

Combinatorial Properties of 2D Discrete Rigid Transformations under Pixel-Invariance Constraints[★]

Phuc Ngo¹, Yukiko Kenmochi¹, Nicolas Passat^{2,3}, and Hugues Talbot¹

¹ Université Paris-Est, LIGM, UPEMLV-ESIEE-CNRS, France

² Université de Strasbourg, LSIIT, UMR 7005 CNRS, France

³ Université de Reims, CReSTIC, EA 3804, France

Abstract. Rigid transformations are useful in a wide range of digital image processing applications. In this context, they are generally considered as continuous processes, followed by discretization of the results. In recent works, rigid transformations on \mathbb{Z}^2 have been formulated as a fully discrete process. Following this paradigm, we investigate – from a combinatorial point of view – the effects of pixel-invariance constraints on such transformations. In particular we describe the impact of these constraints on both the combinatorial structure of the transformation space and the algorithm leading to its generation.

Keywords: Combinatorial structure, Discrete rigid transformation, Pixel-invariance constraints.

1 Introduction

Rigid transformations, (*i.e.*, transformations based on translations and rotations) are involved in the design of many computer vision and image processing techniques (see, *e.g.*, [8,9]). Such transformations are generally performed by considering the Euclidean space (\mathbb{R}^n) associated to the Eulerian space (\mathbb{Z}^n) of the data. As a consequence, they need to be interfaced with a subsequent digitization process to finally produce results in \mathbb{Z}^n .

In [5], we have recently proposed to study rigid transformations on \mathbb{Z}^2 as a *fully discrete process*. In this context, three main questions were considered: (*i*) How many rigid transformations can be defined on a finite subspace of \mathbb{Z}^2 ? (*ii*) How to generate all of them? (*iii*) What are the topological relationships between them? Some combinatorial and algorithmic answers, inspired by the approaches developed in [3,4], were provided, and then contributed to the state of the art in this research field [3,4,2,7]. In [5], a combinatorial structure – namely a *graph* – is used to represent the 2D discrete rigid transformations. This structure has a polynomial complexity $O(N^9)$ where $N \times N$ is the size of images. However, this high complexity makes it difficult to generate the whole graph for large images, and to further find admissible transformations that best match two given images, namely a template and a target image; the later problem is called *image registration*. Practically in computer vision, some constrained search paradigms are used for registration issues (see, *e.g.*, [9,1]). Indeed the constraints introduce prior knowledge of transformations and contribute to reduce the research space.

[★] The research leading to these results has received funding from the French *Agence Nationale de la Recherche* (Grant Agreement ANR-2010-BLAN-0205 03).

In this article, we extend the study initiated in [5], by investigating the effects of geometric constraints on the proposed graph. In particular, we focus on *pixel-invariance constraints*, which consist of forcing the correspondence between points in an initial (sub)space (of \mathbb{Z}^2) and transformed points – or more generally regions.

This study is organised as follows. We first recall background notions on discrete rigid transformations (Sec. 2), and express pixel-invariance constraints in the associated parameter space (Sec. 3). We then develop an algorithmic process for generating a combinatorial structure modeling all the discrete rigid transformations and their relationships under the given constraints (Sec. 4). A complexity analysis is proposed for this algorithm and the induced structure (Sec. 5). Finally, we conclude the article (Sec. 6).

2 Background Notions of Discrete Rigid Transformations

2.1 Digital Images and Digital Rigid Transformations

In the continuous framework, an image can be formalised as a function $\mathcal{I} : \mathbb{R}^2 \rightarrow V$, where V is any value space. A *digital image* associated to \mathcal{I} can then be defined as $I : \mathbb{Z}^2 \rightarrow V$, by sampling \mathcal{I} on the discrete space \mathbb{Z}^2 . In other words, we have $I = \mathcal{I}|_{\mathbb{Z}^2}$, and for each $\mathbf{p} \in \mathbb{Z}^2$, the value $I(\mathbf{p})$ models the value of \mathcal{I} on the associated pixel $\mathbf{p} + [-\frac{1}{2}, \frac{1}{2}]^2$, namely the Voronoi cell of \mathbb{R}^2 induced by \mathbb{Z}^2 around \mathbf{p} . This paradigm relies on the digitisation function, where $[\cdot]$ is a rounding operator, defined as

$$\left| \begin{array}{l} D : \mathbb{R}^2 \longrightarrow \mathbb{Z}^2 \\ (x, y) \longmapsto ([x], [y]) \end{array} \right. \tag{1}$$

In the continuous framework, a 2D rigid transformation, composed of translation and rotation, is expressed as a bijection $\mathcal{T} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined, for any $\mathbf{x} = (x, y) \in \mathbb{R}^2$, by

$$\mathcal{T}(\mathbf{x}) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix} \tag{2}$$

where $a, b, \theta \in \mathbb{R}$ and $\theta \in [0, 2\pi[$. In particular, such a transformation is unambiguously modeled by the triplet of parameters (a, b, θ) , and will sometimes be noted by $\mathcal{T}_{ab\theta}$. When applied on an image $\mathcal{I} : \mathbb{R}^2 \rightarrow V$, it provides a new transformed image $\mathcal{I} \circ \mathcal{T}^{-1} : \mathbb{R}^2 \rightarrow V$.

Following the digitisation paradigm proposed above, a *digital rigid transformation* $T : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$ associated to \mathcal{T} can be defined, for any $\mathbf{p} = (p, q) \in \mathbb{Z}^2$, by

$$T(\mathbf{p}) = D \circ \mathcal{T}(\mathbf{p}) = \begin{pmatrix} [p \cos \theta - q \sin \theta + a] \\ [p \sin \theta + q \cos \theta + b] \end{pmatrix} \tag{3}$$

In general, this function is not bijective. However, by setting $T^{-1} : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$ as $T^{-1} = D \circ \mathcal{T}^{-1}$, it becomes possible to define the digital transformed image $I \circ T^{-1} : \mathbb{Z}^2 \rightarrow V$ with respect to T . In the sequel of this article, we focus on such digital rigid transformations. From this point on – for the sake of readability and without loss of correctness – we will denote them T instead of T^{-1} , due to the bijectivity of \mathcal{T} and \mathcal{T}^{-1} .

