Registration of 3D Point Clouds and Meshes: A Survey From Rigid to Non-Rigid

Gary K.L. Tam[1], Zhi-Quan Cheng[2], Yu-Kun Lai[1], Frank C. Langbein[1], Yonghuai Liu[3], David Marshall[1], Ralph R. Martin[1], Xian-Fang Sun[1] and Paul L. Rosin[1]

Abstract—3D surface registration transforms multiple 3D datasets into the same coordinate system so as to align overlapping components of these sets. Recent surveys have covered different aspects of either rigid or non-rigid registration, but seldom discuss them as a whole. Our study serves two purposes: (i) to give a comprehensive survey of both types of registration, focusing on 3D point clouds and meshes, and (ii) to provide a better understanding of registration from the perspective of data fitting. Registration is closely related to data fitting in that it comprises three core interwoven components: *model selection*, *correspondences & constraints* and *optimization*. Study of these components (i) provides a basis for comparison of the *novelties* of different techniques, (ii) reveals the similarity of rigid and non-rigid registration in terms of problem representations, and (iii) shows how over-fitting arises in non-rigid registration and the reasons for increasing interest in intrinsic techniques. We further summarise some practical issues of registration which include initializations and evaluations, and discuss some of our own observations, insights and foreseeable research trends.

Index Terms—Deformation modeling, digital geometry processing, surface registration, point clouds, meshes, 3D scanning

1 INTRODUCTION

SURFACE registration transforms multiple 3D datasets into the same coordinate system so as to align overlapping components of these sets. The datasets comprise measured points representing surfaces of 3D objects or scenes. Due to limitations of 3D scanning technology, typically multiple datasets must be captured from different viewpoints, each is associated with a different coordinate system. To allow them to be recombined to reconstruct the surfaces that represent the original objects or scenes [1], these data must be registered. *Surface registration* is thus an essential component of the 3D acquisition pipeline and is fundamental to computer vision, computer graphics and reverse engineering. Registering templates to a set of deforming surfaces provides cross-parametrization, and facilitates texture and skeleton transfer, shape interpolation, and statistical shape analysis. Numerous applications also benefit from the continual research on correspondences and registration (e.g. features and saliency), including symmetry detection and articulated object matching, finding object correspondences, fractured object reassembly, sub-part identification, and skeleton and pose construction.

Surface registration may consider *rigid* or *non-rigid* shapes. The former assumes that two (or more) surfaces are related by a rigid transformation. The latter allows deformation (e.g. morphing, articulation) between them. Rigid registration is a challenging problem. Firstly, the data itself poses many difficulties, which may include noise, outliers, and limited amounts of overlap. Noise may take the form of perturbations of points, or unwanted points close to a 3D surface. Outliers are unwanted points far from the surface, which can seriously affect results if not discarded. Limited overlap arises due to different parts of the object being in view in each scan; typically the number of scans is kept low for efficiency, with few points in common between successive scans. Further problems may arise due to self-occlusion when the object is scanned from certain viewing angles. While such problems can be mitigated by careful scanning, they are hard to avoid completely. Secondly, variations in initial positions and orientations (and what is known about them), as well as resolutions of data, can also affect algorithm performance, and must be taken into account when comparing rates of convergence, methods of correspondence determination, and approaches to optimization.

Non-rigid registration is even more difficult, as it not only faces the above challenges but also needs to account for *deformation*, so the solution space is much larger. Unlike the rigid case, where a few correspondences are sufficient to define one candidate rigid transformation for hypothesis testing, both deformation and alignment in the non-rigid case, without strong prior assumptions, often require a lot more reliable correspondences to define. Establishing meaningful and natural correspondences, however, is a challenging problem in its own right. Choice of appropriate representation for the deformation, and suitable

• *Gary K.L. Tam, E-mail: kwok-leung.tam@cs.cardiff.ac.uk*

[1] *School of Computer Science and Informatics, Cardiff University, UK*
[2] *Computer School, National University of Defense Technology, China*
[3] *Department of Computer Science, Aberystwyth University, UK*

tools for evaluation of non-rigid registration methods are two difficult problems. Recent success in rigid surface registration, coupled with the development of scanning devices that can capture time varying surfaces, have brought non-rigid surface registration into focus.

Over the past two decades, many effective rigid registration techniques have been developed. Many of those aforementioned challenges are being addressed [2]. These include [3] which can handle up to 35% noise, and part-in-whole problem, and [4] which can handle up to 40% outliers, and down to 40% overlap as demonstrated in their experiments. Comparatively, non-rigid registration techniques are still in their infancy. Yet many useful techniques have been developed. These include techniques that target articulation, like [5] that can handle large gaps in 4D sequences, reduce high dimensionality of deformations through automatic construction of consensus skeletons and [6] that can register 4D surfaces and produce an *urshape* template; [7] that supports facial capture and animation of avatar heads in *real-time*; and [8], [9] that can handle near- to approximate-isometric deformation. All these have changed the landscape of digital geometry processing, migrating the focuses to dynamic scenes and motions.

The goal of this survey is to overview significant work for new researchers and potential users of registration, to analyse the similarities and differences of these methods, and to provide up-to-date references to this field. While other surveys exist, our specific contributions include the following:

Rigid and Non-Rigid Registration: Our survey covers registrations over the past two decades. We discuss 1) the research development from rigid to non-rigid, 2) the recent advance of intrinsic techniques, and 3) some reasons of such research progress.

Data-fitting and Registration Components: The connection of registration to data fitting and the implications may not be at all obvious to new researchers. We first point out the connection, leading to the discussion of three core registration components in the light of the overfitting problem: *model selection*, *correspondences & constraints* and *optimization*.

Constraints: We carefully study the constraints used in various registration techniques. These properties or techniques limit the search space or improve registration results.

Practicalities and Future work: We discuss various practical issues for potential users such as the choice of functional models, initialization, weight settings, evaluations and finally, consider future directions and trends guided by our analysis of the core components and data fitting.

1.1 Comparsion to other surveys

Among all the existing surveys (Table 1), [11]–[15] focus on rigid registration only. Our survey tries to connect rigid and non-rigid registration, compare and contrast the novel techniques used, and discuss the underlying background to help readers appreciate the developments in the field. [10], [16] cover both rigid and non-rigid registration. [10] focused on applications in medical imaging. Their classification is similar to ours, but they discussed optimization and correspondence together. We particularly separate constraints from optimization to enable a clearer analysis of recent work and the research focuses of the past decade. [16] discussed registration in the more general context of shape correspondence. Our survey, on the other hand, discusses registration from the viewpoint of data fitting. This allows us to focus on components that are more specific to registration techniques. In particular, we highlight the novelties in different components and their uses. We further discuss the problem of over-fitting and provides up-to-date references beyond those in existing surveys.

TABLE 1
Existing Surveys of Surface Registration

Ref	Rigid	Non-Rigid	Focus
[10]	✓	✓	surface registration for medical imaging
[11]	✓		systematic breakdown of ICP and its variants
[12]	✓		registration and fusion of range images
[13]	✓		comparison of several Improved ICPs
[14]	✓		comparison of quadratic approximants and ICP
[15]	✓		coarse vs fine, pairwise vs multi-view alignments
[16]	✓	✓	techniques for shape correspondence

1.2 Scope

This study focuses on registration techniques for 3D point clouds, meshes (representing surfaces) and sparse 3D point data. To limit our scope, various related techniques that require additional information (e.g. silhouettes extracted from images or video sequences, or assumptions concerning viewpoints for the input data) are excluded, as are those which require a specific representation (e.g. range images, 2D arrays of depth points). However, we do include techniques originally designed for range images where they can be applied to point clouds without modification or extra assumptions.

The literature contains many techniques for the closely related problem of 3D *volume image registration*, particularly for medical data. 2D *image registration* is also widely studied. Many kinds of features or landmarks have been proposed, often tailored to the nature of the images, and have inspired the development of surface registration methods (e.g. based on SIFT [17]). Recently, non-parametric image registration techniques have been developed for medical images. These take into account physical properties (e.g. elasticity or viscosity) of tissues, and may also prove useful in surface registration. Readers are referred to [18]–[20] for surveys of such techniques.

1.3 Organization

Our survey is organized as follows. Section 2 establishes the connection between data fitting and registration, and its implications, leading to the discussion of the three key components of registration: *model selection*, *correspondences & constraints* and *optimization*. Section 3 discusses various transformation models. Section 4 classifies various constraints for rigid and non-rigid registration. Section 5 discusses how registration and correspondence determination can be cast in terms of optimization. Section 6 discusses evaluations of registration. Finally, Section 7 concludes with discussions and future possibilities.

1.4 Notation

The following notation is used in the paper. $P = \{\mathbf{p}_1 \dots \mathbf{p}_{N_1}\}$ and $Q = \{\mathbf{q}_1 \dots \mathbf{q}_{N_2}\}$ are two point sets representing overlapping surface portions of a scanned 3D object. $\Sigma \subset \{\mathbf{p}_1 \dots \mathbf{p}_{N_1}\} \times \{\mathbf{q}_1 \dots \mathbf{q}_{N_2}\}$ is a matching relation that denotes the set of all correspondences. We use $a_{ij} = \{\mathbf{p}_i, \mathbf{q}_j\}$ to denote a corresponding pair of points. Sometimes, for illustration purpose, we simplify the notation to $a_i = \{\mathbf{p}_i, \mathbf{q}_i\}$ to denote a one-to-one correspondence, and N is the number of points. $\hat{\mathbf{p}}_i = \chi(\mathbf{a}, \mathbf{p}_i)$ describes a functional model χ with parameter vector $\mathbf{a}$ that maps a point $\mathbf{p}_i$ to $\hat{\mathbf{p}}_i$. $\mathbf{T} = (\mathbf{R}, \mathbf{t})$ is a rigid transformation comprising a rigid rotation and translation, while $\mathbf{A}$ is an affine transformation. $\mathbf{u}$ is a deformation vector. $E(P, Q, \Sigma, \chi, \mathbf{a})$ is an energy functional to be optimized. $p(\cdot)$ is a probability distribution of a random variable. τ is time. DoF means the degree of freedom. χ can be further classified into extrinsic f and intrinsic $g^{-1} \circ h \circ g$ transformations, which will be defined and discussed in the next section.

2 Registration and Data Fitting

Registration has much in common with data fitting. Here, we first provide a general overview of both rigid and non-rigid registration techniques as a whole, in terms of a generalized version of their objective functions. Then we discuss their similarities and differences to data fitting, which leads to our classification scheme and the structure of this survey.

Registration is very often cast as an optimization problem with the objective function of the form:

$$E = E_{\text{data}} + E_{\text{reg}} \tag{1}$$

The data term E_{data} measures the alignment error, and is related to the transformation function χ:

$$E_{\text{data}} = \sum_i \|\mathbf{q}_i - \chi(\mathbf{a}, \mathbf{p}_i)\|^2 \quad \text{s.t.} \quad \{\mathbf{p}_i, \mathbf{q}_i\} \in \Sigma \tag{2}$$

where $\mathbf{p}_i$ belongs to one surface and $\mathbf{q}_i$ to another. Here, we assume χ transforms $\mathbf{p}_i$ to $\hat{\mathbf{p}}_i \in \hat{\mathbf{P}}$ such that the residual sum of squares $\|\mathbf{q}_i - \hat{\mathbf{p}}_i\|^2$ is minimal; $\{\mathbf{p}_i, \mathbf{q}_i\} \in \Sigma$ is the set of all correspondences. The goal of registration is to obtain the best transformation parameters $\mathbf{a}$. E_{reg} is a regularization term that provides additional information to constrain an arbitrary transformation so that the solution $\mathbf{a}$ is a reasonable one. Registration can be further classified into extrinsic and intrinsic techniques (Figure 1).

$$\begin{tikzcd}
P \arrow[r, "f \in F"] \arrow[d, "g \in G"] & Q \arrow[d, "g \in G"] & \text{Extrinsic domain } S \\
P' \arrow[r, "h \in H"'] & Q' & \text{Intrinsic domain } \Psi
\end{tikzcd}$$

Fig. 1. Registration can be classified into extrinsic ($\chi = f$) and intrinsic ($\chi = g^{-1} \circ h \circ g$) technique based on their transformation f, g and mapping function h, see Section 2.

Extrinsic techniques consider a surface lying in Euclidean space. They make use of external properties (e.g. transformations: rotation(s) and/or translation(s)) that can be applied onto the surface. Often, these techniques look for a function $\chi = f \in F : S \to S$ such that surface $P \in S$ can be transformed to align and correspond with surface $Q \in S$, where S is the space of all possible surfaces. We refer to f and S as the *extrinsic* transformation and domain in our context.

Intrinsic techniques, in contrast, consider a surface as free standing. They make use of properties like surface distances and angles which are internal within a surface. Often, these techniques transform/embed the two surfaces P and Q into their canonical forms $P' \in \Psi$ and $Q' \in \Psi$, where Ψ denotes the space of all possible canonical forms, with a function $g \in G : S \to \Psi$, and look for a mapping function $h \in H : \Psi \to \Psi$ so that the surface point $\hat{\mathbf{p}_i} = g^{-1} \circ h \circ g(\mathbf{p}_i)$ matches closely to $\mathbf{q}_i$. We refer to g and Ψ as the *intrinsic* transformation and domain in our context. The objective function may be seen as letting $\chi = g^{-1} \circ h \circ g$:

$$E_{\text{data}} = \sum_i \|\mathbf{q}_i - g^{-1} \circ h \circ g(\mathbf{a}, \mathbf{p}_i)\|^2 \quad \text{s.t.} \quad \{\mathbf{p}_i, \mathbf{q}_i\} \in \Sigma \tag{3}$$

Some techniques may by-pass the embedding process and optimize an objective function which is formulated using intrinsic measures for correspondence establishment. The intuition is similar.

