

Classification of Image Registration Problems Using Support Vector Machines

Steve Oldridge, Sidney Fels, Gregor Miller
University of British Columbia
Department of Electrical and Computer Engineering
steveo, ssfels, gregor@ece.ubc.ca

Abstract

This paper introduces a system that automatically classifies image pairs based on the type of registration required to align them. The system uses support vector machines to classify between panoramas, high-dynamic-range images, focal stacks, super-resolution, and unrelated image pairs. A feature vector was developed to describe the images, and 1100 pairs were used to train and test the system with 5-fold cross validation. The system is able to classify the desired registration application using a 1:Many classifier with an accuracy of 91.18%. Similarly 1:1 classifiers were developed for each class with classification rates as follows: Panorama image pairs are classified at 93.15%, high-dynamic-range pairs at 97.56%, focal stack pairs at 95.68%, super-resolution pairs at 99.25%, and finally unrelated image pairs at 95.79%. An investigation into feature importance outlines the utility of each feature individually. In addition, the invariance of the classification system towards the size of the image used to calculate the feature vector was explored. The classification of our system remains level at ~91% until the image size is scaled to 10% (150 x 100 pixels), suggesting that our feature vector is image size invariant within this range.

1. Introduction

Image registration is the process of calculating spatial transforms which align a set of images to a common observational frame of reference, most frequently one of the images in the set. Registration is a crucial step in any image analysis or understanding task where multiple sources of data are combined. It is commonly used in computational photography [1], remote sensing [10, 4], medical image processing [11, 14], and many other computer vision tasks.

Although the field is rapidly moving towards automatic image registration, algorithms and systems are most often limited to a single application domain, such as stitching panoramas [3], super-resolution imaging [9, 19], high-

dynamic-range (HDR) imaging [16, 15], or focal stacking [1]. The algorithms used vary significantly depending on the type of registration being performed: When images vary by more than just alignment the proper selection of appropriate algorithm is critical in calculating the correct spatial transform. For example feature based methods useful in panorama stitching perform poorly at focal stacking problems because the overlapping high frequency regions from which features are found and matched do not exist. Techniques can often be used on a limited subset of problems from other domains, however no single algorithm exists that will solve all types of registration. Figure 1 presents example image pairs from each type of registration and demonstrates the differences between these images.

This paper introduces a system that attempts to automatically classify registration problems into common computational photography application domains. Support Vector Machines (SVMs) are used to classify image pairs, determining whether a given pair is from a panorama, a high-dynamic-range image, a focal stack, a super-resolution set, or is unrelated. Using this approach it becomes possible to select an appropriate method of solution based on the classification. This begins to approach what Zitová and Flusser [18] term ‘*the ultimate registration method*,’ which is ‘*able to recognize the type of given task and to decide by itself about the most appropriate solution*.’

In order to classify amongst the different categories of registration application using SVMs a feature vector is needed which describes the image pairs. In the case of image registration the variation between the images is one of the most distinguishing features between types. We have evaluated a wide range of features for this description and present them below in Section 3.1. Many of the features examined are general representations of the image as a whole, in the form of histograms or differences of histograms. These representational features are less likely to be affected by the size of the images and as such we have examined the effect of the size of images used to calculate the feature vector on the classification rate of the system. In addition we look at feature importance, evaluating which features