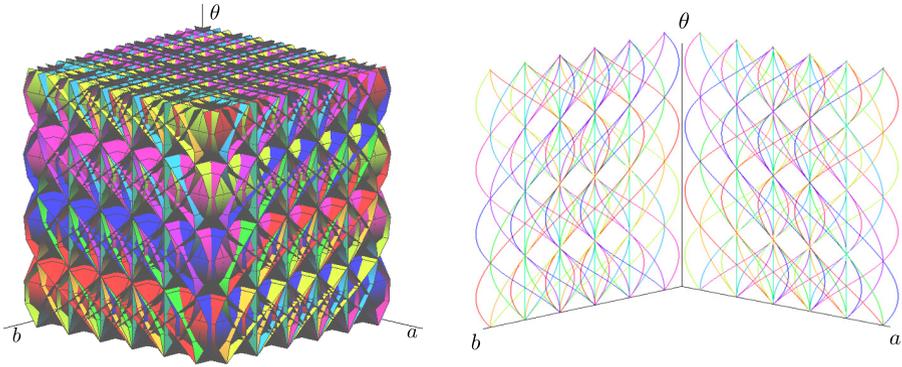


Fig. 1. Tipping surfaces in the 3D parameter space (a, b, θ) (left) and their cross-sections, namely tipping curves, in the 2D planes (a, θ) and (b, θ) (right)

2.2 Discontinuities of Digital Rigid Transformations

For any \mathbf{x} in Eq. (2), if we change the value of (a, b, θ) slightly, then the new point $\mathcal{T}(\mathbf{p})$ also changes slightly. More formally, the function $(a, b, \theta) \mapsto \mathcal{T}_{ab\theta}$ is continuous. Contrariwise, in Eq. (3), an infinitesimal variation of (a, b, θ) may lead to a variation of $T(\mathbf{p})$ from a point of \mathbb{Z}^2 to another one. More precisely, the parameter space \mathbb{R}^3 of (a, b, θ) is divided into 3D open cells where the function $(a, b, \theta) \mapsto T_{ab\theta} = D \circ \mathcal{T}_{ab\theta}$ is piecewise constant, bounded by 2D surfaces where it is discontinuous.

We focus in particular on the triplets (a, b, θ) and their associated transformations $\mathcal{T}_{ab\theta}$, which lead to such discontinuities in the space of digital rigid transformations. Such *critical* transformations are those that map at least one discrete point onto the discrete half-grid $\mathcal{H} = (\mathbb{R} \times (\mathbb{Z} + 1/2)) \cup ((\mathbb{Z} + 1/2) \times \mathbb{R})$ (i.e., the boundaries of the Voronoi cells of \mathbb{R}^2 induced by \mathbb{Z}^2).

Definition 1 ([5]). Let $(a, b, \theta) \in \mathbb{R}^3$, and $\mathcal{T}_{ab\theta} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be its associated rigid transformation. We say that $\mathcal{T}_{ab\theta}$ is a critical transformation if $\exists \mathbf{p} \in \mathbb{Z}^2$ s.t. $\mathcal{T}_{ab\theta}(\mathbf{p}) \in \mathcal{H}$.

It is plain that in the parameter space (a, b, θ) , the critical transformations are modeled by 2D surfaces analytically defined, for any $\mathbf{p} = (p, q) \in \mathbb{Z}^2$ and $k, l \in \mathbb{Z}$, by

$$\left| \begin{array}{l} \Phi_{pqk} : \mathbb{R}^2 \longrightarrow \mathbb{R} \\ (b, \theta) \longmapsto a = \phi_{pqk}(\theta) = k + \frac{1}{2} + q \sin \theta - p \cos \theta, \end{array} \right. \quad (4)$$

$$\left| \begin{array}{l} \Psi_{pql} : \mathbb{R}^2 \longrightarrow \mathbb{R} \\ (a, \theta) \longmapsto b = \psi_{pql}(\theta) = l + \frac{1}{2} - p \sin \theta - q \cos \theta. \end{array} \right. \quad (5)$$

The surfaces Φ_{pqk} (resp. Ψ_{pql}) are termed *tipping surfaces* [5]. Their intersection ϕ_{pqk} (resp. ψ_{pql}) on the 2D plane (a, θ) (resp. (b, θ)) are called *tipping curves*. These tipping surfaces/curves, which correspond to the discontinuities of digital rigid transformations expressed in the parameter space (a, b, θ) , are illustrated in Fig. 1.

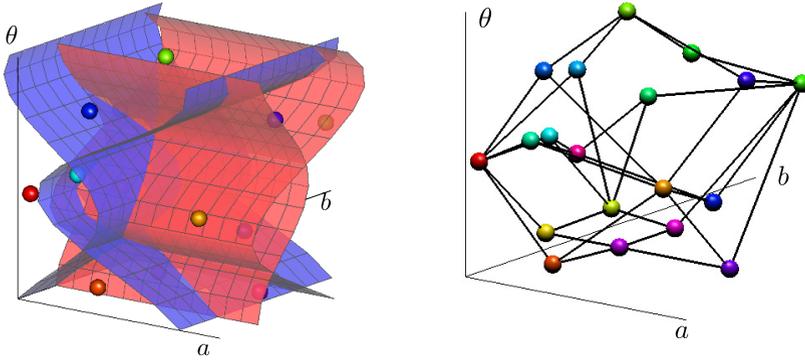


Fig. 2. Parameter space subdivided by four tipping surfaces (left) and its DRT graph (right)

2.3 Partition of Parameter Space and Discrete Rigid Transformation Graph

The digitisation process (Eq. (1)) generally maps distinct rigid transformations (Eq. (2)) onto the same digital rigid transformation (Eq. (3)). More precisely, the set of all the *non-critical transformations* can be partitioned into equivalence classes induced by the equivalence relation \sim defined by $(\mathcal{T}_{ab\theta} \sim \mathcal{T}_{a'b'\theta'}) \iff (T_{ab\theta} = T_{a'b'\theta'})$. This also leads to the straightforward definition of an equivalence relation on the parameters (a, b, θ) associated to these transformations. In this isomorphic framework, each equivalence class is called a *discrete rigid transformation (DRT)*, and is modeled by 3D open cells bounded by 2D tipping surfaces, which subdivide the parameter space (a, b, θ) (see Fig. 1(a)).

In [5], we have shown that this subdivision of the parameter space could be modeled using a dual combinatorial structure, namely a graph. In particular, each 3D open cell (*i.e.*, each DRT) is associated to a vertex, and each tipping-surface segment (linked to a critical transformation) shared by two adjacent 3D open cells, is associated to an edge. The resulting graph is called a *DRT graph* [5] (see Fig. 2).

From a theoretical point of view, the notions introduced above are correctly defined for images and transformations on \mathbb{Z}^2 and \mathbb{R}^2 . Practically, our purpose is to study such transformations on images of *finite sizes*. Under this hypothesis, only a finite subset of digital rigid transformations are relevant, namely those which actually affect such finite images. From this point on, we focus on this finite case and assume that the digital images are defined on subsets of \mathbb{Z}^2 of size $N \times N$ ($N \in \mathbb{N}$). We have the next property.

Property 2 ([5]). *The DRT graph associated to a digital image of size $N \times N$ has a space complexity of $O(N^9)$.*

In [5], an exact computation algorithm was also proposed to build this graph in linear time w.r.t. its size. The DRT graph models a kind of “neighbouring” relationship between DRTs. Indeed, the existence of an edge between two vertices indicates that the associated transformations differ in one pixel among the N^2 ones. This property opens a way of involving this combinatorial structure in image processing tasks.

We now study the effects of forcing the correspondence between points in the initial and transformed spaces. We focus in particular on the subdivision of the parameter space and the induced graph from the algorithmic and combinatorial points of view.

