R' , if ℓ^* is a dual point in R' and p^* is a dual line that crosses either R' or one of the small rectangles directly above or below R' (in the y -direction, if they exist), then the vertical distance $\text{dist}_v(\ell^*, p^*)$ (which is the same as $\text{dist}_v(p, \ell)$) is at most $5\delta_1\delta_2 = 5\varepsilon$.*

Proof. Indeed, if p^* crosses a small rectangle R'' , which is either R' or one of the two adjacent rectangles, as above, then, since the slope of p^* is in $[-\delta_1/2, \delta_1/2]$, its maximum vertical deviation from R'' is at most $2\delta_2 \cdot (\delta_1/2) = \delta_1\delta_2$. Adding the heights $2\delta_1\delta_2$ of R'' , and of R' when $R'' \neq R'$, the claim follows. ◀

► **Lemma 2.**

- (a) *Let $(p, \ell) \in P \times L$ be such that $\text{dist}(p, \ell) \leq \varepsilon$. Let S_i be the small square containing p . If $\delta_1 \geq \varepsilon\sqrt{2}$, then ℓ must cross either S_i or one of the two squares directly above and below S_i . In other words, there exists a j such that $(p, \ell) \in P_j \times L_j$.*
- (b) *Continue to assume that $\text{dist}(p, \ell) \leq \varepsilon$, let i be such that $(p, \ell) \in P_i \times L_i$, and let R' be the dual small rectangle (that arises in the dual processing of S_i) that contains ℓ^* . Then the dual line p^* must cross either R' or one of the two small rectangles lying directly above and below R' (in the y -direction, if they exist).*

Proof. Both claims are obvious; in (a) we use the fact that $\text{dist}_v(p, \ell) \leq \varepsilon\sqrt{2}$, and the assumption that $\varepsilon\sqrt{2} \leq \delta_1$; see below how this is enforced. In (b) we use the fact that $\text{dist}_v(p, \ell) = \text{dist}_v(\ell^*, p^*)$ and that the height of a small rectangle is $2\delta_1\delta_2 = 2\varepsilon > \varepsilon\sqrt{2}$. ◀

The algorithm. We first compute, for each point $p \in P$, the square S_i it belongs to; this can be done in $O(1)$ time, assuming a model of computation in which we can compute the floor function in constant time. Similarly, we find, for each line $\ell \in L$ the squares that it crosses, in $O(1/\delta_1)$ time. This gives us all the sets P_i, L_i , in overall $O(m + n/\delta_1)$ time.

We then iterate over the small squares in the partition of S . For each such square S_i , we construct the dual partitioning of the resulting dual rectangle R into the smaller rectangles R' . As above, we find, for each dual point ℓ^* , for $\ell \in L_i$, the small rectangle that contains it, and, for each dual line p^* , for $p \in P_i$, the small rectangles that it crosses. This takes $O(n_i + m_i/\delta_2)$ time.

We now report, for each small rectangle R' , all the pairs $(p, \ell) \in P_i \times L_i$ for which ℓ^* lies in R' and p^* crosses either R' or one of the small rectangles lying directly above or below R' (if they exist). We repeat this over all small squares S_i and all respective small rectangles R' . Note that a pair (p, ℓ) may be reported more than once in this procedure, but its multiplicity is at most some small absolute constant. The running time of this algorithm is

$$O\left(m + \frac{n}{\delta_1} + \sum_{i=1}^{1/\delta_1^2} \left(n_i + \frac{m_i}{\delta_2}\right) + k\right) = O\left(\frac{n}{\delta_1} + \frac{m}{\delta_2} + k\right),$$

where k is the number of pairs that we report. Lemma 1 guarantees that each reported pair is at distance $\leq 5\varepsilon$ and Lemma 2 guarantees that every pair (p, ℓ) at distance at most ε is reported.

