StOCaMo: Online Calibration Monitoring for Stereo Cameras

Jaroslav Moravec

The Department of Cybernetics, Faculty of Electrical Engineering
Czech Technical University in Prague

June 26, 2023

Problem Introduction

Sensor calibration

- Sensors are not colocated and have their own internal parameters

  $\Rightarrow$ We need to know these for proper sensor fusion

- Calibration room or infrastructure-based

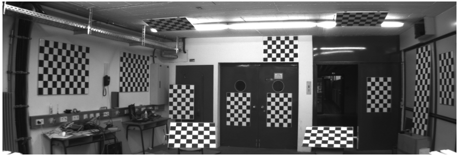

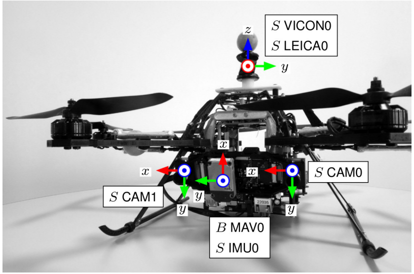

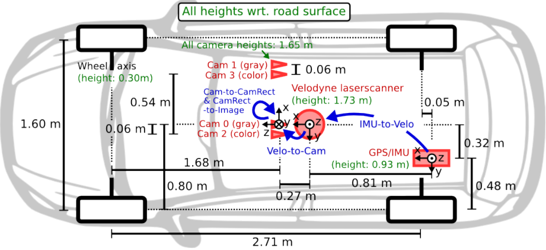

Problem Introduction

Calibration monitoring

- The calibration is essential for all subsequent parts of autonomous operation
- But it is not stable due to vehicle twisting or thermal dilations

$\Rightarrow$ The calibration monitoring could be neccessary

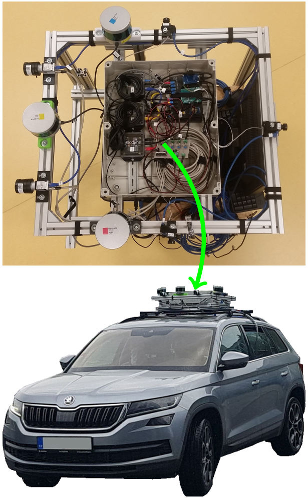

Problem Introduction

Taxonomy of Calibration

- Off-line calibration

  Batch of data B → Construct $f(\theta|B)$ → $\text{argmin}_\theta f(\theta|B)$

  - time-consuming
  - large computational overhead
  - high precision
- On-line calibration

  Small batch of data B_0 → Construct $f(\theta|B_0)$ → $\text{argmin}_\theta \mathbb{E}[f(\theta|B_0)]$

  - random $f(\theta|B_0) \rightarrow$ large variance
  - $\mathbb{E}[f(\theta|B_0)]$ provides higher precision
  - fast response

Problem Introduction

Taxonomy of Calibration

- Off-line calibration

  Batch of data B → Construct $f(\theta|B)$ → $\text{argmin}_\theta f(\theta|B)$

  - time-consuming
  - large computational overhead
  - high precision
- On-line calibration

  Small batch of data B_0 → Construct $f(\theta|B_0)$ → $\text{argmin}_\theta \mathbb{E}\left[f(\theta|B_0)\right]$

  - random $f(\theta|B_0) \rightarrow$ large variance
  - $\mathbb{E}\left[f(\theta|B_0)\right]$ provides higher precision
  - fast response
- On-line calibration monitoring (OCaMo)

  Reference calibration θ^{ref} → Small batch of data B_0 → Calibration validity $P(\theta^{\text{ref}}|B_0)$

