ISSN 2186-7437

NII Shonan Meeting Report

No. 140

Optimisation Methods in Geometric Vision

Tat-Jun Chin
Anders Eriksson
Yasuyuki Matsushita

January 27–31, 2019

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Optimisation Methods in Geometric Vision

Organizers:
Tat-Jun Chin (The University of Adelaide)
Anders Eriksson (The University of Queensland)
Yasuyuki Matsushita (Osaka University)

January 27–31, 2019

1 Meeting information

1.1 Technical abstract

Computer vision is concerned with inferring properties of the world from observations in the form of visual images. Such inverse problems typically take shape as optimization problems, that aim to find the best explanation, for the complex visual phenomenon that gave rise to a set of noisy and incomplete visual measurements. For computer vision applications to be successful, the underlying optimization problems must be supported by efficient and dependable solution methods.

The meeting focuses on a broad subclass of computer vision problems called geometric vision problems. Roughly, these are problems that exploit fundamental geometrical constraints arising from the image formation process or physical properties of the scene (e.g., lighting conditions, characteristics of motions), to extract information of the scene (e.g., depth, 3D shape, camera trajectory, object identities) from the given visual data. Example geometric vision problems include structure-from-motion (SfM), simultaneous localization and mapping (SLAM), pose averaging [15], photometric stereo [23], and motion segmentation [8]. Methods for solving geometric vision problems underpin many useful applications, such as 3D reconstruction, robot navigation, object recognition/tracking, and computational photography [21].

Geometric vision is replete with hard optimization problems. By "hard", we mean that the time needed to solve the optimization problems grows quickly with the size of the input data. Take, for example, the task of robustly estimating the planar perspective transformation (a.k.a. homography) from outlier-contaminated point correspondences between two images. Due to the inherent intractability of robust homography estimation [7], practitioners often rely on simple randomized heuristics to find rough approximate solutions, which neither guarantee optimality nor provide bounds on the approximation error.

The computational difficulty of geometric vision problems is also often compounded by the extremely large size of the input. Take, for example, the task of bundle adjustment, i.e., calculate 3D points and camera poses that are consistent with a set of images of a scene. In the age of big data, the input image set is often obtained by "scraping" Internet photo collections, or by conducting long-term surveillance of a scene using a robot. Such input sizes easily overwhelm traditional computing architectures, and distributed or parallel versions of bundle adjustment must be used [11].

1.2 Technical themes

The overall theme for the proposed meeting is recent theoretical and algorithmic advances on optimization problems in geometric vision. More specific themes include:

- Solvability and approximability of geometric vision problems, e.g., [7, 31];
- Duality in geometric vision problems, e.g., [28, 5, 12];
- Global optimization algorithms for geometric vision, e.g., [7, 12];
- Approximate algorithms including randomized methods, e.g., [22, 20];
- Distributed algorithms for geometric vision problems, e.g., [11];
- Incremental algorithms for online geometric optimization, e.g., [31, 20].

Apart from discussing recent progress in geometric optimization through the above themes, we also aim to chart future research directions and novel application areas. For example:

- The role of machine learning in geometric optimization;
- Geometric optimization on constrained computing platforms (e.g., smartphones, sensor networks);
- Geometric optimization for novel imaging devices (e.g., RGBD cameras, light field cameras); and
- Geometric vision problems from new industries (e.g., self-driving cars, UAVs).

Following the spirit of Shonan Meetings, we will also consider other related topics, based on the interest of the attendees and the trajectory of the discussions.

1.3 List of participants and program

Please see Appendix A (page 11) and Appendix B (page 12). See also the meeting website [1] for more information.

2 Meeting report

The report is structured around two main themes that were discussed in the meeting (Section 2.1), and an overview of several talks (Section 2.2). In Section 3, a summary of the lessons learned on the organisation of the meeting for the computer vision community will be presented.

2.1 Main themes of discussion

The main themes received the most attention during the meeting, though bearing in mind that the list of topics presented and discussed is more diverse than these main discussion points (again, an overview of other topics will be given in Section 2.2).