Data Fitting: Given a set of N 3D data points $\mathbf{p}_1, ..., \mathbf{p}_N$, where $\mathbf{p}_i = (x_i, y_i, z_i)$, various data fitting problems may be posed. The simplest is to find the vector of parameters $\mathbf{a}$ of a functional model $z = f(\mathbf{a}, x, y)$ that best describes the data. To measure the fitting error, the usual formulation considers the sum of squared residuals E:

$$E = \sum_{i=1}^{N} \|z_i - \chi(\mathbf{a}, x_i, y_i)\|^2 + E_{\text{reg}}. \tag{4}$$

In general, data fitting involves both selection of χ, and optimization to find the best $\mathbf{a}$. E_{reg} is a regularization term to avoid overfitting.

Connections and Implications: Comparison of Eqns. 1–4 allows us to further analyze registration as three core components:
Model selection (Section 3): The model χ in both cases is often guessed or provided by the users with *a priori* knowledge of the data. In data-fitting, specific models are given depending on the data, e.g., polynomial equations or splines. In registration, χ is usually based on an assumed transformation model, and is highly dependent on the data and applications. We will discuss the choice of these model χ. These cover rigid transformation (Section 3.1), and both extrinsic (Sections 3.2-3.4) and intrinsic non-rigid transformations (Section 3.5).

One of the major concerns in data-fitting is to avoid *overfitting*, when χ fits noise rather than the data. The same applies to registration where the models f of extrinsic techniques usually have high degrees of freedom. This potentially leads to over-fitting problems (in which there are too many parameters to describe the transformation). In these techniques, regularization becomes highly essential, not only to obtain a smooth surface, but also to constrain arbitrary transformation to obtain a reasonable solution. In intrinsic techniques, both P and Q are transformed into a canonical form and the mapping function h is usually low-dimensional. The intrinsic transformation g provides the most important constraint to avoid arbitrary transformation. This is one reason why these techniques have become the recent research focus.
Correspondences and constraints (Section 4): In data fitting, the *correspondences* are incorporated within the data point (x_i, y_i, z_i), but in registration, correspondence information $\{\mathbf{p}_i, \mathbf{q}_i\} \in \Sigma$ *must* be determined. We will discuss different *constraints and techniques* to restrict the solution space $(\mathbf{a}, \Sigma)$. They may be induced by the functional model χ (Sections 4.1.1, 4.2.3) or related to the regularizations (Sections 4.1.4, 4.2.6), correspondences Σ (like features and saliency, Sections 4.1.2, 4.1.3, 4.2.4, 4.2.5), assumptions on data (like template or space-time surface, Section 4.2.2, 4.2.7) and optimization (like search constraint) (Sections 4.1.5, 4.2.8).
Optimization (Section 5): Both registration and data fitting are formulated as *optimization* problems which can be solved in a similar fashion. We will discuss techniques to find the best parameter $\mathbf{a}$ and possibly Σ. When registration is cast as a continuous optimization problem, $(\mathbf{a}, \Sigma)$ are determined iteratively and simultaneously. Some methods focus on correspondence determination only, and Σ is estimated iteratively. The implication here is that many techniques which were originally designed for data fitting (like information-based criteria) may be applicable in registration.

TABLE 2
Transformation Models

Assumed Transformations	DoF	Model Formulation	Examples
Rigid Transformation	6	Euclidean Transformation	all rigid cases
Rigid Alignment with Non-Rigid Correction	7-18	Displacement Field + Space Time Surface	[21]
		Rigid ICP + Thin Plate Splines	[22]
Piecewise Rigid Transformation	$\ll$ N	Predefined Skeleton	[23]
		Bones and Joints	[5], [24]–[26]
General Deformation & Fine Details	3N-12N	Displacement Fields	[6], [27]
		Local Rigid Transformation	[2], [28]
		Local Affine Transformation	[29]–[35]
(Nearly) Isometric Deformation	$\ll$ N	Intrinsic Transformation	[8], [36]–[41]

Our classification is based on these three components. The focus on constraints, in particular, allows us to compare the novelty of various approaches. Since each component interacts with the others, efforts have been made to avoid repeating information, but we do so where necessary for clarity. While not every registration technique fits well into this framework derived from data fitting, these components and constraints do exist in most surveyed methods.

3 Functional Models

We classify the models χ used in registration from the viewpoint of the *assumed transformations*. Table 2 provides a quick look-up for users who need to register surfaces undergoing a specific type of transformations and their DoFs. Four main types are common: *rigid transformation*, *piecewise rigid deformation*, *general deformation* and *(nearly) isometric deformation*. *Rigid alignment with non-rigid correction* is a transitional case between the rigid and non-rigid models.

3.1 Rigid Transformation (6 DoF)

This is the most important assumption of rigid registration: the surfaces are assumed to be aligned by a *Euclidean transformation* involving a rotation and translation $(\mathbf{R}, \mathbf{t})$: $f(\mathbf{p}_i) = \mathbf{R}\mathbf{p}_i + \mathbf{t}$. This transformation (i) is global—the *same* for every point $\mathbf{p}_i$, (ii) can be uniquely defined by three non-collinear pairs of correspondences, and (iii) has low dimension (6 DoF). Using the first two properties, if both point sets contain at least one pair of triplets that match exactly and provide the correct transformation, it takes $O(N^6)$ time to test all possible transformations for sets of N points. By considering all three properties, registration can be cast as a voting problem [17].

3.2 Rigid Alignment with Non-Rigid Correction

Displacement Field of Space-Time Surfaces (7 DoF): To allow non-rigid correction due to subtle movements of objects during acquisition, [21] proposes a space-time displacement model of the form $f(\mathbf{p}_i^j) = (\mathbf{R}_j\mathbf{p}_i^j + \mathbf{t}_j, \tau^j)$. The transformation consists of the unknown $(\mathbf{R}_j, \mathbf{t}_j)$ that aligns surface in consecutive frame j in the space domain, and a translation τ^j along the time axis. This technique requires dense sampling in both the spatial and temporal domain.

Rigid ICP + Thin Plate Spline (18 DoF): Calibration errors or device non-linearity can lead to imprecise alignment of rigid scans. [22] uses a locally weighted ICP to establish reliable correspondences across multiple scans. Then they optimize the locations of these correspondences across multiple scans in a global manner. These correspondences define the anchors such that those misalignments can be corrected by thin-plate spline (TPS) with $3n + 12$ DoF, where n is the number of anchor correspondences, resulting in a sharper final model. Since n is fixed once the rigid transformation is computed, it requires overall 18 DoF.

3.3 Piecewise Rigid Transformation

Some methods, e.g. for human modelling, assume mainly articulated changes, where bones undergo large rigid transformations and local non-rigid surface deformation (bending or stretching) occurs near joints.

Predefined Skeleton: [23] uses a skeleton as a basis for modeling deformation. Each joint is given some DoFs (e.g. joint angles) and is related to other joints by rigid transformations (e.g. a translation).

Bones and Joints: More recent techniques do not use an explicit skeleton. [24] uses predefined bone information to track bone transformations. [25] searches in a finite set of plausible clustered rigid transforms. The small deformations of joints are obtained by blending the transformations of two adjacent bones in the overlap regions [25]: $f(\mathbf{p}_i) = \sum_j w_j(\mathbf{p}_i)(\mathbf{R}_j\mathbf{p}_i + \mathbf{t}_j)$ where $w_j(\mathbf{p}_i)$ is the weight for bone j. Recently, [5] constructs a consensus skeleton automatically from 4D sequence and [26] recovers a dynamic graph—a poseable skeleton with bones, ball and hinge joints.

The advantage of skeletal representation is the low DoF ($\ll N$), which depends on the available bones or joints. It also allow new poses to be created.

3.4 General Deformation and Fine Surface Details

Further techniques consider more generic deformations which include articulation, non-rigid deformation and fine surface details.

Local Displacement Field (3N DoF): [27] allows points to displace freely, and encourages the displacement vectors to point in similar directions by regularizing high-frequency components of the displacement field. [6] uses *meshless finite element deformation*: $f(\mathbf{p}_i, \tau) = \mathbf{p}_i + \sum_j \psi_j(\mathbf{p}_i)\mathbf{u}_{j,\tau}$ which depends on certain nodes j scattered *near* (but not on) the surface, where $\mathbf{u}_{j,\tau}$ are displacement vectors; $\psi_j(\mathbf{p}_i)$ is the influence of node j on $\mathbf{p}_i$. This technique expects as input a sequence of frames, allows adaptive sampling, reduces DoF to $3n$, where n is the number of nodes. [35] uses thin plate spline (TPS) to define a transformation model (DoF 12 + 3N): $f(\mathbf{p}_i) = \mathbf{p}_i \cdot \mathbf{a} + \mathbf{w} \cdot \omega(\mathbf{p}_i)$, where $\mathbf{a}$ represents the affine transformation parameters and $\mathbf{w}$ is the parameter for the TPS smoothness kernel $\omega(\mathbf{p}_i)$.

Local Rigid Transformation (6N DoF): [2] defines a model: $f(\mathbf{p}_i) = \mathbf{R}_i\mathbf{p}_i + \mathbf{t}_i$ using local rigid transformations $(\mathbf{R}_i, \mathbf{t}_i)$ at each point. Nearby points with the same $(\mathbf{R}_i, \mathbf{t}_i)$ are clustered into patches and are updated every few optimization steps. It is adaptive to the deformation; the DoF depends on the number of clusters. [28] uses local rigid transformations, an iterative as-rigid-as-possible deformation technique [42] and space-time surface constraints to align surfaces from densely temporally sampled sequences.

Local Affine Transformation (12N DoF): Local affine transformations are frequently used in non-rigid registration [29], [31]. Higher DoFs allow more freedom to capture fine surface detail changes (e.g. body fat [43], wrinkles [32]). [34] defines a model: $f(\mathbf{p}_i) = \mathbf{A}_i\mathbf{p}_i$ and uses stiffness to ensure adjacent transformations are similar. [33] uses differential coordinates and can be considered as a local affine transformation with a smoothness constraint. [32] uses an *embedded graph node* model: $f(\mathbf{p}_i) = \sum_{\mathbf{x}_j} w(\mathbf{p}_i, \mathbf{x}_j)[\mathbf{A}_j(\mathbf{p}_i, \mathbf{x}_j) + \mathbf{x}_j + \mathbf{b}_j]$ where $\mathbf{b}_j$ is the translation vector of certain nodes $\mathbf{x}_j$, with weighting function w. These $\mathbf{x}_j$ lie on the surface.

3.5 (Nearly) Isometric Deformation

Intrinsic Transformation: Recent techniques emphasize *intrinsic* models, which assume (nearly-) isometric (distance-preserving) deformation. Intrinsic geometry concerns properties like surface distance and angle. Instead of studying $\hat{\mathbf{p}}_i = f(\mathbf{p}_i)$, they study $\hat{\mathbf{p}}_i = g^{-1} \circ h \circ g(\mathbf{p}_i)$ so that $\hat{\mathbf{p}}_i$ matches closely to $\mathbf{q}_i$ (Figure 1). Here g is the assumed intrinsic transformation and h is often a low-dimensional ($\ll N$) mapping function that establishes the point correspondences (Section 2).

Notable and pioneering works include: [36] seeks a low dimensional embedding that preserves all pairwise *geodesic distances*, [37] uses *generalized multidimensional scaling* to embed one mesh in another for partial matching, [38] uses *diffusion distance* and *Gromov-Hausdorff distance* to handle topological noise, and [40] shows that a single correspondence can establish correspondences for all points using the *heat kernel*. [39] pioneers the use of the *Möbius transformation* (isometry is a subgroup of the Mobiüs group). This has stimulated much subsequent work including [8], [41].

3.6 Observations and Discussion

Extrinsic Models: A general observation made by all extrinsic models of non-rigid techniques is that the more high frequency information on the surface, the more DoF required. In the extreme case, non-rigid registration methods that handle large deformation involving articulation or fine detail preservation, use 12 DoF (an affine transformation) per point or node.

It is well-known that using too many DoF can overfit noise. In rigid registration, the overfitting problem is generally ignored because there are 6 DoF only. When non-rigid registrations assume high DoFs, overfitting can become a problem (though not frequently discussed). To alleviate it, [2], [6], [32] reduce the DoF explicitly by computing local transformations with respect to clusters or nodes, and most non-rigid registration methods use regularization to further constrain the solutions.

Although many of these non-rigid deformation models are different in terms of formulation, they all have one common characteristic. Whether they use adjacency similarity, gradient or Laplacian fields [28] (first-order), thin-plate splines [35] (second-order), or various regularizers (Section 4.2.6), they all try to preserve some kind of smoothness. In term of data fitting, this ensures that the registered surface represents the surface details but not noise.

