2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

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

Jongwon Lee[1], David Hanley[2], and Timothy Bretl[1]

***Abstract*— This paper addresses the problem of choosing a sparse subset of measurements for quick calibration parameter estimation. A standard solution to this is selecting a measurement only if its utility—the difference between posterior (with the measurement) and prior information (without the measurement)—exceeds some threshold. Theoretically, utility, a function of the parameter estimate, should be evaluated at the estimate obtained with all measurements selected so far, hence necessitating a recalibration with each new measurement. However, we hypothesize that utility is insensitive to changes in the parameter estimate for many systems of interest, suggesting that evaluating utility at some initial parameter guess would yield equivalent results in practice. We provide evidence supporting this hypothesis for extrinsic calibration of multiple inertial measurement units (IMUs), showing the reduction in calibration time by two orders of magnitude by forgoing recalibration for each measurement.**

I. Introduction

Inertial measurement units (IMUs) are a common component in mobile robot navigation, comprising accelerometers and gyroscopes to measure specific forces and angular rates along three orthogonal axes. While single IMUs are commonly used, deploying multiple IMUs can enhance measurement accuracy, bandwidth, and fault tolerance without increasing the overall size, weight, power, and cost [1-6].

To realize these benefits, accurate *extrinsic calibration* is crucial to determine the relative pose of each IMU to others. Various extrinsic calibration methods for multi-IMU systems have been developed, including those utilizing prescribed trajectories [7-9] and aiding sensors like cameras [10]. Notably, *self*-calibration methods [11-13], which rely solely on IMU measurements themselves, offer significant benefits, especially in scenarios requiring recalibration due to intentional or unintentional changes in sensor configuration during robot operation, such as part loosening or thermal expansion.

Self-calibration, not bound to a predetermined trajectory, naturally experiences measurements with varying "utility". In cases like calibrating multi-IMU systems on non-rotating bodies such as spacecraft in orbit or cars on straight paths, much of the data can be low in utility. This is especially critical for optimization-based calibration, favored for its accuracy compared to filter-based methods, as it involves processing a significant amount of low-utility (uninformative) data alongside with high-utility (informative) data during optimization. Therefore, identifying subsets of measurements with high utility would offer benefits such as reduced runtime and memory usage for optimization. However, this remains beneficial only if the selection process is done efficiently.

In this study, we present an approach to multi-IMU extrinsic calibration by selecting informative measurement subsets with high utility, significantly streamlining our previous work [13]. This selection is primarily based on method from existing studies [14, 15] but is executed more efficiently. Specifically, we hypothesize that the measure of utility is insensitive to changes in the parameter estimate and equivalent calibration results are obtained by evaluating utility only at some initial guess of the calibration parameters. A valuable consequence of this is eliminating the need for recalibration with each new measurement, thereby reducing total computation time for calibration.

Our paper is structured as follows: Section II reviews existing literature on informative measurement selection for self-calibration. Section III describes an approach with our hypothesis for efficient measurement subset selection. Section IV delineates the solution approach for multi-IMU extrinsic self-calibration using these subsets, building on our prior work [13]. We then proceed to assess its runtime and accuracy against benchmarks in simulation (Section V) and using real-world data (Section VI). Section VII summarizes our findings, provides conclusions, and discuss the implications of this work.

II. Related Works

Selecting measurements with high utility for self-calibration involves optimizing the Fisher information matrix $\boldsymbol{\mathcal{I}}$, whose inverse sets the Cramer-Rao lower bound [16] for parameter estimation. Imagine a full measurement set $\mathcal{D}$ divided into L segments $\mathcal{D}^1, \ldots, \mathcal{D}^L$, with each segment $\mathcal{D}^l$ containing measurements across K consecutive timesteps. The goal is to identify a subset of $\mathcal{D}$

$$\begin{aligned} &\underset{\mathcal{D}^{info} \subset \mathcal{D},\, \mathrm{card}(\mathcal{D}^{info})=\kappa}{\arg\max} && f\left[\boldsymbol{\mathcal{I}}(\mathcal{D}^{info})\right] \\ &\text{subject to} && \kappa < \mathrm{card}(\mathcal{D}), \end{aligned} \qquad \text{(original)}$$

where f maps the input matrix to a scalar value (termed "information" in what follows), with $\mathrm{card}(\cdot)$ denoting set cardinality and κ a predetermined subset size. This combinatorial optimization problem, aimed at maximizing information, is typically addressed through two alternative ways in the existing literature. One approach involves choosing segments with the highest information:

[1]Jongwon Lee and Timothy Bretl are with the Department of Aerospace Engineering, University of Illinois Urbana-Champaign, Urbana, IL 61801, USA (Email: {jongwon5, tbretl}@illinois.edu).

[2]David Hanley is with the School of Informatics and the Translational Healthcare Technologies Group in Centre for Inflammation Research, Institute for Regeneration and Repair, University of Edinburgh, UK (Email: dhanley@ed.ac.uk).

This work was supported by the NASA Grant No. STTR-80NSSC20C0020.

1) Evaluate $f[\boldsymbol{\mathcal{I}}(\cdot)]$ for segments $\mathcal{D}^1$ through $\mathcal{D}^L$.
2) Choose M segments showing the largest $f[\boldsymbol{\mathcal{I}}(\cdot)]$.

(M-largest selection)

Schneider *et al.* [17] and Lv *et al.* [18] used this approach in camera-IMU and LiDAR-IMU calibration, respectively. The second approach uses a greedy algorithm, which selectively adds candidate segments $\mathcal{D}^{new}$ in the informative subset $\mathcal{D}^{info}$ if their utility exceeds threshold λ:

1) Initialize $\mathcal{D}^{info}$ as an empty set.
2) Iterate over candidate segments $\mathcal{D}^{new}$ from $\mathcal{D}^1$ to $\mathcal{D}^L$.
3) Add $\mathcal{D}^{new}$ in $\mathcal{D}^{info}$ if $f[\boldsymbol{\mathcal{I}}(\mathcal{D}^{info}, \mathcal{D}^{new})] - f[\boldsymbol{\mathcal{I}}(\mathcal{D}^{info})] > \lambda$.
4) Continue until all segments are processed.

(greedy algorithm)

Maye *et al.* initially applied this approach in rangefinder-odometer [14] and later in camera-IMU calibration [15].

Both approaches succeed in calibration only if they correctly identify informative subsets. We show in Section V-B an edge case where M-largest selection fails. Therefore, outside of that section, we refrain from discussing M-largest selection, focusing instead solely on greedy algorithm.

However, the usefulness of the greedy algorithm depends on the selection efficiency. We show that the original greedy algorithm by Maye *et al.* [14, 15] may not significantly improve or might even worsen runtime for our multi-IMU calibration scenario in concern. This led us to seek improvements detailed in Section III, where we apply our hypothesis for a more efficient execution of the greedy algorithm.

III. Efficient Identification of Informative Measurement Subsets for Self-Calibration

In this section, we first present the greedy algorithm by Maye *et al.* [14, 15] (Alg. 1) *as it is, without modifications.* Following that, we introduce our adaptation, which evaluates utility at initial calibration parameters (Alg. 2). Additionally, we provide a theoretical analysis showing its computational complexity reduction, expressed in Big O notation.

A. Original Greedy Algorithm

We provide details of Alg. 1 here. Upon receiving new measurements $\mathcal{D}^{new}$, the algorithm obtains the calibration parameter estimate $\widehat{\boldsymbol{\theta}}^{+}$ by minimizing the sum of squared residuals associated with both $\mathcal{D}^{info}$ and $\mathcal{D}^{new}$, starting from an initial estimate $\widehat{\boldsymbol{\theta}}^{-}$ (Line 6). Next, the Jacobian — stacks of derivatives of the residuals (see Section IV-E) with respect to the estimated parameters $\boldsymbol{\Theta}$ — is evaluated. The residuals corresponding to both $\mathcal{D}^{info}$ and $\mathcal{D}^{new}$, along with the current estimate $\widehat{\boldsymbol{\theta}}^{+}$, are used for evaluating the Jacobian. Then, the Fisher information matrix $\boldsymbol{\mathcal{I}}$ is computed (Line 7), and the block entries corresponding to $\boldsymbol{\Theta}$ from the inverse of $\boldsymbol{\mathcal{I}}$ are taken (Line 8). If the utility of $\mathcal{D}^{new}$, when added to $\mathcal{D}^{info}$, exceeds λ, then $\mathcal{D}^{new}$ is considered a high-utility segment (Line 9), prompting parameter updates (Lines 10-12). Note that this process iteratively refines calibration.