  - needs to run on-line
  - small computational overhead

StOCaMo: Stereo on-line calibration monitoring

Main idea

- Examining epipolar distance between detected keypoints


StOCaMo: Stereo on-line calibration monitoring

Main idea

- Examining epipolar distance between detected keypoints

StOCaMo: Stereo on-line calibration monitoring

Main idea

- Examining epipolar distance between detected keypoints

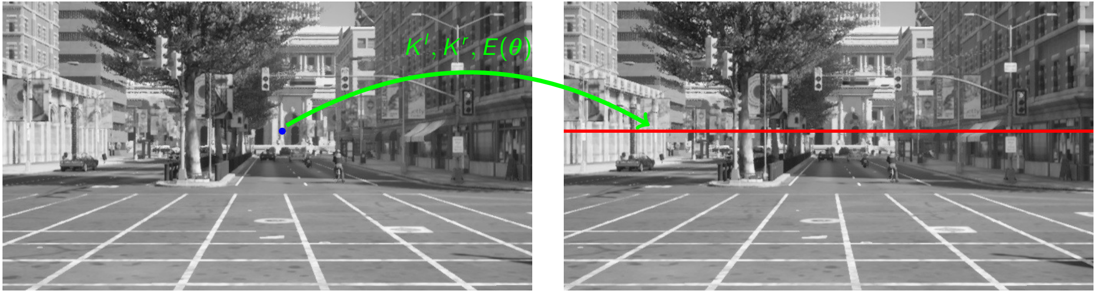

StOCaMo: Stereo on-line calibration monitoring

Main idea

- Examining epipolar distance between detected keypoints

StOCaMo: Stereo on-line calibration monitoring

Main idea

- Examining epipolar distance between detected keypoints
  - single frame estimation, without memory
  - small computational overhead
  - robust, without one-to-one matching

StOCaMo: Stereo on-line calibration monitoring

Robust kernel correlation [5]

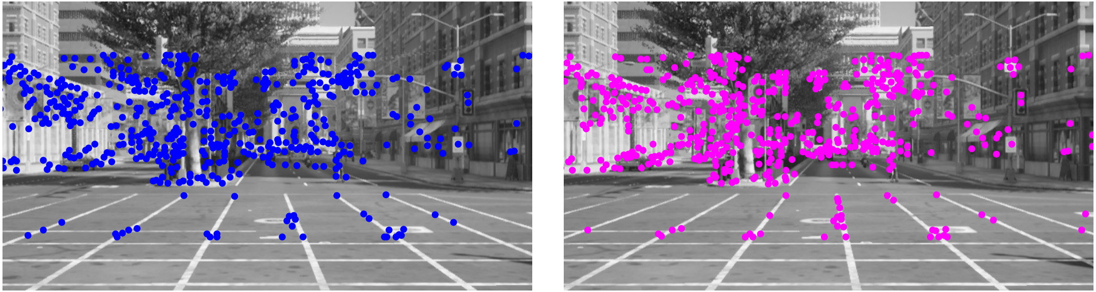

StOCaMo: Stereo on-line calibration monitoring

Robust kernel correlation [5]

StOCaMo: Stereo on-line calibration monitoring

Robust kernel correlation [5]

StOCaMo: Stereo on-line calibration monitoring

Robust kernel correlation [5]

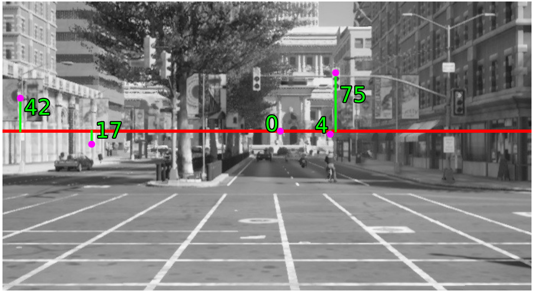

StOCaMo: Stereo on-line calibration moni

Robust kernel correlation [5]

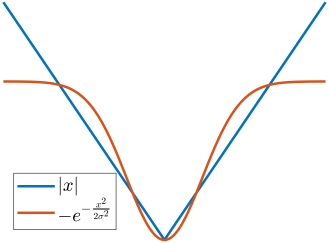

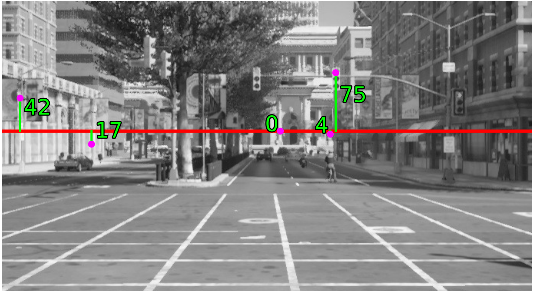

StOCaMo: Stereo on-line calibration monitoring

Robust kernel correlation [5]

