Image alignment for panorama stitching in sparsely structured environments

Giulia Meneghetti, Martin Danelljan, Michael Felsberg and Klas Nordberg

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-121566>

N.B.: When citing this work, cite the original publication.

The original publication is available at www.springerlink.com:

Meneghetti, G., Danelljan, M., Felsberg, M., Nordberg, K., (2015), Image alignment for panorama stitching in sparsely structured environments. In: *Paulsen R., Pedersen K. (eds) Image Analysis. SCIA 2015. Lecture Notes in Computer Science, vol 9127*, pp. 428-439. https://doi.org/10.1007/978-3-319-19665-7_36

Original publication available at:

https://doi.org/10.1007/978-3-319-19665-7_36

Copyright: Springer Verlag (Germany)

<http://www.springerlink.com/?MUD=MP>

Image alignment for panorama stitching in sparsely structured environments

Giulia Meneghetti, Martin Danelljan, Michael Felsberg, and Klas Nordberg

Computer Vision Laboratory, Linköping University, 581 31 Linköping, Sweden
{giulia.meneghetti,martin.danelljan,
michael.felsberg,klas.nordberg}@liu.se

Abstract. Panorama stitching of sparsely structured scenes is an open research problem. In this setting, feature-based image alignment methods often fail due to shortage of distinct image features. Instead, direct image alignment methods, such as those based on phase correlation, can be applied. In this paper we investigate correlation-based image alignment techniques for panorama stitching of sparsely structured scenes. We propose a novel image alignment approach based on discriminative correlation filters (DCF), which has recently been successfully applied to visual tracking.

Two versions of the proposed DCF-based approach are evaluated on two real and one synthetic panorama dataset of sparsely structured indoor environments. All three datasets consist of images taken on a tripod rotating 360 degrees around the vertical axis through the optical center. We show that the proposed DCF-based methods outperform phase correlation based approaches on these datasets.

Keywords: image alignment, panorama stitching, image registration, phase-correlation, discriminative correlation filters

1 Introduction

Image stitching is the problem of constructing a single high resolution image from a set of images taken from the same scene. We consider panorama stitching, the process to merge images taken by a camera fixed on a tripod and rotated about its vertical axis through the optical center. This procedure is visualized in Figure 1. A panorama stitching pipeline usually contains three major steps:

1. Camera calibration: estimate of the camera parameters.
2. Image alignment (or image registration): compute the geometric transformation between the images.
3. Image stitching and blending: transform all the images to a new coordinate system and blend them to eliminate visual artefacts.

Mobile applications and desktop software for panorama images are usually designed to produce visually good results, and therefore focus on the third step. However, the exact estimation of the transformation is required in increasingly

Fig. 1. Left: Visualization of an incorrect image registration output. The blue lines represents the optical axis for each image and the green line represents the vertical axis around which the camera is rotating. In this panorama the top-most image is misaligned with respect to the others. This will greatly reduce the visual quality of the panorama and generate errors in the camera extrinsic parameters. **Center:** Resulting image alignment using POC for an image pair in our *Synthetic dataset*. In this case the images are clearly misaligned. **Right:** Image alignment result of the same image pair using the proposed DFC-CN method. In this case, the alignment is correct and does not contain any visual artefacts.

many fields, including computer graphics (image-based rendering), computer vision (surveillance applications, automatic quality control, vehicular systems applications) and medical imaging (multi-modal MRI merging). In this paper, we therefore investigate the problem of image alignment for panorama stitching.

