ELSEVIER

Information Processing Letters 64 (1997) 107–114

Information Processing Letters

Exact and approximate computational geometry solutions of an unrestricted point set stereo matching problem [1]

Frank Dehne [a,*], Katia Guimarães [b,2]

[a] *School of Computer Science, Carleton University, Ottawa, Canada K1S 5B6*

[b] *Departamento de Informática, Univ. Federal de Pernambuco, 50732-970 Recife, Pe, Brazil*

Received 1 October 1995; revised 1 October 1997
Communicated by S.G. Akl

Abstract

In this paper we study the problem of computing an exact, or arbitrarily close to exact, solution of an unrestricted point set stereo matching problem. Within the context of classical approaches like the Marr–Poggio algorithm, this means that we study how to solve the unrestricted basic subproblems created within such approaches, possibly yielding an improved overall performance of such methods.

We present an $O(n^{2+4k})$ time and $O(n^4)$ space algorithm for *exact* unrestricted stereo matching, where n represents the number of points in each set and k the number of depth levels considered. We generalize the notion of a δ-approximate solution for point set congruence to the stereo matching problem and present an $O((\varepsilon/\delta)^k n^{2+2k})$ time and $O((\varepsilon/\delta)n^2)$ space δ-*approximate* algorithm for unrestricted stereo matching (ε represents measurement inaccuracies in the image). We introduce new computational geometry tools for stereo matching: the *translation square arrangement*, *approximate translation square arrangement* and *approximate stereo matching tree*. © 1997 Published by Elsevier Science B.V.

Keywords: Algorithms; Computational geometry

1. Introduction

Binocular stereo is a technique used in machine vision for creating depth perception from two 2-D images (*views*) recorded from different angles. The difference in location of the same object in the two views, also called *binocular disparity*, determines the depth (i.e. location in the third dimension) of an object. The primary computational problem of binocular stereo, called the *stereo matching problem*, is to identify corresponding objects in the two views (see e.g. [2]). The stereo matching problem has been extensively studied in the *vision* literature, see e.g. [2–7,10–12,14–16,18]. Considerable investigation has been dedicated to the *random-dot stereogram* which consists of two synthetic images of uncorrelated dots that happen to be two views of the same surface. The question studied is to retrieve depth information in such a stereogram which contains no other cues besides location (in contrast to, e.g., knowledge-based approaches which also use shape, intensity, etc.).

* Corresponding author. Email: dehne@scs.carleton.ca. Research partially supported by the Natural Sciences and Engineering Research Council of Canada, CNPq (Brasil) and FACEPE (Pernambuco).

[1] This work was done while the first author was visiting the Universidade Federal de Pernambuco, Brasil.

[2] Email: katia@di.ufpe.br.

PII S0020-0190(97)00168-3

The majority of solutions to the stereo matching problem presented in the literature use order preserving restrictions, heuristics, simulated annealing or other optimization methods. A classical approach is the Marr–Poggio algorithm [10–12]. It uses several versions of the image with increasing resolution, and solves the stereo matching problem for each version by using the results obtained for the previous versions together with order preserving restrictions. The problem is reduced to a large number of unrestricted stereo matching problems for considerably smaller subsets and considerably fewer possible depth levels. We will refer to these problems as the *basic subproblems*. The Marr–Poggio and related algorithms solve the unrestricted basic subproblems by brute force.

In this paper we study the problem of computing an exact, or arbitrarily close to exact, solution for unrestricted point set stereo matching. Within the context of stereo matching approaches like the Marr–Poggio algorithm this means that we study how to solve the unrestricted basic subproblems faster than by brute force and thereby obtain an improved overall performance. We show that an unrestricted basic subproblem can be solved in polynomial time without any loss in accuracy. An arbitrarily close to exact solution can be found with a further improved time complexity.

The following is the formal definition of the unrestricted point set stereo matching problem studied in the remainder.

Let $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ be two sets of points in $\mathbb{R}^2$ representing two views. A *(k, ε)-stereo matching*, Γ, for A and B, where $\varepsilon > 0$ and $1 \leqslant k \leqslant n$, is comprised of (a) a partitioning of A and B into k subsets $A_1, \dots, A_k$, and $B_1, \dots, B_k$, respectively, (b) k bijections $l_i : B_i \to A_i$, $1 \leqslant i \leqslant k$, (called *labelings*) and (c) k *translation vectors* $t_1, \dots, t_k$ with the following property: For all $1 \leqslant i \leqslant k$ and $b \in B_i$, $\mathrm{dist}(b + t_i, l_i(b)) \leqslant \varepsilon$. All distances $\mathrm{dist}(.,.)$ will be measured with respect to the L_∞ metric. The *(k, ε)-stereo matching problem* for A and B consists of finding such a stereo matching, if it exists.

In the above definition, k corresponds to the number of depth levels to be considered, $A_1, \dots, A_k$ and $B_1, \dots, B_k$ represent the points in the different depth layers, and $l_i : B_i \to A_i$, $1 \leqslant i \leqslant k$ the matchings between those points. Due to inaccuracy in measurings, some degree of noise tolerance is necessary, which is represented by the factor ε.

