Reducing Drift in Mosaicing Slit-Lamp Retinal Images

Kristina Prokopetc^{1,2} ¹QuantelMedical 11 Rue du Bois Joli, 63808 Cournon-d'Auvergne, France prokopec.kristina@mail.ru

Abstract

The construction of seamless and accurate mosaics from long slit-lamp retinal video sequences is an important and challenging task in navigated Pan-Retinal Photocoagulation. The main difficulty is accumulated registration drift due to the small number of features away from the optic nerve and the distortion induced by the geometry of the eye and the contact lens. We present a new approach to reduce the drift. Our main idea is to create long-term high precision point correspondences by associating a simple global model with local correction and perform key-frame based Bundle Adjustment. We evaluate the method's performance compared to state-of-the-art. The results obtained with our method show significantly lower accumulated error.

1. Introduction

Pan-Retinal Photocoagulation (PRP) is usually a recommended treatment for patients with advanced proliferative retinopathy, a situation in which abnormal new retinal blood vessels may rupture and bleed inside the eye. During the treatment an ophthalmologist uses a therapeutic laser to make tiny burns on the affected areas of the retina. The conventional PRP through the slit-lamp with manual navigation has been the standard for laser delivery for many decades. However, sequential improvements in the past ten years led to the development of computer guided photocoagulation systems with integrated imaging and automatic navigation [5, 9, 10]. While the fundus camera based system [9] is considered state-of-the-art, the magnification and control offered by the slit-lamp still makes it a very popular choice in the clinical environment. In this context the development of a platform to combine conventional slit lamp based laser delivery with assisted navigation is on demand. Recently, a computer assisted slit-lamp based platform prototype has been developed in QuantelMedical. The prototype combines real-time HD imaging, pre-operative planning and intra-operative navigation. Images acquired with Adrien Bartoli² ²ALCoV-ISIT, / Université d'Auvergne 28 place Henri Dunant, 63001 Clermont-Ferrand, France

adrien.bartoli@gmail.com

the prototype have a narrow field of view (FOV) visualizing only small thin portions of the retina. Obtaining the larger view can greatly facilitate the treatment planning and navigation. This is usually solved using image registration techniques which allow one to stitch multiple images of the same scene to an image mosaic.

Retinal image registration has been well studied, see [1, 4, 27, 26, 23, 13], to name a few. The majority of approaches operates with images obtained with a fundus camera. The quality of this type of images is higher. They have fewer specular reflections, good contrast and almost no blur. The registration implies using 10 consecutive images at most, thus, not dealing with the problem of heavy drift accumulation in a long video sequence. Hence, these methods are not able to provide desirable results when directly applied to mosaicing slit-lamp images, where the length of a sequence spans at least a thousand images. Very few works specifically deal with slit-lamp images [2, 20, 22, 19]. We consider [19] as the state-of-the-art, where a retinal mosaic is built in real-time. This method was integrated into the prototype developed in QuantelMedical. The method is a combination of direct tracking and feature detection approaches, and the mosaic is constructed by registering the images sequentially with pairwise transformations. Despite the low computational cost and simplicity of this approach, due to its 'chaining' nature, alignment errors tend to accumulate, causing images to drift in the mosaic (figure 1). Our main objective in this work is to reduce the drift. The problem of accumulated drift was addressed in a wide variety of image registration applications [25, 16, 3, 7, 12]. The usual approach practiced by many is to perform global Bundle Adjustment (BA). In real-time systems it has been left as a post-processing step for a long time. However, in the past few years a number of real-time local BA-type refinement methods were proposed [17, 11, 6] which allow one to achieve a similar accuracy to conventional BA while reducing computational cost.

