ISE RTS
Institute of Systems Engineering
Real Time Systems Group

Leibniz
Universität
Hannover

Interval-Based Visual-LiDAR Sensor Fusion

Raphael Voges and Bernardo Wagner

2021 IEEE International Conference on
Robotics and Automation (ICRA)

June 3, 2021

DFG Research Training Group (RTG 2159)
i.c.sens - Integrity and Collaboration in Dynamic Sensor Networks

Motivation: Complementary Sensors

- Camera: reidentifiable 2D features

Motivation: Complementary Sensors

- Camera: reidentifiable 2D features
- LiDAR: accurate distance measurements

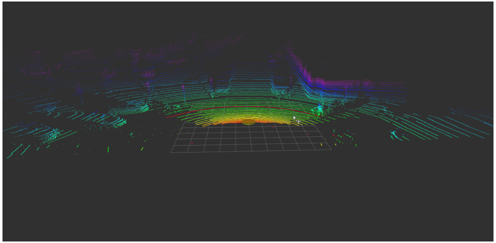

Motivation: Complementary Sensors

- Camera: reidentifiable 2D features
- LiDAR: accurate distance measurements
- Sensor fusion: reidentifiable 3D features

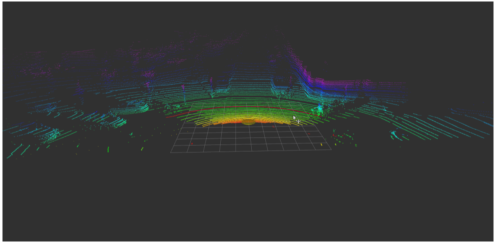

Motivation: Complementary Sensors

- Camera: reidentifiable 2D features
- LiDAR: accurate distance measurements
- Sensor fusion: reidentifiable 3D features
- Uncertainty of fused information must be known for subsequent tasks

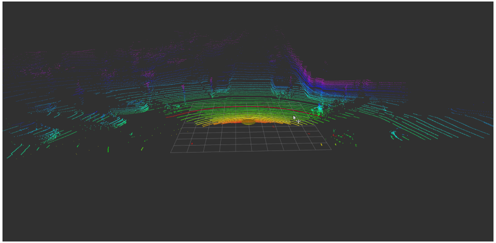

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown

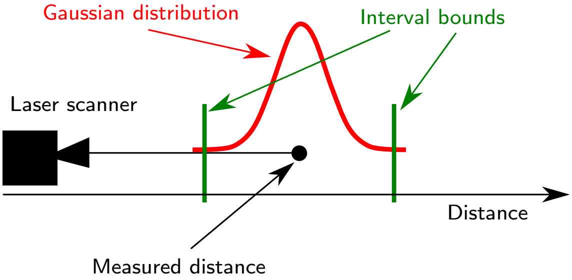

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors

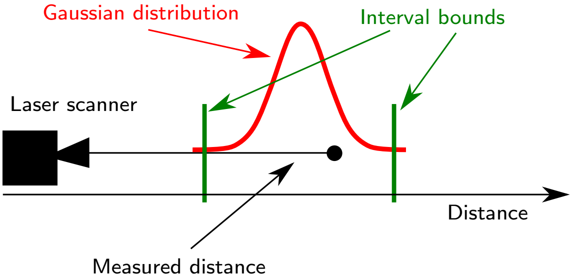

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors
- Results can be guaranteed

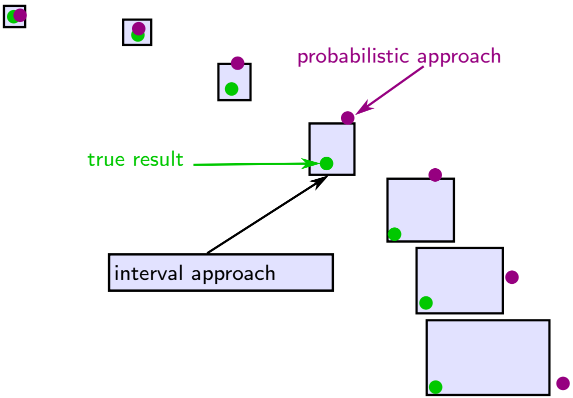

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors
- Results can be guaranteed
  - Important for safety-critical systems