B. Greedy Algorithm with Utility Evaluation at Initial Calibration Parameters

Alg. 2 distinguishes itself from Alg. 1 by eliminating parameter recalibration within its loop, which, along with Fisher information matrix evaluation, is one of the primary bottlenecks in Alg. 1. This change was motivated by the need to improve efficiency without compromising the results. It is based on the hypothesis that utility evaluation for measurements is unaffected by the specific choice of $\boldsymbol{\theta}$, thus not significantly altering the choice of $\mathcal{D}^{info}$ or the calibration results.

Upon receiving new measurements $\mathcal{D}^{new}$, Alg. 2 evaluates the Jacobian *only* for $\mathcal{D}^{new}$ using uncalibrated parameters $\boldsymbol{\theta}^0$, computing the Fisher information matrix $\boldsymbol{\mathcal{I}}$ (Line 5). The repetitive calculation of $\mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}^T \mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}$ is avoided by reusing results from the previous iteration (Line 8). Then, the block entries corresponding to $\boldsymbol{\Theta}$ from the inverse of $\boldsymbol{\mathcal{I}}$ are taken (Line 6). If the utility of $\mathcal{D}^{new}$, when added to $\mathcal{D}^{info}$, exceeds λ, then $\mathcal{D}^{new}$ is considered a high-utility segment (Line 7), prompting parameter updates and retention of $\mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}^T \mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}$ for future iterations (Lines 8-9). Finally, with $\mathcal{D}^{info}$ identified, the calibration is conducted by solving a nonlinear least-squares problem (Line 12).

C. Comparison of Time Complexity

Consequently, Alg. 2 yields reduced time complexity compared to Alg. 1, as shown in Table I. Detailed derivations are provided in Appendices I and II, respectively.

Alg. 2 is notably effective when the high-utility subset selection $\mathcal{D}^{info}$ is not highly sensitive to variations in $\boldsymbol{\theta}$. To validate this, we conduct a sensitivity analysis on the impact of $\boldsymbol{\theta}$ deviations in measured information both in simulations (Section V-C) and hardware experiments (Section VI-B). We also show that calibration results between Alg. 1 and Alg. 2 are comparably insignificant in both the simulation (Section V) and hardware experiment sections (Section VI).

TABLE I: Big O Comparison: Original Greedy vs. Greedy with Utility Evaluation at Initial Calibration Parameters

	Original (Alg. 1)	Init-Param (Alg. 2)
Best	$\mathcal{O}(LKp^2)$	$\mathcal{O}(LKp^2)$
Worst	$\mathcal{O}(L^2Kp^2)$	$\mathcal{O}(LKp^2)$

** L: the number of total segments; K: the number of timesteps (i.e., data) in each segment; p: the number of parameters being estimated*

IV. Self-Calibration for Multiple IMUs

This section presents the solution approach for multi-IMU extrinsic self-calibration with sparse measurement subsets, beginning with a problem overview from Section IV-A to Section IV-C, which serves as a review of our previous work [13]. We detail a method for estimating each IMU's relative orientation in Section IV-D, crucial for Alg. 2's validity, as shown in Section V-C and VI-B. Section IV-E discusses residuals for evaluating the Fisher information matrix and

Algorithm 1 Pseudocode for calibration via the original greedy algorithm [14, 15]

