

Mapping the Problem Space of Image Registration

Steve Oldridge, Gregor Miller, and Sidney Fels
Human Communication Technologies Laboratory
ECE, University of British Columbia
Vancouver, Canada
{steveo, gregor, ssfels}@ece.ubc.ca

Abstract—In this paper we explore a conceptual mapping of the image registration problem into an N-Dimensional problem space based on the properties of the images being registered, in contrast to traditional surveys of image registration which divide the field algorithmically. The five main dimensions of our proposed mapping are variations in: spatial alignment, intensity, focus, sensor type, and structure. Individual algorithms can be thought of as supporting a volume of solutions within the problem domain map, although they are typically designed to solve problems along a single dimension. Existing image registration papers and techniques are taxonomized within this mapping according to these major dimensions. The focus of this paper is threefold. First, an up-to-date survey of image registration techniques is provided, building from previous seminal surveys. Second, a novel taxonomy is presented that organizes the registration problem space based on the variation between the images being registered. Finally, a number of new research areas made possible under this novel taxonomy are examined, and a path is laid out for future research in the field.

Keywords-Image Registration; Literature Survey; Taxonomy; Problem-Centric Computer Vision;

I. INTRODUCTION

Image registration is the process of calculating spatial transforms which align a set of images to a common observational frame of reference, often one of the images in the set. Registration is a key step in any image analysis or understanding task where different sources of data must be combined. It is a critical component of computational photography [1], remote sensing [2], [3], automated manufacturing processes, and medical image processing [4], [5]. More recently it has been used to create navigable models of a scene from a database of photographs [6], to remove unwanted objects from overlapping images, etc.

Previous surveys of image registration [7], [8] divide the field algorithmically, focussing on *how* registration is accomplished. This approach provides a good basis for classifying and comparing algorithmic similarity, however it does little to illuminate the conditions important to the problem of registration. Understanding the image registration problem space and the conditions under which specific algorithms perform best allows for further automation of image registration, and makes the field more accessible to non-experts.

The new taxonomy proposed within this paper maps the registration problem space based on the type of registration problem being solved, focussing on properties of the images or data being registered. The most appropriate method of registration of a set of images is most often determined by how these properties vary in comparison to one another. We provide a mapping of existing papers and techniques into our N-dimensional space of problem variances. Within our mapping individual algorithms can be thought of as supporting a volume of solutions within the problem domain, although they are typically designed to solve problems along a single dimension. By rethinking the problem in this way we change the abstraction of registration from one requiring knowledge and expertise about particular algorithms, both in how they work and when to use them, to one requiring expertise about the registration problem itself, which is much simpler for programmers who are not vision experts to understand and implement. In Section V we further discuss the benefits of this problem centric taxonomy.

Several important image registration techniques and strategies have been developed since Zitová and Flusser first published their survey. The increased computational power of the past decade has made automatic methods common, and under optimal conditions modern algorithms are able to align image pairs more accurately than can be detected by the human eye [9]. Our intention with this paper is to organize these new methods by the conditions under which they perform optimally, and to reorganize the classic or key ideas introduced prior to 2003 that are still in use today. As with Zitová and Flusser we do not examine details of particular algorithms or perform comparative experiments, but instead attempt to map and summarize the main approaches used in registration today. It would be difficult to provide an evaluation of every image registration algorithm, or even every type of registration problem.

The contributions of this paper are threefold. First, an up to date survey of image registration techniques is provided, building off of previous seminal surveys [8], [7], [5]. Second, a novel taxonomy is presented that organizes the registration problem space based on the variation between the images being registered. Finally, a number of new research areas are made possible under this novel taxonomy, and a path is laid out for future research in the field.

In Section II we present an overview of registration methodologies, reexamining the frameworks present by Zitová and Flusser [7], by Pluim, Maintz and Viergever [5], and by Brown [8] in their comprehensive surveys. Section III explores in detail main dimensions of the registration problem space, and outlines the dominant and alternative methods in each category. Methods for evaluating the accuracy of a registration are discussed in Section IV. Finally, Section V summarizes and explores potential future trends and applications.