Intrinsic Models: Intrinsic models assume (nearly) isometric deformation. Fine surface details are usually ignored when establishing correspondences, due to their intrinsic nature. Many methods of this kind require mesh data with consistent topology, and may not work on point cloud data. Large topological errors and holes affect these techniques.

This section has considered many transformations. The appropriate model to use depends on the data itself. Users should determine the kinds of deformations and other properties of their data to select a suitable model. This point may easily be overlooked.

We exclude the discussion of similarity transformation, which concerns an additional scale factor compared to a rigid transformation. The major reason is that, in registration, most of the 3D scanned data are assumed to be multiple views of the same object which share the same scale. When registration is required for establishing correspondences across a set of 3D models from heterogeneous data source, scale becomes an important factor. One example includes the registration of sets of faces from different persons [44]. Procrustes analysis [45] is one such important technique.

4 Constraints for Rigid and Non-Rigid Registration

We next classify various constraints, and *roughly* ordered them in decreasing strength of priors: see Tables 3 (rigid) and 4 (non-rigid), and subsequent subsections. Constraints are the properties that limit the search space for transformation and correspondences.

TABLE 3
Rigid Registration Constraints. T: Transformation; SC / DC: Sparse / Dense Correspondence.

Constraints Classifications		Type	Examples
Transformation-Induced Constraints	Distance and Angle Preserving Properties	SC	[46], [47]
	Lower-Bounding	DC	[48]
	Affine Ratio	T	[4]
	Principal Axes	T	[49]
	Closest Point Criterion	DC	[14], [50], [51]
Features	Spin Images	SC	[52]
	Curvature	SC	[53], [54]
	Moments & Spherical	SC	[54]
	Integral Descriptors	SC	[46], [48], [55]
	DCT-DFT Coefficients	SC	[17]
	Cluster Signature	SC	[46]
Saliency	Differential Properties	SC	[56]
	Rareness	SC	[48]
	Geometry Scale-Space	SC	[17]
	Feature Scale-Space	SC	[48]
	Size & Curvature	SC	[53]
	Multi-Scale Slippage	SC	[57]
	MSER	SC	[58]
Regularization	Equalizing Correspondences	DC	[59], [60]
	Maximizing Correspondences	DC	[13], [61]
Search Constraint	Localization	DC	[62]
	Hierarchical Approaches	DC	[56], [62]

In rigid registration, correspondences assist in further pruning the transformation search space, whilst in the non-rigid case, establishing correspondences is the essential step that drives alignment.

4.1 Constraints for Rigid Registration

We consider transformation-induced constraints, features, saliency, search constraints and regularization.

4.1.1 Transformation-Induced Constraints

Transformation-induced constraints are properties related to rigid transformation that can be used to prune the transformation or correspondence search space.

Distance and Angle Preserving Properties: Distance and angle preserving properties of solids can be used for pruning unreliable sparse correspondences established by other constraints. [46] assumes that the Euclidean distance between the centers of a pair of features on a surface, and the angles between their principal directions and their connecting line, are preserved under a rigid transformation. RANSAC-based DARCES [47] uses distance-preserving constraint to enumerate and test three pairs of correspondences in a sphere to limit the transformation search space.

Lower-Bounding: [48] uses (Euclidean) distance-based lower-bounding and the *branch-and-bound* to reduce the correspondence search space (Section 5.2.1).

Affine Ratio: Given two surfaces P and Q in arbitrary initial poses, [4] randomly selects a 4-point tuple from

P and another 4-point tuple from Q, and compares the affine ratios. If they form a correspondence, their affine ratio will match. They further speed up the enumerate-and-test process using a fast range query search of points that satisfy the affine ratios.

Principal Axes: Principal axes, computed through *principal component analysis* (PCA), are three orthogonal directions (eigenvectors correspond to the largest eigenvalues) where the greatest variance of the datapoints lie. A good review, including their use for coarse registration, is given in [15]. A recent advance [49] combines PCA with the least median of squares, providing resistance to noise and outliers. PCA methods are most useful for whole object registration.

Closest Point Criterion (CPC): CPC constrains the potential dense correspondences. Assuming a rigid transformation, it chooses the closest point $\mathbf{q}_i \in Q$ as the match for $\mathbf{p}_i$. It is used in the *iterative closest point* (standard ICP) algorithm [50], which minimizes:

$$E(\mathbf{R}, \mathbf{t}) = \min_{\mathbf{R},\mathbf{t}} \sum_i \|\mathbf{q}_i - (\mathbf{R}\mathbf{p}_i + \mathbf{t})\|^2 \tag{5}$$

This iterative algorithm must be carefully initialized. At each step a new set of parameters (CPC correspondences, rotation $\mathbf{R}$ and translation $\mathbf{t}$) are computed and updated. [63] recently provided a formal treatment of CPC showing that it guarantees the established point correspondences cannot be *too* wrong at each step: it is impossible that some correspondences are of high quality, while their neighbors are of significantly lower quality, because the relationship of correspondences are implicitly captured by the rigid transformation. Many modifications to ICP have been proposed, to improve speed, range and rate of convergence, and robustness—see the earlier survey in [15]. Recent improvements explicitly model inliers and outliers [64], and confidence in correspondences using graduated assignment [59], [60], [65].

There are alternatives of CPC. The first of these was *point-to-plane* ICP. [51] minimizes the shortest distance between a point and the tangent plane of the closest point on another surface. This allows faster tangential movement of the surfaces. [66] studies motion of one surface within the squared distance field of another, and derives a *local quadratic approximant* for the squared distance function; standard ICP uses a linear approximation. They reveal that this local quadratic approximant gives a hybrid between standard and point-to-plane ICP; [14] derives a registration method with quadratic convergence. These methods give faster and more stable convergence with accuracy similar to that of standard ICP. Independently, [67] proposed a similar method based on a generalized Gauss-Markov model which uses a statistical framework to model noise.

4.1.2 Features

Features are quantities (e.g. principal curvatures) that describe a point. Multiple features can be concatenated to form a feature vector. When features of two points (or clusters) match, a meaningful correspondence may exist. Higher-dimensional features offer greater pruning power because of the low probability of matching all components of the feature, but often take longer to compute. Low-dimensional features are quick to compute and is usually used in conjunction with other constraints like saliency (Section 4.1.3). Features are sometimes referred as *point signatures* if they have high descriptive and pruning power. Such features can be defined in a transformation invariant way, and help to establish sparse correspondences. These features include spin images [52], mean curvature [56], Gaussian curvature [53], moments and spherical harmonics [54], integral descriptors [46], [48], [55], FFT and DCT coefficients of distributions of normals [17], cluster signatures of size and anisotropy [46] and SHOT signature [68].

4.1.3 Saliency

Saliency is a measure of local significance in a surface: salient points/regions are those whose properties are unlike most of their neighbours. They are used for key point/region detection, often in conjunction with features (Section 4.1.2). They reduce the size of correspondence space, potential mismatches and obtain more reliable coarse correspondences. These saliency measures include: differential properties [56], rareness [48], scale space analysis on geometry [17] and features [48], size and curvature based on visual saliency [53], multi-scale slippage [57] and maximally stable extremal regions [58].

4.1.4 Regularization Constraints

Regularization involves adding penalty terms to the objective function. It incorporates *a priori* information, improves search, and avoids local minima during optimization. Two regularization terms are used in rigid registration; details of the optimization processes involved are given in Sections 5.2.3 and 5.2.2.

Equalization: [65] casts registration as a continuous optimization problem that handles both rigid transformation and correspondence determination simultaneously. It uses CPC, and formulates a regularization term based on entropy: $-\sum_i \sum_j \mathbf{M}_{ij} \log \mathbf{M}_{ij}$ where $\mathbf{M}$ is the matching probability matrix of all correspondences. When all point matches are *equally likely*, the entropy reaches a maximum. This term can be interpreted as equalizing the importances of point matches, forcing them to the same distance and thus causing the surfaces to run parallel, rather than intersecting one another [59]. From an optimization point of view, this helps to convexify the error landscape of the objective function. [13] uses a similar argument involving equalization of confidence weightings based on mean field variables.

Maximizing Correspondences: [65] uses the term $-\sum_i \sum_j \mathbf{M}_{ij}$ to maximize correspondences and bias the objective function towards *more matches* (i.e. it favors $\mathbf{M}_{ij} = 1$ over $\mathbf{M}_{ij} = 0$). This maximizes the overlapping area.

4.1.5 Search Constraints

Search constraints are for efficiency. These include:

Localization: [62] gives a fast ICP algorithm (Section 4.1.1) based on *localization*. They restrict the search for a new CPC at each iteration to the local neighborhood of the correspondence in the previous iteration, avoiding a global exhaustive search.

Hierarchical Approaches: [62] further combines the above with a coarse-to-fine *hierarchical* search technique. Use of down-sampled data speeds up the first few iterations of ICP; the resolution is gradually increased to obtain more reliable alignment. [56] uses a hierarchy of surface features in the form of points, curves and finally surfaces. They estimate a Euclidean transformation using a few salient points. The transformation is then refined using curves, and surfaces.

4.2 Constraints for Non-Rigid Registration

This section discusses constraints for non-rigid registration, assigning them to categories of markers, templates, deformation-induced constraints, features, saliency, regularization, envelopes of motion and search constraints (Table 4).

4.2.1 Markers

3D *markers* provide sparse correspondences. For example, early commercial systems used simple reference objects (e.g. spheres) clamped to a surface before data acquisition, which are then identified by fitting methods [69]. Other techniques identify markers from an accompanying video sequence using e.g. Laplacian convolution filters [23] or SIFT [33]. Some approaches simply assume correspondences are provided [34].

4.2.2 Templates

Templates simplify the part-to-part alignment problem to part-to-whole alignment, by providing strong structural information including priors for shape, constraints on arbitrary deformation, and connectivity. These greatly help to handle topological noise and missing data. [6], [28], [79] constructs a template, an *urshape*, by accumulating and filling in missing data from other scans in the same sequence. They have been used extensively in non-rigid registration to constrain dense correspondences [23], [29], [32], [34], [80]. This works for objects with well defined forms, such as the human face and body.

TABLE 4
Non-Rigid Registration Constraints.

Constraints Classification		Type	Examples
Markers	3D Objects	SC	[69]
	Optical/Manual	SC	[23], [33], [34]
Template		DC, D	[6], [23], [28], [34]
Deformation-Induced Constraints	Isometry Consistency	SC, D	[2], [8], [9], [36]–[40], [70]–[72]
	Self-Deformation Distortion	SC, D	[70]
Features	Spin Images	SC	[25], [72]
	Curvature	SC	[71], [72]
	Integral Descriptors	SC	[46]
	Heat Kernel Signature	SC	[73]–[78]
Saliency	Extremities	SC	[70]
	Slippage	SC	[57], [71]
Regularization	Orthonormality	DC,D	[32]
	Handling Holes	DC,D	[31]
	Geometric Coherence	DC,D	[6], [25], [27], [29], [30], [32], [35]
	Articulation	DC,D	[24]
	Temporal Coherence	DC,D	[6], [79]
Envelope of Motion		DC	[21], [28]
Search Constraints	Closest Point Criteria	DC	[2], [6], [25], [29], [34], [46], [79]

T: Transformation; D: Deformation;
SC, DC: Sparse / Dense Correspondence

4.2.3 Deformation-Induced Constraints

The deformation model itself may induce constraints.

Isometry Consistency: Generalising the distance preserving constraint in the rigid case, some methods assume inelastic or *isometric* deformation. We first discuss the various definitions of surface distance and then the use of isometry for correspondence.

Definition of surface distances: Surface distances are intrinsic measures defined solely by a metric tensor. Let $(M, \hat{g})$ be a Riemannian manifold with a metric tensor $\hat{g}$. The intrinsic Laplace-Beltrami Operator Δ is determined by $\hat{g}$, and induces a family of surface distances of the form: $d(x, y)^2 = \sum_k \Omega(\lambda_k)(\Phi_k(x) - \Phi_k(y))^2$ where λ_k and Φ_k are the eigenvalues and eigenfunctions of Δ. These include the diffusion distance [38] ($\Omega(\lambda_k) = e^{-2t\lambda_k}$) which measures the degree of connectivity of two points by paths of length t; the commute-time distance [81] ($\Omega(\lambda_k) = \lambda_k^{-1}$) which is a scale-invariant version of the diffusion distance; and the biharmonic distance [82] ($\Omega(\lambda_k) = \lambda_k^{-2}$) which balances the local and global properties of diffusion distance. The biharmonic distance is more shape-aware and insensitive to topology.

The geodesic distance can be defined via the length of curve γ on M: $d(x, y) = \min_\gamma \int_0^1 \|\gamma'(r)\| dr$, $\gamma(0) = x$ and $\gamma(1) = y$ where $\hat{g}$ provides the inner product to measure length. Geodesic distance is sensitive to noise and holes. [83] proposes fuzzy geodesic distances which trades off between precision and stability. [84] approximates geodesic distances by unfolding boundary of holes. The geodesic distance is often found by solving the eikonal equation and wavefront propagation [85]. [86] recognizes the connection between heat and distance, and develops an efficient scheme to find geodesic distance, which is less sensitive to noise and holes: $d^2(x,y) = \lim_{t\to 0} -4t \log h_t(x,y)$ where $h_t(x,y) = \sum e^{-t\lambda_k}\Phi_k(x)\Phi_k(y)$ is the heat kernel.