Image alignment methods can be divided into two categories, namely *feature based* methods and *direct* (or *global*) methods [19, 24, 4]. The feature-based methods work by first extracting descriptors from a set of image features (e.g. points or edges). These descriptors are then matched between pairs of images to estimate the relative transformation. Direct methods instead estimate the transformation between an image pair by directly comparing the whole images. Feature based methods often provide excellent performance in cases when there are sufficient reliable features in the scene. However, these methods struggle and often fail in sparsely structured scenes, i.e. when not enough distinct features can be detected. We find such cases, for example, in indoors scenarios, where uniform walls, floors and ceilings are common, or in outdoor panoramas, where sky and sea can dominate. In this work, we tackle the problem of image alignment for panorama stitching in sparsely structured scenes, and therefore turn to the direct image alignment methods. Since our camera is rotating of small angle the transformation between two consecutive images can approximated as a translation in the image plane. Given this assumption we restrict our investigation to phase correlation approaches [17, 11]. This approach is usually robust against sparsely structured images. However, it is desirable that an alignment method is robust to the perspective distortions caused by the rotation around

the vertical camera axis. Recently, Discriminative Correlation Filter (DCF) [3, 9, 14, 8] based approaches have successfully been applied to the related field of visual tracking, and have achieved state-of-the-art results on standard tracking benchmarks [23, 16]. These methods have shown robustness to many types of distortions and changes in the target appearance, including illumination variations, in-plane rotations, scale changes and out-of-plane rotations [9, 14, 8]. The multi-channel DCF approaches also provide a consistent method of utilizing more general image features, instead of just relying on grayscale values. We therefore investigate to what extent DCF based methods can be used for the image alignment problem in panorama stitching.

1.1 Contributions

In this paper, we investigate the image alignment problem for panorama stitching in sparsely structured scenes. For this application, we evaluate four different correlation-based techniques in an image alignment pipeline. Among phase correlation approaches, we evaluate the standard POC method and a regularized version of POC developed for surveillance systems [11].

Inspired by the success of DCF-based visual trackers, we propose an image alignment approach based on DCF. Two versions are evaluated: the standard grayscale DCF [3] and a multi-channel extension using color names for image representation, as suggested in [9]. Image alignment results for these four methods are presented on three panorama stitching datasets taken in sparsely structured indoors environments. We provide quantitative and qualitative comparisons on one synthetic and two real datasets. Our results clearly suggest that the proposed DCF-based image alignment methods outperform the POC-based methods.

2 Background

Image alignment is a well-studied problem with applications in many different fields. Image stitching, target localization, automatic quality control, super-resolution images, multi-modal MRI merging, are some of many applications that use registration between images. Many techniques have been proposed [19][24][4] in the last decades. Image alignment methods can be divided in two major categories.

2.1 Feature based methods

The feature-based methods mostly vary in the way of extracting and matching the features in an image pair. After the corresponding features have been found, a process of outlier removal is used to improve robustness to false matching.

The estimation of the geometric transformation is usually computed from the corresponding features using Direct Linear Transformation and then refined using bundle adjustment techniques. A classical example of a feature-based registration approach is Autostitch [5, 6] (a panorama stitching software).

Feature-based methods often fail to perform accurate image alignment in sparsely structured scenarios or when the detected features are unevenly distributed. In such cases, direct methods are preferable since they utilize a global similarity metric and do not depend on local features.

2.2 Direct methods

Direct methods can be divided into intensity-based and correlation-based approaches, depending on the kind of similarity measure (or error function). The sum of square difference (SSD) and the sum of absolute differences (SAD) between the intensity values of the pixels of two images, are two intensity-based similarity metrics. Correlation-based approaches compute a similarity score using a “sliding window” procedure. Among these normalized cross correlation (NCC) computes the scalar product of two image windows and divides it by the product of their norms. Given the error (or score) function, various techniques can be applied to find the optimum, such as, exhaustive search of all possible alignments, which is prohibitively slow, hierarchical coarse-to-fine techniques based on image pyramids, which suffer from the fact that image structures need not be on the same scale as the displacement, and Fourier transform-based techniques. The latter techniques are based on the shift theorem of the Fourier transform, the correlation between two images is computed as the multiplication of the first with the complex conjugate of the second one. If the images are related by a translation, phase correlation technique [17], estimates the shift between them by looking for the peak in the inverse Fourier transform of the normalized cross-power spectrum. The normalization is introduced since, it significantly improves the peak estimation with respect to using the cross-power spectrum [15].