We present a method for drift correction in mosaicing monocular slit-lamp images. The solution is to expand the



Figure 1: (Best viewed electronically) Examples of mosaics obtained with [19]. (a) - registration drift is visible through the mismatched vascular structure, (b) - example with drift induced blurred regions and duplication. Visual assessment was performed by and expert.

span of tracks across the images. Our methodology relies on three main assumptions: (*i*) point correspondences presented in multiple views provide more constraints, (*ii*) a simple global motion model associated with local correction can be used to predict the track location, this may help to obtain tracks longer than short-inter-frames with improved precision, (*iii*) using a simple global model to initialize key-frame based local BA can be as accurate as performing global BA while being less computationally expensive. We perform partial assessment of our approach and evaluate it against state-of-the-art [19].

2. Methodology

Our algorithm consists of the following steps: (1) Initialization with a key-point detector, (2) Motion Estimation based on key-frames selection, (3) Prediction, using a popular tracking algorithm, (4) Track Correction using a simple global model with local adjustment (5) Key-frame Instantiation and Local BA. The detailed description is given in the subsequent paragraphs.

2.1. Initialization

We use latin bold to refer to key-points (such as **p**) and Greek chatacter τ to refer to point tracks. Frame index-

ing is denoted as $f = 1, ..., n_f$ and $k = 1, ..., n_k$ is used for key-frame indexing. We start by obtaining a set of keypoints $\{\mathbf{p}_i\}_{i=1}^{n_i}$ detected on the first frame $I_{f=1}$ and defining an initial set of tracks $\{\tau_j\}_{j=1}^{n_j} = \{\mathbf{p}_i\}_{i=1}^{n_i}$. We also tag the first frame as a key-frame $I_{f=1} \rightarrow I_{k=1}$. Here and in the following steps all the computation and processing is performed on the image where only the visible part of the retina is kept and strong specular reflections have been filtered out. A segmentation mask is produced such that image pixels which do not belong to the retina are assigned to zero (i.e. zero-intensity pixels) and to one otherwise. This is done using thresholding followed by morphological refinement [19]. In the experimental section we assess different types of key-point detectors, SIFT [15], the minimum eigen value algorithm (minEig) [24] and their impact on the performance of the proposed algorithm. We also use a uniform grid of points (UGrid) evenly placed on the area of the visible retina to complement the evaluation.

2.2. Motion Estimation

Inter-frame motion estimation with a simple model as used in [19] seems to be robust but inaccurate, typically up to 5 pixels [18]. We can use this simple global model to create better inter-frame correspondences, and then tracks. The slit-lamp system's optics include several parts moving independently, namely the contact lens and the camera. This complicates the derivation of an accurate, simple and physically valid transformation to relate the images geometrically. We use the affine transformation in our work as a best tradeoff [18]. When the new frame I_f comes we estimate the motion to the last key-frame $A_{f\to k-1}$ by solving a Linear Least Squares (LLS) problem where we minimize the sum of squared transfer discrepancies:

$$\tilde{\theta} = \min_{\theta} \sum_{i=1}^{n_i} \| \mathbf{q}_i - w(\mathbf{p}_i; \theta) \|_2^2$$
(1)

where $\bar{\theta}$ is an estimated (6×1) vector of motion parameters of the last key-frame. The transformation function has the form $w(\mathbf{p}; \theta)$ and \mathbf{p}_i , \mathbf{q}_i are key-point correspondences from the current and the previous frames respectively.

2.3. Prediction

We propagate the existing *query* tracks τ_j using the Kanade-Lucas-Tomasi (KLT) algorithm [21] obtaining *candidate* tracks as:

$$\tau'_{j} = \mathcal{KLT}(\tau_{j}, I_{f-1}, I_{f}) \tag{2}$$

The key-point associated with the *candidate* track is then checked for zero-intensity (i.e. intensity values of all color channels equal to zero). If true it is then rejected as a faulty prediction because the track is considered valid only if it belongs to the visible part of the retina.

We have chosen KLT as it is an appearance based method which uses local search. It is fast and robust just enough to handle changes between consecutive frames. It can cope with a sudden motion better compared to statistical approaches such as the Extended Kalman Filter (EKF) where the redundancy exists in time.