Main theme 1: More meaningful interactions between machine learning and geometry

In line with the broader trends in computer vision, a persistent theme in the meeting is the usage of machine learning, particularly deep neural networks (or "deep learning"), to perform tasks that are typically done using geometric techniques prior to the resurgence of neural networks. These tasks include object pose estimation [18, 13], place and scene recognition [4, 9], 3D reconstruction [24], and simultaneous localisation and mapping (SLAM) [3].

The investigation of deep learning towards the above problems is justified, due to the significant improvements in performance it has enabled in problems such as image segmentation and object recognition [19]. However, for the geometric problems above, it is unclear under what conditions are deep learning approaches able to provide significantly better accuracy and performance than techniques that model the intrinsic geometry of the problems. In fact, some of the attendees remarked that it is "quite difficult" to develop a deep learning method that can provide highly accurate object pose estimates.

On the other hand, "pure" geometric techniques are brittle, especially when presented with noisy visual inputs from unseen before environments. To enable geometric techniques to generalise well to different operating conditions, usually a manual tuning process is required to reselect the algorithm parameters.

During the meeting, Richard Hartley proposed to combine deep learning and geometric techniques in a way that leverages the intrinsic strengths of both approaches: ability to exploit inherent regularity in scenes by training deep networks with large amounts of data, and the usage of mathematically justified models for scene understanding by exploiting the scene geometry.

A potentially fruitful framework along the above lines is to backpropagate geometric errors into the deep network; the general concept is illustrated in the following diagram.

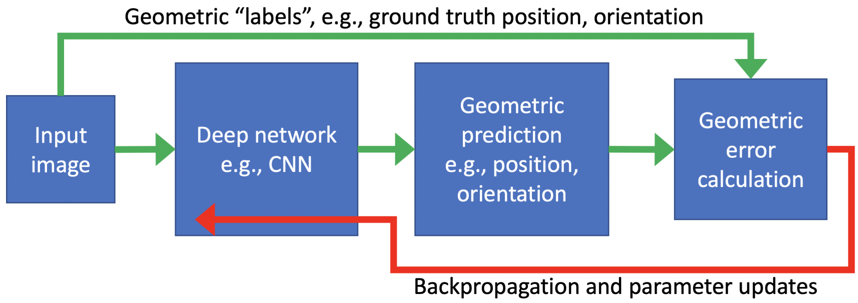

For example, in the problem of object pose estimation, the deep network can be used to predict the corners of the 3D bounding box of the target object in the image, which can then be used to analytically compute the object pose. If a ground truth pose is available, the pose difference can be used to calculate a residual which is then backpropagated to adjust the network parameters. Backpropagating "non-standard" error functions has been done, for e.g., in [14].

Main theme 2: The role of strong duality in computer vision

Computer vision concerns itself with understanding the real world through the analysis of images. The problems that this analysis give rise to often lead to highly complicated optimization problems with a mixture of challenging objective functions and constraints, involving high- or even infinite-dimensional variables in terms of curves and surfaces. Due to the non-convexity of the overwhelming majority of these problems they are typically exceedingly difficult to solve globally.

A prominent theme in this meeting was on the role of duality principles in the field of Computer Vision and related areas. The mathematical concept of duality lies at the core of a great part of the most efficient optimization algorithm currently in use. A key reason behind this is the notion of strong duality, a property in optimization theory establishing the equivalence of a given minimization problem to that of an associated and convex problem, known as the dual problem.

There is a well established theory on Lagrangian duality and the existence of strong duality with respect to convex optimization problems. Unfortunately, much less is currently understood regarding the role of strong duality in non-convex optimization, hence significant portions of the existing theory can not be directly applied to this setting. There are however a few notable exceptions to this rule, problems that are non-convex but for which strong duality does indeed hold. These are problems that provably admit a large number of local minima but can still be solved in a way that guarantees a global optima by invoking duality principles.

In recent years, empirical observations have been made suggesting that, under specific conditions, a broad class of non-convex problems, fundamental to Computer Vision, also belong to this group of exceptions. The consequences of establishing such a proposition would be far-reaching, both in theory and in practice. Such a result would imply that the global solution to these primal non-convex problems are equivalent to the solution of the associated dual convex problem. This allows us to solve exceedingly difficult non-convex problems, with numerous local minima, implicitly through their dual formulations with a guarantee of global optimality and in many instances also in polynomial time.