Input: uncalibrated parameters $\boldsymbol{\theta}^0$,
Output: calibrated parameters $\boldsymbol{\theta}^*$
$\mathcal{D}^{info} \leftarrow \varnothing$
$\widehat{\boldsymbol{\theta}}^- \leftarrow \boldsymbol{\theta}^0$
for $new \leftarrow 1$ to L **do**
  $\widehat{\boldsymbol{\theta}}^+ \leftarrow \texttt{Calibrate}(\mathcal{D}^{info}, \mathcal{D}^{new}; \widehat{\boldsymbol{\theta}}^-)$
  $\boldsymbol{\mathcal{I}}(\mathcal{D}^{info}, \mathcal{D}^{new}) \leftarrow \mathbf{J}(\mathcal{D}^{info})|_{\widehat{\boldsymbol{\theta}}^+}^T \mathbf{J}(\mathcal{D}^{info})|_{\widehat{\boldsymbol{\theta}}^+} + \mathbf{J}(\mathcal{D}^{new})|_{\widehat{\boldsymbol{\theta}}^+}^T \mathbf{J}(\mathcal{D}^{new})|_{\widehat{\boldsymbol{\theta}}^+}$
  $\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}} \leftarrow \left[\boldsymbol{\mathcal{I}}(\mathcal{D}^{info}, \mathcal{D}^{new})^{-1}\right]_{\boldsymbol{\Theta}\boldsymbol{\Theta}}$
  if $\frac{1}{2}\log\frac{|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info}}|}{|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}}|} > \lambda$ or $\mathcal{D}^{info} == \varnothing$ **then**
    $\mathcal{D}^{info} \leftarrow \mathcal{D}^{info} \cup \mathcal{D}^{new}$
    $|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info}}| \leftarrow |\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}}|$
    $\widehat{\boldsymbol{\theta}}^- \leftarrow \widehat{\boldsymbol{\theta}}^+$
  end if
end for
$\boldsymbol{\theta}^* \leftarrow \widehat{\boldsymbol{\theta}}^-$
return $\boldsymbol{\theta}^*$

Algorithm 2 Pseudocode for calibration via greedy algorithm with utility evaluation at initial calibration parameters

Input: uncalibrated parameters $\boldsymbol{\theta}^0$,
Output: calibrated parameters $\boldsymbol{\theta}^*$
$\mathcal{D}^{info} \leftarrow \varnothing$
for $new \leftarrow 1$ to L **do**
  $\boldsymbol{\mathcal{I}}(\mathcal{D}^{info}, \mathcal{D}^{new}) \leftarrow \mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}^T \mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0} + \mathbf{J}(\mathcal{D}^{new})|_{\boldsymbol{\theta}^0}^T \mathbf{J}(\mathcal{D}^{new})|_{\boldsymbol{\theta}^0}$
  $\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}} \leftarrow \left[\boldsymbol{\mathcal{I}}(\mathcal{D}^{info}, \mathcal{D}^{new})^{-1}\right]_{\boldsymbol{\Theta}\boldsymbol{\Theta}}$
  if $\frac{1}{2}\log\frac{|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info}}|}{|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}}|} > \lambda$ or $\mathcal{D}^{info} == \varnothing$ **then**
    $\mathcal{D}^{info} \leftarrow \mathcal{D}^{info} \cup \mathcal{D}^{new}$ (Keep $\mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}^T \mathbf{J}(\mathcal{D}^{info})|_{\boldsymbol{\theta}^0}$ as well)
    $|\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info}}| \leftarrow |\boldsymbol{\Sigma}_{\boldsymbol{\Theta}|\mathcal{D}^{info},\mathcal{D}^{new}}|$
  end if
end for
$\boldsymbol{\theta}^* \leftarrow \texttt{Calibrate}(\mathcal{D}^{info}; \boldsymbol{\theta}^0)$
return $\boldsymbol{\theta}^*$

formulating the calibration's nonlinear least-squares problem. Section IV-F introduces additional residuals for constraining the relationship between the sparse measurement subsets, followed by the calibration solution approach.