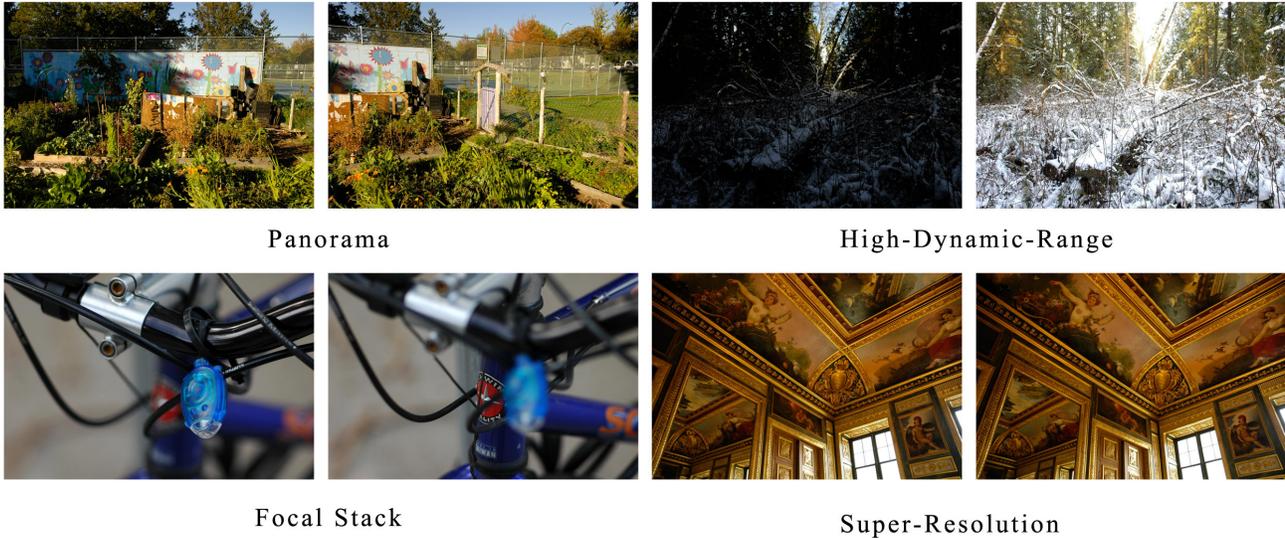


Figure 1. Image pairs representative of the different types of registration problems that occur in computational photography. From the top left: panorama, high-dynamic-range, focal stack, and super-resolution.

are best at distinguishing between the different classes of registration.

Support vector machines are a learning based approach and classify based on a set of training data. A set of 1100 image pairs were created, manually classified according to the type of registration, and used to train and test the system. This set was used to train both 1:many classification of the registration problem domain, as well as 1:1 classification of individual registration problem types, i.e. is this pair likely from a panorama? This 1:1 classification is useful for validating the performance of the 1:many classifier which should see similar results. The set of image pairs used to train the system will eventually be made available online [12] for researchers who wish to improve upon these results.

Section 2 looks at traditional taxonomies of registration and how these can be used to guide in the classification of registration problems. Section 3 gives a brief description of Support Vector Machines and how they are used to classify, outlining the feature vector that is used to classify the image pairs and the set of image pairs used in training and testing. The results of the classification of this set of pairs by our system is shared in Section 4, and the effect that the size of the image has on classification is examined. Finally Section 5 concludes and presents future work.

2. Related Work

Image registration survey papers provide a methodological understanding of the different algorithms used to solve the registration problem. Brown [2] divides registration into four components: feature space, search space, search

strategy, and similarity metric. Within Brown’s framework, knowledge of the types of variation that occur in image sets is 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 which may be of interest once the images have been aligned. These are labeled by Brown ‘corrected distortions’, ‘uncorrected distortions’, and ‘variations of interest’ respectively. This paper outlines a system that attempts to automatically detect both ‘corrected’ and ‘uncorrected distortions’ and use them as a basis for classification of the type of registration problem to be solved, which in turn guides in the selection of suitable algorithms. We focus on Brown’s ‘corrected distortions,’ which are easier to detect and provide significant guidance in the selection process.

More recently, Zitová and Flusser [18] have differentiated the field of registration into area and feature based methods. The four basic steps of image registration under their model are: feature detection, feature matching, mapping function design, and image transformation and resampling. While they do not provide a model of variation equivalent to Brown’s, they discuss in detail the advantages and drawbacks of each method, allowing a similar mapping of methodology from situation. In the conclusion of their survey of registration techniques they propose the creation of ‘the ultimate registration method,’ which is ‘able to recognize the type of given task and to decide by itself about the most appropriate solution.’ This paper is an attempt at rec-

ognizing different types of registration tasks.