For image alignment, this latter technique is preferable, since it uses all the information available in the image. Given two images that differ just by a translation, POC is a simple and robust technique for retrieving the displacement between them. Therefore, we consider this class of techniques in the present paper and suggest a novel approach to image alignment. We propose to use MOSSE [3] and its colornamespace-extension, that have recently been show to perform best on tracking of local regions [9].

2.3 Related Work

The phase correlation method is a frequency domain technique used to estimate the delay or shift between two signals that differ only by a translation. It is based on the shift theorem of the Fourier transform and determine the location of the peak in the inverse Fourier transform of the normalized cross-power spectrum. It is usually applied in motion estimation applications, since it is very accurate and robust to illumination variation and noise in the images. Many versions have been proposed during the years [4, 19, 24]. However, phase correlation for image alignment was first introduced by Kuglin and Hines [17], who compute the displacement as the maximum of the inverse Fourier transform of the normalized

cross-power spectrum between the two images. The drawback was that the recovered translation has pixel accuracy only. Subpixel precision techniques, were later introduced for improving the peak estimation, using fitting functions [12, 20], or finding approximate zeros of the gradient of the inverse Fourier transform of the normalized cross-power spectrum [1], which is more robust against border effect and multiple motions. Foorosh [12] suggested to prefilter the phase difference matrix to remove aliased components (generally at high spatial frequencies), but filtering must be adjusted to each image and sensor. Phase correlation can also be used to estimate other image transformations than pure translation, such as in-plane rotation and scale between two images [10]. Other techniques are based on the Log-polar transformation, since it is invariant to rotation and scaling. Used in combination with correlation, it can robustly estimate scale and in-plane rotations [22]. Chen *et al.* [7] propose a solution for rotated and translated images that computes Fourier-Mellin invariant (FMI) descriptors. Among all the variants that use phase correlation, Eisenbach *et al.* [11] implemented a regularized version based on the noise of the image for surveillance system application.

Phase correlation does not decrease the efficiency depending on the baseline between the images (if they consistently overlap), moreover it is robust against sparsely structured images where feature-based methods fail to detect features. Among all the available phase correlation techniques, we choose the original phase correlation algorithm (POC) [17] and the regularized phase-only correlation version [11].

Recently, discriminative correlation filter based methods [3, 9, 14, 8] have successfully been applied to visual tracking. They have shown to provide excellent results on benchmark datasets [23, 16], while operating above real-time frame rates. These methods work by learning an optimal correlation filter from a set of training samples of the target. Bolme *et al.* proposed the MOSSE tracker [3]. Like standard POC, it only employs grayscale images to estimate the displacement of the target. This method has however been extended to multi-channel features (e.g. RGB) [9, 14, 8]. Danelljan *et al.* [9] performed an extensive evaluation of several color representations for visual tracking in the DCF-based Adaptive Color Tracker (ACT). In their evaluation it was shown that the Color Names (CN) representation [21] provided the best performance. In this work, we therefore investigate two DCF-based methods for the problem of image alignment. First, we evaluate the standard grayscale DCF method (MOSSE) and second, we employ the color name representation used in the ACT tracker in the multi-channel extension of the standard DCF.

3 Method

In this section we present the the evaluation pipeline for the image alignment and the methods investigated. The images are assumed to have been taken from a camera fixed on a tripod. Between each subsequent image, the camera is assumed to have been rotated around the vertical axis through the optical center. The

Fig. 2. Representation of the geometrical relation between the angle α and the displacement d .

rotations are further assumed to be small enough such that two subsequent images have a significantly overlapping view of the scene, at least 50%. We also assume known camera calibration and we, therefore, work on rectified images to compensate for any lens distortion.

3.1 Image Alignment Pipeline

Our image alignment pipeline contains three basic steps that are performed in an iterative scheme. Given a pair of subsequent images u and v , the following procedure is used:

1. Estimate the displacement d between the images u and v using an image alignment method.
2. Using the displacement d , estimate the 3×3 homography matrix H , that maps the image plane of v to the image plane of u .
3. Warp image v to the image plane of u using the homography H .
4. Iterate from 1. using the warped image as v .