2.4. Track Correction

We proceed with the refinement procedure to correct the position of the predicted *candidates* (figure 2). We first warp the new image using the previously estimated affine transformation as:

$$I_f^{\omega} = \omega(I_f, \mathbf{A}_{f \to k-1}) \tag{3}$$

We perform an exhaustive search in a 5×5 neighborhood w around the *query* tracks locations on the warped image I_f^{ω} to find a possible update $\tilde{\tau}_j$ by minimizing a similarity metric. We search on the warped image because it allows us to find an estimate in a local area which is directly related to the perceived misalignment. We evaluate several metrics in this study, namely the Sum of Squared Distances (SSD), Normalized Cross Correlation (NCC) and Sum of Hamming Distances (SHD). Finally the corrected position

of the predicted tracks locations is computed using the previously estimated motion as:

$$\bar{\tau}_j = \phi(\tilde{\tau}_j, \mathbf{A}_{f \to k-1}) \tag{4}$$

where ϕ is the back-warping function.



Figure 2: Schematic illustration of track prediction and correction on a sample track τ_i .

2.5. Key-frame Instantiation and Local BA

We compute the tracking loss L in the current frame as the percentage of lost tracks from the last key-frame to provide the condition for inclusion of new key-points and then tracks $L = \frac{sizeof(\tau \in I_f)100}{sizeof(\tau \in I_k)}$. This does not indicate re-initialization of the tracking process in case of full occlusion. It rather allows us to assure that sufficiently many points are tracked at all times. Thus, if L > 50%, we detect new key-points τ_f^{new} as in the *Initialization* step. We then filter out those new tracks which fall in the predefined local neighborhood (7 × 7 pixels in our experiments) and join the two sets of tracks. This is done to keep new tracks not too close to the existing ones and avoid populating new tracks with redundant locations. Finally, the current frame is tagged as new key-frame $I_f \rightarrow I_{k+1}$.

We then invoke a local BA-type routine. The idea is to minimize the reprojection error. An unknown 2D point **g** is associated with each track $\tau_{k,j}$ and an affine transform $w(\mathbf{g}; \theta)$ with each key-frame. The presence/absence of a track in a key-frame is given by an indicator variable $v_{k,j} \in \{0, 1\}$. The reprojection error to minimize is:

$$\min_{\mathbf{g}_{j},\theta} \sum_{k=1}^{n_{k}} \sum_{j=1}^{n_{j}} v_{k,j} \parallel \tau_{k,j} - w(\mathbf{g};\theta) \parallel_{2}^{2}$$
(5)

we solve this with matrix factorization in the LLS sense [8]. We repeat from *Motion Estimation* step for the rest of the sequence.

3. Dataset Acquisition

The datasets used for evaluation were obtained from four retinal examination videos (figure 3). The examination was performed on volunteers in the University Hospital of Saint-Étienne, France. The navigated PRP system developed at QuantelMedical was used. The videos were captured with a CCD camera at 60fps. Typical videos are 2-3 minutes long. We took every 5th frame to produce images sequences to simplify the evaluation routine. The imaging set-up implies a fixed camera on the moving base, controlled by the ophthalmologist, which undergoes only translation. Small rotations caused by patient's head tilts occasionally occur. The examination videos are populated with uneven illumination caused by changing exposure settings to adjust to the patient's comfort. The retina is illuminated with a narrow light beam focused using a direct contact lens. The standard way of retinal examination is to perform a closed loop motion starting from the optic nerve, moving to the periphery and coming back. Full occlusion may occur due to a patient's sudden move and/or specular reflections induced by the contact lens. Dealing explicitly with such challenging conditions is out of the scope of this work. Thus, our video samples were chosen in such a way that no full occlusion occurred in a sequence.