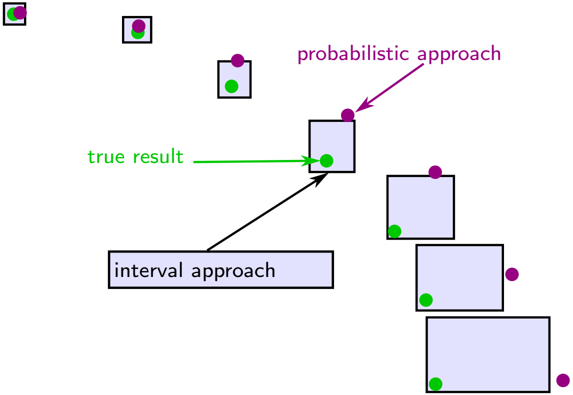

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors
- Results can be guaranteed
  - Important for safety-critical systems
- Reverse view:
  - Dismiss infeasible solutions

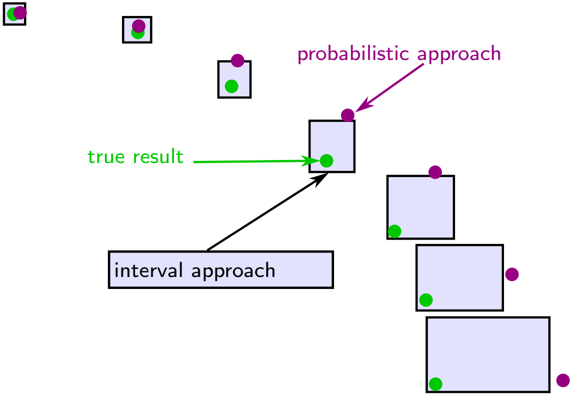

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors
- Results can be guaranteed
  - Important for safety-critical systems
- Reverse view:
  - Dismiss infeasible solutions
- Computations are deterministic

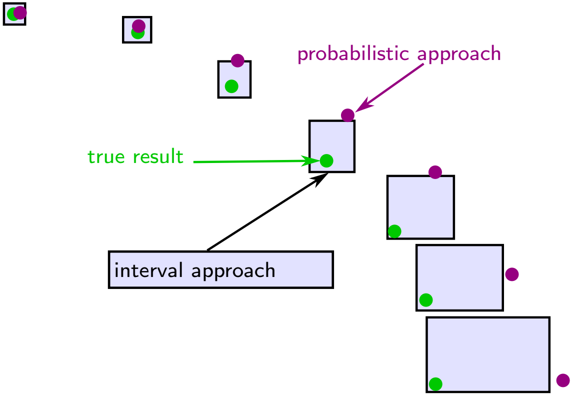

Motivation: Why Unknown But Bounded Errors?

- Error distribution is often unknown
  - Unknown systematic errors
- Results can be guaranteed
  - Important for safety-critical systems
- Reverse view:
  - Dismiss infeasible solutions
- Computations are deterministic
  - Proofs become possible

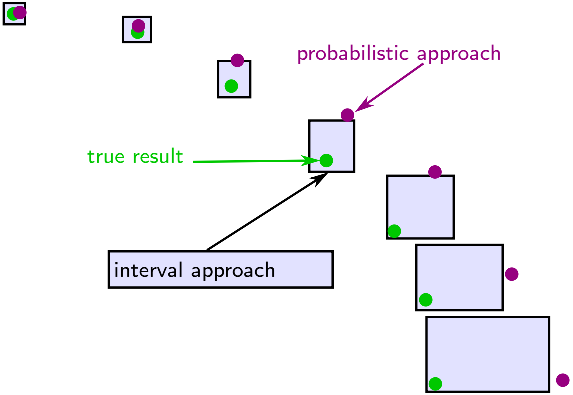

Bounded Sensor Error Models: Laser Scanner

- Unknown but bounded errors for:

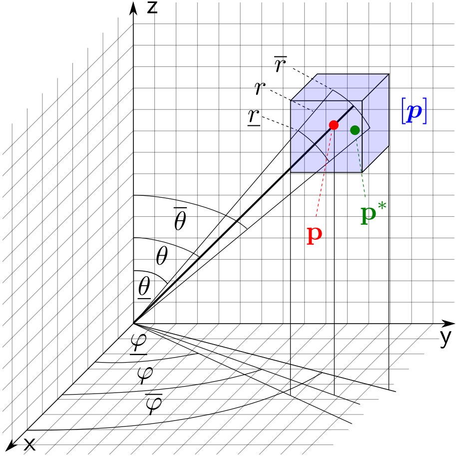

Bounded Sensor Error Models: Laser Scanner

- Unknown but bounded errors for:
  - Distance measurement $[r] = [\underline{r}, \overline{r}]$

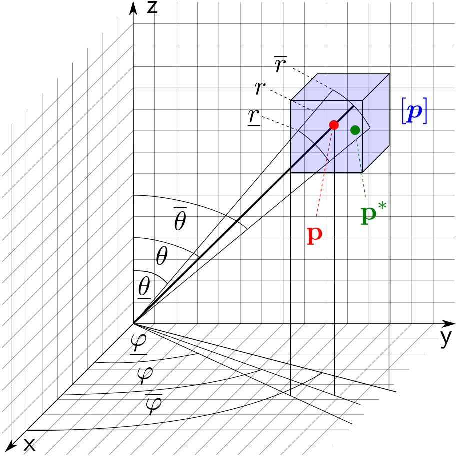

Bounded Sensor Error Models: Laser Scanner

- Unknown but bounded errors for:
  - Distance measurement $[r] = [\underline{r}, \overline{r}]$
  - Angular components $[\theta] = [\underline{\theta}, \overline{\theta}]$ and $[\varphi] = [\underline{\varphi}, \overline{\varphi}]$

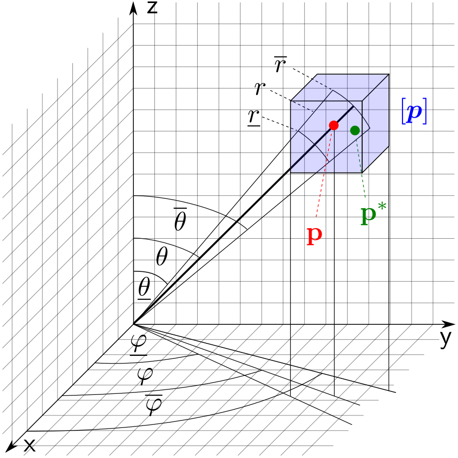

Bounded Sensor Error Models: Laser Scanner

- Unknown but bounded errors for:
  - Distance measurement $[r] = [\underline{r}, \overline{r}]$
  - Angular components $[\theta] = [\underline{\theta}, \overline{\theta}]$ and $[\varphi] = [\underline{\varphi}, \overline{\varphi}]$
- Conversion into Cartesian coordinates