Systems for automatic image registration exist, however they are most often limited to a single application such as stitching panoramas [3], super-resolution [9, 19], high dynamic range (HDR) imaging [16, 15], or focal stacking [1]. These techniques can be used on a limited subset of problems from other domains, however no single algorithm exists that will solve all types of registration. Yang et al. [17] extend the flexibility of their algorithm within other problem domains by analyzing the input image pairs and setting parameters accordingly, however the single underlying algorithm still fails in a number of their test cases.

Drozd et al. [8] propose the creation of an expert system based tool for autonomous registration of remote sensing data, and outline a plan to use information derived from image metadata and user tags to select from amongst correlation based, mutual information based, feature based, and wavelet based methods. Unfortunately their description is more of a preliminary proposal and doesn't provide results of the performance of their expert system or of how appropriate the registration techniques selected were at solving the problems they were chosen for.

Oldridge et al. [13] implemented an expert system that uses a series of rules applied to measured variations of the image pair in order to classify image registration problems. Their system used a two step process: a 1:1 classification to determine whether a pair was likely to be part of a given class, and a normalized cross correlation step which ran all candidate algorithms and compared the results of their transforms in each corresponding error space, choosing the transform that was minimum across all spaces. Using this system they were able to achieve an 87% classification rate for classification into the three categories of purely spatial varying (panorama), intensity varying (high-dynamic-range), and focus varying (focal stacks). Adding an unrelated category lowered this classification to 78%, however their test set of only 60 image pairs provide a limited basis from which to draw significant conclusions. Our proposed system uses a 1:Many classification scheme which does not require the registration algorithms' results, dramatically reducing the required computation. In addition we train and test our system using a set of 1100 image pairs, providing a much more robust evaluation of our methodology.

Finally, Support Vector Machines have been used similarly to classify the aesthetics of photographs by Datta et al. [7]. Here a feature vector was developed to describe individual images, consisting of measures of: colorfulness, saturation, hue, rule-of-thirds, familiarity, wavelets-based texture, size, aspect ratio, region composition, depth of field, and shape convexity. These features were calculated for a set of images appropriated from photo.net, which ranks the aesthetics of photos on a scale of 7, and a support vector machine was trained to classify images as either high quality

(> 5.8) or low quality (< 4.2), achieving a 77% accuracy.

3. Classification of Image Registration Problems Using Support Vector Machines

Support Vector Machines (SVMs) are a supervised learning technique similar to neural networks, which can be used for classification. SVMs work by projecting a description of the element to be classified, known as a feature vector, into a higher dimensional data space where the different classes become linearly separable. LibSVM [5] and the accompanying scripts were used to find the appropriate SVM parameters for our problem, normalize the feature vectors, train the classification models, test the models, cross validate the results, and explore the relevance of features from our feature vector. We utilize the recommended Radial Basis Function (RBF) based kernel, and linearly scale our feature vector data to [-1,1]. A grid search of the two RBF parameters C and γ was performed to find the best settings for our problem independently for each classifier. To prevent over-fitting a 5-fold cross-validation scheme was utilized. This method breaks the test set into five equal subsets, testing each subset on a classifier trained from the other four subsets.

The development of a feature vector that appropriately describes the different aspects of the image pairs to be classified is a critical part of the system. In order to classify image registration we have developed a feature vector to describe both the images themselves, as well as the variation between the image pairs.

Additionally, to train and test the system a significant dataset of image pairs is needed. We have created a set consisting of 1100 image pairs, representing the five classifications of: panorama, high-dynamic-range image, focal stack, super-resolution image, and finally unrelated images. The set is made available to other researchers on our website [12] in order to allow for direct comparison when improving upon our classification results.

3.1. Feature Vector

The description of the image pairs by the feature vector is critical in the system's ability to classify. In order to ensure that the SVM has enough information about the pair to correctly classify it we include information about each image, measures of the global relationships between aspects of the image pair, and also local relationship (pixel-wise) between these aspects. In Section 4.1 we examine the importance of each feature in our feature vector.

For each image we calculate the average intensity, hue, and saturation of the image. The power of each image, which serves as a measure of how in-focus the image is, is calculated by applying a 5x5 Laplace filter to detect edges, squaring the result, and averaging across the entire image.