The discussion were led primarily by Fredrik Kahl, Frank Dellaert, Luca Carlone, Anders Eriksson and Robert Mahony. The topics were mainly centered around problems involving pose, with a particular focus on two specific applications, namely rotation Averaging and SLAM, see Figure 1, both fundamental problems in Computer Vision and Robotics.

Two primary goals of this particular theme and line of research were identified and discussed at the Shonan meeting. They were, firstly to establish new mathematical foundations detailing the role of strong duality in Computer Vision and to further the understanding of the function of strong duality in a broad class of related optimization problems. Secondly, to develop efficient numerical algorithms dedicated to solving the resulting dual optimization problems.

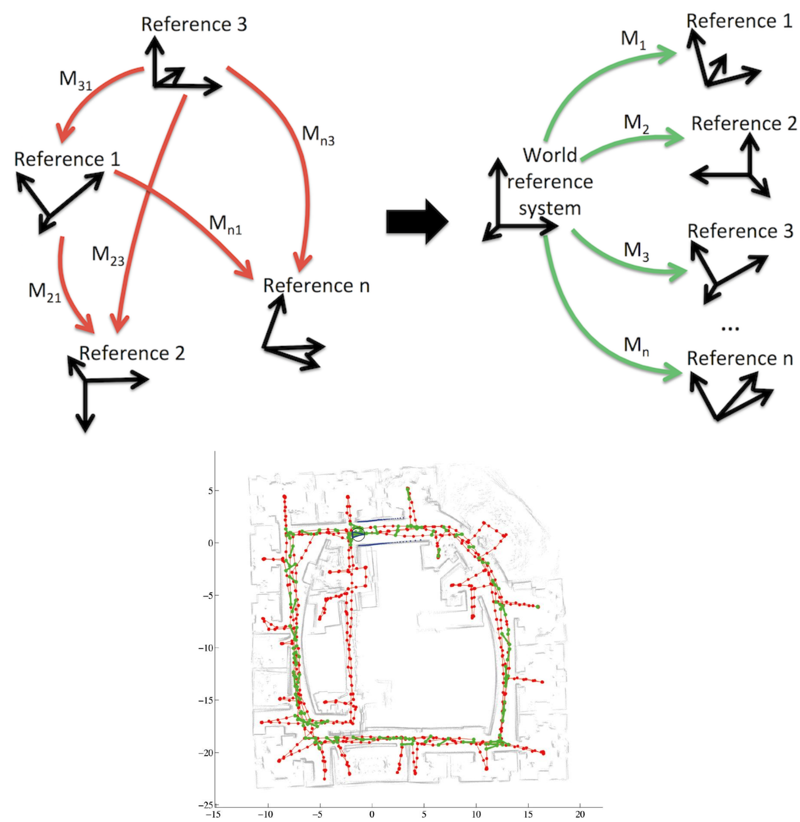

Figure 1: Rotation averaging and SLAM, two fundamental problems related to 3D reconstruction, analysis and navigation.

2.2 Overview of selected talks

Discrete-continuous graphical models for robust perception

Assistant Professor Luca Carlone, MIT

Pose graph optimisation (PGO) lies at the core of a robotic perception module. Specifically, PGO is required to enable a robot to navigate in an unknown environment by iteratively reconciling a sequence of visual motion observations and keeping track of the trajectory of the robot. Computationally, PGO is inherently an intractable problem due to the manifold structure of the variables of interest (robot positions and orientations). Moreover, in situations where long-term operations are essential, there is a large number of measurements to be processed by the PGO algorithm.

This talk surveys a class of PGO algorithms that apply convex relaxation and solve the problem using convex solvers, specifically Semi Definite Programming (SDP) [25]. It has been established that, there is a nonzero bound, such that if the noise the measurements are no greater than the said bound, then the convex relaxation is tight. Moreover, good empirical performances of convex relaxation methods have been demonstrated. Notably, the SDP technique is faster than standard iterative optimisation algorithms [25].