4. Evaluation

In retinal imaging it is difficult to evaluate mosaicing methods objectively due to the lack of ground-truth for alignment. A method of generating simulated retinal image sets by modeling geometric distortions and the image acquisition process have been proposed for the case of fundus images [14]. However, in slit-lamp imaging this option is not directly applicable and the adjustment of this technique to our case is out of the scope of this work. Simulation of the imaging process with a virtual camera becomes problematic likewise due to the complexity of the optical set-up. We provide objective quantitative partial performance evaluation of our method in two stages. First, the assessment of the steps of the method which potentially have strong influence on the result evaluated. This is followed by a comparison of the best performing combination to [19].

4.1. Does the Metric Matter?

To assess the impact of the chosen local similarity metric on the precision of the track correction we compare different metrics namely SSD, NCC and SHD. Both SSD and NCC metrics were considered as the popular choice in realtime tracking algorithms and due to the simplicity of the



Figure 3: (Best viewed electronically) Sample images from different slit-lamp datasets. (*a*) - dataset#1, 253 images, (*b*) - dataset#2, 242 images, (*c*) - dataset#3, 169 images, (*d*) - dataset#4, 309 images.

computation. These are correlation based metrics which rely only on the intensity information. The SHD metric on the other hand is well known for its usage in binary feature matching. It is very fast to compute and it captures structural information, which is a favorable feature in case of slit-lamp imaging where illumination variations are often present. We compute the Forward-Backward Consistency (FBC) error. The idea is to track the $\bar{\tau}_{i,j}$ backward continuously performing *Prediction* and *Track Correction* steps. The FBC error is defined as the distance in pixels from the original location of the track to the final location after the backward tracking. We define the acceptance threshold as 3 pixels. Table 1 shows the computed FBC across datasets.

We calculate FBC every time when the correction step is invoked and take an average among all measurements. We show results for different key-point detectors used to initialize the tracks. As one can see, SHD generally provides a lower error among the datasets while SSD comes second and NCC turned out to be the inferior one.



Figure 4: Number of tracks versus frames. Results show performance with UGrid key-points (*red*), minEig (*green*) and SIFT key-points (*blue*) respectively on the experiment *without* track correction.

		dataset#1	dataset#2	dataset#3	dataset#4
UGrid	SSD	3.88	2.71	4.26	3.53
	NCC	4.29	3.41	5.34	5.23
	SHD	2.82	2.70	3.82	2.36
	SSD	3.47	2.64	4.71	3.15
minEig	NCC	4.82	3.70	6.97	5.27
	SHD	2.64	2.05	3.53	2.65
SIFT	SSD	3.56	2.70	4.21	3.95
	NCC	3.62	3.82	5.28	5.18
	SHD	2.81	2.68	3.18	2.89

Table 1: Forward-Backward Consistency for similarity metrics evaluation.

4.2. How Long the Tracks are?

Long-term tracks is a fundamental part of BA-type refinement. Thus, the quality of the method is directly related to the average length of the tracks, the longer the better. We assess the length of the tracks with and without the correction step of our method. To evaluate this we compute the average length of the tracks across different subsets of frames which were established each time a new key-frame was defined. We call it the span, denoted S. We also check the average number of tracks per frame for a given dataset, denoted μ , as it has a heavy impact on the propagation of local alignment errors. Finally, we analyze the number of key-frames instantiated for a given dataset, denoted κ , as an additional indicator of track accuracy, the lower the better. The graph plots given in figure 4 show the number of tracks per frame for tracking without correction using three options to define the key-points. The tracks obtained with UGrid are shown as red curve, minEig is in green and SIFT was used to obtain the tracks which are shown in blue. As one can see, defining the uniform grid of points to initialize the tracks gives higher track/frame rate for the datasets #1 and #2. However, minEig produces more tracks for datasets #3 and #4. One can also see that the second dataset seems to be an easy example due to the the lower amount of spikes presented on the graph. In fact, the spikes on the graph are the events when the new key-frame was instantiated and new tracks were added to the existing ones. Similarly, one can conclude that dataset#4 is the most difficult case for evaluation. This is not only because it has the longest sequence but also because the retina was not prop-