$$KC(\boldsymbol{\theta}) = -\frac{1}{n}\sum_{i\in\mathcal{I}^l}\sum_{j\in \mathrm{kNN}_i^r}\exp\left[-\frac{d^2(\mathbf{x}_j^r \mid \mathbf{x}_i^l, \boldsymbol{\theta})}{2\sigma^2}\right] - \frac{1}{n}\sum_{j\in\mathcal{I}^r}\sum_{i\in \mathrm{kNN}_j^l}\exp\left[-\frac{d^2(\mathbf{x}_i^l \mid \mathbf{x}_j^r, \boldsymbol{\theta})}{2\sigma^2}\right]$$

$\mathcal{I}^l$

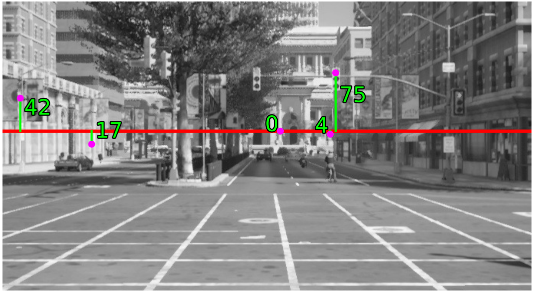

$\mathcal{I}^r$

StOCaMo: Stereo on-line calibration monitoring

Single frame estimation [4]

$$KC(\boldsymbol{\theta}) = -\frac{1}{n}\sum_{i\in\mathcal{I}^l}\sum_{j\in\mathrm{kNN}_i^r}\exp\left[-\frac{d^2(\mathbf{x}_j^r \mid \mathbf{x}_i^l,\boldsymbol{\theta})}{2\sigma^2}\right] - \frac{1}{n}\sum_{j\in\mathcal{I}^r}\sum_{i\in\mathrm{kNN}_j^l}\exp\left[-\frac{d^2(\mathbf{x}_i^l \mid \mathbf{x}_j^r,\boldsymbol{\theta})}{2\sigma^2}\right]$$

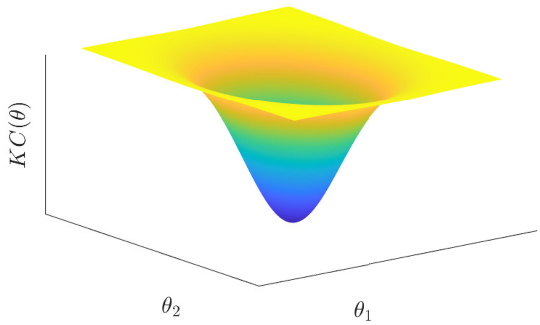

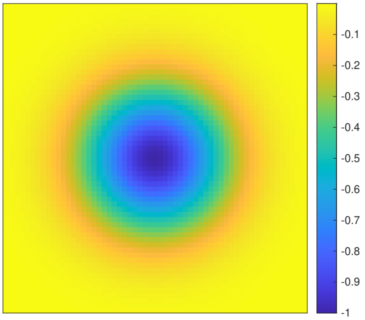

StOCaMo: Stereo on-line calibration monitoring

Single frame estimation [4]

$$F(\boldsymbol{\theta}^{\mathrm{ref}}) = \frac{1}{|\mathrm{grid}|} \sum_{\boldsymbol{\theta} \in \mathrm{grid}} \mathbb{1}\left[KC(\boldsymbol{\theta}) \geq KC(\boldsymbol{\theta}^{\mathrm{ref}})\right]$$