Note that, in general, the solution is not unique. In many real life applications, the preferred matching is indicated by some additional constraints. One such constraint could be, for example, to minimize k.

In Section 2 of this paper we present an algorithm for solving the (k, ε)-stereo matching problem in time $O(n^{2+4k})$ and space $O(n^4)$. Our polynomial time solution is without any loss in accuracy.

The (k, ε)-stereo matching problem for the special case of $k = 1$ is also called the *point set congruence* problem, for which an $O(n^6)$ time algorithm was presented in [1].

Let ε_0 be the smallest ε such that A and B have a (k, ε)-stereo matching. Define the *δ-approximate (k, ε)-stereo matching problem*, $0 < \delta < \varepsilon$, as follows: If $|\varepsilon_0 - \varepsilon| > \delta$, find a (k, ε)-stereo matching for A and B (if exists), otherwise return "don't know". The δ-approximate (k, ε)-stereo matching problem is a relaxed version of stereo matching, where our algorithm has to produce a correct result only for values of ε outside the δ-range of the threshold value ε_0. In general, this relaxation is not critical in practice, but it will allow for a considerable speedup of the algorithm.

In Section 3 of this paper we present an algorithm for solving the δ-approximate (k, ε)-stereo matching problem in time $O((\varepsilon/\delta)^k n^{2+2k})$ and space $O((\varepsilon/\delta) n^2))$. Again, for $|\varepsilon_0 - \varepsilon| > \delta$, our solution is without any loss in accuracy.

The special case $k = 1$, called *approximate point set congruence*, was introduced in [17] and subsequently also studied in [8]. $O((\varepsilon/\delta)^2 n^{2.5})$ and $O((\varepsilon/\delta)^6 n^3)$ time algorithms, respectively, were presented.

Our methods for the general (k, ε)-stereo matching problem are a nontrivial generalization of [8]. The algorithms in [17] and [8] are based on using the centroids and lower left corners, respectively, of the two sets A and B. This is not possible for the general (k, ε)-stereo matching problem, as the partitioning into subsets $A_1, \dots, A_k$, and $B_1, \dots, B_k$ is not given a priori (while for $k = 1$ the partitioning is trivial). Computing this partitioning of A and B is the major additional problem for stereo matching. We present a new tool, the *translation square arrangement*, and a new approximation scheme yielding an *approximate translation square arrangement*. Based on this arrangement, we build an *approximate stereo matching tree* which is the main data structure guiding our algorithm. Our approximation scheme has subtle

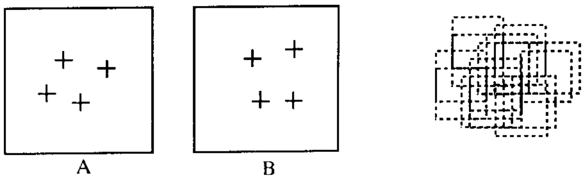

Fig. 1. Two points sets A and B and their translation square arrangement.

but important differences from [8]. It allows us to prove some combinatorial properties of the approximate stereo matching tree and relationships between approximations and exact solutions which yield the obtained results.

2. (k,ε)-stereo matching

For

$$a = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \in A \quad \text{and} \quad b = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} \in B,$$

the square with side length 2ε centered at

$$a - b = \begin{pmatrix} a_1 - b_1 \\ a_2 - b_2 \end{pmatrix}$$

is called the *translation square* $s_\varepsilon(a,b)$. Notice that the square $s_\varepsilon(a,b)$ represents all translation vectors t_i such that $\text{dist}(b + t_i, a) \leqslant \varepsilon$. We call $S_\varepsilon = \{s_\varepsilon(a,b) \mid a \in A, b \in B\}$ the *translation square arrangement* (see Fig. 1). Note that $|S_\varepsilon| = n^2$.

Let R_{S_ε} denote the set of regions in the plane created by S_ε, where all regions not intersecting any $s \in S_\varepsilon$ are identified as one single region called the *external region*.

Observation 1. $|R_{S_\varepsilon}| = O(|S_\varepsilon|^2) = O(n^4)$.

Let $G^{\text{adj}}_{S_\varepsilon} = (R_{S_\varepsilon}, E^{\text{adj}}_{S_\varepsilon})$ be the *adjacency graph* of R_{S_ε} with vertex set R_{S_ε} and edge set $E^{\text{adj}}_{S_\varepsilon}$ connecting all pairs (r_1, r_2) of adjacent regions of R_{S_ε} that share a common boundary edge (not just a vertex). Since $G^{\text{adj}}_{S_\varepsilon}$ is planar (except for the subgraph induced by the node representing the external region) and $G^{\text{adj}}_{S_\varepsilon}$ can be constructed by a standard plane sweep (see e.g. [13]), we observe the following:

Observation 2. *$|G^{\text{adj}}_{S_\varepsilon}| = O(|R_{S_\varepsilon}|) = O(n^4)$, and $G^{\text{adj}}_{S_\varepsilon}$ can be constructed in time $O(n^4 \log n)$.*