Fig. 2 shows a geometrical illustration of the displacement estimation d . O is the common optical center of images u and v , which is projected, respectively, in O_u and O_v . O_{uv} consists in the intersection of the optical axis of u with the image v . The translation between the images is identified as the distance between O_{uv} and O_v . The evaluated image alignment methods used to estimate this displacement, are described in Section 3.2 and 3.3. Given an estimate of the displacement d between the image pair, we can calculate a homography transformation H between the images using the geometry of the problem. Since the camera rotates about the vertical axis through the optical center, the angle α of rotation can be computed as:

$$\alpha = \tan^{-1} \left(\frac{d}{f} \right). \quad (1)$$

Here, f is the focal length of the camera. The homography H between the two images can then be computed as

$$H = K R_\alpha K^{-1}. \quad (2)$$

Here, K is the intrinsic parameter matrix and R_α is the rotation matrix corresponding to the rotation α about the vertical axis

$$R_\alpha = \begin{pmatrix} \cos(\alpha) & 0 & -\sin(\alpha) \\ 0 & 1 & 0 \\ \sin(\alpha) & 0 & \cos(\alpha) \end{pmatrix}. \quad (3)$$

The presented iteration scheme is employed for two reasons. First, it is known that correlation-based methods are biased towards smaller translations due to the fact that a windowing operation and circular correlation is performed. Second, the initial estimate of the displacement is affected by the perspective distortions, since the correlation-based methods assume a pure translation transformation between the image pair. However, as the iterations converge, the translation model will be increasingly correct since the image v is warped according to the current estimate of the displacement. Hence, the estimation of the rotation angle is refined with each iteration. In practice, we noticed that the methods converge already after two iterations in most cases. We therefore restrict the maximum number of iterations to three.

3.2 Phase Only Correlation

The phase correlation method is a frequency domain technique used to estimate the delay or shift between two copies of the same signal. This technique is based on the shift properties of the Fourier transform and determines the location of the peak of the inverse Fourier transform of the normalized cross-power spectrum. Consider two images u and v such that v is translated with a displacement $[x_0, y_0]$ relative to u :

$$v(x, y) = u(x + x_0, y + y_0) \quad (4)$$

Given their corresponding Fourier transforms U and V , the shift theorem of the Fourier transform states that U and V differ only by a linear phase factor

$$U(\omega_x, \omega_y) = V(\omega_x, \omega_y) \cdot e^{i(\omega_x x_0 + \omega_y y_0)} \quad (5)$$

where ω_x and ω_y are the frequency component of the columns and the row of the image. The correlation response s of the normalized cross-power spectrum of U and V is computed from its inverse Fourier transform:

$$s = \mathcal{F}^{-1} \left\{ \frac{U^* \cdot V}{|U^* \cdot V|} \right\} \quad (6)$$

where U^* represents the complex conjugate of U and $\cdot$ denotes the point-wise multiplication. The displacement is then computed as the maximum of the response function. In the ideal case, the inverse Fourier transform of the normalized cross-power spectrum is a delta function centered at the displacement between the two images. A regularized phase correlation version can be found in Eisenbach *et al.* [11], where the response is computed by regularizing the phase correlation using a constant λ . This parameter should be in the order of magnitude

of the noise variance in the individual components of the cross spectrum $V \cdot U^*$:

$$s = \mathcal{F}^{-1} \left\{ \frac{U^* \cdot V}{|U^* \cdot V| + \lambda} \right\} \quad (7)$$

3.3 Discriminative Correlation Filters

Recently, Discriminative Correlation Filters (DCF) based approaches have successfully been applied to visual tracking and have obtained state-of-the-art performance on benchmark datasets [16, 14]. The idea is to learn an optimal correlation filter given a number of training samples of the target appearance. The target is then localized in a new frame by maximizing the correlation response of the learned filter. By considering circular correlation, the learning and detection tasks can be performed efficiently using the Fast Fourier transform (FFT). The tracker implemented by Bolme et al. [3], called MOSSE, uses grayscale patches for learning and detection, and thus only considers luminance information. Multidimensional feature maps (e.g. RGB) [13, 2] generalize this approach, where the learned filter contains one set of coefficients for every feature dimension.