	UGrid				minEig				SIFT			
	μ	κ	S_{mean}	S_{max}	μ	κ	S_{mean}	Smax	μ	κ	S_{mean}	S_{max}
dataset#1	322	19	8	49	128	29	5	39	233	19	9	48
dataset#2	358	19	7	39	106	25	5	37	350	17	9	41
dataset#3	493	30	3	19	1398	30	3	20	372	28	3	18
dataset#4	321	95	1	11	525	87	2	12	205	92	2	11

Table 2: Tracking statistics without track correction. μ - average number of tracks per frame, κ - number of key-frames, S_{mean} - average span, S_{max} - maximum span.

	UGrid			minEig				SIFT				
	μ	κ	S_{mean}	S_{max}	μ	κ	S_{mean}	S_{max}	μ	κ	S_{mean}	S_{max}
dataset#1	308	15	11	52	121	28	6	39	233	16	12	52
dataset#2	359	16	12	44	124	20	7	35	320	12	10	45
dataset#3	485	24	5	26	1351	26	5	24	370	27	4	21
dataset#4	320	90	2	12	550	80	4	15	205	86	3	13

Table 3: Tracking statistics with track correction. μ - average number of tracks per frame, κ - number of key-frames, S_{mean} - average span, S_{max} - maximum span.



Figure 5: Number of tracks versus frames. Results show performance with UGrid key-points (*red*), minEig (*green*) and SIFT key-points (*blue*) respectively on the experiment *with* track correction.

erly illuminated during the examination, thus, not providing sufficient reliable information. The tracking statistics across datasets for this experiment are shown in table 2. One can see that for dataset#1 the maximum span was achieved using UGrid from initialization. However, SIFT shows more consistent tracks for dataset#2. Finally, minEig appears to perform better on datasets #3 and #4. The average number of tracks per frame follows a similar behavior result-

ing in more tracks for datasets #1 and #2 with a uniform grid while for datasets #3 and #4 more tracks are given by minEig. Overall it can be concluded that initializing with UGrid seems to be a tradeoff when we do not incorporate track correction. What happens once the correction step is included in the method? The results of this settings are given in figure 5. The graphs demonstrate that the number of tracks per frame slightly increased for all the datasets. This is supported by the statistics provided in table 3. Indeed, using the result of the evaluation of the similarity metrics, namely SHD, we obtain improvement for all the numbers and on the span and number of key-frames especially. This indicates that the track correction step using a simple global model with local neighborhood based adjustment is an efficient way to obtain longer tracks with better precision.

4.3. Are We Reducing Drift?

As resulted from the previous experiments, tracking initialized with uniform grid and the SHD based track correction scheme provides long, consistent tracks. Now this gives us a solid base for BA initialization. Thus, in this section we evaluate the proposed method with its best performing settings. We compare the method implemented with and without local BA to the baseline method [19]. We use a Loop Closure Error (LCE) metric [18]. This shows how the composition of estimated transformations affects the global registration and accumulated drift. The idea is to initialize a uniform grid of points $g_1, ..., g_{n_l}$ at the first frame of the sequence and use the set of pairwise estimated transformations applied sequentially to transfer the grid through the sequence. The metric computes the discrepancy between the initial and resulting sets of points as:

$$\xi_{LCE} = \sqrt{\frac{1}{l} \sum_{i=1}^{n_l} \|g_i - \zeta\|_2^2}$$
(6)

where $\zeta = w(...(w(g, \theta_{1,2}))..., \theta_{i,1})$. The comparison of our method and [19] are shown in table 4. Results show that the proposed method outperform the baseline method. A significant improvement can be observed on the version of the proposed method where the local BA step was used.

	dataset#1	dataset#2	dataset#3	dataset#4
proposed (1)	30.43	21.75	48.02	49.12
proposed (2)	11.36	5.48	32.16	38.56
baseline [19]	34.18	28.64	48.15	50.72

Table 4: LCE computed across datasets. The proposed (1) is the proposed method with UGrid used for tracks initialization and SHD based local correction step. The proposed (2) is the proposed (1) + local BA.