For a region $r \in R_{S_\varepsilon}$, let

$$E^{\text{cand}}_{S_\varepsilon}(r) = \{(a,b) \in A \times B \mid r \subset s_\varepsilon(a,b)\}$$

denote the *candidate edges* induced by r. For k regions $r_1, \ldots, r_k$ define the *translation square graph* $G^{\text{tsg}}_{S_\varepsilon}(r_1, \ldots, r_k) = (A \cup B, E^{\text{cand}}_{S_\varepsilon}(r_1) \cup E^{\text{cand}}_{S_\varepsilon}(r_2) \cup \cdots \cup E^{\text{cand}}_{S_\varepsilon}(r_k))$. Note that $G^{\text{tsg}}_{S_\varepsilon}(r_1, \ldots, r_k)$ is a bipartite graph.

We now discuss the relationship between the translation square arrangement and (k,ε)-stereo matchings. Consider k regions $r_1, \ldots, r_k$ of R_{S_ε} and assume that the translation square graph $G^{\text{tsg}}_{S_\varepsilon}(r_1, \ldots, r_k)$ has a perfect matching π. In such a case, the following defines the (k,ε)-stereo matching *induced* by $(r_1, \ldots, r_k)$ and π, referred to as $\Gamma^\pi_{S_\varepsilon}(r_1, \ldots, r_k)$: The subsets A_i and B_i $(1 \leqslant i \leqslant k)$ are the sets of point $a \in A$ and $b \in B$, respectively, incident to those edges of the perfect matching π that are contained in the edge set $E^{\text{cand}}_{S_\varepsilon}(r_i)$. The edges of $E^{\text{cand}}_{S_\varepsilon}(r_i)$ selected in the perfect matching π define labeling l_i, and translation vector t_i is the centroid (center of gravity) of r_i, $1 \leqslant i \leqslant k$. Note that, any point inside r_i could be selected as a possible translation vector t_i.

For the remainder, we will omit the subscript ε if it is obvious from the context.

Lemma 3. *Every (k,ε)-stereo matching of A and B corresponds to a k-tuple $r_1, \ldots, r_k$ of regions of R_S such that the translation square graph $G^{\text{tsg}}_S(r_1, \ldots, r_k)$ has a perfect matching, and vice versa.*

Proof. Given k regions $r_1, \ldots, r_k$ of R_S such that $G^{\text{tsg}}_S(r_1, \ldots, r_k)$ has a perfect matching π, then $\Gamma^\pi_S(r_1, \ldots, r_k)$ is a (k,ε)-stereo matching. On the other hand, consider a (k,ε)-stereo matching Γ. Denote with $r_1, r_2, \ldots, r_k$ the k regions of R_S which contain the k translation vectors of Γ. All labelings l_i of Γ, taken together, correspond to a subset of edges (a,b) in $G^{\text{tsg}}_S(r_1, \ldots, r_k)$ which is a perfect matching for $G^{\text{tsg}}_S(r_1, \ldots, r_k)$. □

Recall the following results from [9] on computing maximum matchings.

Lemma 4 (Hopcroft and Karp [9]).

(a) *A maximum matching in a bipartite graph with n vertices and e edges can be found in time* $O(n^{2.5})$.

(b) *Given a maximum matching in a graph G with n vertices and e edges, and another graph G′ with the same vertices which differs from G by at most one edge, then a maximum matching for G′ can be found in time* $O(e)$.

After the following definition, we are ready to present an algorithm for computing a (k,ε)-stereo matching for two point sets A and B.

For any sequence $\phi = (x_1,\dots,x_M)$ let $\phi^{[k]}$ be the sequence of k-tuples of elements of ϕ in lexicographic order, i.e. $\phi^{[k]} = ((x_1,\dots,x_1,x_1,x_1), (x_1,\dots,x_1,x_1,x_2), \dots, (x_1,\dots,x_1,x_1,x_M), (x_1,\dots,x_1,x_2,x_1), \dots, (x_M,\dots,x_M,x_M,x_M))$.

Algorithm 1. *(k,ε)-stereo matching*

1. Compute S, R_S, and the *adjacency graph* G_S^{adj}.
2. Consider any (e.g. depth-first) traversal of $G_{\mathrm{adj}}(R_S)$ starting at some vertex v, traversing the entire graph $G_{\mathrm{adj}}(R_S)$, and returning to the same vertex v. This induces a sequence ϕ of regions $r_i \in R_S$ which contains each region at least once. For each k-tuple of regions $(r_1,\dots,r_k) \in \phi^{[k]}$ determine if the respective translation square graph $G_S^{\mathrm{tsg}}(r_1,\dots,r_k)$ has a perfect matching. Compute a maximum matching for the first k-tuple by using Lemma 4(a) and for all subsequent k-tuples by using Lemma 4(b).
3. If, in Step 2, a graph $G_S^{\mathrm{tsg}}(r_1,\dots,r_k)$ with perfect matching π has been found, report the (k,ε)-stereo matching $\Gamma_S^{\pi}(r_1,\dots,r_k)$ induced by $(r_1,\dots,r_k)$ and π; otherwise report that no (k,ε)-stereo matching exists.