3 Constraints and Feasible Rigid Transformation Set

3.1 Pixel-Invariance Constraints and Interpretation in the Parameter Space

In the context of rigid transformations in \mathbb{R}^2 , forcing the correspondence between points \mathbf{p} in the initial space and \mathbf{p}' in the transformed one leads to restricting the transformations \mathcal{T} . Moreover, forcing the correspondence between $k \geq 2$ distinct pairs of points $(\mathbf{p}_i, \mathbf{p}'_i)$ restricts the number of feasible transformations \mathcal{T} to at most one. Indeed, from the relation expressed in Eq. (2) we obtain for every pair of corresponding points two equations representing the trigonometric surfaces. We then need at least two pairs of corresponding points to obtain a rigid transformation as the feasible transformation. As illustrated in Fig. 3(a), each pair of surfaces represents a pair of corresponding points, and the intersection of two pairs of surfaces determines the feasible transformation.

Restricting discrete rigid transformations, under similar constraints, is more permissive. Indeed, when forcing the correspondence between one or several pairs of pixels $(\mathbf{p}_i, \mathbf{p}'_i)$ of \mathbb{Z}^2 , a larger space of transformations may remain valid (see Fig. 3(b–g)).

Definition 3. Let $\mathbf{p} = (p, q) \in A \subset \mathbb{Z}^2$ and $\mathbf{p}' = (p', q') \in B \subset \mathbb{Z}^2$, such that A, B are of size $N \times N$. There exists a pixel-invariance constraint between \mathbf{p} and \mathbf{p}' if the authorised digital rigid transformations T between A and B satisfy the equality $T(\mathbf{p}) = \mathbf{p}'$, i.e., if

$$p' - 1/2 < p \cos \theta - q \sin \theta + a < p' + 1/2, \tag{6}$$

$$q' - 1/2 < p \sin \theta + q \cos \theta + b < q' + 1/2. \tag{7}$$

More generally, there exist pixel-invariance constraints between two sets $\{\mathbf{p}_i\}_{i=1}^m$ and $\{\mathbf{p}'_i\}_{i=1}^m$ ($m \geq 1$) if $T(\mathbf{p}_i) = \mathbf{p}'_i$ (i.e., if Eqs. (6)–(7) are satisfied) for every $i \in \llbracket 1, m \rrbracket$.

In absence of constraints, the 3D parameter space (a, b, θ) , induced by the subset of size $N \times N$ where the image is defined, is divided into cells whose boundaries are all the tipping surfaces Φ_{pqk} and Ψ_{pql} , with $p, q \in \llbracket 0, N - 1 \rrbracket$ and $k, l \in \llbracket 0, N \rrbracket$. In this context, the whole volume of the parameter space models adequate rigid transformations.

In contrast, under a pixel-invariance constraint, some discrete rigid transformations may become irrelevant. Equivalently, only a part of the parameter space – namely the subspace of parameters (a, b, θ) that satisfy this constraint – remains valid. From Eqs. (6)–(7), this parameter subspace is defined by the intersection of 4 half-spaces associated to 4 tipping surfaces for one pixel correspondence (see Fig. 3(b)).

3.2 Feasible Rigid Transformation Set

More generally, if a set \mathcal{P} of m pixel correspondences is provided, the parameter subspace of relevant transformations is defined as the intersection of m regions induced by

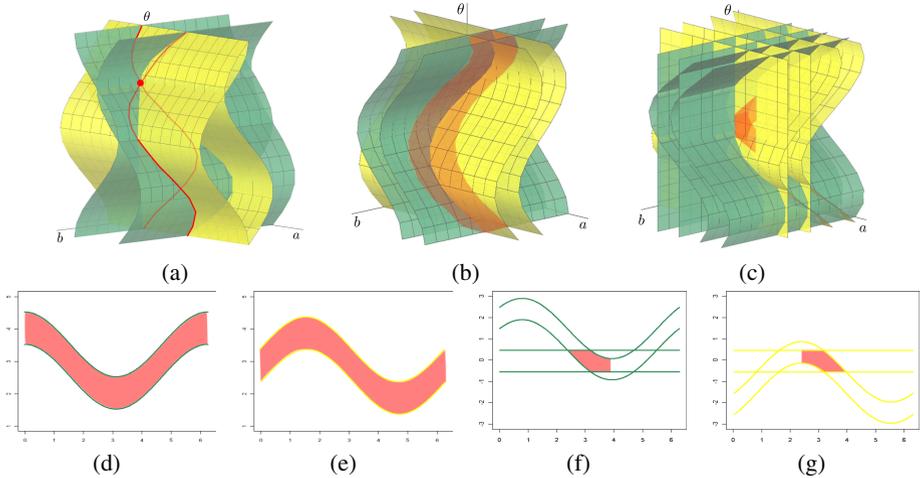


Fig. 3. Rigid transformation sets induced by geometric constraints in continuous (a) and discrete (b,c) frameworks. (a) Transformations with one-point correspondence (red line) and two-point correspondences (red dot). (b,c) Transformations with one-pixel (b) and two-pixels (c) correspondences (red volumes). (d,e) (resp. (f,g)) Cross-sections of (b) (resp. (c)) on the planes (a, θ) and (b, θ) via the use of tipping curves.

these constraints, *i.e.*, as the intersection of $4m$ half-spaces defined by Eqs. (6)–(7). Let us first define the half-spaces induced by tipping surfaces Φ_{pqk} and Ψ_{pql} :

$$H_{pqk}^+ = \{(a, b, \theta) \mid a > \Phi_{pqk}(b, \theta)\} \text{ and } H_{pqk}^- = \{(a, b, \theta) \mid a < \Phi_{pqk}(b, \theta)\}, \quad (8)$$

$$V_{pql}^+ = \{(a, b, \theta) \mid b > \Psi_{pql}(a, \theta)\} \text{ and } V_{pql}^- = \{(a, b, \theta) \mid b < \Psi_{pql}(a, \theta)\}. \quad (9)$$

The subspace of interest, called *feasible rigid transformation set*, is defined as follows.

Definition 4. Let $\mathcal{P} = \{(\mathbf{p}_i, \mathbf{p}'_i)\}_{i=1}^m$ ($m \geq 1$) be a set of corresponding pixel pairs. The feasible rigid transformation set (FRTS) associated to \mathcal{P} is the subspace $\mathcal{R} \subset \mathbb{R}^3$ of the parameter space (a, b, θ) , defined as

$$\mathcal{R} = \bigcap_{i \in \llbracket 1, m \rrbracket} \left(H_{p_i q_i p'_i}^+ \cap H_{p_i q_i p'_i + 1}^- \cap V_{p_i q_i q'_i}^+ \cap V_{p_i q_i q'_i + 1}^- \right).$$

Note that a constraint of one pixel pair (namely, $m = 1$), the FRTS is observed as a “tube” in the parameter space (a, b, θ) (see Fig. 3(b)). We can consider that a pixel-invariance constraint leads to all rotations with a center $\mathbf{x} \in \mathbb{R}^2$ inside of pixel $\mathbf{p} \in \mathbb{Z}^2$ (*i.e.*, $\mathbf{x} \in \mathbf{p} + [-1/2, 1/2]^2$). For constraints of two pixel pairs (namely, $m = 2$), the FRTS becomes a “bounded region” (see Fig. 3(c)).