[75] extends the diffusion distance on 2D manifold to 3D volume. It targets volume-preserving deformation and is found useful in medical imaging. Different Riemannian metrics can be induced by appropriately selecting the arc length. [76], for example, develops an affine-invariant version of diffusion distance by constructing a new Riemannian metric tensor. It is useful for equi-affine-invariant deformation.

Use of isometry: There are many ways to use the isometry constraint, which indicates that the surface distance between two points is the same before and after non-rigid deformation. The direct use of geodesic distance can be seen in [70], where they define *geodesic distortion* in terms of maximum differences of geodesic distances, and in [2], where they define *isometric consistency* in terms of ratios of geodesic distances. [5] extends the latter to work on 1D skeletons for space-time surface registration. There are techniques that formulate correspondences as a minimization of overall geodesic distortion [9]. Isometry can also be formulated as a transformation model (Section 3.5).

Self-deformation Distortion: [70] uses a *self-deformation* distortion measure to avoid unrealistic sparse correspondences arising due to large stretching (e.g. a limb point being mapped to the torso) or repeated structures (e.g. a left leg point being mapped to the right leg). This approach is based on mesh differential deformation techniques and the use of salient points (Section 4.2.5). This work discusses the establishment of meaningful correspondences across deforming objects of quite different kinds.

4.2.4 Features & Signature

Features designed for rigid registration like spin images, mean curvature and integral descriptors (Table 4) have been used to find correspondences in non-rigid registration.

There are isometry invariant features. [73] proposes the *heat kernel signature* (HKS), which is based on heat diffusion at a point. It is intrinsic (invariant to isometric deformation), multi-scale, informative (contains all information about the intrinsic geometry of a shape) and stable. It is used in [40] to define isometric mapping. [87] relates and reinterprets HKS with other spectral shape matching techniques through the *Gromov-Wasserstein distance*.

Over the past years, variants of HKS have been proposed. [74] uses Fourier transform to avoid scale differences and defines a scale-invariant version of HKS. [75] extends HKS to volumetric data, using volumetric distance. [76] defines a new first fundamental form and uses a finite-element technique to define an affine-invariant version of HKS. [77] proposed the Wave Heat Kernel Signature (WKS) which uses a quantum mechanical approach to capture multi-scale details. WKS is shown to be more descriptive than HKS. There is an emerging approach [78] to automatically learn the optimal feature descriptor, in terms of localization, discriminativity and various invariances.

4.2.5 Saliency

Slippage features designed for rigid registration are used in [57] to identify correspondences between non-rigid shapes. [8], [70] use extrema of a deformation invariant function, the *integral geodesic distance* $G(\mathbf{p}) = \int_{\mathbf{s}\in P} d_g(\mathbf{p},\mathbf{s})dP$, to locate salient points. $G(\mathbf{p})$ is large if $\mathbf{p}$ is far away from the rest of the mesh (e.g. at the tip of a finger for a whole body model). This integral is stable under changes of surface detail.

The heat kernel signature (HKS, Section 4.2.4) provides the basis for the definition of salient points and important components for non-rigid shapes. [88] defines "Persistent Heat Signature" for points using persistent homology. The idea is to track and identify critical points (maxima, minima) when the topology of HKS changes over the surface. [89], [90] apply the concept of maximally stable extremal regions (MSER) from computer vision to detect stable regions/components in non-rigid shapes.

4.2.6 Regularization

In Section 3.6, we pointed out the importance of regularization for extrinsic non-rigid registration techniques. Here we survey and classify these techniques.

Orthonormality: [32] uses an affine transformation $\mathbf{A} = [\tilde{\mathbf{a}}_1\tilde{\mathbf{a}}_2\tilde{\mathbf{a}}_3]$ for each node to model fine surface details. In order to encourage the overall deformation model to respect articulation, they maximize local *rigidity*:

$$\sum_{\mathbf{x}_i}\begin{pmatrix}(\tilde{\mathbf{a}}_1^T\tilde{\mathbf{a}}_2)^2 + (\tilde{\mathbf{a}}_1^T\tilde{\mathbf{a}}_3)^2 + (\tilde{\mathbf{a}}_2^T\tilde{\mathbf{a}}_3)^2 + \\ (1-\tilde{\mathbf{a}}_1^T\tilde{\mathbf{a}}_1)^2 + (1-\tilde{\mathbf{a}}_2^T\tilde{\mathbf{a}}_2)^2 + (1-\tilde{\mathbf{a}}_3^T\tilde{\mathbf{a}}_3)^2\end{pmatrix}$$

Handling Holes: [31] avoids the effects of *holes* due to occlusion by minimizing $\sum_i \omega_i^2\|\mathbf{p}_i^j - \hat{\mathbf{z}}_i^j\|_2^2 + \sum_i(1-\omega_i^2)^2$ where $\hat{\mathbf{z}}_i^j$ are depth values. Holes usually have large depth values, causing a high error. The optimizer gives holes small weights ω_i in the first term. The second term prevents a trivial solution with all $\omega_i = 0$.

Geometric Coherence: Various methods encourage geometric properties to vary as *smoothly* as possible, or non-rigid transformation fields to be *coherent*. [28], [33] use *differential coordinates* (first order) and [35] uses *thin plate* splines (TPS) (second order) in non-rigid deformation to provide smoothness. [27] encourages displacement vectors to point in similar directions by analyzing the low-frequency components of the displacement field. [30] regularizes the local affine transformation $\mathbf{A}_i$ of each triangle i using $\sum_i\sum_{j\in\mathcal{N}(i)}\|\mathbf{A}_i - \mathbf{A}_j\|_F^2 + \sum_i\|\mathbf{A}_i - \mathbf{I}\|_F^2$ where $\mathcal{N}(i)$

is the set of neighbouring triangles to i, and $\|.\|_F$ is the Frobenius norm. The first term ensures adjacent transformations are as similar as possible [29], whilst the second term penalizes over-smoothing which leads to unwanted changes in shape. [6] encourages the normals of adjacent points to vary as smoothly as possible using $\sum_i \sum_{j\in\mathcal{N}(\hat{\mathbf{p}}_i)} \left((\hat{\mathbf{p}}_i - \mathbf{p}'_j)^T \mathbf{n}_i\right)^2 + \sum_i \sum_{j\in\mathcal{N}(\hat{\mathbf{p}}_i)} \left(\mathbf{n}'_i - \mathbf{n}'_j\right)^2$ where $\mathcal{N}(\mathbf{p}_i)$ and $\mathbf{n}'$ are the neighborhood and the normal of $\mathbf{p}'$ on the reconstructed shape. The first term encourages points to lie in the tangent plane of their neighbors and the second term aligns adjacent normals.

Articulation Structure: In articulation, points belonging to the same *bone* have the same transformation because of the bone's rigidity, except near *joints*. To ensure smooth transition of geometry of a point $\mathbf{p}$ at joint or to fill in missing data between two bones b_i, b_j, a regularization term $\sum_{\mathbf{p}} \|\mathbf{T}_{b_i}(\mathbf{p}) - \mathbf{T}_{b_j}(\mathbf{p})\|^2$ would be used. [24] further assumes skeletal information is available and [25] finds bone transformations from clustering and uses geodesic consistency.

Temporal Coherence: This is used for regularizing non-rigid registration when capturing sequences of 3D scans. [6], [79] employ *conservation of momentum* to avoid high frequency tangential deformation noise and remove temporal jittering by minimizing $\sum_\tau \sum_i (\mathbf{u}(\mathbf{p}_i, \tau+1) - 2\mathbf{u}(\mathbf{p}_i, \tau) + \mathbf{u}(\mathbf{p}_i, \tau-1))^2$ where $\mathbf{u}(\mathbf{p}_i, \tau)$ is the deformation field of a point $\mathbf{p}_i$ at time τ. They further suggest two regularization terms, $\sum_\tau \sum_i (\det(\nabla \mathbf{u}(\mathbf{p}_i, \tau)) - 1)^2$ and $\sum_\tau \sum_i (\mathbf{u}(\mathbf{p}_i, \tau+1) - \mathbf{u}(\mathbf{p}_i, \tau))^2$ to preserve *volume* and *temporal smoothness*. A deformation field is locally volume preserving if $\det(\nabla \mathbf{u}) = 1$. The temporal term imposes a small penalty on velocity in order to avoid fluttering artifacts in regions that are not well-constrained by data points, such as boundaries.

4.2.7 Envelope of Motion

Correspondences can be computed from *space-time surfaces*: the envelope of motion is given by trajectories of points perpendicular to their normal field. [21] was the first to consider space-time surfaces for both rigid registration and non-rigid correction. Instantaneous kinematics are used to define and minimize the dot product of the velocity field and the normal field. For non-rigid correction, they sample points, compute local instantaneous space-time velocities, and propagate results to neighboring points using regularization. [28] extends the method to larger deformations by extracting an urshape [79], and using a non-linear *as-rigid-as-possible* model [42] for deformation. Space-time surfaces require very dense spatial and temporal sampling of data.

4.2.8 Search Constraints

The Closest Point Criterion (CPC) has been used extensively in non-rigid registration, often in an optimization process similar to that of rigid registration. CPC is applied to find tentative correspondences at each step. Practical optimization issues are discussed later in Sections 5.1.1, 5.5.

4.3 Observations and Discussion

Our classification (Tables 3 and 4) shows that several similar high-level constraints are used in both rigid and non-rigid registration, namely transformation-(deformation-) induced constraints, features, saliency, regularization and search constraints. In particular, many features and saliency have been *reused* in non-rigid registration. They are usually defined on local surface geometry and are (at least, assumed to be) invariant under transformation. Some constraints have been adopted by non-rigid registration, e.g., [46] defines an isometry constraint based on Euclidean distance, which is replaced by geodesic distance in [2].

A significant difference is that while no rigid registration methods use templates, many non-rigid methods do. The template helps to control the extra DoF in the non-rigid case. This works because non-rigid registration is typically applied to specific objects with a well-known form, such as human faces and bodies. Rigid registration is often applied to unknown or much more general objects and scenes, and templates are less useful.

Isometry is a very powerful constraint which is sufficient to define a class of intrinsic transformation models (Section 3.5), and is used in optimization in various ways. From a high level viewpoint, recent intrinsic techniques *enforce* isometric consistency. Extrinsic non-rigid registration allows fairly arbitrary deformation (due to the high dimensionality of the models (an over-fitting problem)) but *encourages* isometry through the heavy use of regularization. In general, most non-rigid constraints (both extrinsic and intrinsic) focus on modelling the relationship between correspondences, whilst this relationship is implicitly captured by the rigid transformation in rigid registration.

For rigid registration, the closest point criterion (CPC) is effective, as the rigidity constraint is imposed by the assumption of rigid transformation. In non-rigid registration, most deformation models allow arbitrary local transformation, and the results may not adequately respect the material properties of the underlying surfaces. CPC *alone* is not likely to be very effective in the non-rigid case. That is why we classify CPC as transformation-induced constraint in rigid case but search constraint in non-rigid case.

Choice of features for feature matching is not well discussed in the literature. Our own experiences and those reported elsewhere can be summarised as follows. Curvatures are fast to compute, but easily affected by noise and incomplete surfaces. Integral invariants are robust to noise and give similar information to curvature. Users must define a suitable ball and grid size for feature matching to get good performance. Slippage analysis is unsuited to smooth surfaces due to its definition. SHOT signatures perform better than spin images [68]. A comparison of several recent features and saliency measures for non-rigid shapes can be found in [91].

Discussion on priors and space-time registration: Here we discuss the constraints of non-rigid registrations in terms of the strength of prior. Markers and templates are explicit prior knowledge provided by users. Transformation (deformation)-induced constraints and *reliable* features & saliency measures are important priors. The transformation-induced constraints are stronger and can often be used for evaluating and pruning correspondences established by the latter. These constraints allow registration to be carried out on highly sparse data, or data from a single view, and often in a pairwise manner.

The focus of space-time surface registration is the use of weaker priors, e.g. various regularizations and envelop of motion. The idea is that more temporal data helps constrain an arbitrary deformation. These techniques usually target sequences, require denser temporal sampling, assume small deformation and large overlap of regions for feature correspondences between consecutive frames. Perhaps the weakest prior is the envelop of motion which requires both dense spatial and temporal sampling. (Search constraints are mostly for efficiency.)

Recent space-time registrations focus on reducing the dimension of the deformation (alleviating the over-fitting problem) by constructing a low-dimensional representation (e.g. skeletons [5], bones and joints [26]). These enable registration techniques to handle sparser data and multiple or all frames at a time, and generate more accurate results. There are methods that focus on global registration in terms of the alignment order of (multiple) sequences to reduce accumulation errors [92].

TABLE 5
Optimization Methods.