Theorem 5. *The (k,ε)-stereo matching problem for two point sets A and B with n points each, $\varepsilon > 0$, can be solved in time $O(n^{2+4k})$ and space $O(n^4)$.*

Proof. The correctness of Algorithm 1 follows from Lemma 3. The time complexity of Step 1 is $O(n^4 \log n)$, see Observation 2. Note that, when traversing G_S^{adj} in Step 2, any particular region might be traversed several times. However, the total number of regions traversed in ϕ is $O(|R_S|) = O(n^4)$, see Observation 2. Hence, the number of k-tuples of regions enumerated in sequence $\phi^{[k]}$ of Step 2 is $O(n^{4k})$. The translation square graphs for two consecutive k-tuples of regions differ in one edge, and the size of each translation square graph is $O(n^2)$. Therefore, the time for computing/maintaining the maximum matchings for all k-tuples of regions is $O(n^{2+4k})$. The time for Step 3 is dominated by the time for Step 2. Thus, the claimed time complexity for Algorithm 1 follows. The space requirement of Algorithm 1 is determined by the space for the adjacency graph G_S^{adj} which is $O(n^4)$, see Observation 2. □

3. Approximate (k,ε)-stereo matching

Let ε_0 be the smallest ε such that A and B have a (k,ε)-stereo matching. We recall the definition of the *δ-approximate (k,ε)-stereo matching problem*, $0 < \delta < \varepsilon$, given in Section 1: If $|\varepsilon_0 - \varepsilon| > \delta$, find a (k,ε)-stereo matching for A and B (if exists), otherwise return "don't know".

For ease of description we assume that ε/δ is an integer. However, all results presented in the remainder hold for any $0 < \delta < \varepsilon$.

We define the *approximate translation square arrangement*, S_ε^* as follows: Let the *δ-grid* be a grid of horizontal and vertical lines of distance δ, respectively, covering the entire plane. The line crossings are called *gridpoints*. For each translation square $s_\varepsilon(a,b) \in S_\varepsilon$ define as its *approximation* $s_\varepsilon^*(a,b)$ the square obtained from $s_\varepsilon(a,b)$ by moving its center to the closest gridpoint. Let $S_\varepsilon^* = \{s_\varepsilon^*(a,b) \mid s_\varepsilon(a,b) \in S_\varepsilon\}$ be the *approximate translation square arrangement*.

Note that several translation squares $s_\varepsilon(a,b)$ might have the same approximation. For each $s^* \in S_\varepsilon^*$ we define its *multiplicity* $m(s^*) = |\{s_\varepsilon(a,b) \in S_\varepsilon \mid s_\varepsilon^*(a,b) = s^*\}|$.

Analogously to Section 2 we define for S_ε^* (instead of S_ε) its set of regions $R_{S_\varepsilon^*}$, adjacency graph $G_{S_\varepsilon^*}^{\mathrm{adj}}$, the candidate edges $E_{S_\varepsilon^*}^{\mathrm{cand}}(r)$ induced by $r \in R_{S_\varepsilon^*}$, the translation square graph $G_{S_\varepsilon^*}^{\mathrm{tsg}}(r_1,\dots,r_k)$ for $r_1, r_2, \dots, r_k \in R_{S_\varepsilon^*}$, and the *induced* (k,ε)-stereo matching $\Gamma_{S_\varepsilon^*}^{\pi}(r_1,\dots,r_k)$ if $G_{S_\varepsilon^*}^{\mathrm{tsg}}(r_1,\dots,r_k)$ has a perfect matching π.

The main advantage of the approximate translation square arrangement S_ε^* is that it has a considerably smaller number of regions.

Lemma 6. $|R_{S_\varepsilon^*}| = O((\varepsilon/\delta)|S|) = O((\varepsilon/\delta)n^2)$.

Proof. Let $\mathcal{B}$ be the set of regions of the δ-grid adjacent to the border of some $s \in S_\varepsilon^*$. Since each region $r \in R_{S_\varepsilon^*}$ contains a (distinct) region of $\mathcal{B}$, it follows that $|R_{S_\varepsilon^*}| \leqslant |\mathcal{B}|$. On the other hand, each $s \in S_\varepsilon^*$ has at most $O(\varepsilon/\delta)$ regions of the δ-grid adjacent to its border, which implies that $|\mathcal{B}| = O((\varepsilon/\delta)n^2)$. □

If $G_{S_\varepsilon^*}^{\mathrm{tsg}}(r_1,\dots,r_k)$ has a perfect matching π for some k regions $r_1,\dots,r_k \in R_{S_\varepsilon^*}$ of the approximate translation square arrangement S_ε^*, then $\Gamma_{S_\varepsilon^*}^{\pi}(r_1,\dots,r_k)$ is called a *(k,ε,δ)-stereo matching approximation* for A and B.