The FRTS generated by m pixel correspondences is divided into DRTs induced from the $(N^2 - m)$ remaining pixels of the given image of size $N \times N$. In particular, the combinatorial structure of DRTs in a FRTS, modeling this subdivision, is a subgraph of the DRT graph. We introduce, in the following, a notion of *directional convexity* of a region R , such that R is “convex along an axis”, and show that any FRTS \mathcal{R} is directionally convex. This property is used in the next sections to study the combinatorial structure of DRTs in \mathcal{R} .

Definition 5. We say that a region $R \subseteq \mathbb{R}^n$ in an n -variable space (x_1, \dots, x_n) is x_k -convex if, for any two points $\mathbf{p}_1, \mathbf{p}_2 \in R$ such that the segment $[\mathbf{p}_1\mathbf{p}_2] = \{\alpha \cdot \mathbf{p}_1 + (1-\alpha) \cdot \mathbf{p}_2 \mid \alpha \in [0, 1]\}$ is parallel to the x_k -axis, $[\mathbf{p}_1\mathbf{p}_2]$ is included in R .

Property 6. The FRTS \mathcal{R} is both a - and b -convex.

4 Combinatorial Structure of Discrete Rigid Transformations in a Feasible Rigid Transformation Set

So far, we know that a FRTS contains all rigid transformations satisfying all constraints, and is subdivided into DRTs (see Sec. 2.3). This section presents a method for constructing a combinatorial structure of DRTs in a FRTS by following the three stages: (1) finding the boundaries of a FRTS, (2) finding tipping surfaces passing a FRTS and (3) constructing a DRT graph in a FRTS. Before describing these stages, we first explain an algorithm for building a graph modeling a subdivision of the parameter space from a given set of tipping surfaces, which is used later in the first and third stages.

4.1 Sweeping Algorithm for Incremental Partition Graph Construction

We may generalize the problem of subdivision of the parameter space by tipping surfaces as follows: given a set of tipping surfaces S , we would like to construct a graph modeling the subdivision of the parameter space (a, b, θ) induced by S . Such a graph is called a *partition graph* and denoted by G . In G each vertex is associated to a 3D open cell of the subdivision, and each tipping-surface segment shared by two adjacent 3D open cells, is associated to an edge. This problem can be answered with the help of 3D arrangements of surfaces [6]. In [5], we have proposed the *sweeping algorithm* for constructing a DRT graph of a given image (see Section 2.3), which is a specific case of generating a partition graph. Such a method has a complexity $\mathcal{O}(n^3)$, where n is the number of surfaces. Here we present a similar method that builds the partition graph G based on the relations that link tipping surfaces and tipping curves (see Eqs (4)–(5) and Fig. 1). Such a subdivision can be fully described from its two cross-sections on the planes (a, θ) and (b, θ) , respectively expressed by two sets of tipping curves. Therefore, instead of constructing directly the partition graph in the 3D parameter space (a, b, θ) , we will first build the structures of the graphs in the 2D planes (namely, (a, θ) and (b, θ) planes), and then combine them to build the complete partition graph. Note that G becomes a DRT graph if we consider all tipping surfaces for a given image.

We first define a *cut* for a plane – either (a, θ) or (b, θ) – denoted by γ , as a monotonic line intersecting exactly once for each tipping curve in the plane. A cut is then represented by its sequence of intersecting tipping curves (see Fig. 4(a)). Such a cut can be modeled by a directed graph according to their sequences of tipping curves.

Definition 7. Let $\gamma = (\phi_1, \phi_2, \dots)$ be a cut. A graph $G_\gamma = (V_\gamma, E_\gamma)$ w.r.t. γ consists of
– a set of vertices $V_\gamma = \{v_0, v_1, \dots\}$, and
– an ordered set of labelled edges $E_\gamma = ((v_0, v_1, \phi_1), (v_2, v_1, \phi_2), \dots)$ and each edge $(u, w, f) \in E_\gamma$ connects two vertices $u, w \in V_\gamma$ separated by the tipping curve f , which is considered as an edge label.

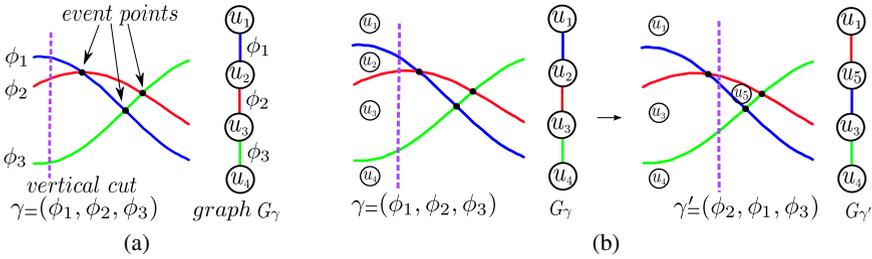


Fig. 4. (a) Example of a cut and its graph. (b) Progress of the cut at an event point by which the cut is updated and the corresponding graph is modified.

In practice, the elements of E_γ are also ordered in the same way as γ (see Fig. 4(a)).

The main idea of the sweeping method in 2D is that a cut is swept through all tipping curves on the plane in finite time, allowing us to construct the graph afterwards. We assume that γ starts at $\theta = 0$ and ends at $\theta = 2\pi$. While sweeping the cut, its sequence changes only at intersections of tipping curves, called *event points*. The moment at which a cut reaches an event point, the algorithm performs an update of its sequence, and generates new vertices and edges in the graph (see Fig. 4(b)). We call this an *elementary step* of the algorithm. The set of event points forms a series of elementary steps. Therefore, instead of moving the cut continuously, we need to maintain a set of sorted event points w.r.t. θ , and progress the cut in their increasing order to build the graph incrementally.

For building a graph G in 3D, two *cuts* are used such that each cut sweeps on either the plane (a, θ) or (b, θ) . We denote those cuts by γ_a and γ_b respectively. For each update of the cuts, γ_a and γ_b , the associated graphs, G_{γ_a} and G_{γ_b} , are respectively modified, so that a part of G is generated. We call such a part of G a *partial graph*, denoted by δG . In fact, δG is a combination of the two graphs G_{γ_a} and G_{γ_b} as follows (see Fig. 5).

Definition 8. The partial graph $\delta G = (\delta V, \delta E)$ is generated from $G_{\gamma_a} = (V_{\gamma_a}, E_{\gamma_a})$ and $G_{\gamma_b} = (V_{\gamma_b}, E_{\gamma_b})$, such that

- $\delta V = \{(v_a, v_b) \mid v_a \in V_{\gamma_a}, v_b \in V_{\gamma_b}\}$, and
- $\delta E = \{((u_1, v), (u_2, v), \phi_u) \mid u_1, u_2 \in V_{\gamma_a}, v \in V_{\gamma_b}, (u_1, u_2, \phi_u) \in E_{\gamma_a}\} \cup \{((u, v_1), (u, v_2), \phi_v) \mid v_1, v_2 \in V_{\gamma_b}, u \in V_{\gamma_a}, (v_1, v_2, \phi_v) \in E_{\gamma_b}\}$.