We optimize the running time by choosing δ_1, δ_2 to satisfy $m/\delta_2 = n/\delta_1$ and $\delta_1\delta_2 = \varepsilon$. That is, we want to choose $\delta_1 = \sqrt{n\varepsilon/m}$ and $\delta_2 = \sqrt{m\varepsilon/n}$. These choices are effective, provided that both δ_1, δ_2 are at most 1, for otherwise the primal partition or the dual partitions does not exist. If $\delta_2 > 1$, that is, if $n < m\varepsilon$, we simply choose $\delta_1 = \varepsilon$, and run only the primal part of the algorithm, outputting all the pairs in $\bigcup_i P_i \times L_i$. The cost is now $O(m + n/\varepsilon + k) = O(m + k)$. (This is the naive implementation, which is now efficient since n is so small.) If $\delta_1 > 1$, we pass directly to the dual plane, flip the roles of P and L , and solve the problem in the naive manner just described, at the cost of $O(n + k)$. Otherwise (when both δ_1 and δ_2 are ≤ 1), the cost is $O(\sqrt{mn}/\sqrt{\varepsilon} + k)$. The cost of the algorithm is therefore always bounded by $O(n + m + \sqrt{mn}/\sqrt{\varepsilon} + k)$.

Recall also that in the proof of Lemma 2 we needed the inequality $\varepsilon\sqrt{2} \leq \delta_1$. This will hold when $m \leq n$ (and $\varepsilon \leq 1/2$, as we assume). In the complementary case $m > n$, we simply flip the roles of points and lines (that is, we start the analysis in the dual plane).

In conclusion, we have obtained the following main result of this section.

► **Theorem 3.** *Let P be a set of m points in the unit disk B in the plane, let L be a set of n lines that cross B , and let $0 < \varepsilon \leq 1/2$ be a prescribed parameter. We can report all pairs $(p, \ell) \in P \times L$, for which $\text{dist}(p, \ell) \leq \varepsilon$, in time $O(n + m + \sqrt{mn}/\sqrt{\varepsilon} + k)$, where k is the actual number of pairs that we report; all pairs at distance at most ε are reported, and every reported pair lies at distance at most 5ε .*

3 Review of the other algorithms

Near neighbors in point-plane configurations. Here we are given a set P of m points in the unit ball B in \mathbb{R}^3 , a set Π of n planes crossing B , and a prescribed error parameter $0 < \varepsilon \leq 1/2$. We solve the approximate incidences problem for P and Π with accuracy ε . As in the planar case, we approximate the point-plane distance $\text{dist}(p, \pi)$ by the z -vertical distance $\text{dist}_v(p, \pi)$. We partition Π into $O(1)$ subfamilies, according to the directions of the

normals, and treat each family separately, assuming that all the normal directions in this family are close to the z -direction, making the distance approximation behave well.

We assume that $P \subset S = [0, 1]^3$, and apply a two-stage partitioning, one in the primal space and one in the dual space, with a suitable choices for the corresponding parameters δ_1 , δ_2 , similar to the way it was done in the plane. We obtain an approximate incidence reporting algorithm that runs in $O(n + m + \sqrt{mn}/\varepsilon + k)$ time, where k is the actual number of pairs that we report.

Nearly congruent pairs in the plane. We are given two point sets P , Q , of respective sizes m and n , and parameters r , ε , and present an algorithm that reports all pairs $(p, q) \in P \times Q$ in the unit disk B , such that $|pq| \in [r - \varepsilon, r + \varepsilon]$, and each pair that it reports lie at distance in $[r - \alpha\varepsilon, r + \alpha\varepsilon]$, for some constant $\alpha > 1$. The problem is equivalent to an approximate incidences problem between P and the set of congruent circles $C := \{c_q \mid q \in Q\}$ where c_q is the circle of radius r centered at a point q . We assume that r is bounded away from 0 and that $\varepsilon \ll r \leq 1/2$.

We present two different solutions. The first one, inspired by an idea of Indyk et al. [13], does not use duality, so it is insensitive to cases where m and n differ significantly. The second solution does use duality, and is sensitive to such differences; it is more similar to the preceding solutions for the point-line and point-plane approximate incidences problems. We review here only the first solution.