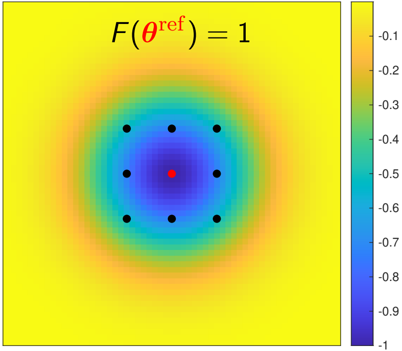

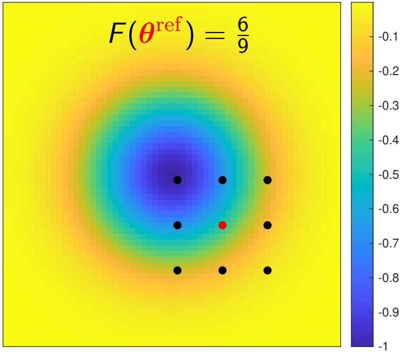

StOCaMo: Stereo on-line calibration monitoring

Single frame estimation [4]

$$F(\boldsymbol{\theta}^{\mathrm{ref}}) = \frac{1}{|\mathrm{grid}|} \sum_{\boldsymbol{\theta} \in \mathrm{grid}} \mathbb{1} \left[KC(\boldsymbol{\theta}) \geq KC(\boldsymbol{\theta}^{\mathrm{ref}}) \right]$$

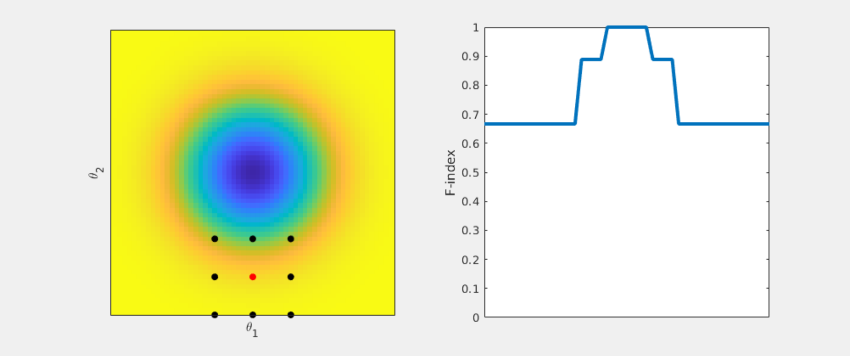

StOCaMo: Stereo on-line calibration monitoring

Single frame estimation [4]

$$F(\boldsymbol{\theta}^{\mathrm{ref}}) = \frac{1}{|\mathrm{grid}|} \sum_{\boldsymbol{\theta} \in \mathrm{grid}} \mathbb{1}\left[KC(\boldsymbol{\theta}) \geq KC(\boldsymbol{\theta}^{\mathrm{ref}})\right]$$

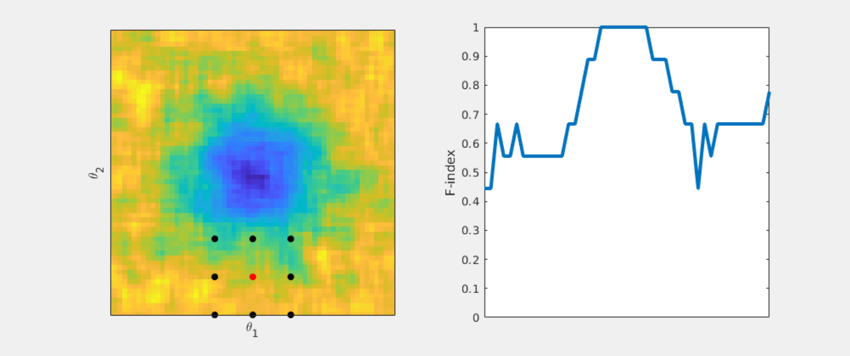

StOCaMo: Stereo on-line calibration monitoring

Single frame estimation [4]

$$V(\boldsymbol{\theta}^{\text{ref}}) = \frac{p_c(\boldsymbol{\theta}^{\text{ref}})}{p_c(\boldsymbol{\theta}^{\text{ref}}) + p_d(\boldsymbol{\theta}^{\text{ref}})}$$