In the application of image alignment, we are only interested in finding the translation between a pair of images. The first image is set as the reference training image used for learning the correlation filter. We consider the D -dimensional feature map component u_j with $j \in \{1, \dots, D\}$. The goal is to learn an optimal correlation filter f_j per feature dimension that minimizes the following cost:

$$\varepsilon = \left\| \sum_{j=1}^D f_j \star u_j - g \right\|^2 + \lambda \sum_{j=1}^D \|f_j\|^2. \quad (8)$$

Here, the star $\star$ denotes circular correlation. The first term is the L^2 -error of the actual correlation output on the training image compared to the desired correlation output g . In this case, g is a Gaussian function with the peak on the displacement. The second term is a regularization with a weight λ . The considered signals f_j , u_j and g are all of the same size, corresponding to the image size in our case. The filter that minimizes the cost (8) is given by

$$F_j = \frac{G^* \cdot U_j}{\sum_{k=1}^D U_k^* \cdot U_k + \lambda}. \quad (9)$$

Here, capital letters denote the discrete Fourier transform (DFT) of the corresponding signals.

To estimate the displacement, the correlation filter is applied to the feature map v extracted from the second image. The correlation response is computed in the Fourier domain as:

$$s = \mathcal{F}^{-1} \left\{ \sum_{j=1}^D F_j^* \cdot V_j \right\} = \mathcal{F}^{-1} \left\{ G \cdot \frac{\sum_{j=1}^D U_j^* \cdot V_j}{\sum_{j=1}^D U_j^* \cdot U_j + \lambda} \right\} \quad (10)$$

Fig. 3. **Left:** Sample image from the *Synthetic dataset*. **Center:** Sample images from the *Lunch Room Blue* dataset. **Right:** Sample image from the *Lunch Room* dataset.

The displacement can then be found by maximizing s .

The multi-channel DCF provides a general framework for incorporating any kind of pixel-dense features. Danelljan et al. [9] recently performed an evaluation of several color features in a DCF based framework for visual tracking. In their work it was shown that the Color Names (CN) [21] representation concatenated with the grayscale channel provides the best result compared to several other color features. We therefore evaluate this feature combination in the presented DCF approach. We refer to this method as DCF-CN.

Eq. 10 resembles the procedure (6) used for computing the POC response. However, two major distinctions exist. First, DCF employs the desired correlation output g , which is usually set to a Gaussian function with a narrow and centered peak. In standard POC the desired response is implicitly considered to be the Dirac function. In the DCF approach g acts as a lowpass filter, providing a smoother correlation response. The second difference is that the cross-correlation is divided by the cross power spectrum in the POC approach. In the DCF method, the cross-correlation is instead divided by the power spectrum of the reference image. For this reason, DCF is not symmetric but depends on which image that is considered the reference.

4 Experiments

In this section we present the datasets we use and the evaluation of the image alignment methods in our pipeline.

4.1 Datasets

As far as we know, no dataset with sparsely structured scene were public available, then we acquired the following three datasets¹.

¹ <http://www.cvl.isy.liu.se/research/datasets/panorama-dataset/index.html>

Synthetic dataset: consists of 72 images of a room rendered with Blender ² with a resolution of 1280×1920 px. Intrinsic parameters were retrieved from Blender and the camera is rotating by 5 degrees between consecutive images. This dataset depicts an almost sparsely structured scene.

Lunch Room Blue: consists of 72 images acquired with a Canon DS50 and perspective lenses with a resolution of 1280×1920 px at poor light condition.

Lunch Room: consists of 72 images acquired with a Canon DS70 and wide angle lenses Samyang 2.8/10mm (about 105 degree of field of view), with a resolution of 5740×3780 px.