In the remainder of this section we first show how to efficiently compute a (k,ε,δ)-stereo matching approximation (if exists), and then study how a stereo matching approximation can be used to solve the δ-approximate stereo matching problem. Note that the latter problem requires an "*exact*" (k,ε)-stereo matching to be reported if $|\varepsilon_0 - \varepsilon| > \delta$.

When it is obvious from the context, we will omit the subscript ε.

3.1. Computing a (k,ε,δ)-stereo matching approximation

Analogously to Section 2 we make the following observations.

Observation 7.
(a) $|G_{S^*}^{\mathrm{adj}}| = O(|R_{S^*}|) = O((\varepsilon/\delta)n^2)$.
(b) *$G_{S^*}^{\mathrm{adj}}$ can be constructed in time* $O((\varepsilon/\delta)n^2 \log n)$.

Let $\mathcal{T}^*$ be a spanning tree of the adjacency graph $G_{S^*}^{\mathrm{adj}}$, rooted at the node representing the external region. For each node of $\mathcal{T}^*$ representing a region $r \in R_{S^*}$, we will store the centroid of r and the set $E_{S_\varepsilon^*}^{\mathrm{cand}}(r)$ of candidate edges induced by r. Storing all sets $E_{S_\varepsilon^*}^{\mathrm{cand}}(r)$ explicitly requires space $O((\varepsilon/\delta)n^4)$. In order to reduce this memory size we will represent all sets $E_{S_\varepsilon^*}^{\mathrm{cand}}(r)$ in an incremental way by storing $E_{S_\varepsilon^*}^{\mathrm{cand}}(r)$ explicitly only for the root of $\mathcal{T}^*$, and for each edge $e = (r_1,r_2)$ of $\mathcal{T}^*$ the change between $E_{S_\varepsilon^*}^{\mathrm{cand}}(r_1)$ and $E_{S_\varepsilon^*}^{\mathrm{cand}}(r_2)$, referred to as $\Delta(e) = \Delta(r_1,r_2)$. The tree $\mathcal{T}^*$ together with the centroids and incremental representation of the candidate edges is called the *approximate stereo matching tree $\mathcal{T}^*$*. (See Fig. 2.) It will be our main tool for computing a stereo matching approximation.

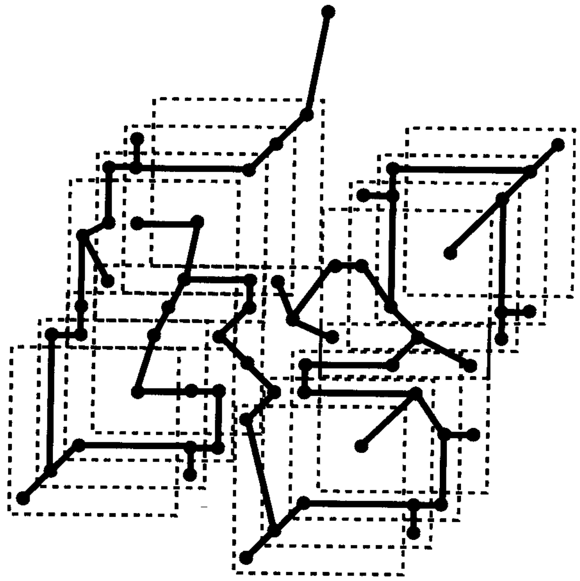

Fig. 2. An approximate translation square arrangement and its approximate stereo matching tree.

Lemma 8.

$$\sum_{e\in\mathcal{T}^*} |\Delta(e)| = O\left(\frac{\varepsilon}{\delta}n^2\right).$$

Proof. For each edge $e = (r_1,r_2)$ of $\mathcal{T}^*$ let its *weight* $w(e)$ be the sum of all multiplicities $m(s)$, $s \in S^*$, such that e crosses the boundary of s (i.e., r_1 and r_2 are on different sides of the boundary of s). The weight $w(e)$ is equal to the number of candidate edge changes, $|\Delta(e)|$. Let

$$w(\mathcal{T}^*) = \sum_{[e\in\mathcal{T}^*]} w(e) = \sum_{[e\in\mathcal{T}^*]} |\Delta(e)|$$

denote the total weight of $\mathcal{T}^*$. We will show that $w(\mathcal{T}^*) = O((\varepsilon/\delta)n^2)$.

For each region $r \in R_{S^*}$, let

$$w(r) = \sum\nolimits_{[e \text{ edge of } \mathcal{T}^* \text{ incident with } r]} w(e),$$

and for each square $s \in S^*$, let $w(s)$ be the number of all edges $e \in \mathcal{T}^*$ that cross the boundary of s. It follows that

$$w(\mathcal{T}^*) = O\Big(\sum_{r\in R_{S^*}} w(r)\Big) = O\Big(\sum_{s\in S^*} m(s)w(s)\Big).$$