	Formulations	Rigid Registration & Correspondence Examples	Non-Rigid Registration & Correspondence Examples
Local Deterministic	Gradient Descent & related methods	[11], [50]–[52], [54], [62], [93]	[24], [28], [30], [34]
	Newton, Gauss-Newton, L-M method	[14], [46], [66], [94]	[2], [31], [32], [79], [80]
	Quasi-Newton		[6], [23], [29], [30]
	Expectation Maximization	[27]	[9], [27]
Global Deterministic	Branch and Bound & Tree Search	[48]	[70]
	Graduated Assignment	[59], [60]	[35]
	Mean Field Annealing	[13]	
	Gaussian Field Framework	[95]	
	Game Theoretical Framework	[96]	
	Spectral Method		[2], [5], [97]
Stochastic	Genetic Algorithm & Simulated Annealing	[15]	
	Particle Filtering	[3]	
	Hough Transform	[17]	[39]
	RANSAC	[39]	[71], [98]
	Belief Propagation		[72]
Con-strained Search	Prune and Search	Sections 4.1.2, 4.1.3	
	Geometric Hashing	[53]	
	Embedding Techniques		[36]–[38], [99]

5 Optimization Methods

Rigid and non-rigid registration techniques are typically cast as non-linear optimization (see Table 5). Many techniques optimize both transformations and correspondences whilst some optimize transformations or find correspondences only.

5.1 Local Deterministic Optimization

Local deterministic methods look for a solution that maximizes/minimizes an objective function locally. These techniques are efficient but also depend heavily on initialization and often converge to local minimum.

5.1.1 Gradient Descent, Newton and related methods

Gradient-descent, *Newton*, *(damped) Gauss-Newton*, *quasi-Newton* and *Levenberg-Marquardt* (L-M) are the most frequently used optimization techniques in rigid and non-rigid registration.

Rigid Registration: The optimal transformation $\mathbf{a}$ at each iteration is updated by (a) taking derivatives of the objective function, (b) linearizing the rigid motion (rotation and translation) [94] or kinematic (helical motion) parameters [66], (c) setting the derivatives to zero and (d) solving the resulting system of equations. [14] points out that standard ICP is similar to a gradient-descent method because it always finds the best transformation locally at each step. In rigid transformation estimation, a closed-form least-squares solution (preferable for efficiency and robustness) exists if three or more distinct non-collinear pairs of correspondences are given [100], which can be solved by SVD, quaternions, orthonormal matrices or dual quaternions.

Non-Rigid Registration: Most methods based on CPC formulate an energy functional with data and regularization terms. We summarize some key implementation issues here. [29] suggests a multi-resolution approach, optimizing a low resolution mesh to find correspondences, then optimizing again on a high resolution mesh. Methods handling large deformations generally start the optimization with small emphasis on the data term because CPC tends to be invalid.

After a few steps, more emphasis is put on the data term because the original shape has been deformed to lie close enough to the aligning shape [29], [30]. This is similar to graduated convexity (Section 5.2.2) except that weights are assigned empirically.

5.1.2 Expectation Maximization (EM)

Some methods solve for groups of parameters in alternate steps, e.g. first fixing correspondences and optimizing transformation parameters, then fixing transformation parameters and optimizing the correspondences. These are EM approaches.

Rigid Registration: [27] represent the centroids of one point set using a Gaussian mixture model (GMM) and align them with another set of data points. The EM algorithm alternates between two steps: (1) it refines the parameters (the transformations and the covariances of GMM), and computes the posterior probability distributions of the GMM centroids through Bayes theorem, (2) computes the rigid transformation parameters that maximize the likelihood.

Non-Rigid Registration: [27] extends the above method to non-rigid registration by regularizing the displacement fields using coherence.

Non-Rigid Correspondence: [9] formulates the problem of shape correspondences as a combinatorial optimization, solved using EM approach, with the goal of minimizing the overall isometric distortion of correspondences all over the surface.

5.2 Global Deterministic Optimization

Local optimization can become stuck in local minima if the initial solution is not close to a global solution. Global optimization tries to find a global solution and to avoid local minima. Registration, when cast as global optimization, does a complete search with bounds (exact global solution), or approximates / relaxes the problem into continuous settings so that a near global solution can be found.

5.2.1 Branch and Bound & Tree Search

A decision tree can be formed where each node is a possible correspondence $\{\mathbf{p}_i, \mathbf{q}_j\}$. The root is $\{\emptyset, \emptyset\}$; the set of nodes forming a path from root to a leaf node is one set of possible correspondences between two surfaces. *Branch-and-bound* is a complete search technique, which is based on a lower bound on the cost function. Its efficiency depends on how tight the bound is, and how quickly it can be computed.

Rigid Registration: [48] uses distance root mean squared error:

$$\mathrm{dRMS}^2(P,Q) = \frac{1}{n^2}\sum_i\sum_j(\|\mathbf{p}_i - \mathbf{p}_j\| - \|\mathbf{q}_i - \mathbf{q}_j\|)^2$$

as a bound on CPC error (Eqn 5); it is both tight and fast to compute.

Non-Rigid Correspondence: [70] represents potential correspondences in a *tree*, and uses a self-deformation distortion measure for a set of correspondences to prune whole branches of a tree search. This method allows *natural correspondences* to be found between two non-rigid 3D shapes of quite different kinds.

5.2.2 Graduated Assignment

The *graduated non-convexity*, which is a deterministic annealing technique, seeks a convex function $E = E_D + \beta E_S$ to approximate the non-linear objective function E_D by adding a regularization term E_S. β controls convergence of the system in a similar way to temperature in simulated annealing.

Rigid Registration: [65] initiated the application of graduated non-convexity to registration. They cast the combinatorial search as a constrained continuous non-linear optimization, with cost function:

$$\begin{aligned} E(\mathbf{a}) = &\sum_j\sum_i \mathbf{M}_{ij}\|\mathbf{q}_j - \mathbf{R}\mathbf{p}_i - \mathbf{t}\|^2 \\ &+ \frac{1}{\beta}\sum_j\sum_i \mathbf{M}_{ij}\log\mathbf{M}_{ij} - \alpha\sum_j\sum_i \mathbf{M}_{ij}, \end{aligned} \tag{6}$$

where α, β are weights, $\mathbf{M}$ is the matching probability matrix, and $(\mathbf{R}, \mathbf{t})$ is a rigid transformation. The second and third terms are regularization terms (see Section 4.1.4). β is the inverse computational temperature of the system. This method can be sped up by k nearest neighbors [59] and CPC [60].

Non-Rigid Registration: [35] defines a transformation model using thin plate spline, and uses graduated assignment for non-rigid registration and optimization.

5.2.3 Mean Field Annealing

Mean field annealing is a deterministic approximation to the simulated annealing technique.

Rigid Registration: [13] minimizes the cost function:

$$\begin{aligned} E(\mathbf{M},\mathbf{R},\mathbf{t}) = &\sum_i m_i\|\mathbf{q}_i - \mathbf{R}\mathbf{p}_i - \mathbf{t}\|^2 \\ &- \frac{1}{\beta}\prod_i m_i - \frac{1}{\beta}\frac{\gamma-1}{\gamma}\sum_i m_i^\gamma \end{aligned}$$

where m_i is the mean field variable and is a vector in $\mathbf{M}$ (the matching probability matrix). The first term is the data term and the others are regularization terms (see Section 4.1.4).

5.2.4 Gaussian Fields

The *Gaussian field* framework replaces a non-linear, non-smooth error function by a convex smooth approximation, with the help of mollification: a summation of a large number of Gaussian functions which, by the Central Limit Theorem, leads to a single Gaussian. Tracking the global extremum is much easier as the error is differentiable even for noisy input data.

Rigid Registration: [95] uses the objective function:

$$E_\sigma(\mathbf{a}) = \max_{\mathbf{R},\mathbf{t}}\sum_i\sum_j \exp\left(-\frac{d^2(\mathbf{p}_i,\mathbf{q}_j)}{\sigma^2}\right)$$

where $d(\mathbf{p}_i, \mathbf{q}_j) = \|\mathbf{q}_i - \mathbf{R}\mathbf{p}_j - \mathbf{t}\|$. σ works like the temperature in annealing technique. This method shows robustness in the presence of strong noise.

5.2.5 Game Theoretical Framework

Finding a large set of correspondences that express a high level of mutual compatibility can be formulated as an inlier selection process of a non-cooperative game [101]. The inlier selection is modelled as two players competing with one another.

Rigid Correspondences: [96] uses game theory to recognize rigid objects in a cluttered scene. They use the consistency of Euclidean distances between point pairs to model a competitive game. To obtain reliable correspondences, they define two steps: 1) to run the competitive game to obtain a set of reliable sparse correspondences, 2) to propagate the sparse set to a dense set using neighbouring information.

5.2.6 Spectral Correspondence

A spectral method [97] takes into account the relationship between points and correspondences, and formulates it as an eigen-decomposition problem.

Non-Rigid Correspondence: [2] extends the concept to the non-rigid case. Two relationships are considered: (i) the similarities of two points (as difference of principal curvatures), and (ii) the isometric consistency (as difference of geodesic distances) of two pairs of correspondences. Having defined an affinity matrix of correspondences, they compute its principal eigenvector. The eigenvector entries define the confidence values of correspondences that belong to a coherent cluster. This helps to prune unreliable ones.

5.3 Stochastic Optimization

Many techniques model registration and the correspondence determination problem with statistics and probabilistic approaches in order to handle datasets with noise, outliers and missing data.

5.3.1 Genetic Algorithm and Simulated Annealing

Genetic algorithm adapts the idea of evolution and natural selection [102]. *Simulated annealing* is inspired by thermodynamics. These rigid techniques are discussed in [15].

5.3.2 Particle Filtering

Particle filtering is a recursive Bayesian estimation technique to estimate the posterior distribution of a state given past observations. Its strength lies in its ability to handle non-linear functions and non-Gaussian noise. It is often used in object tracking.

Rigid Registration: [3] sees the optimization of rigid registration as a sequence of steps to arrive at a global solution. They use particle filtering to *track* the optimal state (the rigid transformation). The method handles noise, arbitrary initial estimates, missing data, and both sparse and dense point sets. At each iterative step, the algorithm draws N_s particles (here, initial transformations) from the proposed distribution, and feeds them into a local optimizer (e.g. ICP) to produce candidate solutions. Importance weights are adjusted. New particles are predicted, and the process repeats. Multiple particles help to avoid local minima.

5.3.3 Hough Transform (Voting)

The *Hough transform* was originally developed for line detection in images, and was later developed to detect other shapes and find transformation parameters. It operates on a quantized parameter space using an accumulator.

Rigid Registration: [17] identifies salient points, generates candidate transformations, votes for their parameters in the accumulator space, and picks the peaks after all votes have been cast. The accuracy of the method depends on the accumulator resolution, so ICP is subsequently applied for fine alignment.

Non-Rigid Correspondence: [39] uses voting for non-rigid correspondences. The key novelty is to project a genus-zero topology mesh in a mid-edge parametrization so that a global Möbius transformation can be used, as for a rigid transformation. The method repeatedly samples three points, applies the Möbius transformation and tests if the transformation is valid or not using isometric consistency. This method votes for correspondences.

5.3.4 RANSAC

Random sample consensus (RANSAC) is a robust algorithm for fitting models in the presence of many outliers. It follows a hypothesis-and-test paradigm.

Rigid Registration: [4] considers rigid registration as a largest common pointset (LCP) problem. Given two point sets P and Q, RANSAC finds the largest subset $P' \subset P$ such that the error between $\mathbf{T}(P')$ and Q is less than some predefined threshold, where $\mathbf{T}$ is a rigid transformation. The affine ratio constraint reduces the number of possible correspondences.

Non-Rigid Correspondence: [71] introduces importance sampling into RANSAC. Potential correspondences are drawn from a probability density function. Isometric consistency is formulated as a likelihood to adjust the posterior probability such that more reliable correspondences are determined in the next iteration. Results are improved using a tangent-space optimization technique that moves points along the mesh. [98] uses an entropy-based planning strategy to select important matches and reduce sampling cost.

5.3.5 Message Passing and Belief Propagation

Random variables, which have a Markov property, can be represented as a graph in a *Markov random field* (MRF), where the undirected edges model the *interdependencies* between random variables.

Non-Rigid Correspondence: [72] uses MRF to model the inter-dependency of correspondences, aiming to find correspondences that maximize a certain joint probability distribution. Two probability potentials are considered: (i) discrepancies between the compressed spin images of two points, and (ii) isometric consistency between two pairs of correspondences. The simplifying assumption of MRF is that all probabilities are independent of others and can be multiplied together. They apply *loopy belief propagation* to find a (local) maximum of probability.

5.4 Constrained Search

Constrained search establishes special constraints to limit the search space. We summarize them as follows.

5.4.1 Enumerate and Prune

Enumerate and prune is the usual strategy to obtain sparse (or initial) correspondences. In general, it is impossible to enumerate all correspondences. Saliency, features (Sections 4.1.2, 4.1.3) and other constraints are usually employed to prune the huge search space.

5.4.2 Geometric Hashing

Geometric hashing pre-computes and stores a transformation invariant representation of a patch in a hash structure. The hash bins are used to restrict the number of potential matches.

Patch Correspondences: [53] extracts patches using salient geometric features and hashes the transformation invariant representations. The representation has complexity $\binom{|P_s|}{3}$ as it takes every point-triple from the patch P_s containing $|P_s|$ points. Given a new patch Q_s, a new representation is formed and found in the hash structure in $\binom{|Q_s|}{3}$ time. The key is to break the correspondence problem of complexity $\binom{|P_s|}{3} \cdot \binom{|Q_s|}{3}$ into two sub-problems; this allows simultaneous processing of all hashed patches. The method has huge memory requirements.