II. IMAGE REGISTRATION METHODOLOGIES

Past image registration surveys have provided a methodological taxonomy for understanding the different algorithms used to solve the registration problem. Brown [8] divides registration into four components: feature space, search space, search strategy, and similarity metric. The later work of Zitová and Flusser [7] divides the field into area and feature based methods, and their model reflects the shift towards feature based methods that occurred between the two papers. The four basic steps of image registration under their model are: feature detection, feature matching, mapping function design, and image transformation and resampling. Like Brown we have chosen to leave image transformation and resampling out of our taxonomy; these steps, though important for most computational photography applications, can be considered rendering problems, and are independent of spatial alignment. Maintz [4] provides insight into the use of registration in medical imaging, providing important methods and variations relevant to that field. The taxonomy divides both algorithmically and based on the modality of the data, again providing a similar mapping. Pluim et al. [5] also survey medical imaging, focusing on Mutual-Information-Based registration techniques. Their taxonomy classifies algorithms into two main categories: methodological aspects and aspects of application.

Within her framework Brown talks about how knowledge of the types of variation that occur in image sets can be used to guide selection of the most suitable components for a specific problem. Variations are divided into three classes: variations due to differences in acquisition that cause the images to be misaligned, variations due to differences in acquisition that cannot be easily modeled (such as lighting or camera extrinsics), and finally variations due to movement of objects within the scene. These are labeled by Brown “corrected distortions”, “uncorrected distortions”, and “variations of interest” respectively. Zitová and Flusser provide a model of variation according to the manner in which the images were acquired: different viewpoints, different times / conditions, different sensors, and finally scene to model registration. Within their survey they do not use this mapping directly, however in many cases they discuss the type of problem each method has been designed to solve, allowing a similar mapping of methodology from situation. Similarly

Pluim et al.’s “aspects of the application” entail image modalities, subject of the registration, and the object of registration. This delineation provides an excellent starting point for variations that are important within the medical imaging community.

It is this concept of variations that we have chosen to base our taxonomy on, extending these initial ideas into specific variations common in computational photography and remote sensing and exploring the successful algorithms in each categorization. Our five main categories are: Purely spatial variation, variations in intensity, variations in focus, variations in sensor type, and finally variations in structure. These categories could be differentiated further, and others certainly exist, however the majority of registration methods encountered target one of these types of image variation.

III. CLASSIFICATION OF REGISTRATION PROBLEMS

Although the field is rapidly moving towards automatic image registration, algorithms and systems are most often limited to a single application area such as stitching panoramas, super-resolution, high-dynamic-range (HDR) imaging, focal stacking, multimodal imaging, etc. In most cases these techniques will perform reasonably on a limited subset of problems from other domains, however no single algorithm exists that will accurately solve all types of registration problems.

Examining the different applications that rely on image registration from a data centric perspective provides insight into the relative success of algorithms on particular types of data. Each of our main categories: purely spatial variation, significant variations in intensity, significant variations in focus, variations in sensor type, and finally variations in structure are examined below in detail.

A. Purely Spatial Variations

Image pairs that differ *purely* spatially are the most common type of image registration problem. Applications that require registration of images that vary spatially include panorama stitching, super resolution, and remote sensing. Although area-based methods derivative of Lucas and Kanade [10], [11] are capable of solving these types of registration problems, feature based methods are the most common technique applied and are generally considered much more accurate unless the image pairs contain little high-frequency information from which to find and match features. Feature based methods work by matching features detected in each image and finding a transform which maximizes the number of features which correctly correspond once the images are aligned. Methods vary in their selection of interest points from which features are calculated, in their feature descriptors, in their method of matching, and finally in the method and error function used to solve for the aligning transform.