For the image acquisition of the real datasets, a panorama head was used to approximate a fixed rotation of 5 degrees around the vertical axis about the optical center of the camera. These datasets were acquired in the same room with different light conditions. They naturally contain more structure than the synthetic images. We have tested all methods on rectified images to remove lenses distortion effects. Figure 3 shows sample images for the three dataset used.

4.2 Results and discussion

We compare four different correlation-based methods: phase correlation (POC) [17], regularized phase correlation (RPOC) [11] and discriminative correlation filter (DCF-CN) with and without (DCF) the color names representation [3, 9]. For reference, a state-of-the-art feature-based approach for panoramic image stitching has been included [18]. The results are shown using three different evaluation metrics. Table 1 I shows the standard deviation of the estimated angles compared from the reference angle of 5 degrees. Table 1 II shows the success rate of the four methods on the three dataset. An estimated angle is considered to be an inlier (and therefore a success) if the error is smaller than a threshold. The value for the threshold has been computed as the 95th percentile of the absolute error on each dataset for all four methods. However, the threshold is set to 2 degrees for the synthetic dataset. This is due to the high failure rate of the POC methods on this particular dataset. Finally, Table 1 III shows the average estimated angle in the three datasets when only considering the successful estimates (inliers).

We observe that the proposed DCF based methods outperform the POC methods in all three datasets. The achieved success rates (Table 1 II) in the synthetic dataset clearly demonstrate that POC-based methods fail on the majority of cases. This is most probably due to the sparsity of visual structure in the scene. In the same scenario, both DCF based methods provide a 100% inlier rate and below 0.07 degrees in standard deviation. Among the successful estimates on the synthetic dataset, the DCF based approaches still outperform the evaluated POC methods. The average angle (Table 1 III) is correct within 0.01 degrees for the DCF method. For the Lunch Room Blue dataset DCF and DCF-CN achieve significantly lower standard deviations of 0.61 and 0.62 degrees respectively compared to 1.41 for POC and 1.44 degrees for RPOC. Similarly, there is a clear

² <http://www.blender.org/>

Table 1. Results of each method for all three datasets. **I:** Standard Deviation of the estimated angles from the reference angle (degrees). **II:** Inlier rate for the four methods (threshold set at 95 percentile). **III:** Average inter-frame rotation in degrees (successful cases).

	Synthetic dataset			Lunch Room Blue			Lunch Room		
	I	II	III	I	II	III	I	II	III
Feature-based	0.95	98.63%	4.97	4.68	84.93%	5.04	0.86	94.52%	4.70
POC	2.52	31.94%	5.20	1.41	90.41%	5.41	0.56	97.22%	5.29
RPOC	2.47	41.67%	4.98	1.44	91.78%	5.19	1.57	87.50%	5.08
DCF	0.07	100.00%	5.00	0.62	98.63%	5.18	0.51	97.22%	4.97
DCF-CN	0.06	100.00%	4.99	0.61	98.63%	5.17	0.50	97.22%	4.98

difference in the inlier rate. On the Lunch Room dataset, the standard DCF and DCF-CN achieve a slight improvement over normal POC, while RPOC provides inferior results. Table 1 II shows that the DCF-based methods provide the same inlier rate as POC. However, they perform better in terms of accuracy both in standard deviation (Table 1 I) and mean angle estimation (Table 1 III). Table 1 II and Table 1 III show that the feature-based method performs well when is able to retrieve reliable features. Nevertheless, we can notice that it is not as strong as the DCF based methods. The success of the DCF based approaches is likely due to its robustness to geometric distortions, which has previously been demonstrated in the application of visual tracking. This property is largely attributed to the desired correlation output g , which regularizes the correlation response as discussed in Section 3.3. Moreover, our results indicate an improvement in precision and robustness when using the color names representation instead of only grayscale images in our DCF-based framework. Figure 1 Center and Right show a comparison between the standard POC and the DCF using color names representation on an image pair.

5 Conclusions

In this paper, we tackle the problem of image alignment for panorama stitching in sparsely structured scenes. We propose an image alignment pipeline based on discriminative correlation filters. Two DCF-based versions are evaluated on three panorama datasets of sparsely structured indoor environments. We show that the proposed methods are able to perform robust and accurate image alignment in this scenario. Additionally, DCF-based methods are shown to outperform the standard and the regularized phase-correlation approaches.