Feature	Panorama	HDR
Intensity(1)	105.415	91.5064
Intensity(2)	109.824	177.419
Hue(1)	23.3616	39.0234
Hue(2)	23.656	46.2149
Saturation(1)	154.964	115.74
Saturation(2)	153.859	74.4806
Power(1)	73.6296	38.6469
Power(2)	80.5668	37.3
Num Features(1)	26403	4982
Num Features(2)	27092	5395
Num Features Per Pixel(1)	0.0175552	0.0033125
Num Features Per Pixel(2)	0.0180133	0.0035871
Intensity Overlap	0.966408	0.40814
Hue Saturation Overlap	0.963122	0.41148
Focus Overlap	0.0596722	0.0180572
Matched Features	4467	594
Matched Features Per Pixel	0.002970	3.94947e-4
Matched-Feat Centroid (1)	(0.455,0.594)	(0.481,0.419)
Matched-Feat Centroid (2)	(0.487,0.548)	(0.423,0.427)
Intensity Overlap UL	0.0598358	0.0178404
Intensity Overlap UM	0.0597852	0.0178797
Intensity Overlap UR	0.0599069	0.0179156
Intensity Overlap ML	0.0598524	0.0179641
Intensity Overlap MM	0.0597819	0.0179501
Intensity Overlap MR	0.059732	0.0179681
Intensity Overlap LL	0.0598517	0.018018
Intensity Overlap LM	0.0596769	0.0180459
Intensity Overlap LR	0.059611	0.0180106

Table 1. Example of the feature vector and its corresponding values for a panorama image pair and a high-dynamic-range (HDR) image pair.

Finally the number of SIFT features in each image is calculated. In [13] these values and their differences across image pairs were the basis for their rule based expert system classifier.

These differences between images provides us with significant insight regarding which category a pair belongs to. Rather than simply take the difference between average values we compare the overlap of the images’ histograms, both globally and locally. Intensity and Power histograms are calculated using 64 bins, and their overlap is calculated as the of intersection between the two dense histograms. Our joint two dimensional Hue / Saturation histogram has 30x32 bins and is compared similarly. The number of matched SIFT features, as well as the number matched per pixel is calculated for the vector. In [3], Brown uses the number of matched features as a basis for selecting the next image to combine into a panorama, performing a similar classification, motivating this inclusion. In addition the centroid of the matched features is calculated for each image.

Finally, to make local image comparisons we divide the images into nine equal regions, comparing the overlap of local intensity histograms for each section of the image. Table 1 demonstrates the complete feature vector, and provides example values for a pair of panorama images and a pair of high-dynamic-range images.

3.2. Data Set

To train and test the SVM a set of image pairs which are representative of each of our classes under various conditions possible is necessary. In our creation of this test set we have attempted to include image pairs taken in different lighting conditions and settings so as not to unwittingly bias the learning system. Photographs were taken in pairs, with a specific single application in mind, and categorized accordingly. Images were scaled from 3008x2000 pixels to our base size of 1504x1000 pixels in size in order to accommodate memory limitations of our SIFT feature implementation. As we will see in Section 4.2, which looks at the feature vector’s invariance to scale, this is unlikely to affect performance of the classification of full sized images. 1100 pairs of images were taken in total, divided evenly amongst the five possible groupings: 220 Panorama pairs, 220 High-Dynamic-Range pairs, 220 Focal pairs, 220 Super-Resolution pairs, and finally 220 ‘unrelated’ pairs. Unrelated pairs consist of images taken from within the same category, for example both images are from a focal stack, and are distributed evenly amongst the four application based categories.

The same set of image pairs was used both in the training of our 1:many classifier, which labels across all classes, and our 1:1 classifiers, which attempt to classify whether an image is a part of a given class or not using all other classes as negative training cases.

4. Results

As mentioned in Section 3 we have trained and tested our SVM using 5-fold cross validation in order to prevent overfitting of our data set. The 1:many classification rate of the system for full sized images is 91.18%. This rate of classification makes the automation of image registration tools feasible, and would allow photographers to reliably group sets of photos automatically by type and apply an appropriate registration algorithm.