Therefore, when an elementary step is applied, the sweep progresses as the partial graph δG is generated and integrated in G for constructing the final graph as well. The following proposition has been originally proposed in [5] for constructing a DRT graph, and this is valid for a partition graph as well.

Proposition 9. Let S be a set of tipping surfaces and G be a partition graph modeling the subdivision of the parameter space by S . We have

$$G = \bigcup_{i \in \llbracket 1, e \rrbracket} \delta G_i,$$

where δG_i is a partial graph at the i -th elementary step and e is the number of ordered event points.

More details about the sweeping algorithm for tipping surfaces can be found in [5].

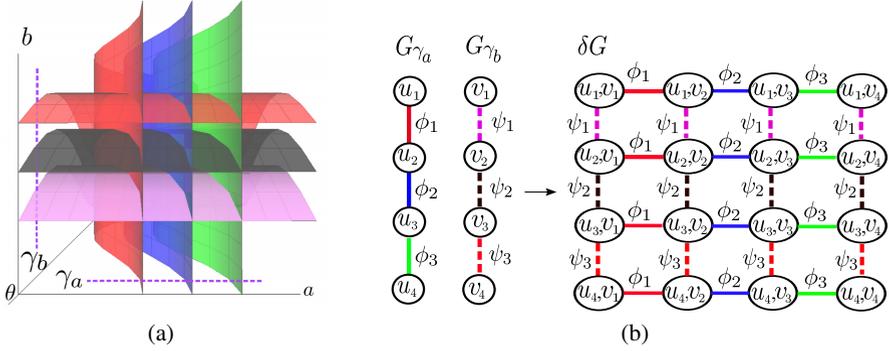


Fig. 5. Construction of a partial graph δG by combining two graphs G_{γ_a} and G_{γ_b}

4.2 Finding the Feasible Rigid Transformation Set Boundary

It is possible to describe a FRTS \mathcal{R} by a set of half-spaces constituting the boundary of \mathcal{R} , instead of using all the half-spaces of \mathcal{R} as described in Definition 4. This section explains how to find such a set using the above sweeping algorithm.

A FRTS \mathcal{R} in the parameter space (a, b, θ) can be fully described from its two cross-sections \mathcal{R}_H and \mathcal{R}_V on the planes (a, θ) and (b, θ) , defined as

$$\mathcal{R}_H = \bigcap_{i \in \llbracket 1, m \rrbracket} \left(h_{p_i, q_i}^+ \cap h_{p_i, q_i'}^- \right) \text{ and } \mathcal{R}_V = \bigcap_{i \in \llbracket 1, m \rrbracket} \left(v_{p_i, q_i}^+ \cap v_{p_i, q_i'}^- \right) \quad (10)$$

where h_{p_i, q_i}^+ and $h_{p_i, q_i'}^-$ (resp. v_{p_i, q_i}^+ and $v_{p_i, q_i'}^-$) are the cross-sections in the plane (a, θ) (resp. (b, θ)) of H_{p_i, q_i}^+ and $H_{p_i, q_i'}^-$ (resp. V_{p_i, q_i}^+ and $V_{p_i, q_i'}^-$). This is illustrated in Fig. 3(d–g). They are expressed as shown in Eq. (8) (resp. Eq. (9)) by replacing the tipping surfaces Φ_{pqk} and Ψ_{pql} by the tipping curves ϕ_{pqk} and ψ_{pql} respectively. We call $h_{p_i, q_i}^+, v_{p_i, q_i}^+$ *upper half-planes* and $h_{p_i, q_i'}^-, v_{p_i, q_i'}^-$ *lower half-planes*.

Relying on the similarity of \mathcal{R}_V and \mathcal{R}_H , hereafter we consider only \mathcal{R}_H . Our problem is then specified as follows: given a constraint set of half-planes of \mathcal{R}_H defined from m corresponding pixel pairs, $\mathcal{P} = \{(\mathbf{p}_i, \mathbf{p}'_i)\}_{i=1}^m$, report the boundary half-planes of \mathcal{R}_H . From Property 6, it is obvious that \mathcal{R}_H contains two sets of boundary half-planes:

- a *upper boundary sequence* $U = (h_{p_i, q_i}^+, \dots)$ contains only the upper half-planes;
- a *lower boundary sequence* $L = (h_{p_i, q_i'}^-, \dots)$ contains only the lower half-planes.

The 2D sweeping algorithm, presented in Sec. 4.1, is used to find such U and L of \mathcal{R}_H , such that the cut γ is now represented as a sequence of half-planes intersecting it. Note that no partition graph is built in this stage, and we only need to observe the sequence of the cut γ during its update in order to obtain U and L . Indeed, while sweeping γ its sequence changes at event points. We remark that γ is in \mathcal{R}_H when its sequence of half-planes is separated into two successive sequences of γ^+ and γ^- , namely $\gamma = \gamma^+ \gamma^-$, where γ^+ contains only the upper half-planes and γ^- contains only the lower half-planes. Moreover, we see that the last element of γ^+ and the first element of γ^- determine the upper and lower boundaries of \mathcal{R}_H respectively. The cut is located out of \mathcal{R}_H when there is no more such a separation. According to the change of γ in \mathcal{R}_H , the

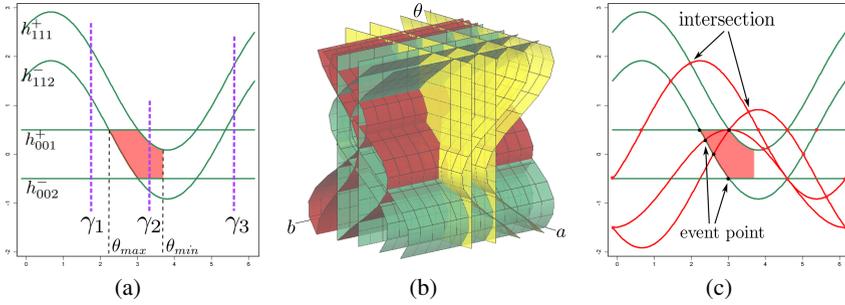


Fig. 6. (a) Progression of the cut γ in the cross-section \mathcal{R}_H of FRTS \mathcal{R} on the plane (a, θ) . The initial cut is $\gamma_1 = (h_{111}^+, h_{112}^-, h_{001}^+, h_{002}^-)$; when it goes in \mathcal{R}_H it has become $\gamma_2 = (h_{001}^+, h_{111}^+, h_{002}^-, h_{112}^-)$; when it goes out of \mathcal{R}_H it becomes $\gamma_3 = (h_{111}^+, h_{112}^-, h_{001}^+, h_{002}^-)$. (b) Example of tipping surfaces (in red) passing \mathcal{R} in the parameter space (a, b, θ) and (c) its cross-section \mathcal{R}_H on the plane (a, θ) .

upper and lower boundaries are added progressively in U and L at each event point. See Fig. 6(a) for the illustration.