A. Notation

We denote the rotation matrix that represents frame $\mathcal{F}_A$'s orientation in frame $\mathcal{F}_B$'s coordinates by ${}^B_A\mathbf{R} \in SO(3)$. The operator $\mathbf{C}(\cdot)$ converts a unit quaternion ${}^B_A\mathbf{q}$ to its corresponding rotation matrix ${}^B_A\mathbf{R}$. Coordinate transformations are thus defined:

$$ {}^B\mathbf{p} = {}^B_A\mathbf{R}\,{}^A\mathbf{p} = \mathbf{C}\left({}^B_A\mathbf{q}\right){}^A\mathbf{p}. $$

We distinguish quantities using these symbols:

- a tilde $\widetilde{(\cdot)}$ denotes a sensor measurement,
- a hat $\widehat{(\cdot)}$ denotes an estimate,
- a zero superscript $(\cdot)^0$ denotes an initial guess (pre-optimization),
- and an asterisk superscript $(\cdot)^*$ denotes a final value (post-optimization).

B. Inertial Sensor Model

We model accelerometer measurements from each IMU as

$$ {}^I\widetilde{\mathbf{a}} = {}^I\mathbf{a}_{WI} - {}^I\mathbf{g} + \mathbf{b}_a + \mathbf{n}_a \tag{1} $$

where

- ${}^I\mathbf{a}_{WI}$ is the IMU's linear acceleration relative to the world frame, expressed in the IMU frame,
- ${}^I\mathbf{g}$ is gravitational acceleration in the IMU frame,
- $\mathbf{b}_a$ is a time-varying bias modeled by a random walk
$$ \mathbf{b}_{a,k+1} - \mathbf{b}_{a,k} \sim \sigma_{\mathbf{b}_a}\sqrt{\Delta t}\cdot\mathcal{N}\left(\mathbf{0},\,\mathbf{1}\right), $$
- and $\mathbf{n}_a$ is stochastic noise
$$ \mathbf{n}_a \sim \sigma_a/\sqrt{\Delta t}\cdot\mathcal{N}\left(\mathbf{0},\,\mathbf{1}\right) $$

with Δt as the IMU's sampling interval. We model gyroscope measurements from each IMU as

$$ {}^g\widetilde{\boldsymbol{\omega}} = \mathbf{C}\left({}^g_I\mathbf{q}\right){}^I\boldsymbol{\omega}_{WI} + \mathbf{b}_g + \mathbf{n}_g \tag{2} $$

where

- $\mathbf{C}\left({}^g_I\mathbf{q}\right)$ is the rotation matrix (written in terms of the corresponding quaternion) that models the IMU frame's orientation relative to the gyroscope frame, accounting for gyroscope misalignment,
- ${}^I\boldsymbol{\omega}_{WI}$ is the IMU's angular velocity relative to the world frame, expressed in the IMU frame,
- $\mathbf{b}_g$ is a time-varying bias modeled a random walk
$$ \mathbf{b}_{g,k+1} - \mathbf{b}_{g,k} \sim \sigma_{\mathbf{b}_g}\sqrt{\Delta t}\cdot\mathcal{N}\left(\mathbf{0},\,\mathbf{1}\right), $$
- and $\mathbf{n}_g$ is stochastic noise
$$ \mathbf{n}_g \sim \sigma_g/\sqrt{\Delta t}\cdot\mathcal{N}\left(\mathbf{0},\,\mathbf{1}\right). $$

We assume that noise densities σ_a, σ_g and bias instabilities $\sigma_{\mathbf{b}_a}$, $\sigma_{\mathbf{b}_g}$ in continuous-time are identified for all sensors before extrinsic calibration. Additionally, we assume intrinsic parameters like scale factors and axis non-orthogonality are calibrated, and all IMUs are temporally synchronized.

C. Problem Statement (Multi-IMU Extrinsic Calibration)

Supposing $N+1$ IMUs whose frames are indexed as $I_0, \ldots, I_N$, we define:

- ${}^{I_n}\mathbf{p}_{I_nI_0} := \mathbf{p}_{I_n}$ for the frame I_0's position relative to frame I_n, expressed in frame I_n,
- ${}^{I_n}_{I_0}\mathbf{q} := \mathbf{q}_{I_n}$ for the quaternion describing frame I_0's orientation relative to frame I_n, assuming IMU frames are accelerometer-aligned,
- and ${}^{I_n}_{g_n}\mathbf{q} := \mathbf{q}_{g_n}$ for the quaternion describing frame g_n's orientation relative to frame I_n (gyroscope misalignment).