We take the circle c_o of radius r centered at the origin o , and partition it into $2\pi/\sqrt{\varepsilon}$ equal canonical arcs, each with a central angle $\sqrt{\varepsilon}$. We replace each arc γ by a sector of an annulus A_γ of radii $r \pm \varepsilon$ that has γ as its ‘midline’, and enclose A_γ by a rectangle R_γ . Simple calculations show that the sides of R_γ are at most $\sqrt{\varepsilon} \times 3\varepsilon$.

We fix γ , and for each $q \in Q$ we translate R_γ to $R_\gamma(q) := q + R_\gamma$. We get a collection of n isothetic rectangles, and the m points of P . We tile up the unit disk by a grid whose cells are isothetic to R_γ , partition the points of P among the grid cells, and, for each $R_\gamma(q)$, report all pairs (p, q) such that p lies in one of the at most four cells that $R_\gamma(q)$ overlaps. We repeat this for each of the $O(1/\sqrt{\varepsilon})$ canonical arcs. The resulting algorithm runs in time $O((m + n)/\sqrt{\varepsilon} + k)$, where k is the actual number of pairs that we report. Our second approach, which uses duality, yields runs in $O(m + n + \sqrt{mn}/\sqrt{\varepsilon} + k)$ time, which is an improvement when m and n differ significantly.

Near-neighbor point-circle configurations. The duality-based approach can be extended to handle the approximate incidence reporting problem for points and arbitrary (rather than congruent) circles in the plane. The main difference is that general circles can be dualized into points in three dimensions, so our algorithm uses a standard grid decomposition in the primal plane, as in the cases of lines and congruent circles, but the dual partitionings take place in three dimensions, as in the case of planes.

To facilitate the second dual decomposition step, we replace the standard distance between points and circles by the *power* of a point with respect to a disk. We show that the distortion caused by this change is small, and our gain is that in the dual setup the points of P become planes (and the circles become points), so the machinery used for points and planes can be easily adapted to handle this case too. The algorithm runs in time $O(m + n + m^{1/3}n^{2/3}/\varepsilon^{2/3} + k)$, where k is the actual number of pairs that we report.

Reporting all nearly congruent pairs in three dimensions. Here we consider the three-dimensional version of the problem of nearly congruent pairs, where we are given sets P and

Q of m and n points, respectively, in the unit ball B in \mathbb{R}^3 , and parameters $0 < \varepsilon \ll r \leq 1/2$, and wish to report all pairs $(p, q) \in P \times Q$ such that $\text{dist}(p, q) \in [r - \varepsilon, r + \varepsilon]$, ensuring that each pair (p, q) that we report satisfies $\text{dist}(p, q) \in [r - \alpha\varepsilon, r + \alpha\varepsilon]$, for some absolute constant $\alpha > 1$. This is an approximate incidence reporting problem between P and spheres of radius r centered at the points of Q .

As before, we have two alternative solutions, one using the technique of Indyk et al. [13], and one using duality. Both extensions are reasonably routine, although some nontrivial technical issues have to be faced when extending the techniques to three dimensions. The first approach, runs in time $O((m+n)/\varepsilon + k)$. By using duality one can get a better bound (replacing $(m+n)/\varepsilon$ by \sqrt{mn}/ε) when the sizes of P and Q differ substantially.

Reporting all point-line neighbors in three dimensions. Let P be a set of m points in the unit ball B in three dimensions, let L be a set of n lines that cross B , and let $\varepsilon > 0$ be a given error parameter. We present an algorithm for the approximate incidence reporting problem involving P and L .