By using two cuts γ_a and γ_b sweeping on the two planes (a, θ) and (b, θ) , we can find the boundary of a FRTS \mathcal{R} . At each event point either on (a, θ) or (b, θ) , the algorithm updates and checks the sequences of both cuts. From this, we obtain the boundary of \mathcal{R} , and the first θ at which both sequences of γ_a and γ_b are separated in two parts (resp. have no more separation), denoted by θ_{min} (resp. θ_{max}) of \mathcal{R} that need in the next stage.

4.3 Finding Tipping Surfaces Passing by a Feasible Rigid Transformation Set

As the subdivision of a FRTS \mathcal{R} is induced by the tipping surfaces existing in \mathcal{R} (see Fig. 6(b)), we need to determine such tipping surfaces among all Φ_{pqk} and Ψ_{pql} for $p, q \in \llbracket 0, N-1 \rrbracket$ and $k, l \in \llbracket 0, N \rrbracket$, where $N \times N$ is the image size. Now looking at the cross-sections of \mathcal{R} , this problem is equivalent to finding tipping curves ϕ_{pqk} (resp. ψ_{pql}) passing \mathcal{R}_H (resp. \mathcal{R}_V) (see Fig. 6(c)).

A tipping curve ϕ_{pqk} passes \mathcal{R}_H if it intersects one of the boundary segments of $\phi_{p'q'k'}$ of \mathcal{R}_H . This is easily verified by the following steps:

- (i) verify if ϕ_{pqk} and $\phi_{p'q'k'}$ intersect, *i.e.*, if the following relation is satisfied [5, Property 2]: $\Delta_1 + \Delta_2 > 0$ and $|KP \pm \sqrt{\Delta_1}| \leq P^2 + Q^2$ and $|KQ \pm \sqrt{\Delta_2}| \leq P^2 + Q^2$, where $P = p - p'$, $Q = q - q'$, $K = k - k'$, $\Delta_1 = P^2(P^2 + Q^2 - K^2)$ and $\Delta_2 = Q^2(P^2 + Q^2 - K^2)$;
- (ii) if they intersect, then calculate the following values at the intersection [5, Corollary 1]: $\sin \theta = \frac{KQ \pm \sqrt{\Delta_1}}{P^2 + Q^2}$ and $\cos \theta = \frac{KP \pm \sqrt{\Delta_2}}{P^2 + Q^2}$, and verify if $\theta_{min} \leq \theta \leq \theta_{max}$, where θ_{min} and θ_{max} are obtained in Sec. 4.2;
- (iii) if (ii) is verified, then calculate $a_{upper} = \max_{h_{pqk}^+ \in U} \phi_{pqk}(\theta)$ and $a_{lower} = \min_{h_{pqk}^- \in L} \phi_{pqk}(\theta)$, and verify if $a_{upper} \leq a \leq a_{lower}$, where the value a at the above intersection θ is calculated from Eq. (4).

Note that the values $\cos \theta$ and $\sin \theta$ are used to represent θ . All $\cos \theta$, $\sin \theta$, $\cos \theta_{min}$, $\sin \theta_{min}$, $\cos \theta_{max}$, $\sin \theta_{max}$, a , a_{min} and a_{max} are quadratic irrationals¹. As shown in [5], their comparisons can be achieved exactly in constant time.

¹ A quadratic irrational is an irrational number that is a solution of some quadratic equations.

We define *tipping surfaces of interest* as a set of tipping surfaces that bound or pass a FRTS \mathcal{R} . Similarly, *tipping curves of interest* of $\mathcal{R}_{\mathcal{H}}$ (resp. $\mathcal{R}_{\mathcal{V}}$) is the set of the boundary tipping curves of $\mathcal{R}_{\mathcal{H}}$ (resp. $\mathcal{R}_{\mathcal{V}}$) and the tipping curves passing $\mathcal{R}_{\mathcal{H}}$ (resp. $\mathcal{R}_{\mathcal{V}}$).

4.4 DRT-Graph Construction in a Feasible Rigid Transformation Set

In order to build the DRT graph in a FRTS \mathcal{R} , we use the sweeping algorithm described in Sec. 4.1. However the cut γ in this part sweeps from θ_{min} to θ_{max} instead of $[0, 2\pi]$, and contains only tipping surfaces existing between the upper and lower boundaries of \mathcal{R} . The following question then arises: how can we detect event points in \mathcal{R} ? Or when is an elementary step applied? Because of the similarity of $\mathcal{R}_{\mathcal{V}}$ and $\mathcal{R}_{\mathcal{H}}$, in the following we consider only the cross-section $\mathcal{R}_{\mathcal{H}}$ of \mathcal{R} . Event points in $\mathcal{R}_{\mathcal{H}}$ are now defined as intersections of tipping curves of interest being either on a boundary segment or inside of $\mathcal{R}_{\mathcal{H}}$, as illustrated in Fig. 6(c). According to its nature, it is called either a *boundary event point* or an *inside event point*. Similarly to the method in Sec. 4.3 if an intersection coordinate (θ, a) satisfies $\theta_{min} \leq \theta \leq \theta_{max}$ and $a_{min} \leq a \leq a_{max}$, then it is verified to be an event point. The algorithm described in Sec. 4.1 deals with any inside event points. In contrary, the boundary event points must be treated separately as follows.

As described in Sec. 4.1, an elementary step at each event point consists of (i) updating the graphs G_{γ_a} and G_{γ_b} according to the change of γ_a and γ_b respectively (explained below) and (ii) building the partial graph δG from G_{γ_a} and G_{γ_b} (see Definition 7).

In [5], we classified inside event points into two cases: simple intersections and degeneracies (see Fig. 7). Figure 8 shows an elementary step at a simple intersection. In [5], the degeneracies are processed by modifying this simple case.

Regarding boundary event points, they can be classified into the following six cases (Fig. 9), which are easily detected by checking the tipping curves intersecting at the event point with the tipping curves in γ and the upper and lower bound sequences U and L . The procedure for handling them in simple cases is explained below, while the degenerate cases are treated similarly to [5] and omitted in this paper due to the page number limitation. As illustrated in Fig. 9, an event point:

- changes the boundary, which is either upper (a) or lower (b);
- does not change the boundary, such that one of the tipping curves
 - goes in (resp. out) \mathcal{R} by the upper boundary (c) (resp. (d));
 - goes in (resp. out) \mathcal{R} by the lower boundary (e) (resp. (f)).

We first explain how to update the cut for (a) and (b). Without loss of generality, let $\mathbf{q} = \{\phi_u, \phi_v\}$ be a boundary event point generated by two tipping curves ϕ_u, ϕ_v and γ, γ' be the cuts before and after \mathbf{q} respectively. Assuming $\gamma = (\phi_1, \phi_2, \dots, \phi_{n-1}, \phi_n)$, if \mathbf{q} is

- the upper boundary, i.e., $\phi_u = \phi_1$ and $\phi_v = \phi_2$, then $\gamma' = (\phi_v, \phi_2, \dots, \phi_{n-1}, \phi_n)$,
- the lower boundary, i.e., $\phi_u = \phi_n$ and $\phi_v = \phi_{n-1}$, then $\gamma' = (\phi_2, \phi_3, \dots, \phi_{n-1}, \phi_v)$.