The choice of feature descriptor used for image registration is a subject of debate within the community. The descriptors must be distinctive and at the same time robust to changes in viewing conditions as well as to errors of the detector. The comparison of different feature descriptors has been recently made by Mikolajczyk and Schmid [12]. In their survey they compare shape context [13], steerable filters [14], PCA-SIFT [15], differential invariants [16], spin images [17], SIFT [18], complex filters [19], moment invariants [20], and cross-correlation for different types of interest regions, concluding that the ranking of the descriptors is mostly independent of the interest region detector and that the SIFT-based descriptors perform best.

Once features have been detected they must be matched across images. Approximate nearest neighbor matching using k-d trees is the dominant method of matching and is widely used by a number of researchers [21]. Exact nearest neighbor approaches are computationally intractable for large numbers of features and dimensions, and provide a limited advantage over approximate methods [22].

Images are aligned by solving the overconstrained system created by the feature pairs. Random Sample and Consensus (RANSAC) is commonly used [21] because of its robustness and efficiency, however other solvers such as nonlinear least squares fitting have also been applied with success. Yang et al. propose using an Iterative Closest Point Dual Bootstrap method [23] which was reported to perform favorably to RANSAC for difficult image pairs.

B. Intensity Variations

Image pairs that contain significant intensity variations include those used for high dynamic range (HDR) imaging, pairs with varying illumination, and even panorama sequences taken in direct sunlight.

Feature based methods listed in Section III-A work for small variations in intensity, particularly if their feature descriptors are gradient based. However for image pairs with significant intensity variation, the interest regions of feature based methods do not occur in the same location. Sand and Teller [24] attempt to handle intensity varying pairs by only selecting features from parts of the image that can be more easily matched while avoiding parts that are difficult. Their technique was designed for matching two video sequences, achieving good results with the limited spatial variation that entails, however it was not tested on still photographs. More recently Tomaszewska and Mantiuk [25] presented a similar idea, reporting a high quality alignment such that the "photographs were aligned with sufficient accuracy so that there are no visible artifacts in the final HDR image" by using only features that occur across all images in the set.

Schechner and Nayer [26] presented an alignment method based on pyramids of maximum likelihood as a part of their approach to generalize panorama images to incorporate

HDR. Kang et al. [27] also described a technique for creating high dynamic range video from a sequence of alternating high variable intensity exposures. Their sophisticated HDR stitching process uses local alignment and motion estimation to compensate for camera movement and object motion within the scene, a technique tailored to their input data.

Finally, Ward [28], [29] introduced a method specifically designed to align images with significant variations in intensity. The technique thresholds image pairs into pyramidal bitmaps. Bitmaps are analyzed and aligned for translation errors using shift and difference operations over each image. With this method 3 megapixel image sets are aligned in a fraction of a second. Their method deals solely with translation errors, although they discuss the possibility of rotation errors, suggesting that 10% of their data set failed as a result.

C. Focus Variations

Focus variations can occur in image pairs deliberately as is the case with focus stacking, or through motion blur due to movement of the camera or objects in the scene as can sometimes be the case in super resolution imaging. Here again feature based methods surveyed in Section III-A have utility to a point, however image pairs with no overlapping regions of focus do not work; the same interest points are not detected across images with different focal planes. Flusser, Zitová and Suk attempt to overcome this by finding blur invariant interest regions [30], [31]. Their evaluation of a single 128x128 pixel image using control points makes it difficult to provide significant conclusion about the quality of registration this provides.

In focal stacking applications, classic intensity based area methods are typically used to find the correct alignment. These methods work by attempting minimizing the difference in actual pixel values between the two images, most often through some form of gradient descent. Lucas and Kanade [10], [11] based sum of square difference area methods are reasonably successful assuming the spatial overlap between images is significant. This is often the case for focus stacking problems, particularly those composed of microscope data where sensor movement is minimal between images. Bradley et al. [32] make use of normalized cross correlation in their virtual microscopy system requiring an overlap of at least 45% between image pairs and ignoring results that fall outside their expected solution area.