We represent each line in \mathbb{R}^3 by the pair of equations $y = ax + b, z = cx + d$. Let ℓ be the line $y = ax + b, z = cx + d$, and let $p = (\xi, \eta, \zeta) \in \mathbb{R}^3$. We approximate $\text{dist}(p, \ell)$ by slicing space by the plane $\pi_p : x = \xi$, and by computing the distance between the points p and $\ell_p := \ell \cap \pi_p = (\xi, a\xi + b, c\xi + d)$. As before, for this approximation to be good, the angle between ℓ and the x -direction should not be too large, say at most $\pi/4$, and we ensure this by partitioning L into $O(1)$ subfamilies, such that each subfamily has this property with respect to some direction u' . We focus on a single family, keep calling it L , and assume that u' is the x -axis. We show that $\text{dist}(p, \ell_p) \leq \sqrt{2}\text{dist}(p, \ell)$ for all $p \in P$ and $\ell \in L$.

We assume that P is contained in the unit cube $S = [0, 1]^3$, and apply the following two-stage partitioning procedure. For a pair of parameters δ_1, δ_2 , whose values will be set later we partition S into $1/\delta_1^3$ pairwise openly disjoint smaller cubes, each of side length δ_1 . For each small cube S_i , let P_i denote the set of all points of P that lie in S_i or in one of the (at most) eight cubes that surround S_i and have the same x -projection as S_i , and let L_i denote the set of all the lines of L that cross S_i . For each such small cube S_i , we pass to a parametric dual four-dimensional space, in which we represent each line $\ell \in L_i$, given by $y = ax + b, z = cx + d$, by the point $\ell^* = (a, b, c, d)$, and represent each point $p = (\xi, \eta, \zeta) \in P_i$ by the 2-plane (in \mathbb{R}^4) $p^* = \{(a, b, c, d) \mid a\xi + b = \eta, c\xi + d = \zeta\}$; p^* is the locus of all points dual to lines that pass through p .

We define the distance in the dual space between a point $\ell^* = (a, b, c, d)$ and a plane p^* , for a primal point $p = (\xi, \eta, \zeta)$, to be the distance between ℓ^* and the point $(a, \eta - a\xi, c, \zeta - c\xi)$, which is the intersection of p^* with the plane defined by $x = a$ and $z = c$. It follows that the distance between ℓ^* and p^* , as defined above, is equal to $\text{dist}(p, \ell_p)$ in the primal space.

Fix a small cube S_i , and assume without loss of generality that $S_i = [0, \delta_1]^3$. Let ℓ be a line in L_i , given by $y = ax + b, z = cx + d$. One can show that ℓ^* lies in the box R given by $-1 \leq a, c \leq 1$ and $-\delta_1 \leq b, d \leq 2\delta_1$. We now partition R into $1/\delta_2^4$ smaller boxes, each of which is a homothetic copy of R scaled down by δ_2 . Concretely, each smaller box R' is congruent to the box $[0, 2\delta_2] \times [0, 3\delta_1\delta_2] \times [0, 2\delta_2] \times [0, 3\delta_1\delta_2]$.

We then show that, for each small box R' , if $\ell^* = (a_\ell, b_\ell, c_\ell, d_\ell)$ is a dual point (of some $\ell \in L_i$) in R' and p^* is a dual plane (of some point $p = (\xi, \eta, \zeta) \in P_i$) that crosses R' or one of its surrounding boxes of the same xz -range, then $\text{dist}(p, \ell) \leq 8\sqrt{2}\delta_1\delta_2$. Conversely, if $\text{dist}(p, \ell) \leq \delta_1\delta_2$ then (p, ℓ) belong to some subproblem $P_j \times L_j$, and p^* crosses the small dual region R' containing ℓ^* or one of its nearby regions.

The algorithm is now immediate: We compute the sets P_i, L_i , for $i = 1, \dots, 1/\delta_1^3$, in overall $O(m + n/\delta_1)$ time. Then, for each small cube S_i , we consider the partitioning of the

resulting dual box R into the smaller boxes R' . As above, we find, for each $\ell \in L_i$, the small region that contains the dual point ℓ^* , and, for each $p \in P_i$, the small regions that the dual plane p^* crosses. We report, for each small region R' , all the pairs $(p, \ell) \in P_i \times L_i$ for which ℓ^* lies in R' and p^* crosses either R' or one of the at most eight small regions that surround R' and have the same xz -range. We repeat this over all small cubes S_i and all respective small regions R' . With a suitable optimization of the values of δ_1 and δ_2 , the running time is $O\left(m + n + m^{1/3}n^{2/3}/\varepsilon^{2/3} + k\right)$.