Similarly, the procedures for updating the cut for (c) and (d) are given as follows. Let $\mathbf{q} = \{\phi_u, \phi_v\}$ be an event point on the upper boundary, i.e., $\phi_u = \phi_1$. We have two cases:

- when ϕ_v goes in $\mathcal{R}_{\mathcal{H}}$, i.e., $\phi_v \neq \phi_2$, then $\gamma' = (\phi_1, \phi_v, \phi_2, \dots, \phi_n)$;
- when the curve ϕ_v goes out $\mathcal{R}_{\mathcal{H}}$, i.e., $\phi_v = \phi_2$, then $\gamma' = (\phi_1, \phi_3, \dots, \phi_n)$.

The procedures for (e) and (f) can be considered in the same way. Fig. 10 illustrates he elementary steps for those boundary event points.

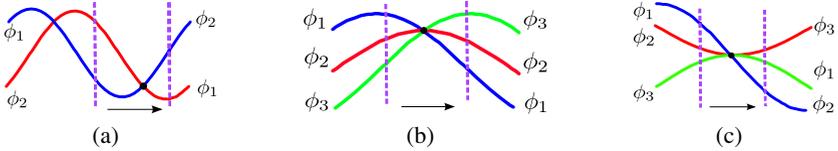


Fig. 7. Inside event point classification: a *simple intersection* if it is generated by only two tipping curves (a), otherwise it is a *degeneracy* (b,c), i.e., when there are more than two tipping curves

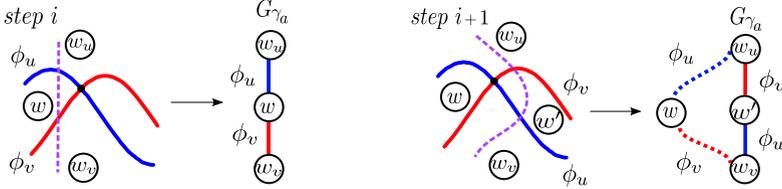


Fig. 8. Updating the graph G_{γ_a} w.r.t. the change of the cut γ_a at a simple intersection

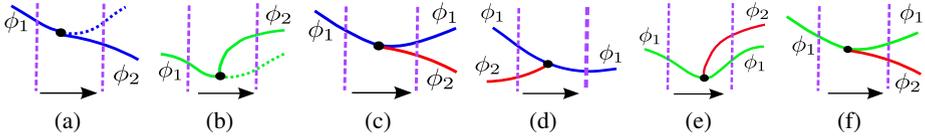


Fig. 9. Classification of simple boundary event points, an event point having a tipping curve that: changes a boundary (a,b), goes in and out by an upper boundary (c,d), goes in and out by a lower boundary (e,f). Upper and lower boundaries are colored in blue and red respectively.

5 Complexity and Experiments

5.1 Complexity Analysis

The space complexity of a DRT graph is proportional to its number of vertices and edges. In absence of constraints, the DRT graph G for an image of size $N \times N$, has a $O(N^9)$ complexity (Property 2). This results from the fact that (i) the number of event points of the whole space is $O(N^6)$, and (ii) at each elementary step (i.e., for each event point), there are $O(N^3)$ vertices generated in the partial graph δG of G [5]. Given one pixel-invariance constraint, some of the potential DRTs become irrelevant. Following the similar proof scheme as above, we can show that the number of event points (i) decreases from $O(N^6)$ to $O(N^5)$ (due to Property 4 in [5] on tipping curves periodicity), and (ii) at each elementary step, $O(N^2)$ vertices are generated instead of $O(N^3)$ since, as explained in Sec. 4, δG of G is generated from two graphs G_{γ_a} and G_{γ_b} of two cuts γ_a and γ_b sweeping on the planes (a, θ) and (b, θ) respectively. Each of the cut intersects at most $O(N^2)$ tipping curves on the plane. Then at each intersection, there are $O(N^2)$ vertices generated in δG . This leads to the following property.

Property 10. *The DRT graph G associated to a digital image of size $N \times N$ under one pixel-invariance constraint has a space complexity of $O(N^7)$.*

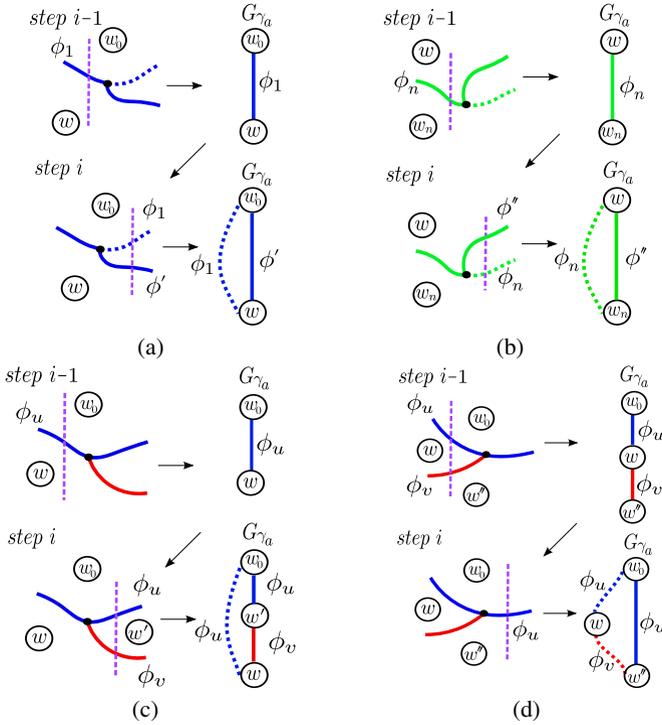


Fig. 10. Illustrations of elementary steps –update γ and generate its graph G_γ – for a tipping curve changing (a) an upper or (b) a lower boundary, and going (c) in or (d) out of an upper boundary

Regarding more complex pixel invariance constraints, we cannot use a similar approach to obtain the theoretical upper bound complexity of G in a FRTS. This is due to the space complexity of G depending on the distance between the involved pixels. Thus, we construct G using our method and we investigate its complexity in practice.

5.2 Computational Experiments

Experiments were carried out on 2D digital images I of size $N \times N$, for 1, 2, 3, 5 and 10 constraints, to investigate how these constraints affect the complexity of the DRT graph. The first experiment (Fig. 11(a)) validates the theoretical $\mathcal{O}(N^7)$ space complexity for one given pixel-invariance constraint. Previous works in [2,7] on discrete rotations provide a complexity of only $\mathcal{O}(N^3)$. However, they consider only a rotation center at a pixel center, while our approach makes no such assumption. In other words, we consider any discrete rotation whose rotation center is located inside a pixel region, due to one pixel-invariance constraint, instead of a pixel center. For this reason, the complexity is increased from $\mathcal{O}(N^3)$ to $\mathcal{O}(N^7)$. Given two pixel-invariance constraints, the bounded FRTS \mathcal{R} varies randomly with the selected corresponding pixels. Results obtained for an image of size 5×5 are shown in Fig. 11(b) for two random pixel choices with some fixed distances. Results for different image sizes are shown in Fig. 11(c). By taking into account the largest complexity for each image size we obtain a worst case complexity $\mathcal{O}(N^{5.5})$ of a DRT graph G in \mathcal{R} .