The 1:1 classifiers were similarly trained and provide a point of comparison to the 1:many classifier. Classification rates for our 1:1 classifiers are as follows: Panorama image pairs are classified at 93.15%, high-dynamic-range pairs at 97.56%, focal stack pairs at 95.68%, super-resolution pairs at 99.25%, and finally unrelated image pairs at 95.79%. Table 2 summarizes these classification rates.

Type	Classification Rate
1:1 Panorama	93.15%
1:1 High-Dynamic-Range	97.56%
1:1 Focal Stack	95.68%
1:1 Super-Resolution	99.25%
1:1 Unrelated pairs	95.79%
1:Many Overall	91.18%

Table 2. Summary of classification rates for our 1:Many and 1:1 classifiers.

4.1. Feature Importance

Examining the importance of the individual features with regards to classification provides insight into how the classifications are taking place. Chen et al. [6] developed a measure of feature importance known as an 'FScore' which measures the discrimination of two sets of real numbers. Table 3 presents the FScores of our feature vector for our 1:Many classifier.

Ordered features	FScore
Hue / Saturation Overlap	2.524745
Intensity Overlap	2.022537
Matched Features	1.269260
Matched Features Per Pixel	1.269115
Intensity Overlap UR	1.199512
Intensity Overlap UM	1.197718
Intensity Overlap UL	1.196403
Intensity Overlap LL	1.195831
Intensity Overlap MR	1.193893
Focus Overlap	1.193237
Intensity Overlap MM	1.192491
Intensity Overlap ML	1.190727
Intensity Overlap LR	1.189334
Intensity Overlap LM	1.187563
Power(1)	0.364715
Power(2)	0.352534
Num Features Per Pixel(1)	0.266387
Num Features(1)	0.260805
Num Features Per Pixel(2)	0.234179
Num Features(2)	0.229223
Intensity(1)	0.227655
Saturation(2)	0.170803
Hue(2)	0.126434
Hue(1)	0.121199
Intensity(2)	0.069913
Saturation(1)	0.066153
Matched-Feat Centroid (2 X)	0.060326
Matched-Feat Centroid (1 X)	0.042694
Matched-Feat Centroid (2 Y)	0.006689
Matched-Feat Centroid (1 Y)	0.002026

Table 3. Ordered list of feature importance and corresponding FScore measure.

Using the FScore ranking as a basis, features can also be removed from the system and the impact on classification measured, providing a basis for improvement of the speed of the system at a cost of accuracy. With our complete feature vector we achieved a classification rate of 91.18%. Reducing the number of features to: 30 results in a 90.90% classification rate, 15 results in a 90.52% classification rate, 7 results in 86.59% accuracy, and finally using only the top 3 features results in 85.83% classification rate. Figure 2 summarizes these results.

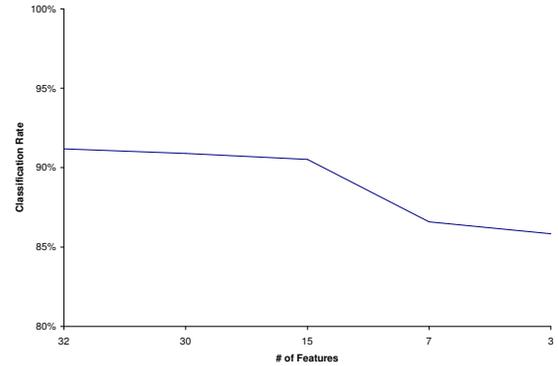


Figure 2. Summary of classification rates based on reducing the number of features in the feature vector. For 32, 30, 15, 7, and 3 features respectively.

Also of interest is the top three features from each 1:1 classifier. These provide insight into how classification of each type of registration is being performed. The top three features from each category are summarized in Table 4. As expected the most important features in each class relate directly to the common forms of variation that are indicative of that class.

