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

IROS '24 ABU DHABI

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

- For multi-IMU systems, estimating the relative position and orientation (pose) of IMUs – *extrinsic calibration* – is essential
- Extrinsic calibration only with the measurement of IMUs themselves, requiring neither prescribed trajectories nor aiding sensors (e.g., cameras), is called *self*-calibration
- Self-calibration should be performed efficiently on large datasets, which are commonly generated during data collection (e.g., spacecraft in orbit, cars on the road)

We propose an efficient $\boxed{\text{self-calibration method for multiple IMUs}}$ $\boxed{\text{by identifying high-utility data}}$

Based on our prior work [1]
Based on an existing approach [2,3]

PROBLEM STATEMENT

Our calibration aims to estimate extrinsic parameters:
- $\mathbf{p}$: relative position
- $\mathbf{q}$: relative orientation
- ${}^{g}_{I}\mathbf{q}$: gyroscope misalignment

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

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

EXISTING APPROACH (GREEDY ALGORITHM)

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

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

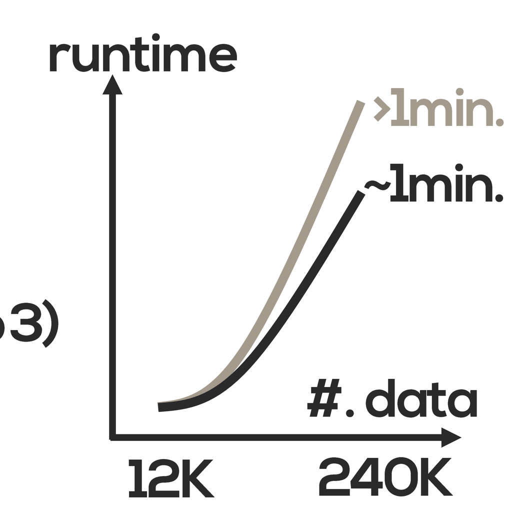

MODIFIED APPROACH

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

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

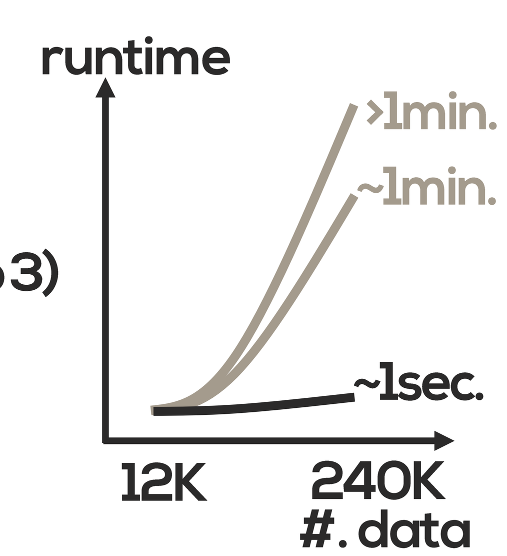

SENSITIVITY ANALYSIS OF UTILITY

- A sensitivity analysis of the measured utility, $f[\mathcal{I}(\cdot)|_{\boldsymbol{\theta}}]$, was performed by introducing $\delta\mathbf{p}, \delta\mathbf{q}$ from the true $\boldsymbol{\theta}$

Deviation in relative position and orientation ($\delta\mathbf{p}, \delta\mathbf{q}$) from their true values ($\mathbf{p}, \mathbf{q}$)

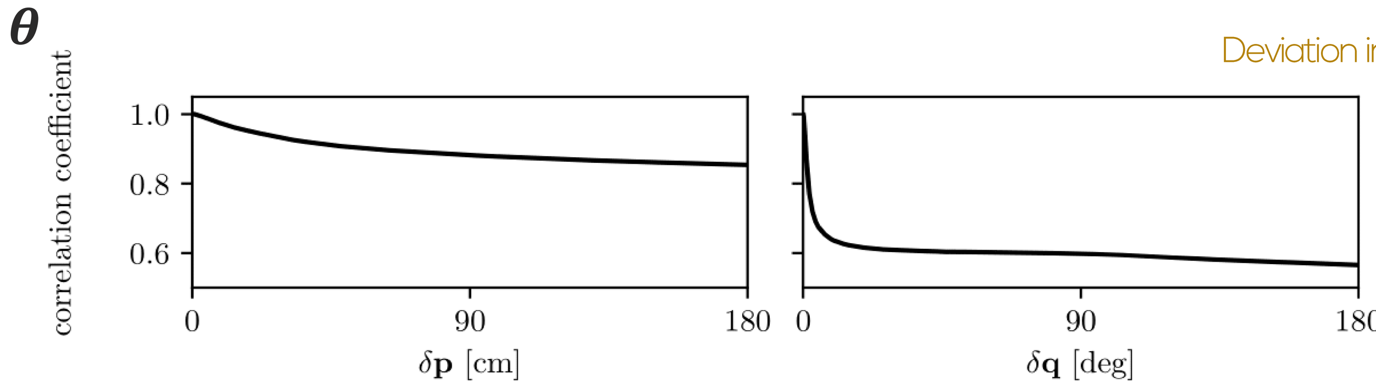

- Information is more sensitive to $\mathbf{q}$ than to $\mathbf{p}$
  → Good initial guesses for $\mathbf{q}$ is needed for successful measurement subset selection

