IROS '24
ABU DHABI

Efficient Extrinsic Self-Calibration of Multiple IMUs using Measurement Subset Selection

Jongwon Lee[1], David Hanley[2], and Timothy Bretl[1]
[1] Department of Aerospace Engineering, University of Illinois Urbana-Champaign, USA
[2] School of Informatics, University of Edinburgh, UK

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN
THE UNIVERSITY of EDINBURGH

Introduction

- Extrinsic calibration, estimating the relative pose of IMUs, is essential for multi-IMU systems
- Self-calibration uses only IMU data, without prescribed trajectories or external sensors (e.g., cameras)

Our contribution

- Self-calibration must be efficient for large datasets, which are common outcomes in data collection scenarios (e.g., spacecraft, vehicles)

Jacobian: Function of the measurements and parameter estimates

$\mathcal{J} = \mathbf{J}^T \mathbf{J}$

: Measurements with high-utility

We propose an efficient self-calibration method for multiple IMUs by identifying high-utility data

Based on our prior work (RAL'22)

Based on an existing approach ("Greedy algorithm")

Paper LinkedIn

Problem statement

- **Given:** Measurements $\mathcal{D} = \bigcup_{l=1}^{L} \mathcal{D}^l$
- **To find:** Parameter estimate $\hat{\boldsymbol{\theta}}$ while identifying an informative subset $\mathcal{D}^{\text{info}} \subseteq \mathcal{D}$

Our calibration aims to estimate extrinsic parameters:

- $\mathbf{p}$: relative position
- $\mathbf{q}$: relative orientation
- ${}^{g}_{I}\mathbf{q}$: gyroscope misalignment

IROS '24
ABU DHABI

Paper

LinkedIn

Existing approach (Greedy algorithm)

Given a candidate segment $\mathcal{D}^{\text{new}}$ during iteration over the measurement segments...

Step 1: State Initialization

Informative subset: $\mathcal{D}^{\text{info}}\ (\subset \{\mathcal{D}^1, \ldots, \mathcal{D}^{\text{new}-1}\})$

Parameter estimate: $\hat{\boldsymbol{\theta}}^-$

Step 2: Calibration

$$\hat{\boldsymbol{\theta}}^+ = \text{Calibrate}(\mathcal{D}^{\text{info}}, \mathcal{D}^{\text{new}}; \hat{\boldsymbol{\theta}}^-)$$

Step 3: Utility Gain Evaluation

$$f\left[\mathcal{I}(\mathcal{D}^{\text{info}}, \mathcal{D}^{\text{new}})\big|_{\hat{\boldsymbol{\theta}}^+}\right] - f\left[\mathcal{I}(\mathcal{D}^{\text{info}})\big|_{\hat{\boldsymbol{\theta}}^+}\right] > \lambda?$$

Step 4: State Update (If true in Step 3)

$$\mathcal{D}^{\text{info}} \leftarrow \mathcal{D}^{\text{info}} \cup \mathcal{D}^{\text{new}}$$

Parameter estimate: $\hat{\boldsymbol{\theta}}^- \leftarrow \hat{\boldsymbol{\theta}}^+$

IROS '24 ABU DHABI

Paper LinkedIn

Modified approach

Given a candidate segment $\mathcal{D}^{\text{new}}$ during iteration over the measurement segments...

Step 1: State Initialization

Informative subset: $\mathcal{D}^{\text{info}} (\subset \{\mathcal{D}^1, \ldots, \mathcal{D}^{\text{new}-1}\})$

Parameter estimate: $\boldsymbol{\theta}^0$

Step 2: Calibration

$\hat{\boldsymbol{\theta}}^+ = \text{Calibrate}(\mathcal{D}^{\text{info}}, \mathcal{D}^{\text{new}}; \hat{\boldsymbol{\theta}}^-)$

Step 3: Utility Gain Evaluation

$f\left[\mathcal{I}(\mathcal{D}^{\text{info}}, \mathcal{D}^{\text{new}})\big|_{\boldsymbol{\theta}^0}\right] - f\left[\mathcal{I}(\mathcal{D}^{\text{info}})\big|_{\boldsymbol{\theta}^0}\right] > \lambda?$

Step 4: State Update (If true in Step 3)

$\mathcal{D}^{\text{info}} \leftarrow \mathcal{D}^{\text{info}} \cup \mathcal{D}^{\text{new}}$

Step 5: Final Calibration (After completing iterations)

$\hat{\boldsymbol{\theta}}^+ = \text{Calibrate}(\mathcal{D}^{\text{info}}; \boldsymbol{\theta}^0)$

Evaluation: Comparison against benchmarks

- We compared the multi-IMU extrinsic self-calibration in three different calibration modes:
  - The full set (**Baseline**)
  - Subset selected by the existing approach (**Greedy (original)**)
  - Our modified approach (**Greedy (init-param)**)
- The methods were evaluated over three trajectories, each lasting over 20 minutes