4.2. Invariance to Scale

Computation time of the feature vector is exponentially (n^2) related to the size of the images from which it is calculated. Many of the elements in our feature vector are global calculations such as histograms or averages which are unlikely to be significantly affected by the size of the image. As such, we have investigated the impact of image size on the classification rate of our system. To test this feature vectors were generated from 1504x1000 pixel images calculated at: 100%, 80%, 60%, 50%, 40%, 30%, 25%, 20%, 15%, 10%, 5%, 4%, 3%, and 2% scale, and a 1:many classifier was trained using 5-fold cross validation. As we see in Figure 3 classification remains level at ~91% until the image is scaled down to 10% of its original size (150 x 100

Type	# 1 Feature	# 2 Feature	# 3 Feature
1:1 Panorama	Power(1)	Power(2)	Matched Features
1:1 High-Dynamic-Range	Intensity Overlap	Intensity (1)	Intensity Overlap UR
1:1 Focal Stack	Power(1)	Power(2)	Matched Features
1:1 Super-Resolution	Matched Features	Matched Features Per Pixel	Intensity Overlap UR
1:1 Unrelated pairs	Hue / Saturation Overlap	Matched Features	Matched Features Per Pixel

Table 4. Top three features of each 1:1 classifier. As expected the most important features in each class relate directly to the common forms of variation that are indicative of that class.

pixels). Decreasing the size of the images to 2% of their original size (30 x 20 pixels) results in a classification rate of 79.7%.

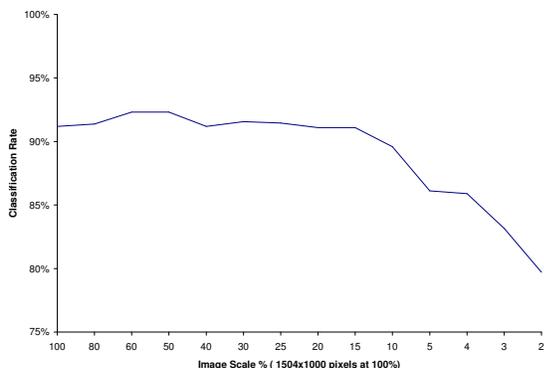


Figure 3. Degradation of classification as the size of the input pairs decreases. Classification remains level around 91% until the image is decreased to 10% of its original size (150 x 100 pixels). Decreasing the size of the images further to 2% of its original size (30 x 20 pixels) still results in a classification rate of 79.7%

This decrease of image size, in combination with the selection of features based on importance, begins to reduce the computation necessary for classification to the point where it becomes possible to do on-camera. On-camera classification would allow photographers to automatically organize panoramas, high-dynamic-range images, focal stacks, and super-resolution images as they are taken, significantly reducing the manual labor currently involved in their creation.

Investigating feature importance at smaller scales we see that many features are unaffected by the change in scale. The number of SIFT features, both matched and total per image, is a notable exception: as image size decreases the overall number of SIFT features and the number of matches drops significantly, even for super-resolution or panorama images, particularly when images are smaller than 10%. Additionally the focus overlap becomes much more impor-

tant in distinguishing between classes, jumping in importance from #10 to #3.

5. Conclusions and Future Work

In this paper we have described a novel classification of image pairs according to the category of registration they belong to using support vector machines. 1100 pairs of images was collected, divided evenly amongst the five possible groupings: 220 Panorama pairs, 220 High-Dynamic-Range pairs, 220 Focal pairs, 220 Super-Resolution pairs, and finally 220 'unrelated' pairs, and is made available online to support future research [12].

A 1:many classifier was trained which is able to classify between panoramas, high-dynamic-range-images, focal stacks, super-resolution images, and unrelated image pairs with a 91.18% accuracy. 1:1 classifiers were also developed to classify each of the categories individually. Classification rates for our 1:1 classifiers are as follows: Panorama image pairs are classified at 93.15%, high-dynamic-range pairs at 97.56%, focal stack pairs at 95.68%, super-resolution pairs at 99.25%, and finally unrelated image pairs at 95.79%.

The importance of features was investigated and the 1:many classification rate was measured for feature vectors of various size, taken from a feature vector ordered by FS-core. Classification was somewhat affected by the reduction in features, and use of the full feature vector is recommended for maximum accuracy.