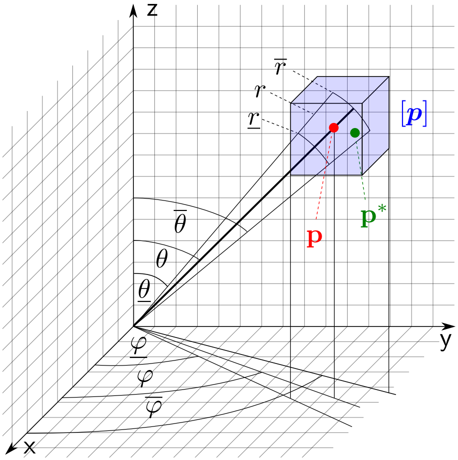

Bounded Sensor Error Models: Laser Scanner

- Unknown but bounded errors for:
  - Distance measurement $[r] = [\underline{r}, \overline{r}]$
  - Angular components $[\theta] = [\underline{\theta}, \overline{\theta}]$ and $[\varphi] = [\underline{\varphi}, \overline{\varphi}]$
- Conversion into Cartesian coordinates
  - 3D box that is guaranteed to contain the true 3D point: $\mathbf{p}^* \in [\boldsymbol{p}]$

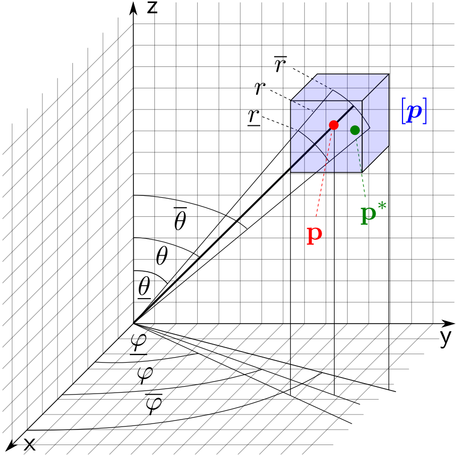

Bounded Sensor Error Models: Camera

- Unknown but bounded errors:
  - Detection and matching of image features
- Error bounds:
  - Empirically under consideration of outliers

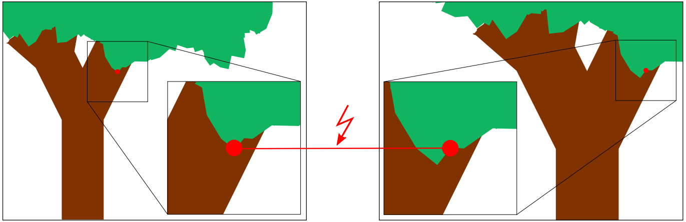

Bounded Sensor Error Models: Camera

- Unknown but bounded errors:
  - Detection and matching of image features
- Error bounds:
  - Empirically under consideration of outliers

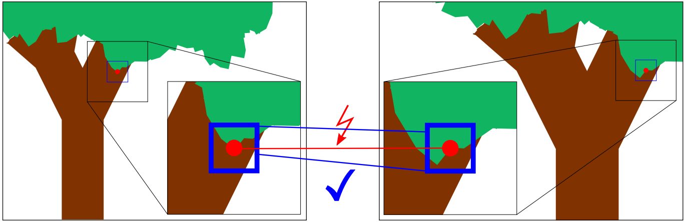

Visual-LiDAR Sensor Fusion

- Goal: Assign distance information to 2D image features

Visual-LiDAR Sensor Fusion

- Goal: Assign distance information to 2D image features
- Project 3D laser scan boxes onto the 2D image plane

Traditional projection without considering uncertainties.

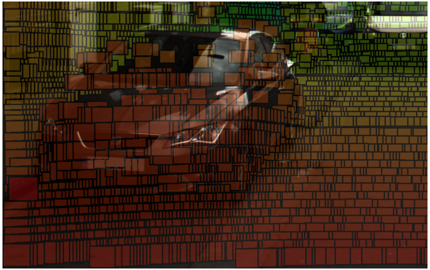

Interval boxes enclosing the projected laser scan points. The color of each box corresponds to the midpoint of the depth interval.

Visual-LiDAR Sensor Fusion

- Goal: Assign distance information to 2D image features
- Project 3D laser scan boxes onto the 2D image plane
- Extrinsic transformation between camera and laser scanner required

Traditional projection without considering uncertainties.