Reporting all point-circle neighbors in three dimensions. In preparation for the final algorithm, that finds all nearly congruent copies of a given triangle in a set of n points in \mathbb{R}^3 , we first solve the following problem. Let P be a set of m points in the unit ball B in \mathbb{R}^3 , let C be a set of n congruent circles in \mathbb{R}^3 of radius $r \leq 1/2$ that cross B , and let $\varepsilon \ll r$ be a prescribed error parameter. We present an efficient algorithm for the approximate incidence reporting problem for P and C .

This is perhaps the most complex algorithm in our collection. We slice each circle into canonical arcs, replace each arc by a sector of a torus of width ε around it, enclose each torus sector by a suitable (bounded) cylinder, and reduce our problem to that of reporting point-cylinder containments. We further reduce the problem by cutting space by parallel slabs of width $\sqrt{\varepsilon}$ in some suitable direction, say the x -direction, by partitioning the points of P among the slabs, and by considering only those toric/cylindrical pieces that form sufficiently small angle with the x -direction. For each such slab σ , we take the points in σ , replace each cylinder that intersects σ , or a nearby slab, by the full line that supports its axis, and run the approximate incidence reporting algorithm involving the points in the slab and the lines associated with the slab, repeating this over all slabs and tori sectors. The resulting algorithm runs in time $O\left((m+n)/\varepsilon^{1/2} + m^{1/3}n^{2/3}/\varepsilon^{7/6} + k\right)$, where k is the number of (distinct) reported pairs.

4 Reporting all nearly congruent triangles

In this section we put to work the algorithms in (e) and (g) (see Section 1), to obtain an efficient solution of the first step in solving the approximate point pattern matching problem in \mathbb{R}^3 (see its review in the introduction), where we are given a sampled “reference” triangle Δabc , for a triple of points a, b, c in the first set A , and a prescribed error parameter $\varepsilon > 0$. Our goal is to report all triples p, q, o in the second set B that span a triangle “nearly congruent” to Δ ; that is, triples that satisfy

$$||pq| - |ab|| \leq \varepsilon, \quad ||po| - |ac|| \leq \varepsilon, \quad \text{and} \quad ||qo| - |bc|| \leq \varepsilon. \quad (1)$$

We allow to report triples that satisfy (1) with $\alpha\varepsilon$ on the right-hand sides rather than ε , for some fixed constant α . Let ab be the longest edge of Δ . We require that $\beta \leq |ab| \leq 1/2$ for some fixed constant β . We also require that the height h of Δ from c (perpendicular to ab) is larger than some fixed constant s . We assume that $\beta, s \gg \varepsilon$. Our approximation guarantee α increases as β and s decrease.

We first report all pairs $(p, q) \in B^2$ such that $||pq| - |ab|| \leq \varepsilon$, using the algorithm specified in (e) (incidences between congruent spheres and points). This takes $O(n/\varepsilon + N)$ time, where N is the number of pairs that we report. Let Π denote the set of reported pairs. We know that all the desired pairs are included in Π , and that every pair (p, q) in Π satisfies $||pq| - |ab|| \leq \alpha'\varepsilon$, for some absolute constant α' . We prune Π , leaving in it only pairs (p, q) satisfying $||pq| - |ab|| \leq \varepsilon$. We continue to denote the resulting set as Π , and its size by N .

Let (p, q) be a pair in Π . Any point o that satisfies $||po| - |ac|| \leq \varepsilon$ and $||qo| - |bc|| \leq \varepsilon$ lies in the intersection $K = K_{p,q}$ of two spherical shells, one centered at p with radii $|ac| \pm \varepsilon$, and one centered at q with radii $|bc| \pm \varepsilon$. The following lemma allows us to replace K by a torus that is congruent to a fixed torus that depends only on Δ . See Figure 2.