Finally the invariance of the classification system towards the scale of the image used to calculate the feature vector was explored. Feature vectors were generated at: 100%, 80%, 60%, 50%, 40%, 30%, 25%, 20%, 15%, 10%, 5%, 4%, 3%, and 2% scale, and a 1:many classifier was trained using 5-fold cross validation. The classification of our system remains level at ~91% until the image is scaled to 10% of its original size (scaled to 150 x 100 pixels), suggesting that our feature vector is image size invariant within that range. Decreasing the size of the images to 2% of their original size (30 x 20 pixels) results in a classification rate of 79.7%.

Improvement of the classifier through the development of a system which deals with sets of images, rather than

image pairs, is left as future work. Image sets of a particular class are often taken in sequence, allowing the sequential use of our classifier to combine pairs into sets, however thought must be put into the system to prevent classification errors from compounding. This set based system would replace Brown's "Recognizing Panoramas" [3] providing a solution that is capable of "Recognizing Panoramas, High-Dynamic-Range Images, Focal Stacks, and Super-Resolution images."

Further, the development of a system which uses this classifier in conjunction with image registration and compositing algorithms would allow for fully automatic image registration, and remains an important area of future research. Such a system would be Zitová and Flusser's 'ultimate registration method' [18].

Finally, to improve upon this 'ultimate registration method,' a deep understanding of which registration algorithms perform best under specific conditions is desirable. Such an understanding would allow for the creation of a data set and subsequent classifiers that are tuned to not only registration applications, but also to particular algorithms. Comparative research within the image registration field is currently lacking, and conclusive evaluations of different algorithms' performance in comparison to another, particularly across different problem domains, would be a significant contribution to the field.

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*, pages 294–302, New York, NY, USA, 2004. ACM Press. 1, 3
- [2] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24:325–376, 1992. 2
- [3] M. Brown and D. Lowe. Recognising panoramas. *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, 2:1218–1225, 16–16 Oct. 2003. 1, 3, 4, 7
- [4] J. B. Campbell. *Introduction to remote sensing*. Guildford Press, 4th edition, 2008. 1
- [5] C.-C. Chang and C.-J. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>. 3
- [6] Y. W. Chen and C. J. Lin. *Combining SVMs with various feature selection strategies*. Springer, 2006. 5
- [7] R. Datta, D. Joshi, J. Li, and J. Wang. Studying aesthetics in photographic images using a computational approach. pages III: 288–301, 2006. 3
- [8] A. L. Drozd, A. C. Blackburn, I. P. Kasperovich, P. K. Varshney, M. Xu, and B. Kumar. A preprocessing and automated algorithm selection system for image registration. volume 6242, page 62420T. SPIE, 2006. 3
- [9] 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*, pages 1262–1264. IEEE Computer Society, 1999. 1, 3
- [10] T. M. Lillesand and R. W. Kiefer. *Remote Sensing and Image Interpretation*. Wiley, 6th edition, 2007. 1
- [11] J. Maintz and M. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998. 1
- [12] S. Oldridge, S. Fels, and G. Miller. Image registration classification dataset. <http://www.ece.ubc.ca/~hct/research>, January 2011. 2, 3, 6
- [13] S. Oldridge, G. Miller, and S. Fels. Automatic classification of image registration techniques. In *ICVS 09: Proceedings of the International Conference on Computer Vision*, October 2009. 3
- [14] J. Pluim, J. Maintz, and M. Viergever. Mutual-information-based registration of medical images: a survey. *Medical Imaging, IEEE Transactions on*, 22(8):986–1004, Aug. 2003. 1
- [15] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: High dynamic range in a wide field of view. *Int. J. Comput. Vision*, 53(3):245–267, 2003. 1, 3
- [16] G. Ward. Robust image registration for compositing high dynamic range photographs from handheld exposures. *Journal of Graphics Tools*, 8:17–30, 2003. 1, 3
- [17] 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*, 29(11):1973–1989, Nov. 2007. 3
- [18] B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21:977–1000, 2003. 1, 2, 7
- [19] B. Zitová, J. Kautsky, G. Peters, and J. Flusser. Robust detection of significant points in multiframe images. *Pattern Recogn. Lett.*, 20(2):199–206, 1999. 1, 3