D. Sensor Variations

Image pairs taken using different sensors are commonly referred to as multimodal image pairs. They are common to both medical imaging and to remote sensing applications, where proper alignment of two or more modalities provides significant additional information. When sensors differ there is no guarantee that intensities, gradients, or edges will be

similar, and both feature based and intensity based methods fail to find alignments.

Sharma and Paval [33] propose to overcome this by making the images as similar as possible, transforming images into representations invariant to polarity reversals before applying traditional area based techniques. Irani and Anandan [34] similarly transform images into high-pass energy images which are significantly less sensitive to sensor variations. These methods have been further developed by Liu et al. [35], [36] who use Gabor filtering as their local frequency representation. More recently, Henn and Witsch [37] define two nonlinear distance functions and minimize these to find the optimal alignment.

In the field of medical imaging, maximization of mutual information, developed simultaneously by Viola and Wells [38] and by Collignon et al. [39], is the most common method used and was comprehensively surveyed by Pluim, Maintz and Viergever [5]. Bardera, Feixas and Boada proposed two new similarity measures based off of Jensen's difference applied to Rényi and Tsallis-Havrda-Charvát entropies claiming that their proposed measures are more robust than the normalized mutual information for some modalities and a determined range of the entropy parameter. Gan et al. [40] suggest using Kullback-Leibler distance if a priori knowledge of the joint intensity distribution is available. Makela et al. [41] provide an overview of further methods focusing specifically on cardiac images. While features based on mutual information are being used in pattern recognition [42], [43] they have not to our knowledge been used successfully for registration of multimodal images.

Automatic registration of 3D shapes and volumetric slices of sensor data are common within the medical imaging community, however they fall outside the scope of this survey.

E. Variations in Structure

Image pairs can vary significantly in the structure of the scene they depict, either because objects within the scene have moved, or more commonly in medical imaging because objects have changed over time. Aligning these images in spite of these changes often requires non-rigid transforms that solve for the alignment of regions or at the extreme of individual pixels. Methods discussed here can include a variety of other variations, although sensor variations are most common. This concept of multiple dimensions of variation is discussed further in Section V.

Bookstein proposed one of the seminal concepts of non-rigid models for image registration: using thin plate splines to interpolate affine transforms [44]. Moshfeghi proposed an alternative model based on elasticity [45]. Christensen et al. introduce the idea of viscous fluid representations of deformable registration [46], while Bro-Nielson and Gramkow [47] significantly accelerate this concept using a fast fluid

model. McInerney and Terzopoulos [48] provide a good survey of early nonrigid techniques and their application within the medical imaging community.

Rueckert et al. [49] present a nonrigid method that uses a global affine transform, followed by a local B-spline matching of normalized mutual information voxels, applying their technique to the registration of breast MR images. Rohde et al. [50] make several contributions to the field: use of radially symmetric basis functions rather than B-splines to model the deformation field; a metric to identify regions that are poorly registered and over which the transformation needs to be improved; partitioning of the global registration problem into several smaller ones; and a new constraint scheme that allows us to produce transformations that are topologically correct. They compare the proposed approach to more traditional ones listed above and show that their new algorithm compares favorably to those in current use. More recently D'Agostino et al. [51] propose modeling the registration as a viscous fluid that deforms under the influence of forces derived from the gradient of the mutual information registration criterion, validating their method by matching simulated T1-T1, T1-T2 and T1-PD images.

Crum, Hartkens and Hill present a more recent survey of nonrigid image registration [52] providing a more in depth analysis of this subject. The differentiation of which method of deformation is most applicable to different modalities or problems is a complex question and is ultimately beyond the scope of this survey but as will be discussed in Section V is an important step for the future of the field.