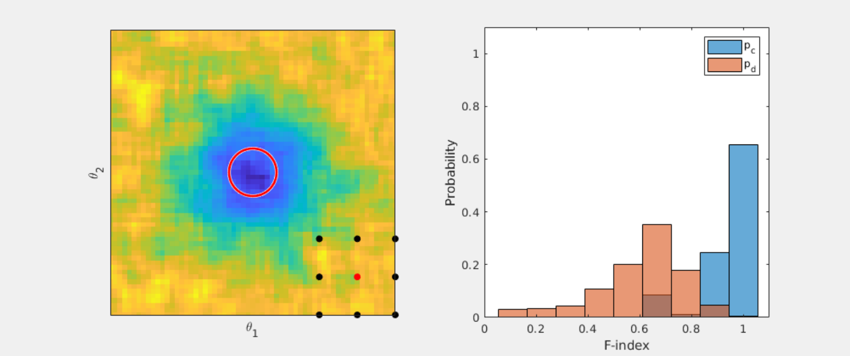

F-index evaluation & V-index parameters learning

- Synthetic dataset for parameter learning:
  - CARLA [2]: 24° VFOV and 1241×376 px (rectified)
  - 155 sequences with 200 frames each
  - Calibration tolerance was set to 0.005 → kernel σ and sampling grid size
- Evaluating F-index on seven decalibration magnitudes:
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0, 0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
- Selecting p_c and p_d based on 0.005 and 0.05 magnitudes, respectively

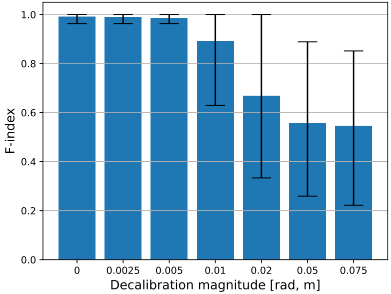

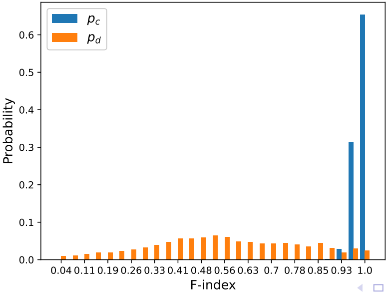

Experiments

Synthetic decalibration on real data

- Two real datasets:
  - KITTI [3]: 29.5° VFOV and 1241×376 px (rectified)
  - EuRoC MAV [1]: 55° VFOV and 752×480 px (unrectified)
- Two magnitudes of synthetic decalibration:
  - Small: $[-0.005, 0.005]$ m or rad
    $\longrightarrow$ should not report a decalibration (examines TN and FP)
  - Large: $[-0.02, -0.01] \cup [0.01, 0.02]$ m or rad
    $\longrightarrow$ should report a decalibration (examines TP and FN)

	TP	FN	TN	FP	Prec.	Recall	Acc.
KITTI	11854	156	11666	344	97.2	98.7	97.9
EuRoC	33240	3580	36321	499	98.5	90.3	94.5

Experiments

Synthetic decalibration on real data

- Two real datasets:
  - KITTI [3]: 29.5° VFOV and 1241×376 px (rectified)
  - EuRoC MAV [1]: 55° VFOV and 752×480 px (unrectified)
- Two magnitudes of synthetic decalibration:
  - Small: $[-0.005, 0.005]$ m or rad
    $\longrightarrow$ should not report a decalibration (examines TN and FP)
  - Large: $[-0.02, -0.01] \cup [0.01, 0.02]$ m or rad
    $\longrightarrow$ should report a decalibration (examines TP and FN)

	TP	FN	TN	FP	Prec.	Recall	Acc.
KITTI	11854	156	11666	344	97.2	98.7	97.9
EuRoC	33240	3580	36321	499	98.5	90.3	94.5

Experiments

Predicting downstream data processing failure

- Tested on KITTI (below) and EuRoC
- Downstream = ORB-SLAM2, failure = RMSE larger than threshold (as in [6])
- 100 decalibrations of 6 decalibration magnitudes
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
  - each decalibration is tested on ten random frames for more informative statistics

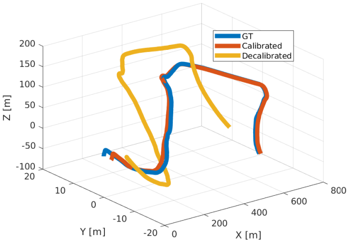

Dec.	TP	FN	TN	FP	Acc.
0.0025	0	0	990	10	99.0
0.005	0	0	959	41	95.9
0.01	262	218	376	144	63.8
0.02	699	191	82	28	78.1
0.05	924	66	3	7	92.7
0.075	928	72	0	0	92.8
Avg.					**87.1**
[6]					62