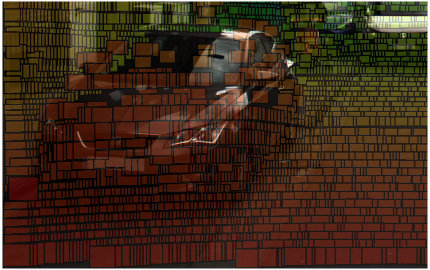

Interval boxes enclosing the projected laser scan points. The color of each box corresponds to the midpoint of the depth interval.

Visual-LiDAR Sensor Fusion

- Goal: Assign distance information to 2D image features
- Project 3D laser scan boxes onto the 2D image plane
- Extrinsic transformation between camera and laser scanner required
- Sensor clocks are assumed to be synchronized

Traditional projection without considering uncertainties.

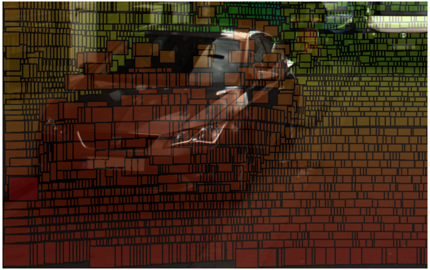

Interval boxes enclosing the projected laser scan points. The color of each box corresponds to the midpoint of the depth interval.

Visual-LiDAR Sensor Fusion

- Set of all projected scan boxes that have a non-empty intersection with the image feature:

$$\mathcal{S}_j = \{\, i \in \{1, \ldots, N_l\} \mid [\tilde{\mathbf{X}}_i^C] \cap [\tilde{\mathbf{X}}_j^C] \neq \varnothing \,\}$$

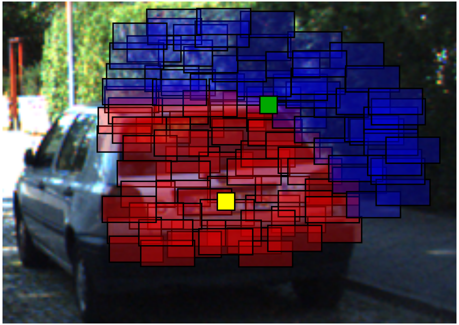

Visual-LiDAR Sensor Fusion

- Set of all projected scan boxes that have a non-empty intersection with the image feature:

$$\mathcal{S}_j = \{\, i \in \{1, \ldots, N_l\} \mid [\tilde{\mathbf{X}}_i^C] \cap [\tilde{\mathbf{X}}_j^C] \neq \varnothing \,\}$$

- Compute the z-coordinate (depth) of the image feature as the union over the z-coordinates of all overlapping scan boxes:

$$[z_j^C] = \bigcup_{i \in \mathcal{S}_j} [z_i^C]$$

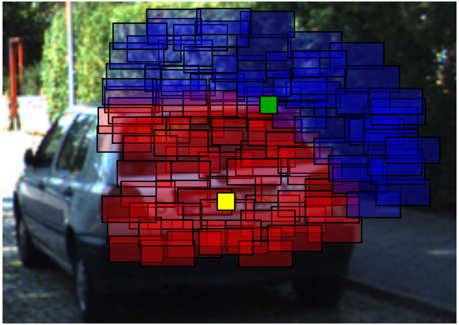

Visual-LiDAR Sensor Fusion

- Set of all projected scan boxes that have a non-empty intersection with the image feature:

$$\mathcal{S}_j = \{\, i \in \{1, \ldots, N_l\} \mid [\tilde{\mathbf{X}}_i^C] \cap [\tilde{\mathbf{X}}_j^C] \neq \varnothing \,\}$$