► **Lemma 4.** *Assume that Δ is sufficiently fat, in the sense that $\beta \leq |ab| \leq 1/2$ and $h \geq s$, for some absolute positive constants β, s that satisfy $\varepsilon \ll \beta, s$. Then there exists a circle $\gamma_{p,q}$ of radius h such that K is contained in the torus $T_{p,q}$ that is the Minkowski sum of $\gamma_{p,q}$ and a ball of radius $\varepsilon' \leq \delta\varepsilon$ around the origin, where the constant δ depends on β and s .*

Proof. Denote the lengths of the edges of the triangle Δabc by $u = |ab|$, $v = |ac|$ and $w = |bc|$. Let g the point where h meets ab and let $z = |ag|$. We have $z^2 + h^2 = v^2$ and $(u - z)^2 + h^2 = w^2$, from which we obtain that $z = \frac{u^2 + v^2 - w^2}{2u}$, and we denote this expression as $z = z(u, v, w)$. Consider an alignment of Δ within the plane of Δpqo , such that a coincides with p and ab overlaps pq . Let g now be a point on pq at distance z from $p = a$. Then c lies on the circle $\gamma_{p,q}$ of radius h , centered at g , and contained in the plane perpendicular to pq through g . See Figure 2(b).

Fix some point $o \in K$. We claim that o must be at distance $\leq \delta\varepsilon$ from $\gamma_{p,q}$, for some fixed constant δ that depends on β and s . Indeed, since $(p, q) \in \Pi$ and $o \in K$, we can write $|pq| = u + \varepsilon_1$, $|po| = v + \varepsilon_2$, and $|qo| = w + \varepsilon_3$, where $|\varepsilon_i| \leq \varepsilon$ for $i = 1, 2, 3$.

Consider the alignment of Δ with Δpqo , as above, and imagine that we perturb the edges ab , ac , and bc of Δ by ε_1 , ε_2 , and ε_3 , respectively, so that Δ is continuously deformed into Δpqo . We claim that o cannot move too far as a result of this deformation so the distance between o and c must be small.

To see this, let h' be the height of Δpqo from o , let g' be the point at which h' meets pq , and let $z' = |pg'|$. We claim that $|z' - z| \leq \delta\varepsilon$ and $|h' - h| \leq \delta\varepsilon$ for some absolute constant δ . To see this, using the function $z = z(u, v, w)$ defined above, we have $z' = z(u + \varepsilon_1, v + \varepsilon_2, w + \varepsilon_3)$, and routine calculations show that, for ε sufficiently small, we have $|z' - z| = O(|\nabla z(u, v, w) \cdot (\varepsilon_1, \varepsilon_2, \varepsilon_3)|) \leq \delta'\varepsilon$, where δ' depends on β .

Similarly, by Heron's formula, we can think of h as a function $h(u, v, w)$, given by

$$h(u, v, w) = \frac{2\text{Area}(\Delta)}{u} = \frac{2\sqrt{\tau(\tau - u)(\tau - v)(\tau - w)}}{u},$$

where $\tau = \frac{1}{2}(u + v + w)$. Then $h' = h(u + \varepsilon_1, v + \varepsilon_2, w + \varepsilon_3)$, and, by another routine calculation, $|h' - h| = O(|\nabla h(u, v, w) \cdot (\varepsilon_1, \varepsilon_2, \varepsilon_3)|) \leq \delta''\varepsilon$, for another constant δ'' that depends on β and s . Take $\delta = \sqrt{(\delta')^2 + (\delta'')^2}$, and the lemma follows. ◀