Experiments

Predicting downstream data processing failure

- Tested on KITTI (below) and EuRoC
- Downstream = ORB-SLAM2, failure = RMSE larger than threshold (as in [6])
- 100 decalibrations of 6 decalibration magnitudes
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
  - each decalibration is tested on ten random frames for more informative statistics

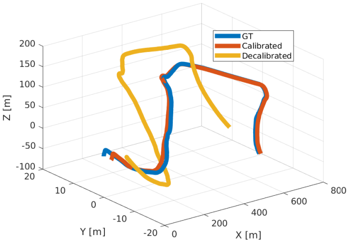

Dec.	TP	FN	TN	FP	Acc.
0.0025	0	0	990	10	99.0
0.005	0	0	959	41	95.9
0.01	262	218	376	144	63.8
0.02	699	191	82	28	78.1
0.05	924	66	3	7	92.7
0.075	928	72	0	0	92.8
Avg.					**87.1**
[6]					62

Experiments

Predicting downstream data processing failure

- Tested on KITTI (below) and EuRoC
- Downstream = ORB-SLAM2, failure = RMSE larger than threshold (as in [6])
- 100 decalibrations of 6 decalibration magnitudes
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
  - each decalibration is tested on ten random frames for more informative statistics

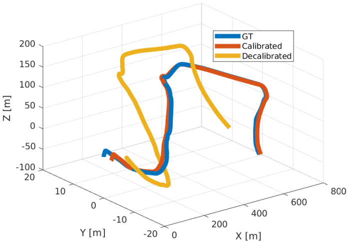

Dec.	TP	FN	TN	FP	Acc.
0.0025	0	0	990	10	99.0
0.005	0	0	959	41	95.9
0.01	262	218	376	144	63.8
0.02	699	191	82	28	78.1
0.05	924	66	3	7	92.7
0.075	928	72	0	0	92.8
Avg.					**87.1**
[6]					62

Experiments

Predicting downstream data processing failure

- Tested on KITTI (below) and EuRoC
- Downstream = ORB-SLAM2, failure = RMSE larger than threshold (as in [6])
- 100 decalibrations of 6 decalibration magnitudes
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
  - each decalibration is tested on ten random frames for more informative statistics

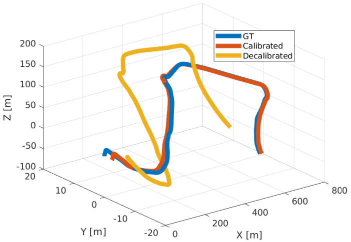

Dec.	TP	FN	TN	FP	Acc.
0.0025	0	0	990	10	99.0
0.005	0	0	959	41	95.9
0.01	262	218	376	144	63.8
0.02	699	191	82	28	78.1
0.05	924	66	3	7	92.7
0.075	928	72	0	0	92.8
Avg.					**87.1**
[6]					62

Experiments

Predicting downstream data processing failure

- Tested on KITTI (below) and EuRoC
- Downstream = ORB-SLAM2, failure = RMSE larger than threshold (as in [6])
- 100 decalibrations of 6 decalibration magnitudes
  - $[-\delta, \delta]$ m or rad, $\delta \in \{0.0025, 0.005, 0.01, 0.02, 0.05, 0.075\}$
  - each decalibration is tested on ten random frames for more informative statistics
- Regression of the RMSE with validity index of StOCaMo

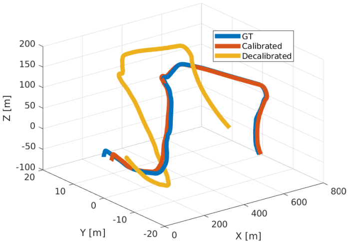

Dec.	TP	FN	TN	FP	Acc.	ρ
0.0025	0	0	990	10	99.0	
0.005	0	0	959	41	95.9	
0.01	262	218	376	144	63.8	
0.02	699	191	82	28	78.1	
0.05	924	66	3	7	92.7	
0.075	928	72	0	0	92.8	
Avg.					**87.1**	**-0.76**
[6]					62	0.44