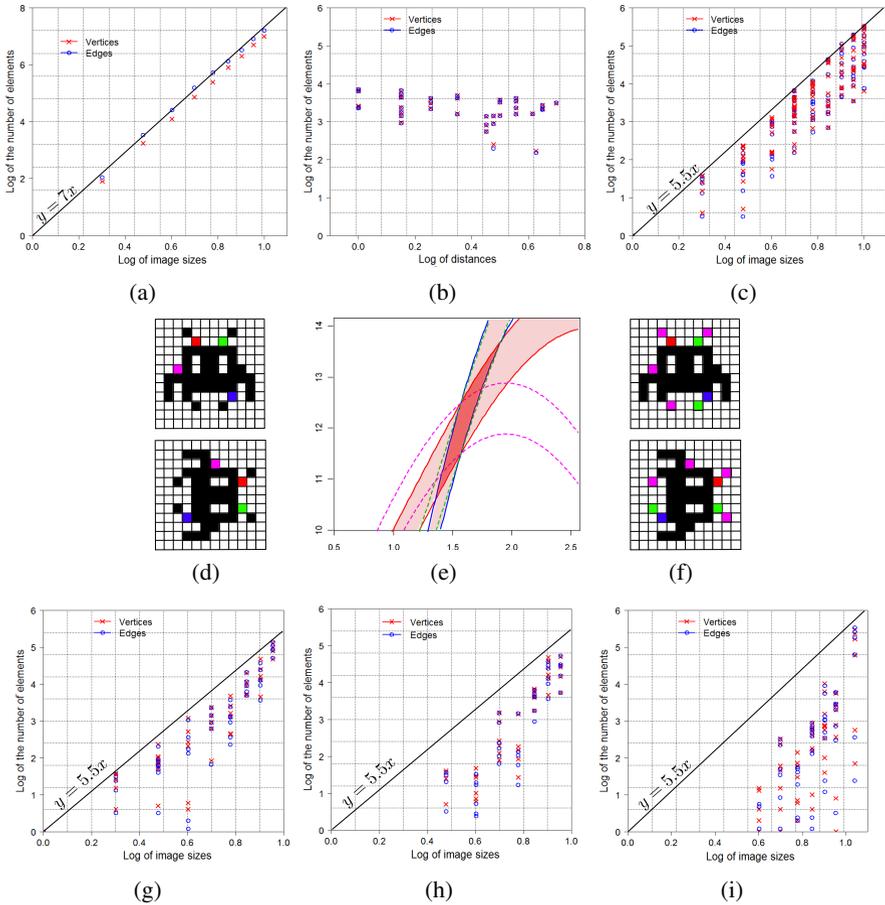


Fig. 11. Experiments for one (a), two (b,c), three (d-g), five (h) and ten (i) pixel-invariance constraints (see Sec. 5.2)

We could expect that \mathcal{R} is strictly reduced when given supplementary pixel-invariance constraints, *i.e.*, more than two constraints. Nevertheless, the third experiment shows that it is not always true. Let us set two pixel-invariance constraints (in red and blue in Fig. 11(d)), and the associated FRTS \mathcal{R} (Fig. 11(e)). A supplementary constraint determined by green pixels in Fig. 11(d) does not reduce \mathcal{R} ; green dotted lines do not pass through \mathcal{R} . Actually (see Fig. 11(f)), there exist two kinds of pixels: one contains some pixels providing no contribution to \mathcal{R} such as the green ones, and another contains pixels allowing to reduce \mathcal{R} such as purple ones. Consequently, the complexity of G depends on given pixel-invariance constraints, but not on the number of constraints.

However, in practice, the more constraints are imposed, the lower the complexity of the DRT graph. This is shown in the experiments for 5 and 10 constraints illustrated in Fig. 11(h) and (i) respectively. Overall there is a downward trend in the numbers of vertices and edges in the DRT graph, though the experimental complexity do not exceed $\mathcal{O}(N^{5.5})$.

6 Conclusion

This article continued the study initiated in [5] by investigating the effects of geometric constraints on rigid transformations of digital images. In this work, we addressed pixel-invariance constraints which consist of specifying the correspondence between points in an initial (sub)space (of \mathbb{Z}^2) and pixels in the transformed space. By enforcing correspondence between one or several pairs of pixels, we consequently restrict the feasible transformations into a parameter subspace, called a feasible rigid transformation set (FRTS), in which all such constraints are satisfied. A proposed algorithm allowed us to build a combinatorial structure (namely a graph) for modeling the subdivision of the FRTS on a subset of \mathbb{Z}^2 of size $N \times N$. We have theoretically analysed the complexity of the graph with one given pixel-invariance constraint to be $O(N^7)$. For more than one constraint, the complexity could not be theoretically calculated. However, using our proposed graph construction method we actually built the graph and experimentally investigated its complexity, which was shown not to exceed $O(N^{5.5})$.

Note that the FRTS is generated from a finite intersection of regions of the imposed constraints. In practice, due to the precision of pixel-invariance constraints, we may obtain an empty set of feasible rigid transformations induced by these constraints. In order to avoid this problem, one solution can be to change the resolution of the images. Namely, we degrade the image resolution as in [7] until we find a non-empty and bounded FRTS.

References

1. Amintoosi, M., Fathy, M., Mozayani, N.: A fast image registration approach based on SIFT key-points applied to super-resolution. *Imaging Science Journal* (2011)
2. Amir, A., Kapah, O., Tsur, D.: Faster two-dimensional pattern matching with rotations. *Theoretical Computer Science* 368(3), 196–204 (2006)
3. Hundt, C., Liškiewicz, M.: Combinatorial Bounds and Algorithmic Aspects of Image Matching under Projective Transformations. In: Ochmański, E., Tyszkiewicz, J. (eds.) *MFCs 2008*. LNCS, vol. 5162, pp. 395–406. Springer, Heidelberg (2008)
4. Hundt, C., Liškiewicz, M., Ragnar, N.: A combinatorial geometrical approach to two-dimensional robust pattern matching with scaling and rotation. *Theoretical Computer Science* 410(51), 5317–5333 (2009)
5. Ngo, P., Kenmochi, Y., Passat, N., Talbot, H.: Combinatorial structure of rigid transformations in 2D digital images. Technical Report, HAL 00643734 (2012)
6. Sharir, M.: Recent Developments in the Theory of Arrangements of Surfaces. In: Pandu Rangan, C., Raman, V., Sarukkai, S. (eds.) *FST TCS 1999*. LNCS, vol. 1738, pp. 1–21. Springer, Heidelberg (1999)
7. Thibault, Y.: Rotations in 2D and 3D discrete spaces. PhD thesis, University Paris-Est (2010)
8. Yilmaz, A., Javed, O., Shah, M.: Object tracking: A survey. *ACM Computing Surveys* 38(4), 1–45 (2006)
9. Zitová, B., Flusser, J.: Image registration methods: A survey. *Image and Vision Computing* 21(11), 977–1000 (2003)