Future work will consider extending our evaluation with other panorama datasets of even more challenging scenarios. We will also look into generalizing our image alignment pipeline for more general image mosaicking problems.

Acknowledgments: The authors thank Andreas Robinson for providing the synthetic dataset. Funding granted by the VPS project.

References

1. Alba, A., Aguilar-Ponce, R., Vigueras-Gmez, J., Arce-Santana, E.: Phase correlation based image alignment with subpixel accuracy. In: *Advances in Artificial Intelligence*, vol. 7629, pp. 171–182. Springer Berlin Heidelberg (2013)
2. Boddeti, V.N., Kanade, T., Kumar, B.V.K.V.: Correlation filters for object alignment. In: *CVPR*. IEEE (2013)
3. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: *CVPR* (2010)
4. Brown, L.G.: A survey of image registration techniques. *ACM Computing Surveys* 24, 325–376 (1992)
5. Brown, M., Lowe, D.: Recognising panoramas. In: *IJCV*. vol. 2, pp. 1218–1225. Nice (October 2003)
6. Brown, M., Lowe, D.G.: Automatic panoramic image stitching using invariant features. *IJCV* 74(1), 59–73 (2007)
7. Chen, Q., Defrise, M., Deconinck, F.: Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition. *TPAMI* 16(12), 1156–1168 (1994)
8. Danelljan, M., Häger, G., Shahbaz Khan, F., Felsberg, M.: Accurate scale estimation for robust visual tracking. In: *BMVC* (2014)
9. Danelljan, M., Shahbaz Khan, F., Felsberg, M., van de Weijer, J.: Adaptive color attributes for real-time visual tracking. In: *CVPR* (2014)
10. De Castro, E., Morandi, C.: Registration of translated and rotated images using finite fourier transforms. *TPAMI* 9(5), 700–703 (1987)
11. Eisenbach, J., Mertz, M., Conrad, C., Mester, R.: Reducing camera vibrations and photometric changes in surveillance video. In: *AVSS*. pp. 69–74 (Aug 2013)
12. Foroosh, H., Zerubia, J., Berthod, M.: Extension of phase correlation to subpixel registration. *Image Processing, IEEE Transactions on* 11(3), 188–200 (Mar 2002)
13. Galoogahi, H., Sim, T., Lucey, S.: Multi-channel correlation filters. In: *ICCV* (2013)
14. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *CoRR* abs/1404.7584 (2014)
15. Horner, J.L., Gianino, P.D.: Phase-only matched filtering. *Applied Optics* 23(6), 812–816 (15 March 1984)
16. Kristan, M., Pflugfelder, R., Leonardis, A., Matas, J., et al.: The visual object tracking vot2014 challenge results. In: *ECCVW* (2014)
17. Kuglin C.D., H.: The phase correlation image alignment method. In: *1975 International Conference on Cybernetics and Society*. (1975)
18. MATLAB: Computer Vision Toolbox - version 8.4.0 (R2014b). The MathWorks Inc., Natick, Massachusetts (2014)
19. Szeliski, R.: Image alignment and stitching: A tutorial. *Found. Trends. Comput. Graph. Vis.* 2(1), 1–104 (Jan 2006)
20. Takita, K.: High-accuracy subpixel image registration based on phase-only correlation. *IEICE transactions on fundamentals of electronics, communications and computer* 86(8), 1925–1934 (2003)
21. van de Weijer, J., Schmid, C., Verbeek, J.J., Larlus, D.: Learning color names for real-world applications. *TIP* 18(7), 1512–1524 (2009)
22. Wolberg, G., Zokai, S.: Robust image registration using log-polar transform. In: *ICIP* (2000)
23. Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: *CVPR* (2013)

24. Zitov, B., Flusser, J.: Image registration methods: a survey. *Image and Vision Computing* 21, 977–1000 (2003)