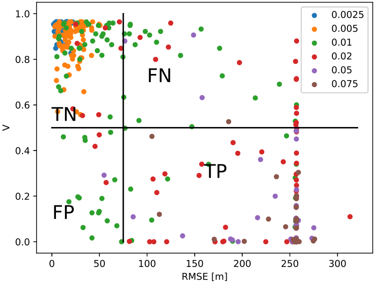

Future Work & Conclusion

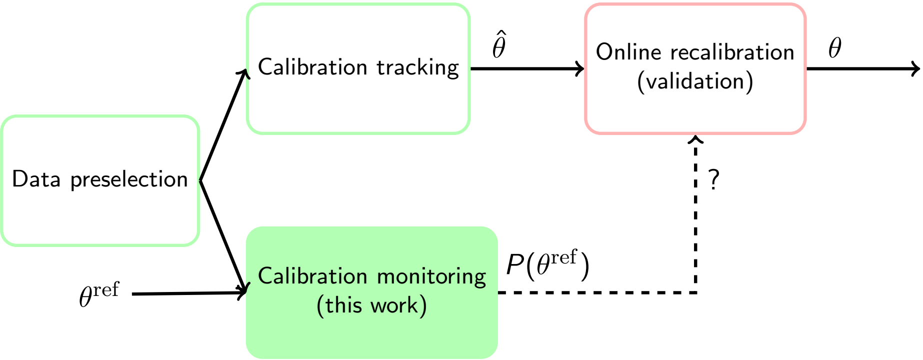

Future Work & Conclusion

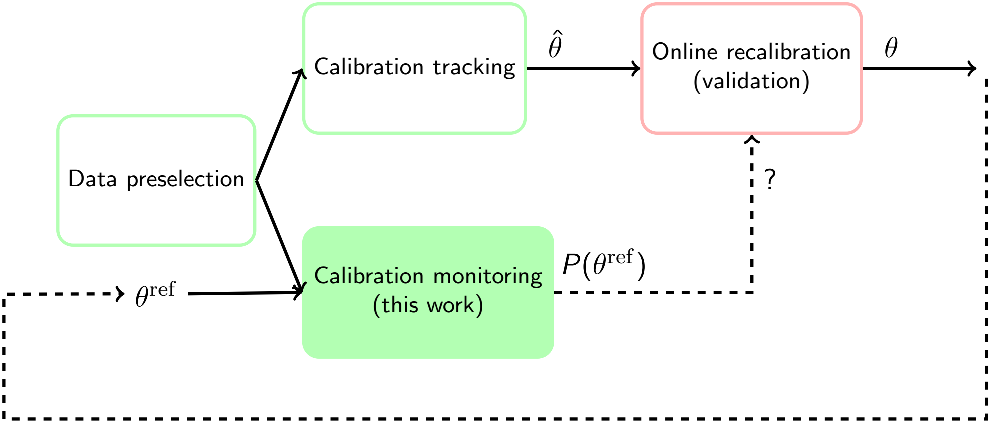

References I

[1] Michael Burri, Janosch Nikolic, Pascal Gohl, et al. "The EuRoC micro aerial vehicle datasets". In: (2016).

[2] Alexey Dosovitskiy, German Ros, Felipe Codevilla, et al. "CARLA: An Open Urban Driving Simulator". In: *Proc. of the Annual Conf. on Robot Learning*. 2017, pp. 1–16.

[3] Andreas Geiger, Philip Lenz, and Raquel Urtasun. "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite". In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2012.

[4] Jesse Levinson and Sebastian Thrun. "Automatic Online Calibration of Cameras and Lasers". In: *Proceedings Robotics: Science and Systems Conference*. Vol. 2. 2013, p. 7.

[5] Yanghai Tsin and Takeo Kanade. "A correlation-based approach to robust point set registration". In: 2004, pp. 558–569.

[6] Jiapeng Zhong, Zheyu Ye, Andrei Cramariuc, et al. "CalQNet-Detection of Calibration Quality for Life-Long Stereo Camera Setups". In: 2021, pp. 1312–1318.