- Compute the z-coordinate (depth) of the image feature as the union over the z-coordinates of all overlapping scan boxes:

$$[z_j^C] = \bigcup_{i \in \mathcal{S}_j} [z_i^C]$$

- Results in re-identifiable 3D features with information about accuracy

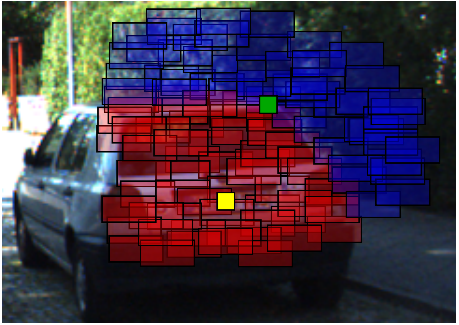

Visual-LiDAR Sensor Fusion

- Exemplary results of the approach for Visual-LiDAR data fusion

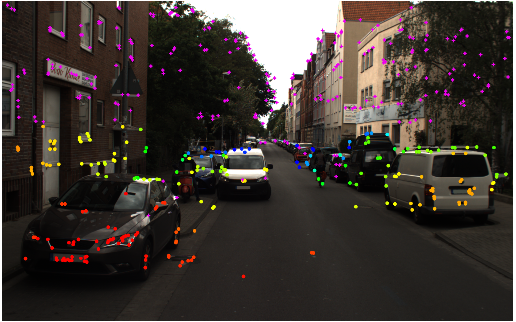

Colored by depth (red: close, blue: distant). Features without depth information are colored pink.

Colored by depth accuracy (red: accurate, blue: inaccurate).

Guaranteed Visual-LiDAR Odometry

General idea

- Dead reckoning

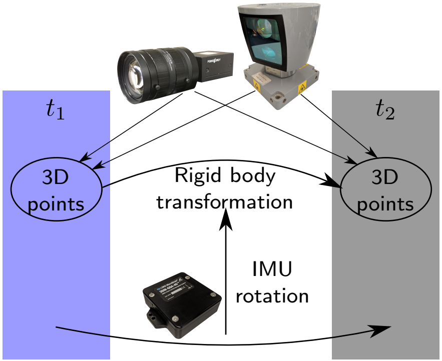

Guaranteed Visual-LiDAR Odometry

General idea

- Dead reckoning
- Visual-LiDAR sensor fusion
    - Corresponding 3D boxes

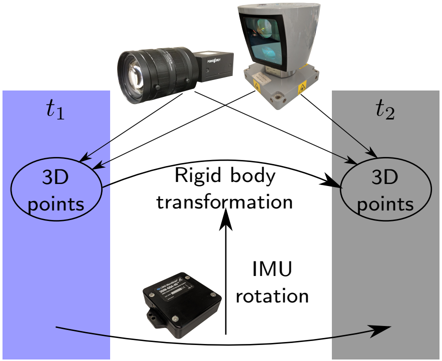

Guaranteed Visual-LiDAR Odometry

General idea

- Dead reckoning
- Visual-LiDAR sensor fusion
  - Corresponding 3D boxes
- Integrate angular velocities to constrain rotation

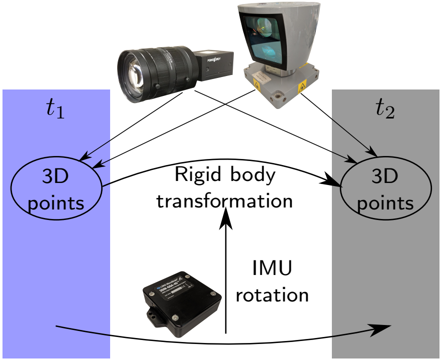

Guaranteed Visual-LiDAR Odometry

General idea

- Dead reckoning
- Visual-LiDAR sensor fusion
  - Corresponding 3D boxes
- Integrate angular velocities to constrain rotation
- Compute 6 DOF rigid body transformation under interval uncertainty
  - Constraint Satisfaction Problem: forward-backward contractors [1]