Our calibration aims to estimate extrinsic parameters $\mathbf{p}_{I_n}$ and $\mathbf{q}_{I_n}$ for each IMU $n \in \{1, \ldots, N\}$ and $\mathbf{q}_{g_n}$ for each IMU $n \in \{0, \ldots, N\}$, given measurements ${}^{I_n}\widetilde{\mathbf{a}}_k$, ${}^{g_n}\widetilde{\boldsymbol{\omega}}_k$ at each time k for each IMU $n \in \{0, \ldots, N\}$. Additionally, estimating time-varying biases $\mathbf{b}_{a_n,k}$, $\mathbf{b}_{g_n,k}$ at each time k for each IMU $n \in \{0, \ldots, N\}$, as well as the angular acceleration of the base IMU ${}^{I_0}\boldsymbol{\alpha}_{I_0,k} := {}^{I_0}\boldsymbol{\alpha}_k$ at each time

k, proves useful in refining sensor parameter estimates in our previous study [13].

D. Relative Orientation Initialization

Prior to subset selection in Alg. 2, we initialize each IMU's relative orientation to the base IMU, ${}^{I_0}_{I_n}\mathbf{q}$, using Yang *et al.*'s method [20]. This preliminary step enhances Alg. 2's validity by providing plausible IMU orientation guesses.

Assuming angular velocity $\boldsymbol{\omega}$ remains constant between t_k and t_{k+1}, with $\Delta t = t_{k+1} - t_k$, the orientation change in the world frame expressed in the IMU frame is

$$ {}^{I}_{W}\mathbf{q}(t_{k+1}) = \begin{bmatrix} \frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|}\sin\left(\frac{|\boldsymbol{\omega}|}{2}\Delta t\right) \\ \cos\left(\frac{|\boldsymbol{\omega}|}{2}\Delta t\right) \end{bmatrix} \otimes {}^{I}_{W}\mathbf{q}(t_k). $$

Hence, the rotation from $\mathcal{F}_{I(t_k)}$ to $\mathcal{F}_{I(t_{k+1})}$ is

$$ {}^{I(t_{k+1})}_{I(t_k)}\mathbf{q} = \begin{bmatrix} \frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|}\sin\left(\frac{|\boldsymbol{\omega}|}{2}\Delta t\right) \\ \cos\left(\frac{|\boldsymbol{\omega}|}{2}\Delta t\right) \end{bmatrix}. $$

Supposing that we have ${}^{I_0(t_{k+1})}_{I_0(t_k)}\mathbf{q}$ and ${}^{I_n(t_{k+1})}_{I_n(t_k)}\mathbf{q}$ from ${}^{I_0}\widetilde{\boldsymbol{\omega}}_k$ and ${}^{I_n}\widetilde{\boldsymbol{\omega}}_k$ respectively, the following relationship holds:

$$ {}^{I_0(t_{k+1})}_{I_0(t_k)}\mathbf{q} \otimes {}^{I_0}_{I_n}\mathbf{q} = {}^{I_0}_{I_n}\mathbf{q} \otimes {}^{I_n(t_{k+1})}_{I_n(t_k)}\mathbf{q}. $$

Or equivalently,

$$ \mathcal{L}({}^{I_0(t_{k+1})}_{I_0(t_k)}\mathbf{q}) \cdot {}^{I_0}_{I_n}\mathbf{q} = \mathcal{R}({}^{I_n(t_{k+1})}_{I_n(t_k)}\mathbf{q}) \cdot {}^{I_0}_{I_n}\mathbf{q} $$

where

$$ \mathcal{L}(\mathbf{q}) = \begin{bmatrix} q_w\mathbf{I}_3 + \lfloor \mathbf{q}_{xyz} \rfloor & \mathbf{q}_{xyz} \\ -\mathbf{q}_{xyz}^T & q_w \end{bmatrix}, $$

$$ \mathcal{R}(\mathbf{q}) = \begin{bmatrix} q_w\mathbf{I}_3 - \lfloor \mathbf{q}_{xyz} \rfloor & \mathbf{q}_{xyz} \\ -\mathbf{q}_{xyz}^T & q_w \end{bmatrix}. $$