The above theoretical results are based on non-robust formulations of PGO. However, in real-life settings, outliers (wrong measurements) are inevitable. If the convex relaxation methods are directly applied on such measurements, the solutions will be biased. The talk proposes a discrete continuous (DC) framework for PGO, by introducing binary "label" variables into the problem that enables the assignment of measurements as inliers or outliers. In more detail, if a measurement is identified as an outlier, its effect on the estimates are discounted. An interesting connection between the DC-PGO problem and Markov Random Fields (MRF) was also pointed out. It is likely that algorithmic developments for the DC-PGO will become a fruitful research direction.

Parametrised complexity in geometric optimisation

Associate Professor Tat-Jun Chin, The University of Adelaide

A number of important and useful geometric problems in computer vision (e.g., robust estimation, point set registration) have been proven to be computationally intractable, not only to solve, but also to approximate [6, 27]. On the other hand, practical applications in real-world settings require these problems to be solved with some form of guarantee so as to avoid unexpected breakdowns. Hence, it is not sufficient to rely on heuristic methods to solve these problems, since the quality and/or performance of heuristic methods cannot be bounded.

This talk proposed to investigate the *structural properties* of the hard geometric problems and develop fixed parameter tractable (FPT) algorithms. Briefly, FPT algorithms exploit additional insights or structures to a problem, and constrain the exponential growth in runtime to parameters that depend on the special structures only. Such an endeavour is encapsulated under the field of parametrised complexity analysis [10].

The talk illustrates an example based on inlier set maximisation or consensus maximisation [6], which is a very common problem in computer vision. It is shown how consensus maximisation has a FPT formulation, if an upper bound on the number of outliers in a problem instance is known. It will be interesting to apply the FPT framework on other intractable problems in computer vision (e.g., point set registration [27]).

Methods for robustified nonlinear least-squares

Professor Christopher Zach, Chalmers University

Given noisy observations of a number of scene points in a number of images, the goal of bundle adjustment (BA) is to estimate the camera poses and 3D coordinates of the scene points that are consistent with the observations. This is typically casted in a least squares objective function, where the sum of squared reprojection errors is minimised. To estimate the variables of interest, non-linear optimisation techniques such as Levenberg-Marquardt are applied.

If the measurements contain outliers, however, the least squares solution will be biased. A more robust norm such as Tukey's Biweight or Geman-McClure must thus be used in the objective function. If the robust norm is smooth, the resulting robustified BA problem can still be solved directly using standard iterative optimisation algorithms. It has been shown, however, that more effective techniques exist, namely, iteratively reweighted least squares (IRLS) [16] and lifted optimisation [2]. Moreover, it can be established that both IRLS and lifted optimisation minimise the same "half-quadratic" objective function, whereby a set of weights (one per measurement) is introduced into the original robustified BA problem [30]. By explicitly introducing weight variables, the half-quadratic objective function suffers from far fewer bad local minima.

This talk proposed new ways of tackling the half-quadratic objective function that can potentially deliver better performance than IRLS and lifted optimisation. The first is an "iterated lifting" technique, where a graduated or annealed optimisation approach is introduced in half-quadratic optimisation. The second is a multi-objective optimisation (MOO) technique, where multiple versions of the original robustified BA problem, each differing from the other by the half-quadratic kernel, are jointly optimised using MOO. Initial empirical results show that iterated lifting and MOO optimisation are able to converge to better solutions than IRLS and regular lifted optimisation.

Semi-calibrated photometric stereo

Professor Yasuyuki Matsushita, Osaka University

High-fidelity shape estimation is a central topic in computer vision. One of the promising approaches is *photometric stereo* that uses *photometric* information to determine the 3D shape. More specifically, photometric stereo estimates surface normal of a static scene in a pixel-wise manner from a set of observations obtained under varying light conditions from a fixed camera. It has been understood that, for a Lambertian surface, its surface normal map can be uniquely determined from three observations.