For an $s \in S^*$ and edge $e = (r_1, r_2) \in \mathcal{T}^*$ crossing the boundary of s, r_1 or r_2 must be adjacent to the boundary of s. For each $s \in S^*$, the number of regions adjacent to the boundary of s is $O(\varepsilon/\delta)$. Hence, $w(s) = O(\varepsilon/\delta)$, and $w(\mathcal{T}^*) = O((\varepsilon/\delta)\sum_{|s\in S^*|} m(s))$. Since $\sum_{|s\in S^*|} m(s) = |S| = O(n^2)$, it follows that $w(\mathcal{T}^*) = O((\varepsilon/\delta)n^2)$. □

Corollary 9. *The approximate stereo matching tree $\mathcal{T}^*$ (together with the incremental representation of candidate edges) is of size $O((\varepsilon/\delta)n^2)$, and can be built in time $O((\varepsilon/\delta)n^2\log(\varepsilon/\delta n))$.*

Algorithm 2. *Computing a (k,ε,δ)-stereo matching approximation*

1. Compute S^*, R_{S^*}, $G^{\text{adj}}_{S^*}$, and the *approximate stereo matching tree* $\mathcal{T}^*$ (together with the incremental representation of candidate edges).
2. Consider a depth-first traversal of $\mathcal{T}^*$, starting at the root, traversing $\mathcal{T}^*$, and returning to the root. This induces a sequence Φ of regions $r_i \in R_{S^*}$ which contains each region at least once. For each k-tuple of regions $(r_1,\ldots,r_k) \in \Phi^{[k]}$ determine if the respective translation square graph $G^{\text{tsg}}_{S^*}(r_1,\ldots,r_k)$ has a perfect matching. Compute a maximum matching for the first k-tuple by using Lemma 4(a) and for all subsequent k-tuples by using Lemma 4(b).
3. If, in Step 2, a graph $G^{\text{tsg}}_{S^*}(r_1,\ldots,r_k)$ with perfect matching π has been found, report $\Gamma^{\pi}_{S^*}(r_1,\ldots,r_k)$; otherwise report that no (k,ε,δ)-stereo matching approximation exists.

Theorem 10. *A (k,ε,δ)-stereo matching approximation for two point sets A and B with n points each, $k \geqslant 2$, $0 < \varepsilon < \delta$, (if exists) can be computed in time $O((\varepsilon/\delta)^k n^{2+2k})$ and space $O((\varepsilon/\delta)n^2)$.*

Proof. The correctness of Algorithm 2 follows immediately from the definition of a (k,ε,δ)-stereo matching approximation. The time complexity of Step 1 is $O((\varepsilon/\delta)n^2)$; see Lemma 6, Observation 7, Lemma 8 and Corollary 9. Note that, when traversing $\mathcal{T}^*$ in Step 2, any particular region might be traversed several times. However,the total number of regions traversed in Φ is $O(|R_S|) = O((\varepsilon/\delta)n^2)$, see Observation 7. Hence, the number of k-tuples of regions enumerated in sequence $\Phi^{[k]}$ of Step 2 is $O((\varepsilon/\delta)^k n^{2k})$. From Lemma 8 it follows that the total number of edge updates for maintaining the translation square graphs for all k-tuples of regions of Φ is also $O((\varepsilon/\delta)^k n^{2k})$. The time per edge update for maintaining the maximum matchings for a translation square graph is $O(n^2)$, see Lemma 4(b). Hence, the time complexity for Step 2 is $O((\varepsilon/\delta)^k n^{2+2k})$. The time for Step 3 is dominated by the time for Step 2. Thus, the claimed time complexity for Algorithm 1 follows. The space requirement of Algorithm 2 is determined by the space for $\mathcal{T}^*$ which is $O((\varepsilon/\delta)n^2)$, see Observation 9. □

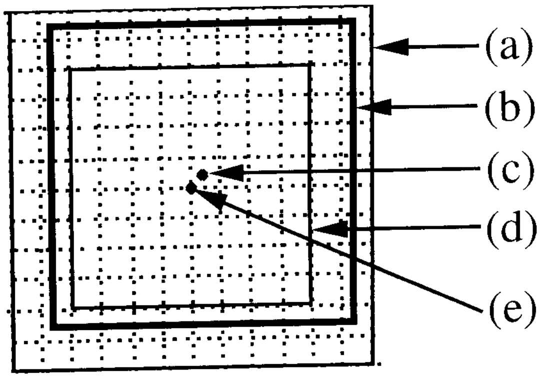

Fig. 3. Containment relationships between translation squares and their approximations. (a) $s^*_{\varepsilon+\delta/2}(a,b)$. (b) $s_\varepsilon(a,b)$. (c) Center of $s_\varepsilon(a,b)$. (d) $s^*_{\varepsilon-\delta/2}(a,b)$. (e) Center of $s^*_{\varepsilon-\delta/2}(a,b)$ and $s^*_{\varepsilon+\delta/2}(a,b)$.

3.2. *Solving the δ-approximate (k,ε)-stereo matching problem*

We will now study the relationship between (k,ε,δ)-stereo matching approximations and the δ-approximate (k,ε)-stereo matching problem. Note that the latter problem requires not a stereo matching approximation but an "*exact*" (k,ε)-stereo matching to be reported if $|\varepsilon_0 - \varepsilon| > \delta$.