IV. EVALUATION OF REGISTRATION ACCURACY

Zitová and Flusser [7] identify three measures of registration accuracy in their survey: localization error, matching error, and alignment error. Localization error represents mistakes in the location of feature based methods' interest regions. Matching error is measured as the number of false matches between features. Finally, alignment error measures the difference between the proposed alignment and the actual between-image geometric distortion. Localization and matching error are specific to feature based methods and can be measures of problems within the steps of those methods, however alignment error is ultimately what researchers must be concerned with when comparing across non-feature based methods.

The unbiased evaluation of registration techniques is a significant problem within the field. Techniques are most often compared using researchers own image pairs, which can vary significantly in resolution, quality, difficulty, and number, making it extremely difficult to understand how well their algorithms perform. Some researchers provide a point of comparison by testing their images with several other methods, however this selection of image pairs may be unknowingly biased to work well with their methods.

Azzari et al. recently propose the use of a set of synthetic data with a known ground truth which they have made available online [9]. This is an important starting point, however their image sets are low resolution (320x240), limited to translation and rotation, and contain no variation in intensity, focus, sensor, or structure. A much more robust and high resolution set of images is necessary for the evaluation of existing and new techniques, and would be a critical contribution to the field of image registration.

The creation of such a test set would allow for a much more detailed mapping of the image registration problem space. By creating image pairs with both ground truth transforms and known variations the performance of algorithms in different parts of the problem space could be calculated and the n-dimensional volume that each supports could be calculated. This would allow developers with knowledge about the forms of variation that their system encounters to select the most appropriate algorithm(s), rather than relying on rough conventions that exist within the field.

V. CONCLUSIONS AND FUTURE WORK

Image registration remains one of the most important tasks in computer vision when combining information from various sources. This paper gives an up-to-date survey of image registration techniques, building from previous surveys [8], [7], [5]. We have presented a mapping of the image registration problem domain that focuses on the types of variation that occur between images to be registered. Conceptualizing these forms of variation as dimensions of the image registration problem, we have an abstraction which allows us to think of algorithms as volumes within the n-dimensional problem space. Existing algorithms have been introduced into this mapping according to the main form of variation that they have been designed to support, and instances where these algorithms support other forms of variation have been outlined, providing an initial mapping of the space that is summarized below in Table I. The reorganization of image registration into our variation-centric taxonomy provides a basis for several opportunities that advance the field of image registration:

First, using variations as a starting point it is possible to create a language model with which to describe registration problems. By describing problems in terms of the variations that exist in their images instead of their algorithmic details it would be possible to create a system that automatically selects the most appropriate algorithm from those available. Problem centric methodologies are preferential to purely algorithmic descriptions in a number of ways. First, they give those programmers who are not vision specialists access to advanced image processing techniques without requiring specific knowledge of the underlying algorithms that implement them. Second, they allow improved algorithms to seamlessly replace older implementations providing programmers using a problem centric software library with an

instantaneous upgrade path without reprogramming or integrating a new implementation. Finally, if as often happens the conditions around the problem change, the programmer can automatically select a more appropriate algorithm simply by changing their description. We liken this to the field of computer graphics where a model describing the world is provided by the programmer, and is interpreted using the available resources into the scene.

This concept can be extended further into the fully automatic classification of registration problems[53], [54]. Many of the variations explored in this paper can be automatically detected from the image pairs, allowing for an estimation of where the pairs exist in problem space. This allows for the selection of an algorithm appropriate to the given conditions of the problem. Such an automated system would come close to the 'ultimate registration method' described by Zitová and Flusser at the end of their survey; a system able to recognize the type of task and to decide by itself about the most appropriate solution.

In addition, most image registration methods are designed to work along a single dimension, however the combination of these is becoming more common, particularly in sensor / structure combinations for multimodal non-rigid medical imaging. Another notable multidimensional example is Schechner and Nayer's HDR panorama stitching method [26]. The examination of other combinations of variation such as focus and structure could prove interesting, although many of these higher dimensional pairings are likely best solved by examining the groups of images as a whole and choosing which pairs to match, a topic not covered within this survey.