In practice, the intensity observations may contain non-Lambertian components, such as specular reflectance and cast/attached shadows, that can be regarded as outliers. To deal with the outliers, robust estimation techniques are employed in the past; for example, ℓ_1-norm minimisation and robust principal component analysis in the context of photometric stereo [29, 17]. The robustness against outliers has been improved in the past decade; however, *calibration* of light sources (both geometric and photometric calibration) remains a practical issue due to its laborious process.

This talk introduced a semi-calibrated photometric stereo method, with which the need for photometric calibration of light sources can be eliminated. It has been believed that for a Lambertian photometric stereo, it is needed to know the light source directions and light source intensities. However, it is shown that the knowledge of light source intensities is unnecessary, but even under the condition, a unique shape can be determined. The talk showed that there exists a linear solution technique to the problem, and it introduced an efficient alternating minimisation strategy to the problem.

Geometric point light source calibration via structure from motion

Dr. Michael Waechter, Osaka University

Estimating the position or direction of a light source accurately is essential for many computer vision tasks, such as shape from shading, photometric stereo, or reflectance and material estimation. In these tasks, inaccurate light positions immediately cause errors in their estimates. Despite the importance of accurate light calibration, it remains laborious as researchers have not yet come up with accurate and easy to use techniques.

Previously, many geometric point light source calibration procedures involved the use of mirror spheres. Unfortunately, the mirror sphere based methods suffer from the difficulty of determining the sphere boundary and the point of specular spike location, resulting in unstable accuracy of the calibration.

This talk introduced a new geometric point light source calibration method [26] that uses a plane and pins that are stuck at unknown locations on the plane. By observing the cast shadows generated by the pinheads, the problem of geometric point light source calibration can be solved in a similar manner to Structure-from-Motion (SfM). Although the pinhead locations are unknown, they can be simultaneously recovered together with the light source positions/directions. It showed the new application domain of geometric computer vision.

3 Lessons for potential future meetings

As a conclusion to this report, we would like to report some of the lessons we have learned as organisers of this Shonan Meeting.

First, it is worth highlighting that this was likely to be the first Shonan Meeting that was organised and attended exclusively by members of the computer vision community[1]. Therefore, there was not a significant body of experience to draw from in the organisation and planning of the meeting.

There are a few aspects of the program (Appendix B) that could be improved. Chiefly, the adopted program was inspired by typical computer vision meetings, where attendees are encouraged give a presentation (oral or poster) to describe his/her own current favourite research topic. While the carefully selected list of invitees ensured strong commonalities in the presentations and discussions (e.g., Lagrangian duality and global optimality in geometric optimisation, interaction between geometry and deep learning), there was not a strong coalescence around a few clear-cut research questions/themes on which unified discussions and progress could be made.

In future instances of Shonan Meetings (or other meeting series in the style of Dagstuhl) by and for the computer vision community, we suggest to

- Identify a handful of “Area Chairs” for key topics, perhaps even during the meeting proposal stage. The role of an Area Chair could include the selection of presenters and the contents of the presentations, and leading the discussions during the breakout sessions.

[1]These are the researchers who regard conferences such as Computer Vision and Pattern Recognition (CVPR) and journals such as IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) to be their flagship publication venues.

- Deemphasise the CVPR tradition to give each participant a speaking slot. Rather, encourage the participants to contribute productively and passionately to the identified research themes via discussions and breakout sessions.

Another aspect for improvement is gender diversity in the list of participants. While the invitation list included female researchers, the gender of the actual participants in the meeting was heavily biased towards males.

Despite the areas of improvements above, there are also very encouraging signs from the meeting. Chiefly, almost all of the attendees agree that the Shonan Meeting series represents an excellent alternative to the traditional conferences and workshops targeted by the computer vision community. With the explositive growth of commercial interest and investment in AI, the research directions in computer vision are often driven by immediate industry demands. In this context, events in the format of Shonan Meetings offer a welcomed opportunity to put longer-term scientific agendas back in the focus of researchers.

On the other hand, we believe that Shonan Meetings could benefit from greater participation from the computer vision community, which has been successful in recent years to attract the next generation of computing researchers, as well as in identifying commercially impactful problems and applications.

References

[1] `https://cs.adelaide.edu.au/~tjchin/shonan140/`.

[2] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *International Journal of Computer Vision*, 19(1):57–91, 1996.