Lemma 11. *If $\varepsilon < \varepsilon_0 - \delta$ then there exists no $(k, \varepsilon + \delta/2, \delta)$-stereo matching approximation for A and B.*

Proof. Assume that $\varepsilon < \varepsilon_0 - \delta$ and there exists a $(k, \varepsilon + \delta/2, \delta)$-stereo matching approximation Γ for A and B. Hence, there exist k regions $r_1, \ldots, r_k \in R_{S^*_{\varepsilon+\delta/2}}$ such that $G^{\text{tsg}}_{S^*_{\varepsilon+\delta/2}}(r_1, \ldots, r_k)$ has a perfect matching π and $\Gamma = \Gamma^{\pi}_{S^*_{\varepsilon+\delta/2}}(r_1, \ldots, r_k)$. For any $(a,b) \in A \times B$ we observe that $s^*_{\varepsilon+\delta/2}(a,b)$ is contained in $s_{\varepsilon+\delta}(a,b)$, see Fig. 3. Thus, all edges of $G^{\text{tsg}}_{S^*_{\varepsilon+\delta/2}}(r_1, \ldots, r_k)$ are also edges of $G^{\text{tsg}}_{S_{\varepsilon+\delta}}(r_1, \ldots, r_k)$, and π is also a perfect matching for $G^{\text{tsg}}_{S_{\varepsilon+\delta}}(r_1, \ldots, r_k)$. Therefore, it follows from Lemma 3 that $\Gamma^{\pi}_{S_{\varepsilon+\delta}}(r_1, \ldots, r_k)$ is a $(k, \varepsilon + \delta)$-stereo matching for A, B. A contradiction, since $\varepsilon + \delta < \varepsilon_0$. □

Lemma 12. *If $\varepsilon > \varepsilon_0 + \delta$ then there exists a $(k, \varepsilon - \delta/2, \delta)$-stereo matching approximation for A and B.*

Proof. Assume that $\varepsilon > \varepsilon_0 + \delta$. Hence there exists a $(k, \varepsilon - \delta)$-stereo matching Γ for A and B. Consider the regions $r_1, \ldots, r_k \in R_{S_{\varepsilon-\delta}}$ defined in the proof of Lemma 3, such that $G^{\text{tsg}}_{S_{\varepsilon-\delta}}(r_1, \ldots, r_k)$ has a perfect matching π and $\Gamma = \Gamma^{\pi}_{S_{\varepsilon-\delta}}(r_1, \ldots, r_k)$. Since $s_{\varepsilon-\delta}(a,b)$ is contained in $s^*_{\varepsilon-\delta/2}(a,b)$ for all $(a,b) \in A \times B$ (see Fig. 3), it follows that all edges of $G^{\text{tsg}}_{S_{\varepsilon-\delta}}(r_1, \ldots, r_k)$ are also edges of $G^{\text{tsg}}_{S^*_{\varepsilon-\delta/2}}(r_1, \ldots, r_k)$ and π is also a perfect matching for $G^{\text{tsg}}_{S^*_{\varepsilon-\delta/2}}(r_1, \ldots, r_k)$. Hence, $\Gamma^{\pi}_{S^*_{\varepsilon-\delta/2}}(r_1, \ldots, r_k)$ is a $(k, \varepsilon - \delta/2, \delta)$-stereo matching approximation for A and B. □

Lemma 13. *If $\varepsilon > \varepsilon_0 + \delta$ then every $(k, \varepsilon - \delta/2, \delta)$-stereo matching approximation for A and B is also a (k, ε)-stereo matching for A and B.*

Proof. Assume that $\varepsilon > \varepsilon_0 + \delta$ and let Γ be a $(k, \varepsilon - \delta/2, \delta)$-stereo matching approximation for A and B. Consider the regions $r^*_1, \ldots, r^*_k \in R_{S^*_{\varepsilon-\delta/2}}$ and perfect matching π for $G^{\text{tsg}}_{S^*_{\varepsilon-\delta/2}}(r^*_1, \ldots, r^*_k)$ such that $\Gamma = \Gamma^{\pi}_{S^*_{\varepsilon-\delta/2}}(r^*_1, \ldots, r^*_k)$. Let Γ be composed of a partitioning into subsets A_i and B_i, labelings l_i, and translation vectors t_i $(1 \leqslant i \leqslant k)$ as given in the definition of $\Gamma^{\pi}_{S^*_{\varepsilon-\delta/2}}(r^*_1, \ldots, r^*_k)$. Recall that t_i is a point inside r^*_i $(1 \leqslant i \leqslant k)$. Let r_i be the region of R_{S_ε} that contains t_i $(1 \leqslant i \leqslant k)$. Since $s^*_{\varepsilon-\delta/2}(a,b)$ is contained in $s_\varepsilon(a,b)$ for all $(a,b) \in A \times B$ (see Fig. 3), all edges in $G^{\text{tsg}}_{S^*_{\varepsilon-\delta/2}}(r^*_1, \ldots, r^*_k)$ are also edges of $G^{\text{tsg}}_{S_\varepsilon}(r_1, \ldots, r_k)$. Hence, π is also a perfect matching for $G^{\text{tsg}}_{S_\varepsilon}(r_1, \ldots, r_k)$, and Γ is also a (k, ε)-stereo matching for A and B; see Lemma 3. □

The above three lemmas lead to the following algorithm for solving the δ-approximate (k, ε)-stereo matching problem.