Figure 6: (Best viewed electronically) Examples of improved areas of the mosaics given in figure 1 with corrected drift using proposed approach. First column - originals, second column - corrected versions.

Improvement on the mosaics is shown in figure 6. The vessel misalignment initially present in 6a (which corresponds to the first mosaic given in figure 1) was corrected and the vessel remains continuous. The blurred vessel in figure 6c and duplicated one from figure 6e (which corresponds to the second mosaic given in figure 1) were also corrected and visual quality has been improved as it is shown in figures 6d and 6f respectively.

5. Conclusion and Future Work

We have presented a method for drift reduction in mosaicing slit-lamp retinal video sequences. We have validated it using a simple global motion model that can efficiently produce long-term tracks with better precision for long video sequences. We also demonstrated that using a grid of points distributed uniformly over the visible part of the retina generally provides a better initialization for tracking. We have proposed a new local refinement procedure which can be applied not only for mosaicing slit-lamp images but also within the scope of other applications such as object tracking in the non-medical domain. The algorithm is planned to be extended and integrated into the interventional planning within the slit-lamp prototype developed at QuantelMedical.

References

- [1] K. M. Adal, R. M. Ensing, R. Couvert, P. van Etten, J. P. Martinez, K. A. Vermeer, and L. van Vliet. A hierarchical coarse-to-fine approach for fundus image registration. In *Biomedical Image Registration*, pages 93–102. 2014. 1
- [2] J. Asmuth, B. Madjarov, P. Sajda, and J. W. Berger. Mosaicking and enhancement of slit lamp biomicroscopic fundus images. *British journal of ophthalmology*, 85(5):563– 565, 2001. 1
- [3] M. Brown and D. G. Lowe. Automatic panoramic image stitching using invariant features. *International journal of computer vision*, 74(1):59–73, 2007. 1
- [4] A. Can, C. V. Stewart, B. Roysam, and H. L. Tanenbaum. A feature-based technique for joint, linear estimation of high-order image-to-mosaic transformations: mosaicing the curved human retina. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(3):412–419, 2002. 1
- [5] J. Chhablani, A. Mathai, P. Rani, V. Gupta, J. F. Arevalo, and I. Kozak. Comparison of Conventional Pattern and Novel Navigated Panretinal Photocoagulation in Proliferative Diabetic RetinopathyComparison of PASCAL and NAVILAS for PRP. *Investigative ophthalmology & visual science*, 55(6):3432–3438, 2014. 1
- [6] J. Civera, A. J. Davison, J. A. Magallón, and J. Montiel. Drift-free real-time sequential mosaicing. *International Journal of Computer Vision*, 81(2):128–137, 2009.
- [7] J. Hannuksela, P. Sangi, J. Heikkilä, X. Liu, and D. Doermann. Document image mosaicing with mobile phones. In *Image Analysis and Processing*, 2007. ICIAP 2007. 14th International Conference on, pages 575–582, 2007. 1
- [8] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. 2003. 4
- [9] M. Kernt, R. Cheuteu, E. Vounotrypidis, C. Haritoglou, A. Kampik, M. W. Ulbig, and A. S. Neubauer. Focal and panretinal photocoagulation with a navigated laser (NAVI-LAS). *Acta ophthalmologica*, 89(8):662–664, 2011. 1
- [10] M. Kernt, R. E. Cheuteu, S. Cserhati, F. Seidensticker, R. G. Liegl, J. Lang, C. Haritoglou, A. Kampik, M. W. Ulbig, and A. S. Neubauer. Pain and accuracy of focal laser treatment for diabetic macular edema using a retinal navigated laser (NAVILAS). *Clin Ophthalmol*, 6(1):289–296. 1
- [11] G. Klein and D. Murray. Parallel tracking and mapping for small AR workspaces. In *Mixed and Augmented Reality*, 2007. ISMAR 2007. 6th IEEE and ACM International Symposium on, pages 225–234, 2007. 1
- [12] K. Konolige and M. Agrawal. FrameSLAM: From bundle adjustment to real-time visual mapping. *Robotics, IEEE Transactions on*, 24(5):1066–1077, 2008. 1
- [13] F. Laliberté, L. Gagnon, and Y. Sheng. Registration and fusion of retinal images-an evaluation study. *Medical Imaging*, *IEEE Transactions on*, 22(5):661–673, 2003.
- [14] S. Lee, M. D. Abràmoff, and J. M. Reinhardt. Validation of retinal image registration algorithms by a projective imaging distortion model. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pages 6471–6474, 2007. 4