[3] M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison. CodeSLAM - learning a compact, optimisable representation for dense visual SLAM. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2018.

[4] S. Brahmbhatt, J. Gu, K. Kim, J. Hays, and J. Kautz. Geometry-aware learning of maps for camera localization. In *CVPR*, 2018.

[5] L. Carlone, G. Calafiore, C. Tommolillo, and F. Dellaert. Planar pose graph optimization: duality, optimal solutions, and verification. *IEEE Trans. on Robotics*, 2016.

[6] T.-J. Chin, Z. Cai, and F. Neumann. Robust fitting in computer vision: easy or hard? In *European Conference on Computer Vision (ECCV)*, 2018.

[7] T.-J. Chin and D. Suter. *The maximum consensus problem: recent algorithmic advances*. Synthesis Lectures on Computer Vision. Morgan & Claypool Publishers, 2017.

[8] T.-J. Chin, H. Wang, and D. Suter. The ordered residual kernel for robust motion subspace clustering. In *Advances on Neural Information Processing Systems (NIPS)*, 2009.

[9] R. Clark, S. Wang, A. Markham, N. Trigoni, and H. Wen. VidLoc: A deep spatio-temporal model for 6-DoF video-clip relocalization. In *CVPR*, 2017.

[10] R. G. Downey and M. R. Fellows. *Paramterized complexity*. Springer, 1999.

[11] A. Eriksson, J. Bastian, T.-J. Chin, and M. Isaksson. A consensus-based framework for distributed bundle adjustment. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[12] A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin. Rotation averaging and strong duality. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[13] M. Feng, S. Hu, M. H. Ang, and G. H. Lee. 2D3D-MatchNet: learning to match keypoints across 2D image and 3D point cloud. In *International Conference on Robotics and Automation (ICRA)*, 2019.

[14] B. Fernando and S. Gould. Discriminatively learned hierarchical rank pooling networks. *International Journal of Computer Vision*, 124(3), 2017.

[15] R. Hartley, K. Aftab, and J. Trumpf. L1 rotation averaging using the Weiszfeld algorithm. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.

[16] P. J. Huber. *Robust statistics*. John Wiley and Sons, 1981.

[17] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2012.

[18] A. Kendall, M. Grimes, and R. Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *CVPR*, 2015.

[19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.

[20] H. M. Le, T.-J. Chin, and D. Suter. An exact penalty method for locally convergent maximum consensus. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[21] Y. Matsushita, E. Ofek, W. Ge, X. Tang, and H.-Y. Shum. Full-frame video stabilization with motion inpainting. In *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2006.

[22] T.-H. Oh, Y. Matsushita, Y.-W. Tai, and I. S. Kweon. Fast randomized singular value thresholding for nuclear norm minimization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[23] J. Park, S. N. Sinha, Y. Matsushita, Y.-W. Tai, and I. S. Kweon. Robust multiview photometric stereo using planar mesh parameterization. In *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2016.

[24] R. Ranftl and V. Koltun. Deep fundamental matrix estimation. In *European Conference on Computer Vision (ECCV)*, 2018.

[25] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard. SE-Sync: a certifiably correct algorithm for synchronization over the special Euclidean group. In *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[26] H. Santo, M. Waechter, M. Samejima, Y. Sugano, and Y. Matsushita. Light structure from pin motion: Simple and accurate point light calibration for physics-based modeling. In *European Conference on Computer Vision (ECCV)*, volume 3, pages 3–19.

[27] V. Tzoumas, P. Antonante, and L. Carlone. Outlier-Robust Spatial Perception: Hardness, General-Purpose Algorithms, and Guarantees. *arXiv e-prints*, page arXiv:1903.11683, Mar 2019.

[28] K. Wilson, D. Bindel, and N. Snavely. When is rotations averaging hard? In *European Conference on Computer Vision (ECCV)*, 2016.

[29] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Asian Conference on Computer Vision (ACCV)*, 2010.

[30] C. Zach and G. Bourmaud. Descending, lifting or smoothing: secrets of robust cost optimization. In *European Conference on Computer Vision (ECCV)*, 2018.