Finally, a more detailed mapping is desirable, but as mentioned in Section IV, in order to create such a mapping a well thought out, significant, robust, high resolution set of image pairs agreed upon by the community is necessary for the evaluation of image registration algorithms.

REFERENCES

- [1] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen, "Interactive digital photomontage," in *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, New York, NY, USA, 2004, pp. 294–302.
- [2] T. M. Lillesand and R. W. Kiefer, *Remote Sensing and Image Interpretation*, 6th ed. Wiley, 2007.
- [3] J. B. Campbell, *Introduction to remote sensing*, 4th ed. Guildford Press, 2008.
- [4] J. Maintz and M. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, no. 1, pp. 1–36, 1998.
- [5] J. Pluim, J. Maintz, and M. Viergever, "Mutual-information-based registration of medical images: a survey," *Medical Imaging, IEEE Transactions on*, vol. 22, no. 8, pp. 986–1004, Aug. 2003.

Variation	Invariant	Capable	Barely Capable
Purely Spatial	[13], [21], [14], [15], [16], [17], [18], [22], [19], [20], [23]	[11], [40], [37], [10], [39], [38]	
Intensity	[27], [29], [24], [26], [25], [28]	[21], [40], [37], [18], [39], [38]	[11], [34], [35], [36], [10], [33]
Focus	[11], [32], [30], [10], [31]	[37], [39], [38], [28]	
Sensor	[37], [34], [35], [36], [39], [33], [38]	[43], [42]	
Structure	[44], [47], [46], [52], [51], [48], [45], [50], [49]		[21], [40], [37], [18], [39], [38]

Table I

A SUMMARY OF THE MAPPING OF IMAGE REGISTRATION ALGORITHMS (BY REFERENCE) ACCORDING TO THE FORMS OF VARIATION THAT THEY SUPPORT. ALGORITHMS ARE PLACED WITHIN THE FIVE MAJOR DIMENSIONS OF OUR TAXONOMY INTO THREE BROAD CATEGORIES: ALGORITHMS INVARIANT TO THAT FORM OF VARIATION, ALGORITHMS THAT ARE CAPABLE AT MOST PROBLEMS WITH THAT FORM OF VARIATION, AND FINALLY ALGORITHMS THAT ARE BARELY CAPABLE OF PROBLEMS WITH THAT FORM OF VARIATION. CAPABLE AND BARELY CAPABLE MAPPINGS ARE MADE ON THE BASIS OF RESULTS MENTIONED EITHER WITHIN THE CORRESPONDING PAPER OR AN EQUIVALENT PAPER WITH SIMILAR PROPERTIES. THIS INITIAL MAPPING, ALTHOUGH COARSE, PROVIDES INSIGHT INTO THE CAPABILITY OF EACH IMAGE REGISTRATION ALGORITHMS WITHIN THE REGISTRATION PROBLEM SPACE.