5.4.3 Embedding Techniques

Non-rigid Correspondence: A class of techniques use *embedding* to establish correspondences. [36] computes and aligns eigenvectors of a geodesic kernel. Correspondence is established through exhaustively pairing (sign-flip, eigen-mode switch) and non-rigid alignment of the leading 6 eigenvectors. These eigenvectors capture the low-frequency structure of the meshes. [99] extends the concept further by defining histograms for these eigenvectors for quick pairing. [37], [38] embed a mesh directly in another, by using *generalized multi-dimensional scaling* (GMDS) and *Gromov-Hausdorff distance* for partial matching [103]. In general, this class of work is inspired by [104] which embeds non-rigid models into some canonical space for surface matching. Once the intrinsic geometry is embedded (through $g \in G$, see Figure 1), the correspondence problem ($h \in H$) reduces to a standard rigid or low-dimensional non-rigid matching and alignment problem ($f \in F$) [36], [103].

5.5 Initialization

Global deterministic optimization, stochastic methods and constrained search generally do not require *initialization*. Here we focus on local iterative optimizers that do need an initial estimate.

Rigid Registration: [67] interactively selects three common points to define an initial rigid transformation. [46], [52] use reliable correspondences to automate this process. [14], [54], [59], [60], [62], [64], [94] test their methods on a set of predefined initial transformations within a certain range, and evaluate the range of convergence. In general, the broader and more stable the convergence funnel, the better the method [94]. Certain methods are designed to obtain coarse alignments (e.g. [17], [49]; see [15]).

Non-Rigid Registration: [22], [32] use rigid registration to find an initial estimate. [2] tries multiple initializations if the error of the first is too high. [25] finds the initial transformation of a rigid sub-part by looking for clusters in transformation space defined by all possible correspondences. For a sequence of 3D scans, adjacent frames are roughly aligned, and correspondences can be tracked [6], [21], [28], [33], [79]. Other approaches [23], [24], [29], [30], [43] assume markers or specific segmentation information are provided.

5.6 Observations and Discussion

Rigid and non-rigid registration are typically cast as *optimization problems*. The error landscape depends on the type of data being registered, outliers, noise, and missing data. We summarize some observations and experiences of our own and from the literature.

Rigid Registration: If the surfaces are relatively clean and there is a good initial estimate of alignment, *local optimizers* (e.g. ICP or gradient descent, Newton's method) are the most efficient choice. However, if there is significant noise, or the initialization is poor, these methods may not converge [94]. *Global deterministic optimization* (e.g. the Gaussian field framework [95]) or *stochastic* techniques (e.g. particle filtering [3]) are more reliable in these situations. Stochastic techniques do not guarantee a globally optimal solution, but the solutions are usually good. In practice, using a fast stochastic technique (e.g. RANSAC [4]) to obtain an initial coarse alignment, and a local optimizer for subsequent fine alignment is a good approach. Features and saliency are important for fast pruning and establishing coarse alignment. Hierarchical methods can also find an initial alignment, and optimize it during fine alignment.

Non-Rigid Registration: Extrinsic non-rigid registration usually uses *local optimization* for efficiency. In particular, the quasi-Newton method can avoid the need to invert a large Jacobian or Hessian matrix. However, local optimization typically requires good initialization. Three techniques are used to make local methods work well (i) regularization, (ii) correspondences and (iii) pre-registration. Regularization is useful because it convexifies the error landscape of the objective function. It is often used with decreasing weights in successive steps to avoid premature convergence at a local minimum. Reliable correspondences drive optimization directly and pre-registration or coarse alignment are used to provide good initialization (Section 5.5).

Recent advances in intrinsic techniques study isometric consistency and intrinsic geometry. Many of them use probabilistic techniques, and the solutions are usually good. Over-fitting is less of an issue in these techniques, but they mainly target (nearly-) isometric deformations. Computing surface distances on a large set of points, and handling topological noise and holes are important issues. Sub-sampling of the surfaces may be required.

Quadratic Assignment Problem: Some optimization techniques are related to the intractable quadratic assignment problem (QAP). We summarize them below. QAP considers allocating a set of facilities to a set of locations, with the cost being a function of the distances d and flows f between the facilities, together with the costs b associated with a facility being placed at a certain location. The objective is to find the mapping ϕ which assigns each facility i to a location $\phi(i)$ such that the total cost is minimised:

$$\arg\min_{\phi} \sum_i \sum_j f_{ij} d_{\phi(i)\phi(j)} + \sum_i b_{i\phi(i)}$$

The intuition is to arrange ϕ so that facilities i, j with high flow f_{ij} are assigned to locations $\phi(i)$ and $\phi(j)$ which have small distance $d_{\phi(i)\phi(j)}$. QAP is NP-hard, and all existing techniques approximate or relax the solution one way or another. [65] approximates the QAP by linearising the cost function via Taylor series and using two-way constraints to convert the QAP into a continuous search. Inspired by this, graduated assignment (e.g., [35], [59], [60]) and mean field annealing [13] techniques are further developments with additional constraints. These techniques use a rigid transformation instead of the quadratic term $f_{ij} d_{\phi(i)\phi(j)}$ to model the pairwise relationship of correspondences (facilities & locations in QAP). To explicitly model such a relationship, spectral techniques [2], [97] and a game theoretical framework [96] consider the objective function: $\mathbf{x}^T \hat{\mathbf{M}} \mathbf{x}$ where $\hat{\mathbf{M}}$ summarizes the QAP cost above. Both techniques relax $\mathbf{x}$, the binary assignment, to a continuous one. The difference is that the former constrains $\|\mathbf{x}\|_2 = 1$ (due to the use of Rayleigh quotient); whilst the latter ensures that $\|\mathbf{x}\|_1 = 1$. Markov random field approaches use probability potentials to model the inter-dependences of correspondences using geodesic distance. Inspired by these pairwise techniques, higher-order matching techniques (e.g. [105], [106] which consider ternary relationships) are emerging.

6 EVALUATION

Here, we discuss common evaluation approaches and survey those datasets that are commonly used.

Rigid Registration: The most common way to evaluate rigid registration is to compare the deviation of the rotation and translation parameters directly because the rigid transformation is low-dimensional (6 DOFs). These parameters can be obtained from synthetic datasets, where additional noise, outliers and different level of overlaps may be added/adjusted. There are several publicly available datasets with ground truth (e.g. [107], [108]). The ground truth parameters are defined during the scanning. These datasets are usually small in size. For large scale datasets (e.g. [109]), ground truth are not available, and visual assessment (e.g. sharpness) is required. The range and rate of convergence are also common measures.

Non-Rigid Registration: Evaluation of non-rigid registration is a significant issue, since non-rigidity is difficult to formulate and can most directly be assessed in terms of correspondences. Earlier techniques use visual assessment. [6], [28], [32] used parametrization to demonstrate low visual distortion. [39] uses colour coding to show sparse correspondences.

Recent techniques use public datasets for evaluation. Several datasets are available: 1) benchmark for 3D mesh segmentation [110] (e.g. Ballerina). These meshes have different numbers of vertices, connectivity and severe non-uniformity of sampling. 2) non-rigid world dataset [111] contains various animal and human mesh models with different poses, where each object has ≈3K vertices with arbitrary connectivity. 3) dataset created by existing techniques, with fixed mesh connectivity, and uniformly sampled at high resolution (9-30K) [112] (e.g. Jumping Man), [72] (e.g. Dancing Man), [30] (e.g. Horse Gallop)—created by [113]. These ground truth correspondences are obtained by fitting a common template on the registered data. 4) TOSCA high-resolution dataset is constructed with ground truth correspondences, with fixed number of vertices and connectivities.

When there is no ground truth correspondences, evaluation of techniques using dataset types 1 and 2 are usually based on visual assessment by connecting spare correspondences with lines [9]. Given ground truth correspondences in dataset types 3 and 4, additional metrics can be defined. [8] uses sum of isometric distortion—difference of geodesic distances between the end-points of one correspondence with respect to all the other correspondences. This can be visualized per vertex, assuming there is a one-to-one mapping. [9] sums these isometric distortions all over the surface to obtain one single number. Note that isometric distortion is meaningful only if the datasets are undergoing a (near-)isometric deformation. Note also that the evaluation of new techniques on dataset type 3 may actually compare their performance against the techniques that create the dataset.

A benchmark database [91] based on TOSCA further allows evaluations on isometry, topology, holes, scale, sampling, and noise of features and saliency. [8] also recruited a volunteer to label 10-35 semantically meaningful and consistent feature points for each class of objects in the watertight meshes dataset [114].

In the computer vision literature, [115] discusses how to evaluate registration without ground truth. One way is to evaluate the results by *generalizing* them to unseen data (a well known method in data-fitting). Given ground truth data, receiver operating characteristic (ROC) curves, which are plots of true positive rate against false positive rate, may be used. In general, obtaining ground truth data is difficult. One idea is to use image-based techniques [116] and compare the results with non-rigid registrations.

7 Challenges and Future Directions

In this section, we discuss challenges and possible future directions in registration.

Deformation Models: Section 3 surveyed various transformation models for registration. Most existing methods assume a *single model* for the two surfaces being aligned. This is practical, enabling optimization to simultaneously find the best transformations for all points. However, a surface comprises different materials and may undergo different kinds of deformation. For example, arms and legs (mainly) undergo articulations, whilst deformation of a skirt leads to folds and wrinkles, and in such cases multiple deformation models might be more useful. The main difficulties of doing so are the need for surface segmentation, constraints between adjacent models, and approaches to merge registration and deformation results.

In general, non-rigid registration is problem- and data-dependent. Users and researchers should carefully consider the nature of the deformation. As well as *blendshape* models [117], physical models relating tissues, bones and muscles may be useful in 3D face registration. This area is yet to explore, possibly due to the complexity of facial physiology. Related techniques [118] are emerging.

Over-Fitting & Model Selection: In data-fitting, one does not want the model to describe the noise (i.e. to over-fit the model). Similarly, non-rigid registration should align 3D surfaces rather than noise or holes. Over-fitting often is not a concern in rigid registration as the rigid transformation is low-dimensional, and the same at all points. Non-rigid registration, however, requires higher DoF. (e.g. local affine transformations have $12n$ DoFs, where n is the number of points or nodes.)

It is well-known theoretically that such high dimensionality can lead to over-fitting. Ultimately, two general surfaces can always be brought into alignment with arbitrarily small error. However, the resulting alignment need not represent physically correct correspondences between homologous structures on the surface. Indeed, simply comparing the residual errors or the rate of convergence for two non-rigid extrinsic registration techniques tells us little about which produces the more meaningfully correct result.

To alleviate the over-fitting problem, existing extrinsic techniques encourage smoothness through (i) formulating a priori knowledge or assumptions as regularization, (ii) introducing additional equations by some linear constraints on neighborhoods or (iii) uses a reduced representation (e.g. embedded graph nodes, meshless element models). Recent intrinsic techniques [36], [38], [39] use low dimensional embedding to by-pass the high-dimensionality problem by assuming a strong inelastic isometry constraint.

A closely related issue is how to automatically determine the most appropriate deformation model: in nearly all existing work, deformation models are predetermined and thus require *a priori* knowledge. However, non-rigid registration is similar to data fitting, where there are many well-established techniques for model selection: *cross-validation*, *early stopping* and *model comparison* (e.g. information-based criteria). They could potentially be used to develop a fully automatic method. In information theory, the *minimum description length* criterion is well-known for determining the best model from data; it has been used in image registration [119], face (texture and depth) registration [120] and medical data registration [121] in computer vision, but is relatively little used in computer graphics and shape analysis.

Constraints and Correspondences: Meaningful correspondences should relate homologous sections of the shapes, i.e., sections that are semantically related. Although many techniques and constraints have been proposed to facilitate correspondence search, human assistance and visual inspection are sometimes required to provide semantic input. As we have noted, there are no precise definitions and methods for evaluating quality of correspondences. Such definitions and measures are needed that match human perception or are based on physical behavior.

More effective pruning constraints and techniques are always beneficial. Recent research in intrinsic geometry and isometric consistency has led to important advances for isometric deformation, especially for face and body modeling. Here, we try to bring the analogies with machine learning and provide insights to why these intrinsic techniques are important.

Connection to Machine Learning (ML): There are two kinds of learning methods in ML [122]: inductive and analytical learning. Inductive learning (e.g. neural network) generalizes hypotheses, in terms of a function, from training examples to unseen data. A hypothesis is assumed to be found/fitted empirically from a large amount of training data. The justification is based on statistical inference. The advantage is that it requires little prior knowledge, but the hypothesis may be incorrect if there is not sufficient training data—overfitting. Analytical learning uses domain theory (e.g. rules) and deductive inference to derive hypothesis. It does not require lots of training data but is seriously affected if the domain theory is incorrect.