[31] Q. Zhang and T.-J. Chin. Coresets for triangulation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2018.

Appendix A — List of participants

1. Alessio Del Bue, Italian Institute of Technology
2. Ali Bab-Hadiashar, RMIT University
3. Anders Eriksson, Queensland University of Technology
4. Carl Olsson, Chalmers Institute of Technology
5. Christopher Zach, Chalmers University of Technology
6. Danda Pani Paudel, ETH Zurich
7. David Suter, Edith Cowan University
8. Florian Bernard, Max-Planck-Institute for Informatics
9. Frank Dellaert, Georgia Tech
10. Fredrik Kahl, Chalmers University of Technology
11. Gim Hee Lee, National University of Singapore
12. Hongdong Li, Australian National University
13. Jamie Sherrah, Digital Animal Interactive Inc. (FTSY)
14. Jesus Briales, Facebook
15. Kenichi Kanatani, Okayama University
16. Laurent Kneip, Shanghai Tech
17. Luca Carlone, MIT
18. Michael Brown, York University
19. Michael Waechter, Osaka University
20. Ping Tan Simon, Fraser University
21. Richard Hartley, Australian National University
22. Robert Mahony, Australian National University
23. Simon Lucey, Carnegie Mellon University
24. Sudipta Sinha, Microsoft Research
25. Tarek Hamel, Universite Cote d'Azur
26. Tat-Jun Chin, The University of Adelaide
27. Viorela Ila, University of Sydney
28. Yasuyuki Matsushita, Osaka university
29. Yinqiang Zheng, NII
30. Yongduek Seo, Sogang University

Appendix B — Meeting program

Sunday 27 January 2019

1500-1900	Check-in
1900-2100	Welcome banquet

Monday 28 January

0730-0900	Breakfast
0900-0915	Welcoming address - Tat-Jun Chin, Anders Eriksson, Yasuyuki Matsushita
0915-1000	Keynote 1 - Fredrik Kahl
1000-1015	Michael Brown
1015-1030	Christopher Zach
1030-1100	Coffee break
1100-1115	Robert Mahony
1115-1130	Yongduek Seo
1130-1200	Breakout session 1
1200-1330	Lunch
1330-1345	Sudipta Sinha
1345-1400	Yinqiang Zheng
1400-1415	Michael Waechter
1415-1430	Viorela Ila
1430-1500	Breakout session 2
1500-1530	Coffee break
1530-1545	Hongdong Li
1545-1600	David Suter
1600-1630	Breakout session 3
1630-1800	Free time
1800-1930	Dinner

Tuesday 29 January

0730-0900	Breakfast
0900-0915	Program briefing
0915-1000	Keynote 2 - Frank Dellaert
1000-1015	Ping Tan
1015-1030	Jamie Sherrah
1030-1100	Cofee break
1100-1115	Ali Bab-Hadiashar
1115-1130	Gim Hee Lee
1130-1200	Breakout session 4
1200-1315	Lunch
1315-1330	Group photo

1330-1345 Kenichi Kanatani
1345-1400 Jesus Briales
1400-1415 Carl Olsson
1415-1445 Breakout session 5
1445-1515 Coffee break
1515-1530 Florian Bernard
1530-1545 Alessio Del Bue
1545-1600 Danda Pani Paudel
1600-1630 Breakout session 6
1630-1800 Free time
1800-1930 Dinner

Wednesday 30 January

0730-0900 Breakfast
0900-0915 Program briefing
0915-1000 Keynote 3 - Richard Hartley
1000-1015 Tarek Hamel
1015-1030 Laurent Kneip
1030-1100 Coffee break
1100-1115 Simon Lucey
1115-1130 Luca Carlone
1130-1200 Breakout session 7
1200-1330 Lunch
1300-2045 Excursion and dinner

Thursday 31 January

0730-0930 Breakfast and check-out
0930-0945 Yasuyuki Matsushita
0945-1000 Anders Eriksson
1000-1015 Tat-Jun Chin
1015-1100 Coffee break
1100-1130 Breakout session 8
1130-1200 Conclusion and wrap-up
1200-1330 Lunch and end of meeting