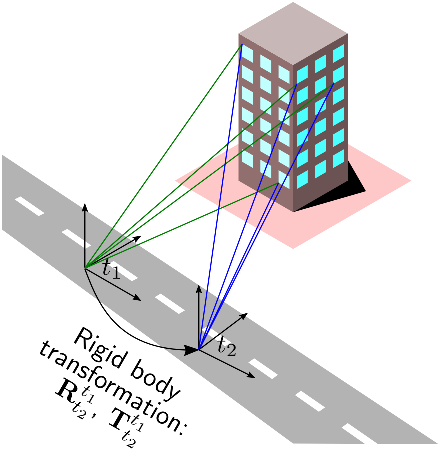

Guaranteed Visual-LiDAR Odometry

Concept of keyframes

- Estimate motion relative to the last keyframe k_1 until we have to insert a new keyframe k_2
- As few keyframes as possible, but as many as necessary

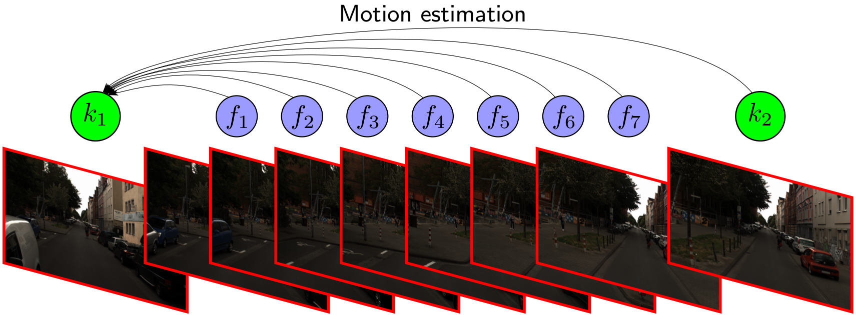

Guaranteed Visual-LiDAR Odometry

Concept of keyframes

- Estimate motion relative to the last keyframe k_1 until we have to insert a new keyframe k_2
- As few keyframes as possible, but as many as necessary

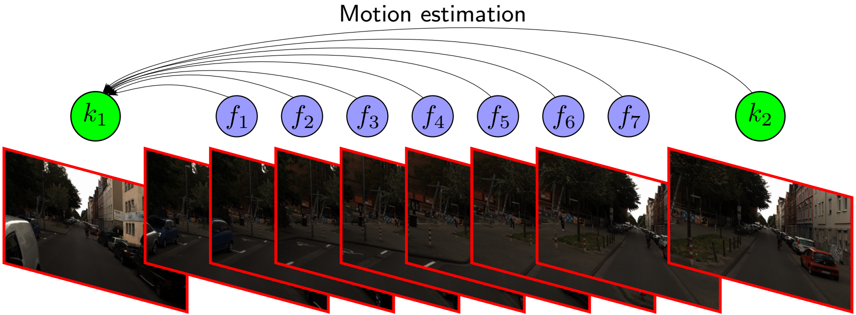

- Dynamic insertion of keyframes depending on the 6D pose uncertainty
- High pose uncertainty $\Leftrightarrow$ insufficient constraints $\Leftrightarrow$ inaccurate 3D features

Experimental Results

- Dead reckoning **without** error propagation at keyframes

Experimental Results

- Dead reckoning **without** error propagation at keyframes
- Quantitative results

Dataset	Correct (%)	Volume (m^3)	Area (m^2)	Rotation accuracy (°)	Features w. depth	Distance per keyframe (m)	Inconsistencies of [2](%)
[3]	100	0.77	1.30	0.39	105	11.35	29.9
0009 [4]	97.8*	1.52	1.81	0.21	201	9.57	16.1
0023 [4]	97.7*	0.91	1.39	0.19	173	11.91	18.2
0106 [4]	98.3*	0.94	1.32	0.22	170	11.48	12.2

Experimental Results

- Dead reckoning **without** error propagation at keyframes
- Quantitative results