A comparison of registration techniques reveals that rigid registration and intrinsic techniques are based on strong prior knowledge (or at least, assumptions). Having a rigid transformation requires *all* points to transform rigidly. Intrinsic transformations require the two surfaces (and *all* points on them) to satisfy isometric consistency. These transformations are similar to the *domain theory* in analytical learning—a global rule, one reason why these techniques can be cast as optimization problems and solved using the Hough transform [17], [39] or RANSAC technique [4], [71] efficiently. This strong *domain theory* leads to two examples which demonstrate the successful application of typical deductive inference: 1) [40] suggests that fixing one correspondence, all other correspondences can be inferred for non-rigid shapes. 2) [98] further answers how many correspondences are required to constrain deformations of non-rigid shapes.

For general extrinsic non-rigid registrations with high DoF, this strong *domain theory* is lost. This is similar to inductive learning with limited prior knowledge. This type of technique assumes the deformation (hypotheses) can be learnt from sufficient training data. Yet, given only two surfaces and a high-dimensional deformation model (like [6], [32]), the solution can only be regularized by other prior knowledge—mostly in the form of regularizations to ensure smoothness. One way of providing more training data is to look into the temporal domain. This explains why recent extrinsic non-rigid techniques focus heavily on registering space-time sequences.

The pitfall of analytical learning is that it depends on a perfect domain theory. For non-rigid shapes, many deformations are only near/approximate-isometric, which may lead to matching ambiguities [8]. In cases with large topological changes, the isometry assumption may fail [38]; handling such issues is a challenging task. Similarly, the rigid transformation is not perfect (i.e. cannot always explain the optimal solution) under noise, outliers and limited amount of overlap. Recent efforts of the community is to model and study these factors (e.g. Section 4.1.4). These approaches may be considered a *combination* of inductive and analytical learning techniques [122]. Learning from rigid registration in this perspective may be a promising direction.

Optimization: Optimization is an important issue in registration as it affects accuracy, robustness and convergence. Many optimization techniques used in rigid registration have yet to be applied to non-rigid cases. A major obstacle to adapting these techniques is the high dimensionality, causing most extrinsic approaches to use quasi-Newton techniques instead of the more stable Newton or L-M methods. The former, however, may not converge if the initial alignment is too far from the optimal solution. Stochastic approaches are promising methods that are not sensitive to initialization. By transforming non-rigid shapes into the intrinsic domain, many stochastic methods have found applications in establishing correspondences, and often near globally optimal solutions are obtained. Adapting these techniques to large scale data, deriving specific objective functions, handling various levels of geometric and topological noise, and extending them to non-isometric and elastic deformation are some of the challenges.

Datasets: Recent research has focused on analyzing sequences of 3D scans: (i) hardware for capturing such sequences at a high frame rate is now readily available, and (ii) sequences assist registration by providing intermediate information, and thereby constraints. However, most techniques can cope with only limited sequences (15-200 frames), mainly due to memory limitations. Faster techniques, and especially ones that can handle more and larger data, are an ever-present goal. Groupwise registration techniques have been developed [33], [123] which are typically more reliable, and reduce the accumulated errors that arise in pairwise registration. Further challenges lie in (i) avoiding the need for markers or templates, and (ii) capturing and processing fine scale details [32]. As has been found in medical imaging, registration utilizing heterogeneous information (e.g. geometry and texture [41], [120]) may give higher accuracy.

Foreseeable Trends:

Rigid Registration is becoming more robust and reliable over the past two decades of researches. Many challenges posed by rigid registration have been addressed [2]. We foresee that rigid registration is becoming application-oriented (e.g., [124]). Developing rigid registration techniques that can handle large datasets with different levels of details and regularity structures is essential. With the advance of technology, handheld devices will soon possess 3D scanning functionality. Real-time techniques that target these ubiquitous devices are new directions.

Non-Rigid Registration, in general, is still at its infant stage, largely because of the large varieties of transformations in the real world and the shortage of our knowledge of these transformations. One direction is to investigate these deformations from a reverse engineering point of view. Many existing non-rigid techniques are still in the stage of proof-of-concept—

showing what can be done. Developing and capturing more ground truth data for each type of these transformations would help mature these techniques.

Evaluation of correspondences and transformations are two important areas. Defining the semantic meaning of correspondences is a hard problem in itself, and currently no work considers the evaluation of transformations. Computation efficiency is another issue. Perhaps, one can employ machine learning to learn the nature or pattern of unreliable or incorrect correspondences. This can speed up optimization to obtain reliable solution.

We also foresee two directions regarding intrinsic techniques: 1) Existing techniques focus mostly on modelling near-isometric deformation and obtaining a global consistent set of correspondences. Relaxing the concept to *local* isometric deformation but ensuring global consistency would be one direction because most real-life deformations are globally non-isometric (e.g, due to deformation at joints) but locally isometric. 2) Most techniques focus on registering surfaces of similar sources. If one wants to register surfaces from heterogeneous sources (e.g., two totally different human faces and their expressions), these techniques do not work. This relates to the morphing problem [30], and is useful for medical applications.

Non-isometric deformations need further consideration. [8] has started in this direction using blended conformal maps. [41] combines intrinsic (conformal map) and extrinsic measures (e.g. texture, Gaussian maps, curvatures) to handle non-isometric deformation. Consideration of elastic and plastic properties of materials is more commonplace in medical imaging, and needs further work in this community, but one foreseeable challenge is the lack of ground truth for evaluation. Designing reliable non-rigid registration for large dataset and for ubiquitous devices (e.g. Kinects) is important.

ACKNOWLEDGMENT

We thank the associate editor and all the anonymous reviewers for their invaluable comments and insights.

REFERENCES

[1] V. Lempitsky and Y. Boykov, "Global optimization for shape fitting," in *Proc. CVPR*, 2007, pp. 1–8.

[2] Q.X. Huang *et al.*, "Non-rigid registration under isometric deformations," *CGF*, vol. 27, pp. 1449–1457, 2008.

[3] R. Sandhu *et al.*, "Particle filtering for registration of 2D and 3D point sets with stochastic dynamics," in *Proc. CVPR*, 2008, pp. 1–8.

[4] D. Aiger *et al.*, "4-points congruent sets for robust pairwise surface registration," *ACM TOG*, vol. 27, pp. 1–10, 2008.

[5] Q. Zheng *et al.*, "Consensus skeleton for non-rigid space-time registration," *CGF*, vol. 29, pp. 635–644, 2010.

[6] M. Wand *et al.*, "Efficient reconstruction of nonrigid shape and motion from real-time 3D scanner data," *ACM TOG*, vol. 28, pp. 1–15, 2009.

[7] T. Weise *et al.*, "Realtime performance-based facial animation," *ACM TOG*, vol. 30, pp. 77:1–77:10, 2011.

[8] V. Kim *et al.*, "Blended intrinsic maps," *ACM SIGGRAPH*, 2011.

[9] Y. Sahillioglu and Y. Yemez, "Minimum-distortion isometric shape correspondence using EM algorithm," *IEEE TPAMI*, vol. PP, p. 1, 2012.

[10] M.A. Audette *et al.*, "An algorithmic overview of surface registration techniques for medical imaging," *Med. Image Analy.*, vol. 4, pp. 201–217, 2000.

[11] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proc. 3DIM*, 2001, pp. 145–152.

[12] M. Rodrigues *et al.*, "Special issue on registration and fusion of range images," *CVIU*, vol. 87, pp. 1–7, 2002.

[13] Y. Liu, "A mean field annealing approach to accurate free form shape matching," *Pat. Recogn.*, vol. 40, pp. 2418–2436, 2007.

[14] H. Pottmann *et al.*, "Geometry and convergence analysis of algorithms for registration of 3D shapes," *IJCV*, vol. 67, pp. 277–296, 2006.

[15] J. Salvi *et al.*, "A review of recent range image registration methods with accuracy evaluation," *IVC*, vol. 25, pp. 578–596, 2007.

[16] O. van Kaick *et al.*, "A survey on shape correspondence," in *Proc. Euro STAR*, 2010.

[17] X. Li and I. Guskov, "Multi-scale features for approximate alignment of point-based surfaces," in *Proc. SGP*, 2005, p. 217.

[18] D.L.G. Hill, "Non-rigid image registration: theory and practice," *British J. Radiology*, vol. 77, pp. S140–S153(1), 2004.

[19] J.P.W. Pluim *et al.*, "Mutual-information-based registration of medical images: a survey," *IEEE T-MI*, vol. 22, pp. 986–1004, 2003.

[20] M. Holden, "A review of geometric transformations for nonrigid body registration," *IEEE T-MI*, vol. 27, pp. 111–128, 2008.

[21] N.J. Mitra *et al.*, "Dynamic geometry registration," in *Proc. SGP*, 2007, pp. 173–182.

[22] B.J. Brown and S. Rusinkiewicz, "Global non-rigid alignment of 3-D scans," *ACM TOG*, vol. 26, p. 21, 2007.

[23] B. Allen *et al.*, "Articulated body deformation from range scan data," *ACM TOG*, vol. 21, pp. 612–619, 2002.

[24] Y. Pekelny and C. Gotsman, "Articulated object reconstruction and markerless motion capture from depth video," *CGF*, vol. 27, pp. 243–253, 2008.

[25] W. Chang and M. Zwicker, "Automatic registration for articulated shapes," *Proc. SGP*, vol. 27, pp. 1459–1468, 2008.

[26] ——, "Global registration of dynamic range scans for articulated model reconstruction," *ACM TOG*, vol. 30, pp. 26:1–26:15, 2011.

[27] A. Myronenko and X. Song, "Point-set registration: Coherent point drift," *IEEE TPAMI*, vol. 32, pp. 2262 –2275, Dec. 2010.

[28] J. Süßmuth *et al.*, "Reconstructing animated meshes from time-varying point clouds," *CGF*, vol. 27, pp. 1469–1476, 2008.

[29] B. Allen *et al.*, "The space of human body shapes: reconstruction and parameterization from range scans," in *ACM SIGGRAPH*, 2003, pp. 587–594.

[30] R.W. Sumner and J. Popović, "Deformation transfer for triangle meshes," in *ACM SIGGRAPH*, 2004, pp. 399–405.

[31] H. Li *et al.*, "Global correspondence optimization for non-rigid registration of depth scans," *Proc. SGP*, vol. 27, pp. 1421–1430, 2008.

[32] ——, "Robust single-view geometry and motion reconstruction," *ACM TOG*, vol. 28, pp. 1–10, 2009.

[33] M. Liao *et al.*, "Modeling deformable objects from a single depth camera," in *Proc. ICCV*, 2009.

[34] B. Amberg *et al.*, "Optimal step nonrigid ICP algorithms for surface registration," in *Proc. CVPR*, 2007, pp. 1–8.

[35] H. Chui and A. Rangarajan, "A new point matching algorithm for non-rigid registration," *CVIU*, vol. 89, pp. 114–141, 2003.

[36] V. Jain and H. Zhang, "Robust 3D shape correspondence in the spectral domain," in *Proc. SMI*, 2006, p. 19.

[37] A.M. Bronstein *et al.*, "Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching," *PNAS*, pp. 1168–1172, 2006.

[38] ——, "A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching," *IJCV*, vol. 89, pp. 266–286, 2010.