We have thus reached the following scenario. We have a set \mathcal{T} of N congruent tori $T_{p,q}$, for $(p, q) \in \Pi$, and a set B (the original one) of n points. By construction, each triple (p, q, o) that defines a triangle for which (1) holds, satisfies $o \in T_{p,q}$. Using our algorithm for point-circle near neighbors in \mathbb{R}^3 , as reviewed in Section 3, we can report all the triples (p, q, o) such that $o \in T_{p,q}$, in time $O(n + N/\varepsilon^{1/2} + n^{1/3}N^{2/3}/\varepsilon^{7/6} + k)$, where k is the number of (distinct) triples that we report; each of the desired triples is reported, and each triple that we report is such that the distance from o to $\gamma_{p,q}$ is at most $\alpha\varepsilon$ for some other fixed constant $\alpha > \delta$. Therefore each triple which we report satisfies (1) with $\alpha\varepsilon$ on the right-hand sides, rather than ε . In summary, we have:

► **Theorem 5.** *Let B be a set of n points in the unit ball in \mathbb{R}^3 . Let Δabc be a fixed reference triangle and let ε an error parameter, so that Δ and ε satisfy the constraints specified in Lemma 4. We can then report all triples $(p, q, o) \in B^3$ that span a triangle nearly congruent to*

Δ , in the sense of (1), in time $\left(n + N/\varepsilon^{1/2} + n^{1/3}N^{2/3}/\varepsilon^{7/6} + k\right)$, where N is the number of pairs reported by our algorithm for approximate congruent pairs in \mathbb{R}^3 (reviewed in Section 3), applied to P with distance $|ab|$, the largest edge length of Δ , and k is the number of (distinct) triples that the algorithm in this section reports; each of the desired triples is reported, and each triple that we report satisfies (1) with $\alpha\varepsilon$ replacing ε , where α is a suitable absolute constant. Each pair is reported at most $O(1)$ times.

References

- 1 P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry, Contemp. Math. 223*, Amer. Math. Soc. Press, Providence, RI, pages 1–56, 1999.
- 2 D. Aiger and K. Kedem. Approximate input sensitive algorithms for point pattern matching. *Pattern Recognition*, 43(1):153–159, 2010.
- 3 D. Aiger, N. J. Mitra, and D. Cohen-Or. 4-points congruent sets for robust pairwise surface registration. *ACM Trans. Graphics*, 27(3):Article 85, 2008.
- 4 P. Brass. Combinatorial geometry problems in pattern recognition. *Discrete Comput. Geom.*, 28(4):495–510, 2002.
- 5 O. Chum and J. Matas. Optimal randomized RANSAC. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(8):1472–1482, 2008.
- 6 *IEEE Int'l Workshop "25 Years of RANSAC" in conjunction with CVPR'06 (RANSAC25'06)*. IEEE Computer Society, 2006.
- 7 G.D. da Fonseca and D.M. Mount. Approximate range searching: The absolute model. *Comput. Geom. Theory Appl.*, 43(4):434–444, 2010.
- 8 J. Erickson. New lower bounds for Hopcroft's problem. *Discrete Comput. Geom.*, 16(4):389–418, 1996.
- 9 M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1982.
- 10 M. Gavrilov, P. Indyk, R. Motwani, and S. Venkatasubramanian. Combinatorial and experimental methods for approximate point pattern matching. *Algorithmica*, 38(1):59–90, 2003.
- 11 P. J. Heffernan and S. Schirra. Approximate decision algorithms for point set congruence. *Comput. Geom. Theory Appl.*, 1(4):137–156, 1994.
- 12 D. P. Huttenlocher and S. Ullman. Recognizing solid objects by alignment with an image. *Internat. J. Computer Vision*, 5(2):195–212, 1990.
- 13 P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric matching under noise: Combinatorial bounds and algorithms. *Proc. 10th ACM–SIAM Sympos. Discrete Algorithms*, pages 457–465, 1994.
- 14 N. Mellado, D. Aiger, and N. J. Mitra. Super 4pcs: fast global pointcloud registration via smart indexing. *Comput. Graphics Forum*, 33(5):205–215, 2014.