Algorithm 3. *δ-approximate (k, ε)-stereo matching*

1. Using Algorithm 2, attempt to compute a $(k, \varepsilon + \delta/2, \delta)$-stereo matching approximation Γ_1 for A and B. If no such stereo matching approximation is found, report "there exists no (k, ε)-stereo matching".
2. Using Algorithm 2 attempt to compute a $(k, \varepsilon - \delta/2, \delta)$-stereo matching approximation Γ_2 for A and B. If such a stereo matching approximation Γ_2 is found, report Γ_2 as a (k, ε)-stereo matching for A and B. Otherwise, report "don't know: $|\varepsilon_0 - \varepsilon| \leqslant \delta$".

Theorem 14. *The δ-approximate (k, ε)-stereo matching problem for two point sets A and B with n points each, $k \geqslant 2$, $0 < \delta < \varepsilon$, can be solved in time $O((\varepsilon/\delta)^k n^{2+2k})$ and space $O((\varepsilon/\delta)n^2)$.*

Proof. The correctness of Algorithm 3 follows from Lemmas 11, 12, and 13, and its time complexity and space requirement from Theorem 10. □

4. Conclusion

We presented algorithms for solving the unrestricted (k, ε)-stereo matching problem in time $O(n^{2+4k})$ and space $O(n^4)$, and the unrestricted δ-approximate (k, ε)-stereo matching problem in time $O((\varepsilon/\delta)^k n^{2+2k})$ and space $O((\varepsilon/\delta)n^2)$. To obtain these results, we introduced the *translation square arrangement*, *approximate translation square arrangement* and *approximate stereo matching tree*. These structures have some interesting combinatorial properties which might make them useful for other applications as well.

Acknowledgements

The authors would like to thank José Ricardo Torreão for introducing them to the stereo vision problem.

References

[1] H. Alt, K. Mehlhorn, H. Wagener, E. Welzl, Congruence, similarity and symmetries of geometric objects, Discrete Comput. Geom. 3 (1988) 237–256.

[2] S.T. Barnard, M.A. Fischler, Computational stereo, Comput. Surveys 14 (4) (1982) 553–572.

[3] S.D. Cochran, G. Medioni, 3-D surface description from binocular stereo, IEEE Trans. Pattern Anal. Machine Intell. 14 (10) (1992) 981–994.

[4] W.E.L. Grimson, A computer implementation of a theory of human stereo vision, in: Proc. Royal Soc. London Ser. B 292 (1981) 217–253.

[5] W.E.L. Grimson, From Images to Surfaces, MIT Press, Cambridge, MA, 1986.

[6] W.E.L. Grimson, Computing stereopsis using feature point contour matching, in: A. Rosenfeld (Ed.), Techniques for 3-D Machine Perception, North-Holland, Amsterdam, 1986, pp. 75–111.

[7] W.E.L. Grimson, D. Marr, A computer implementation of a theory of human stereo vision, in: L.S. Baumann (Ed.), Proc. ARPA Image Understanding Workshop, SRI, 1979, pp. 41–45.

[8] P.J. Heffernan, The translation square map and approximate congruence, Inform. Process. Lett. 39 (1991) 153–159.

[9] J. Hopcroft, R. Karp, An $n^{5/2}$ algorithm for maximal matchings in bipartite graphs, SIAM J. Comput. 2 (4) (1973) 225–231.

[10] D. Marr, Vision, Freeman, New York, 1982.

[11] D. Marr, T. Poggio, Cooperative computation of stereo disparity, Science 194 (1976) 283–287.

[12] D. Marr, T. Poggio, A computational theory of human stereo vision, in: Proc. Royal Soc. London Ser. B 204 (1979) 301–328.

[13] K. Mehlhorn, Data Structures and Algorithms 3: Multidimensional Searching and Computational Geometry, Springer, Heidelberg, 1984.

[14] H.K. Nishihara, Practical real-time imaging stereo matcher, Optical Engineering 23 (5) (1984) 536–545.

[15] M. Okutomi, T. Kanade, A multiple-baseline stereo, IEEE Trans. Pattern Anal. Machine Intell. 15 (4) (1993) 353–363.

[16] K. Prazdny, Detection of binocular disparities, Biological Cybernet. 52 (1985) 93–99.

[17] S. Schirra, Approximate decision algorithms for approximate congruence, Inform. Process. Lett. 43 (1992) 29–34.

[18] C. Stewart, C.R. Dyer, Parallel simulation of a connectionist stereo algorithm on a shared-memory multiprocessor, in: Kumar et al. (Eds.), Parallel Algorithms for Machine Intelligence and Vision, Springer, Heidelberg, 1990.