Dataset	Correct (%)	Volume (m^3)	Area (m^2)	Rotation accuracy (°)	Features w. depth	Distance per keyframe (m)	Inconsistencies of [2](%)
[3]	100	0.77	1.30	0.39	105	11.35	29.9
0009 [4]	97.8*	1.52	1.81	0.21	201	9.57	16.1
0023 [4]	97.7*	0.91	1.39	0.19	173	11.91	18.2
0106 [4]	98.3*	0.94	1.32	0.22	170	11.48	12.2

- Desired guarantees are provided

Experimental Results

- Dead reckoning **without** error propagation at keyframes
- Quantitative results

Dataset	Correct (%)	Volume (m^3)	Area (m^2)	Rotation accuracy (°)	Features w. depth	Distance per keyframe (m)	Inconsistencies of [2](%)
[3]	100	0.77	1.30	0.39	105	11.35	29.9
0009 [4]	97.8*	1.52	1.81	0.21	201	9.57	16.1
0023 [4]	97.7*	0.91	1.39	0.19	173	11.91	18.2
0106 [4]	98.3*	0.94	1.32	0.22	170	11.48	12.2

- Desired guarantees are provided
- Inconsistencies of state-of-the-art approach [2] are reliably detected

Experimental Results

Exemplary image features for a detected inconsistency

Stochastic features colored by depth (red: close, blue: distant, violet: no depth).

Interval features colored by depth midpoint (red: close, blue: distant, violet: no depth).

Experimental Results

Exemplary image features for a detected inconsistency

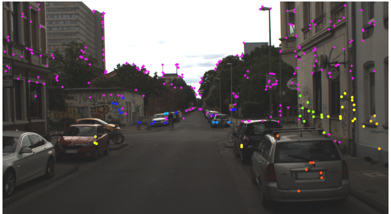

Stochastic features colored by depth
(red: close, blue: distant, violet: no depth).

Interval features colored by depth midpoint
(red: close, blue: distant, violet: no depth).

Interval features colored by uncertainty
(red: certain, blue: uncertain).

Conclusions

- Visual-LiDAR sensor fusion
  - Visual features are augmented by distance information
  - Sensor and transformation errors are considered and propagated

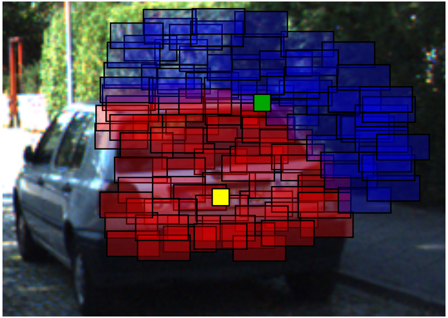

Conclusions

- Visual-LiDAR sensor fusion
  - Visual features are augmented by distance information
  - Sensor and transformation errors are considered and propagated
- Guaranteed visual-LiDAR odometry
  - Takes the feature uncertainty into account
  - Reliably encloses the true localization result
  - Inconsistencies of stochastic approach can be detected

Conclusions

- Visual-LiDAR sensor fusion
  - Visual features are augmented by distance information
  - Sensor and transformation errors are considered and propagated
- Guaranteed visual-LiDAR odometry
  - Takes the feature uncertainty into account
  - Reliably encloses the true localization result
  - Inconsistencies of stochastic approach can be detected
- Future Work
  - Incorporate global constraints (e.g., GNSS, SLAM) to reduce drift
  - Restrict stochastic methods to interval box

References

[1] G. Chabert and L. Jaulin, "Contractor programming," *Artificial Intelligence*, vol. 173, no. 11, pp. 1079–1100, Jul. 2009.

[2] J. Zhang, M. Kaess, and S. Singh, "Real-time Depth Enhanced Monocular Odometry," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Chicago, Illinois, USA, Sep. 2014.

[3] R. Voges, *Dataset: i.c.sens Visual-Inertial-LiDAR Dataset*, `https://doi.org/10.25835/0026408`, 2020.

[4] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, Aug. 2013.

This work was supported by the German Research Foundation (DFG) as part of the Research Training Group i.c.sens [RTG 2159].