- [6] N. Snavely, R. Garg, S. M. Seitz, and R. Szeliski, "Finding paths through the world's photos," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2008)*, vol. 27, no. 3, pp. 11–21, 2008.
- [7] B. Zitová and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, pp. 977–1000, 2003.
- [8] L. G. Brown, "A survey of image registration techniques," *ACM Computing Surveys*, vol. 24, pp. 325–376, 1992.
- [9] P. Azzari, L. Stefano, and S. Mattoccia, "An evaluation methodology for image mosaicing algorithms," in *ACIVS '08: Proceedings of the 10th International Conference on Advanced Concepts for Intelligent Vision Systems*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 89–100.
- [10] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision (darpa)," in *Proceedings of the 1981 DARPA Image Understanding Workshop*, April 1981, pp. 121–130.
- [11] S. Baker and I. Matthews, *Lucas-Kanade 20 Years On: A Unifying Framework*, vol. 56, no. 1, pp. 221 – 255, March 2004.
- [12] K. Mikolajczyk and C. Schmid, "A performance evaluation of local descriptors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 10, pp. 1615–1630, October 2005.
- [13] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 4, pp. 509–522, Apr 2002.
- [14] W. Freeman and E. Adelson, "The design and use of steerable filters," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 13, no. 9, pp. 891–906, Sep 1991.
- [15] Y. Ke and R. Sukthankar, "Pca-sift: a more distinctive representation for local image descriptors," vol. 2, June-2 July 2004, pp. II–506–II–513 Vol.2.
- [16] J. J. Koenderink and A. J. van Doorn, "Representation of local geometry in the visual system," *Biol. Cybern.*, vol. 55, no. 6, pp. 367–375, 1987.
- [17] S. Lazebnik, C. Schmid, and J. Ponce, "A sparse texture representation using affine-invariant regions," vol. 2, June 2003, pp. II–319–II–324 vol.2.
- [18] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, pp. 91–110, 2004.
- [19] F. Schaffalitzky and A. Zisserman, "Multi-view matching for unordered image sets, or "how do i organize my holiday snaps?," in *ECCV '02: Proceedings of the 7th European Conference on Computer Vision*, 2002, pp. 414–431.
- [20] L. Van Gool, T. Moons, and M. Proesmans, "Mirror and point symmetry under perspective skewing," Jun 1996, pp. 285–292.
- [21] M. Brown and D. Lowe, "Recognising panoramas," *Computer Vision, 2003.Proceedings. Ninth IEEE International Conference on*, vol. 2, pp. 1218–1225, 16-16 Oct. 2003.
- [22] S. Maneewongvatana and D. M. Mount, "The analysis of a probabilistic approach to," in *In Proceedings of the 2001 Workshop on Algorithms and Data Structures (WADS01)*, 2001, pp. 276–286.
- [23] G. Yang, C. Stewart, M. Sofka, and C.-L. Tsai, "Registration of challenging image pairs: Initialization, estimation, and decision," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 29, no. 11, pp. 1973–1989, Nov. 2007.
- [24] P. Sand and S. Teller, "Video matching," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 592–599, 2004.
- [25] A. Tomaszewska and R. Mantiuk, "Image registration for multi-exposure high dynamic range image acquisition," in *Proc. Int'l Conf. Central Europe on Computer Graphics, Visualization, and Computer Vision (WSCG)*, 2007.
- [26] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: High dynamic range in a wide field of view," *Int. J. Comput. Vision*, vol. 53, no. 3, pp. 245–267, 2003.
- [27] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski, "High dynamic range video," in *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*. New York, NY, USA: ACM, 2003, pp. 319–325.