RESULTS

- We compared the multi-IMU extrinsic self-calibration in three different calibration modes, each using distinct datasets:
  - 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

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
$\mathbf{p}_{\text{ref}}$: $[100, 100, 0] \pm [25.0, 25.0, 6.3]$ mm, $\mathbf{q}_{\text{ref}}$, ${}^{g}_{I}\mathbf{q}_{\text{ref}}$: $(\mathbf{e}, 0°) \pm (\mathbf{e}, 9.5°)$ for $\forall\, \mathbf{e} \in \mathbb{R}^{3} \setminus \{\mathbf{0}\}$ (angle-axis)
Experiments were performed on an octa-core Intel i7-10700 CPU at 2.90 GHz with 32 GB of RAM

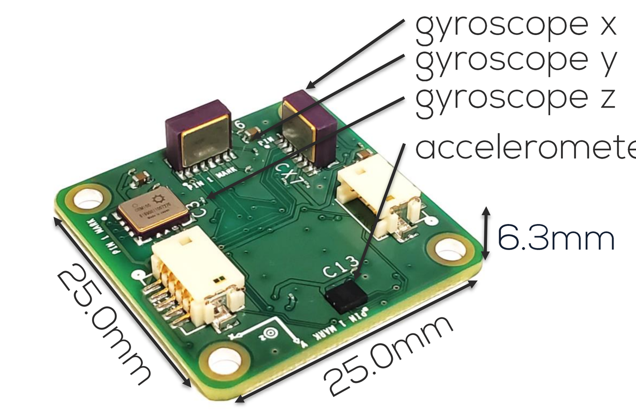

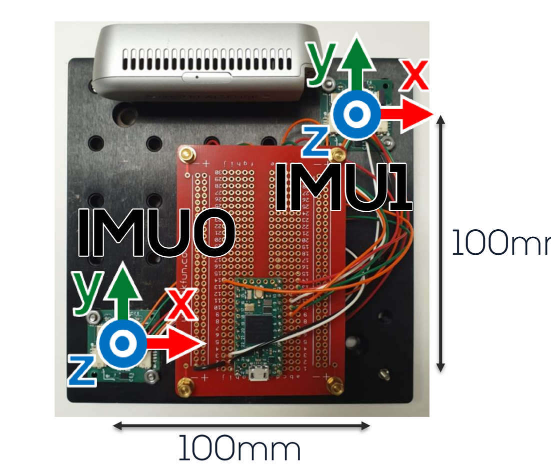

IMU (top) and sensor rig (bottom) used for the experiments

- Greedy (original) and Greedy (init-param) select and use less than 3% of the full set to achieve calibration results that align with reference values
- Greedy (init-param) significantly reduces runtime compared to the baseline (**>1 min → ~1 sec**), while Greedy (original) shows a modest reduction (>1 min → ~1 sec)

CONCLUSIONS AND FUTURE WORK

- We proposed a method for multi-IMU extrinsic calibration by efficiently selecting high-utility measurement subsets
- We hypothesized that in our system, utility – a function of parameter estimates – is largely insensitive to the specific choice of parameters; this allows for evaluation at an initial guess, reducing the need for frequent recalibrations
- To support this hypothesis, we conducted a sensitivity analysis of utility – in future, it may be helpful to provide further evidence by comparing the segments selected by Greedy (original) and Greedy (init-param)

REFERENCES

[1] J. Lee, D. Hanley, and T. Bretl, "Extrinsic calibration of multiple inertial sensors from arbitrary trajectories," IEEE Robot. Autom. Lett., vol. 7, no. 2, pp. 2055–2062, 2022.
[2] J. Maye, P. Furgale, and R. Siegwart, "Self-supervised calibration for robotic systems," in Proc. IEEE Intell. Veh. Symp., 2013, pp. 473–480.
[3] J. Maye, H. Sommer, G. Agamennoni, R. Siegwart, and P. Furgale, "Online self-calibration for robotic systems," Int. J. Robot. Res., vol. 35, no. 4, pp. 357–380, 2016.

ACKNOWLEDGEMENT

- This work was supported by NASA under Grant STTR-80NSSC20C0020.
- Travel funding was provided by Supernal, LLC, with additional support from the author's current internship at Google in approving the travel.

PAPER   LINKEDIN