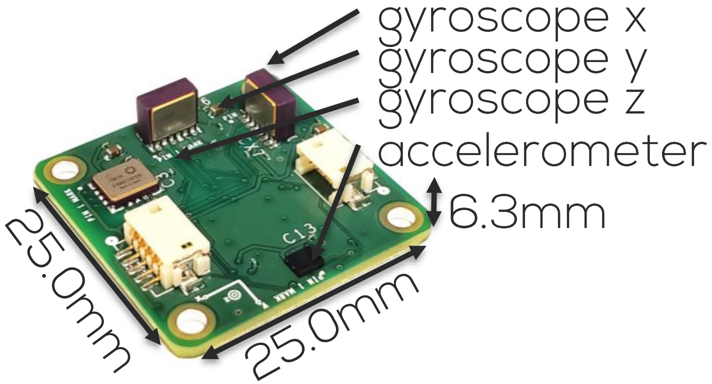

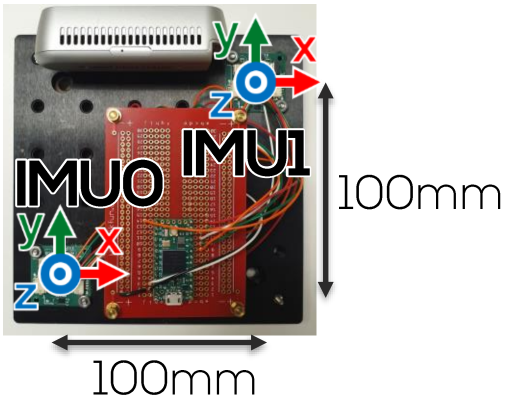

IMU (left) and sensor rig (right) used for the experiments

Results

- Both Greedy algorithms select and use less than 3% of the full set, without compromising the calibration results
- Greedy (init-param) significantly reduces runtime compared to the baseline (**>1 min → ~1 sec**), and even to Greedy (original) **(~1 min → ~1 sec)**

Trajectory		Baseline	Greedy (Original)	Greedy (Init-Param)
`baseline` (1274 [s])	$\mathbf{p}_{I_1}$ [cm]	[-9.80 -9.82 -0.24]	[-9.85 -9.63 -0.34]	[-9.64 -9.58 -0.29]
	$\mathbf{q}_{I_1}$ [deg]	[4.15 3.28 0.27]	[3.39 2.01 -0.42]	[1.98 0.61 -0.40]
	$\mathbf{q}_{g_0}$ [deg]	[-0.74 -0.45 0.04]	[2.47 4.34 -1.08]	[-0.90 1.06 0.21]
	$\mathbf{q}_{g_1}$ [deg]	[3.11 3.31 0.45]	[5.62 6.82 -1.05]	[0.84 2.12 0.09]
	Selected segments [%]	100.00	0.73	1.14
	Runtime [s]	64.09	16.15	**0.98**
`blurry` (1388 [s])	$\mathbf{p}_{I_1}$ [cm]	[-9.81 -9.77 -0.26]	[-9.84 -9.78 -0.30]	[-9.79 -9.72 -0.27]
	$\mathbf{q}_{I_1}$ [deg]	[5.27 4.27 0.25]	[2.22 0.98 -0.25]	[1.96 0.96 -0.65]
	$\mathbf{q}_{g_0}$ [deg]	[-2.96 -2.74 0.00]	[-0.43 -0.06 -0.27]	[-1.76 -1.13 0.45]
	$\mathbf{q}_{g_1}$ [deg]	[2.03 2.11 0.20]	[1.65 1.28 -0.19]	[0.09 0.15 0.05]
	Selected segments [%]	100.00	1.94	0.82
	Runtime [s]	75.81	51.01	**0.90**
`ill-lit` (1276 [s])	$\mathbf{p}_{I_1}$ [cm]	[-9.82 -9.86 -0.21]	[-9.67 -9.96 -0.14]	[-9.65 -9.83 -0.17]
	$\mathbf{q}_{I_1}$ [deg]	[-5.03 -5.79 0.28]	[2.50 1.29 -0.02]	[5.03 3.50 0.10]
	$\mathbf{q}_{g_0}$ [deg]	[-0.36 -0.34 0.02]	[1.70 0.69 -0.18]	[2.55 1.89 -0.40]
	$\mathbf{q}_{g_1}$ [deg]	[-5.79 -5.58 0.39]	[3.82 2.42 0.02]	[7.24 6.20 0.21]
	Selected segments [%]	100.00	2.19	1.05
	Runtime [s]	73.75	59.37	**0.99**

$\mathbf{p}$: relative position, $\mathbf{q}$: relative orientation, ${}^{g}_{I}\mathbf{q}$: gyroscope misalignment

$\mathbf{p}_{\mathrm{ref}}: [100, 100, 0] \pm [250, 250, 63]$ mm, $\mathbf{q}_{\mathrm{ref}}, {}^{g}_{I}\mathbf{q}_{\mathrm{ref}}: (\mathbf{e}, 0°) \pm (\mathbf{e}, 9.5°)$ for $\forall \mathbf{e} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ (angle-axis)

Paper LinkedIn

Conclusion

- We proposed a method for multi-IMU extrinsic calibration by efficiently selecting high-utility measurement subsets
- We hypothesized that utility – a function of parameter estimates – is largely insensitive to specific parameter choices
- This eliminates the need for frequent recalibrations, significantly reducing runtimes compared to existing subset selection methods

Paper LinkedIn