- [28] G. Ward, "Robust image registration for compositing high dynamic range photographs from handheld exposures," *Journal of Graphics Tools*, vol. 8, pp. 17–30, 2003.
- [29] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging. Data Acquisition, Manipulation, and Display*. Morgan Kaufmann, 2005.
- [30] J. Flusser, B. Zitová, and T. Suk, "Invariant-based registration of rotated and blurred images," in *IEEE 1999 International Geoscience and Remote Sensing Symposium. Proceedings*. IEEE Computer Society, 1999, pp. 1262–1264.
- [31] B. Zitová, J. Kautsky, G. Peters, and J. Flusser, "Robust detection of significant points in multiframe images," *Pattern Recogn. Lett.*, vol. 20, no. 2, pp. 199–206, 1999.
- [32] A. P. Bradley, M. Wildermoth, and P. Mills, "Virtual microscopy with extended depth of field," in *DICTA '05: Proceedings of the Digital Image Computing on Techniques and Applications*. Washington, DC, USA: IEEE Computer Society, 2005, p. 35.
- [33] R. K. Sharma and M. Pavel, "Multisensor image registration," *Society for Information Display*, pp. 951–954, 1997.
- [34] M. Irani and P. Anandan, "Robust multi-sensor image alignment," Jan 1998, pp. 959–966.
- [35] J. Liu, B. Vemuri, and F. Bova, "Multimodal image registration using local frequency," 2000, pp. 120–125.
- [36] J. Liu, B. Vemuri, and J. Marroquin, "Local frequency representations for robust multimodal image registration," *Medical Imaging, IEEE Transactions on*, vol. 21, no. 5, pp. 462–469, May 2002.
- [37] S. Henn and K. Witsch, "Multimodal image registration using a variational approach," *SIAM J. Sci. Comput.*, vol. 25, no. 4, pp. 1429–1447, 2003.
- [38] P. Viola and Wells, "Alignment by maximization of mutual information," *International Journal of Computer Vision*, vol. 24, no. 2, pp. 137–154, 1997.
- [39] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *Medical Imaging, IEEE Transactions on*, vol. 16, no. 2, pp. 187–198, April 1997.
- [40] R. Gan, J. Wu, A. C. S. Chung, S. C. H. Yu, and W. M. W. Iii, "Multiresolution image registration based on kullback-leibler distance," in *In The 7th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI04)*. SpringerVerlag, 2004, pp. 599–606.
- [41] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q. C. Pham, T. Katila, and I. Magnin, "A review of cardiac image registration methods," *Medical Imaging, IEEE Transactions on*, vol. 21, no. 9, pp. 1011–1021, Sept. 2002.
- [42] K. Torkkola, "Feature extraction by non parametric mutual information maximization," *J. Mach. Learn. Res.*, vol. 3, pp. 1415–1438, 2003.
- [43] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1226–1238, 2005.
- [44] F. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 11, no. 6, pp. 567–585, Jun 1989.
- [45] M. Moshfeghi, "Elastic matching of multimodality medical images," *CVGIP: Graph. Models Image Process.*, vol. 53, no. 3, pp. 271–282, 1991.
- [46] G. Christensen, R. Rabbitt, and M. Miller, "Deformable templates using large deformation kinematics," *Image Processing, IEEE Transactions on*, vol. 5, no. 10, pp. 1435–1447, Oct 1996.
- [47] M. Bro-Nielsen and C. Gramkow, "Fast fluid registration of medical images," in *VBC '96: Proceedings of the 4th International Conference on Visualization in Biomedical Computing*. London, UK: Springer-Verlag, 1996, pp. 267–276.
- [48] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: a survey," *Medical Image Analysis*, vol. 1, no. 2, pp. 91 – 108, 1996.
- [49] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes, "Nonrigid registration using free-form deformations: application to breast mr images," *Medical Imaging, IEEE Transactions on*, vol. 18, no. 8, pp. 712–721, Aug. 1999.
- [50] G. Rohde, A. Aldroubi, and B. Dawant, "The adaptive bases algorithm for intensity-based nonrigid image registration," *Medical Imaging, IEEE Transactions on*, vol. 22, no. 11, pp. 1470–1479, Nov. 2003.
- [51] E. D'Agostino, F. Maes, D. Vandermeulen, and P. Suetens, "A viscous fluid model for multimodal non-rigid image registration using mutual information," in *MICCAI '02: Proceedings of the 5th International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2002, pp. 541–548.
- [52] W. R. Crum, T. Hartkens, and D. L. G. Hill, "Non-rigid image registration: theory and practice," *British Journal of Radiology*, vol. 77, no. 2, pp. 140–153, 2004.
- [53] S. Oldridge, G. Miller, and S. Fels, "Automatic classification of image registration techniques," in *ICVS 09: Proceedings of the International Conference on Computer Vision*. London, UK: Springer-Verlag, 2009.
- [54] —, "Classification of image registration problems using support vector machines," in *Workshop on the Applications of Computer Vision*, January 2011.