[39] Y. Lipman and T. Funkhouser, "Möbius voting for surface correspondence," in *ACM SIGGRAPH*, 2009, pp. 1–12.
[40] M. Ovsjanikov *et al.*, "One point isometric matching with the heat kernel," *CGF*, vol. 29, pp. 1555–1564, 2010.
[41] Y. Zeng *et al.*, "Intrinsic dense 3d surface tracking," in *Proc. CVPR*, 2011.
[42] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Proc. SGP*, 2007, pp. 109–116.
[43] D. Anguelov *et al.*, "Scape: shape completion and animation of people," *ACM TOG*, vol. 24, pp. 408–416, 2005.
[44] P. Hammond *et al.*, "3D analysis of facial morphology," *Amer. J. Med. Genetics A*, pp. 339–348, 2004.
[45] R. Toldo *et al.*, "Global registration of multiple point clouds embedding the generalized procrustes analysis into an ICP framework," in *Proc. 3DPVT*, 2010.
[46] Q.X. Huang *et al.*, "Reassembling fractured objects by geometric matching," *ACM TOG*, vol. 25, pp. 569–578, 2006.
[47] C.S. Chen *et al.*, "RANSAC-based DARCES: a new approach to fast automatic registration of partially overlapping range images," *IEEE TPAMI*, vol. 21, pp. 1229–1234, nov 1999.
[48] N. Gelfand *et al.*, "Robust global registration," in *Proc. SGP*, 2005, pp. 197–206.
[49] Y.S. Liu and K. Ramani, "Robust principal axes determination for point-based shapes using least median of squares," *CAD*, vol. 41, pp. 293–305, 2009.
[50] P.J. Besl and N.D. McKay, "A method for registration of 3-D shapes," *IEEE TPAMI*, vol. 14, pp. 239–256, 1992.
[51] Y. Chen and G. Medioni, "Object modeling by registration of multiple range images," in *Proc. ICRA*, vol. 3, 1991, pp. 2724–2729.
[52] A.E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes," *IEEE TPAMI*, vol. 21, pp. 433–449, 1999.
[53] R. Gal and D. Cohen-Or, "Salient geometric features for partial shape matching and similarity," *ACM TOG*, vol. 25, pp. 130–150, 2006.
[54] G.C. Sharp *et al.*, "ICP registration using invariant features," *IEEE TPAMI*, vol. 24, pp. 90–102, 2002.
[55] H. Pottmann *et al.*, "Integral invariants for robust geometry processing," *CAGD*, vol. 26, pp. 37–60, 2009.
[56] P. Krsek *et al.*, "Range image registration driven by a hierarchy of surface differential features," in *Proc. OAGM*, 1998, pp. 175–183.
[57] M. Bokeloh *et al.*, "Slippage features," University of Tübingen, Tech. Rep., 2008.
[58] J. Digne *et al.*, "The level set tree on meshes," in *Proc. 3DPVT*, 2010.
[59] Y. Liu, "Automatic 3D free form shape matching using the graduated assignment algorithm," *Pat. Recogn.*, vol. 38, pp. 1615–1631, 2005.
[60] ——, "Automatic registration of overlapping 3D point clouds using closest points," *IVC*, vol. 24, pp. 762–781, 2006.
[61] ——, "Replicator dynamics in the iterative process for accurate range image matching," *IJCV*, vol. 83, pp. 30–56, 2009.
[62] T. Jost and H. Hügli, "A multi-resolution ICP with heuristic closest point search for fast and robust 3D registration of range images," in *Proc. 3DIM*, 2003, p. 427.
[63] Y. Liu, "Constraints for closest point finding," *Pat. Recogn. Letters*, vol. 29, pp. 841–851, 2008.
[64] J.M. Phillips *et al.*, "Outlier robust ICP for minimizing fractional RMSD," in *Proc. 3DIM*, 2007, pp. 427–434.
[65] S. Gold and A. Rangarajan, "A graduated assignment algorithm for graph matching," *IEEE TPAMI*, vol. 18, pp. 377–388, 1996.
[66] H. Pottmann *et al.*, "Registration without ICP," *CVIU*, vol. 95, pp. 54–71, 2004.
[67] A. Gruen and D. Akca, "Least squares 3D surface and curve matching," *J. of Photogrammetry and Remote Sensing*, vol. 59, pp. 151–174, 2005.
[68] F. Tombari *et al.*, "Unique signatures of histograms for local surface description," in *Proc. ECCV*, 2010, vol. 6313, pp. 356–369.
[69] T. Várady *et al.*, "Reverse engineering of geometric models—an introduction," *CAD*, vol. 29, pp. 255–268, 1997.
[70] H. Zhang *et al.*, "Deformation-driven shape correspondence," *CGF*, vol. 27, pp. 1431–1439, 2008.
[71] A. Tevs *et al.*, "Isometric registration of ambiguous and partial data," in *Proc. CVPR*, 2009.
[72] D. Anguelov *et al.*, "The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces," *Proc. NIPS*, vol. 17, pp. 33–40, 2004.
[73] J. Sun *et al.*, "A concise and provably informative multi-scale signature based on heat diffusion," in *Proc. SGP*, 2009, pp. 1383–1392.
[74] M.M. Bronstein and I. Kokkinos, "Scale-invariant heat kernel signatures for non-rigid shape recognition," in *Proc. CVPR*, 2010, pp. 1704–1711.
[75] D. Raviv *et al.*, "Volumetric heat kernel signatures," in *Proc. 3DOR*, 2010, pp. 39–44.
[76] ——, "Affine-invariant diffusion geometry for the analysis of deformable 3d shapes," in *Proc. CVPR*, 2011, pp. 2361–2367.
[77] M. Aubry *et al.*, "The wave kernel signature: A quantum mechanical approach to shape analysis," in *Proc. 4DMOD*, 2011.
[78] A.M. Bronstein, "Spectral descriptors for deformable shapes," Tel Aviv University, Tech. Rep., 2011.
[79] M. Wand *et al.*, "Reconstruction of deforming geometry from time-varying point clouds," in *Proc. SGP*, 2007, pp. 49–58.
[80] L. Zhang *et al.*, "Spacetime faces: high resolution capture for modeling and animation," in *ACM SIGGRAPH*, 2004, pp. 548–558.
[81] F. Fouss *et al.*, "Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation," *IEEE TKDE*, vol. 19, pp. 355–369, 2007.
[82] Y. Lipman *et al.*, "Biharmonic distance," *ACM TOG*, vol. 29, pp. 1–11, 2010.
[83] J. Sun *et al.*, "Fuzzy geodesics and consistent sparse correspondences for deformable shapes," *Proc. SGP*, vol. 29, pp. 1535–1544, 2010.
[84] D.T.P. Quynh *et al.*, "An intrinsic algorithm for computing geodesic distance fields on triangle meshes with holes," *Graph. Models*, vol. 74, pp. 209–220, 2012.
[85] V. Surazhsky *et al.*, "Fast exact and approximate geodesics on meshes," *ACM TOG*, vol. 24, pp. 553–560, 2005.
[86] K. Crane *et al.*, "Geodesics in heat," *CORR*, vol. abs/1204.6216, 2012.
[87] F. Mémoli, "A spectral notion of Gromov-Wasserstein distance and related methods," *Appl. Comp. Harm. Anal.*, vol. 30, pp. 363–401, 2011.
[88] T.K. Dey *et al.*, "Persistent heat signature for pose-oblivious matching of incomplete models," *CGF*, vol. 29, pp. 1545–1554, 2010.
[89] R. Litman *et al.*, "Diffusion-geometric maximally stable component detection in deformable shapes," *Comp. & Graph.*, vol. 35, pp. 549–560, 2011.
[90] ——, "Stable volumetric features in deformable shapes," *Comp. & Graph.*, vol. 36, pp. 569–576, 2012.
[91] E. Boyer *et al.*, "SHREC 2011: robust feature detection and description benchmark," in *Proc. 3DOR*, 2011.
[92] P. Huang *et al.*, "Global temporal registration of multiple non-rigid surface sequences," in *Proc. CVPR*, 2011, pp. 3473–3480.
[93] N. Gelfand *et al.*, "Geometrically stable sampling for the ICP algorithm," in *Proc. 3DIM*, 2003, pp. 260–267.
[94] N.J. Mitra *et al.*, "Registration of point cloud data from a geometric optimization perspective," in *Proc. SGP*, 2004, pp. 22–31.
[95] F. Boughorbel *et al.*, "A new method for the registration of three-dimensional point-sets: The Gaussian fields framework," *IVC*, vol. 28, pp. 124–137, 2010.
[96] A. Albarelli *et al.*, "A non-cooperative game for 3D object recognition in cluttered scenes," in *Intl. Conf. on 3DIMPVT*, 2011, pp. 252–259.
[97] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints," in *Proc. ICCV*, vol. 2, 2005, pp. 1482–1489.
[98] A. Tevs *et al.*, "Intrinsic shape matching by planned landmark sampling," *CGF*, vol. 30, pp. 543–552, 2011.
[99] D. Mateus *et al.*, "Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration," in *Proc. CVPR*, 2008, pp. 1–8.
[100] A. Lorusso *et al.*, "A comparison of four algorithms for estimating 3-D rigid transformations," in *Proc. BMVC*, 1995, pp. 237–246.

[101] A. Torsello *et al.*, "Grouping with asymmetric affinities: A game-theoretic perspective," *Proc. CVPR*, vol. 1, pp. 292–299, 2006.
[102] L. Silva *et al.*, "Precision range image registration using a robust surface interpenetration measure and enhanced genetic algorithms," *IEEE TPAMI*, vol. 27, pp. 762–776, 2005.
[103] A.M. Bronstein *et al.*, *Numerical geometry of non-rigid shapes*. Springer, 2008.
[104] A. Elad and R. Kimmel, "On bending invariant signatures for surfaces," *IEEE TPAMI*, vol. 25, pp. 1285–1295, 2003.
[105] O. Duchenne *et al.*, "A tensor-based algorithm for high-order graph matching," in *Proc. CVPR*, june 2009, pp. 1980 –1987.
[106] Y. Zeng *et al.*, "Dense non-rigid surface registration using high-order graph matching," in *Proc. CVPR*, 2010.
[107] "The Stanford 3D Scanning Repository." [Online] : http://graphics.stanford.edu/data/3Dscanrep/
[108] "Ohio State University/SAMPL Database." [Online] : http://sampl.eng.ohio-state.edu/sampl/data/
[109] "The Digital Michelangelo Project." [Online] : http://graphics.stanford.edu/data/mich/
[110] X. Chen *et al.*, "A benchmark for 3D mesh segmentation," *ACM SIGGRAPH*, vol. 28, 2009.
[111] A.M. Bronstein *et al.*, "Efficient computation of isometry-invariant distances between surfaces," *SIAM J. Sci. Comput.*, vol. 28, pp. 1812–1836, 2006.
[112] P. Sand *et al.*, "Continuous capture of skin deformation," in *ACM SIGGRAPH*, 2003, pp. 578–586.
[113] D. Vlasic *et al.*, "Multilinear models for face synthesis," in *SIGGRAPH Sketches*, 2004, p. 56.
[114] D. Giorgi *et al.*, "SHREC: shape retrieval contest: Watertight models track." [Online] : http://watertight.ge.imati.cnr.it
[115] R.S. Schestowitz *et al.*, "Non-rigid registration assessment without ground truth," in *Med. Image Under. and Analy.*, vol. 2, 2006, pp. 151–155.
[116] D. Bradley *et al.*, "High resolution passive facial performance capture," *ACM TOG*, vol. 29, pp. 1–10, 2010.
[117] H. Li *et al.*, "Example-based facial rigging," *ACM SIGGRAPH*, vol. 29, pp. 32:1–32:6, 2010.
[118] P. Buddharaju *et al.*, "Physiology-based face recognition in the thermal infrared spectrum," *IEEE TPAMI*, vol. 29, pp. 613–626, 2007.
[119] T.F. Cootes *et al.*, "Computing accurate correspondences across groups of images," *IEEE TPAMI*, vol. 32, pp. 1994–2005, 2010.
[120] K.A. Sidorov *et al.*, "Efficient groupwise non-rigid registration of textured surfaces," in *Proc. CVPR*, 2011, pp. 2401–2408.
[121] J.h. Chen *et al.*, "3D point correspondence by minimum description length in feature space," in *Proc. ECCV*, 2010.
[122] T.M. Mitchell, *Machine Learning*. McGraw-Hill, 1997.
[123] F. Wang *et al.*, "Groupwise point pattern registration using a novel CDF-based Jensen-Shannon Divergence," in *Proc. CVPR*, vol. 1, 2006, pp. 1283–1288.
[124] S. Izadi *et al.*, "Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera," in *Proc. UIST*, 2011, pp. 559–568.

Gary K.L. Tam received his BEng (first-class honors), MPhil and PhD degrees from Hong Kong University of Science and Technology, City University of Hong Kong, and Durham University, respectively. He is currently a research associate at Cardiff University, United Kingdom. His research interests include multimedia retrieval & visualization, digital geometry processing and interactive techniques.

Zhi-Quan Cheng received a BSc, MSc, and PhD degree from Computer School at National University of Defense Technology in 2000, 2002 and 2008, respectively. He is a lecturer at the PDL Laboratory, Computer School, National University of Defense Technology (NUDT), and the leader of visual computing team at NUDT. His research interests include computer graphics, and digital geometry processing.

Yu-Kun Lai received his bachelor's degree and PhD degree in computer science from Tsinghua University in 2003 and 2008, respectively. He is currently a lecturer of visual computing in the School of Computer Science, Cardiff University, Wales, UK. His research interests include computer graphics, geometry processing, computer-aided geometric design and computer vision.

Frank C. Langbein received a PhD in 2003 from Cardiff University on "Beautification of Reverse Engineered Geometric Models" and a Diploma in Mathematics from Stuttgart University in 1998. He is currently a lecturer in computer science at Cardiff University, working on modelling, control and machine learning applied to geometry, quantum technologies, chemical synthesis and perception. He is a member of the American Mathematical Society and the IEEE.

Yonghuai Liu is a senior lecturer at Aberystwyth University since 2011. He is currently an associate editor and an editorial board member for a number of international journals, including Pattern Recognition Letters. He has published three books and more than 140 papers in international conference proceedings and journals. His primary research interests lie in 3D computer vision.

Dave Marshall is a Professor in computer science at Cardiff University. He received his BSc and PhD "Three Dimensional Inspection of Manufactured Objects" from University College, Cardiff in 1986 and 1990. He researches in human facial analysis, high dimensional subspace analysis, audio/video image processing, and data/sensor fusion. He has published over 130 papers and one book in these research areas. He is a member of the IEEE and BMVA.

Ralph R. Martin obtained his PhD from Cambridge University. He now leads the Visual Computing Research Group at Cardiff University, and is Director of Scientific Programmes of the One Wales Research Institute of Visual Computing. His publications include over 200 papers and 12 books. He is a Fellow of the Learned Society of Wales, and is on the editorial boards of several leading journals.

Xianfang Sun received his PhD degree in control theory and its applications from the Institute of Automation, Chinese Academy of Sciences. He is a lecturer at Cardiff University. His research interests include computer vision and graphics, pattern recognition and artificial intelligence, and system identification and control. He is on the Editorial Board of Acta Aeronautica et Astronautica Sinica.

Paul L. Rosin is a Professor at the School of Computer Science & Informatics, Cardiff University. Previous posts include Brunel University, Joint Research Centre, Italy and Curtin University of Technology, Australia. His research interests include computer vision, remote sensing, mesh processing, and the analysis of shape in art and architecture.