- [15] D. G. Lowe. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 2
- [16] P. F. McLauchlan and A. Jaenicke. Image mosaicing using sequential bundle adjustment. *Image and Vision computing*, 20(9):751–759, 2002. 1
- [17] E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd. Generic and real-time structure from motion using local bundle adjustment. *Image and Vision Computing*, 27(8):1178–1193, 2009. 1
- [18] K. Prokopetc and A. Bartoli. A comparative study of transformation models for the sequential mosaicing of long retinal sequences of slit-lamp images obtained in a closed loop motion. *Computer Assisted Radiology and Surgery, CARS 2016*, 2016. 2, 3, 7
- [19] R. Richa, R. Linhares, E. Comunello, A. Von Wangenheim, J.-Y. Schnitzler, B. Wassmer, C. Guillemot, G. Thuret, P. Gain, G. Hager, et al. Fundus Image Mosaicking for Information Augmentation in Computer-Assisted Slit-Lamp Imaging. *Medical Imaging, IEEE Transactions on*, 33(6):1304–1312, 2014. 1, 2, 4, 7
- [20] R. Richa, B. Vágvölgyi, M. Balicki, G. Hager, and R. H. Taylor. Hybrid tracking and mosaicking for information augmentation in retinal surgery. In *Medical Image Computing* and Computer-Assisted Intervention–MICCAI 2012, pages 397–404. 2012. 1
- [21] J. Shi and C. Tomasi. Good features to track. In Computer Vision and Pattern Recognition, 1994. Proceedings CVPR'94., 1994 IEEE Computer Society Conference on, pages 593– 600, 1994. 3
- [22] M. Souza, R. Richa, A. Puel, J. Caetano, E. Comunello, and A. Von Wangenheim. Robust Visual Tracking for Retinal Mapping in Computer-Assisted Slit-Lamp Imaging. In *Computer-Based Medical Systems (CBMS), 2014 IEEE 27th International Symposium on*, pages 132–137, 2014. 1
- [23] C. V. Stewart, C.-L. Tsai, and B. Roysam. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. *Medical Imaging, IEEE Transactions on*, 22(11):1379–1394, 2003. 1
- [24] C. Tomasi and T. Kanade. Detection and tracking of point features. 1991. 2
- [25] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustmenta modern synthesis. In *Vision algorithms: theory and practice*, pages 298–372. 1999. 1
- [26] G. Yang and C. V. Stewart. Covariance-driven mosaic formation from sparsely-overlapping image sets with application to retinal image mosaicing. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 1, pages I– 804, 2004. 1
- [27] Y. Zheng, E. Daniel, A. A. Hunter, R. Xiao, J. Gao, H. Li, M. G. Maguire, D. H. Brainard, and J. C. Gee. Landmark matching based retinal image alignment by enforcing sparsity in correspondence matrix. *Medical image analysis*, 18(6):903–913, 2014. 1