Rearranging the equation yields

$$ \left(\mathcal{L}({}^{I_0(t_{k+1})}_{I_0(t_k)}\mathbf{q}) - \mathcal{R}({}^{I_n(t_{k+1})}_{I_n(t_k)}\mathbf{q})\right) \cdot {}^{I_0}_{I_n}\mathbf{q} = \mathbf{0}_{4\times 1}. $$

Stacking the difference of 4-by-4 matrices $\mathcal{L} - \mathcal{R}$ for various timesteps k creates a large matrix $\mathbf{A} \in \mathbb{R}^{4k\times 4}$, leading to $\mathbf{A} \cdot {}^{I_0}_{I_n}\mathbf{q} = \mathbf{0}_{4\times 1}$. We compute ${}^{I_0}_{I_n}\mathbf{q}$ by finding the right-singular vector of $\mathbf{A}$ with the smallest singular value.

E. Intra-Segment Residuals

For each segment $\mathcal{D}^{new}$ with $N+1$ IMU measurements over K timesteps, we define four types of residuals, $\mathbf{r}_a$, $\mathbf{r}_g$, $\mathbf{r}_{\mathbf{b}_a}$, and $\mathbf{r}_{\mathbf{b}_g}$. Residual $\mathbf{r}_a$ relates the nth accelerometer's measurement (${}^{I_n}\widetilde{\mathbf{a}}_k$) to that of base accelerometer's (${}^{I_0}\widetilde{\mathbf{a}}_k$) corrected for bias $\mathbf{b}_{a_0,k}$:

$$ \mathbf{r}_a = {}^{I_0}\widehat{\mathbf{a}}_k - ({}^{I_0}\widetilde{\mathbf{a}}_k - \mathbf{b}_{a_0,k}). $$

${}^{I_0}\widehat{\mathbf{a}}_k$ is derived by transforming ${}^{I_n}\widetilde{\mathbf{a}}_k$ to the base IMU's frame:

$$ {}^{I_0}\widehat{\mathbf{a}}_k = {}^{I_n}_{I_0}\mathbf{R}^{-1}\Big\{({}^{I_n}\widetilde{\mathbf{a}}_k - \mathbf{b}_{a_n,k}) + \lfloor {}^{I_n}_{g_n}\mathbf{R}({}^{g_n}\widetilde{\boldsymbol{\omega}}_k - \mathbf{b}_{g_n,k})\rfloor^2 \mathbf{p}_{I_n} + \lfloor {}^{I_n}_{I_0}\mathbf{R}\,{}^{I_0}\boldsymbol{\alpha}_k\rfloor \mathbf{p}_{I_n}\Big\}, $$

where ${}^{I_n}_{I_0}\mathbf{R} = \mathbf{C}(\mathbf{q}_{I_n})$, ${}^{I_n}_{g_n}\mathbf{R} = \mathbf{C}(\mathbf{q}_{g_n})$, and $\lfloor\cdot\rfloor$ mapping a cross product into skew-symmetric matrix form. Residual $\mathbf{r}_g$ relates the nth gyroscope's measurement (${}^{I_n}\widetilde{\boldsymbol{\omega}}_k$) to that of base gyroscope's (${}^{I_0}\widetilde{\boldsymbol{\omega}}_k$) corrected for biases $\mathbf{b}_{g_n,k}$, $\mathbf{b}_{g_0,k}$:

$$ \mathbf{r}_g = {}^{I_n}_{I_0}\mathbf{R}^{-1}\,{}^{I_n}_{g_n}\mathbf{R}({}^{g_n}\widetilde{\boldsymbol{\omega}}_k - \mathbf{b}_{g_n,k}) - {}^{I_0}_{g_0}\mathbf{R}({}^{g_0}\widetilde{\boldsymbol{\omega}}_k - \mathbf{b}_{g_0,k}), $$

where ${}^{I_n}_{I_0}\mathbf{R} = \mathbf{C}(\mathbf{q}_{I_n})$, ${}^{I_n}_{g_n}\mathbf{R} = \mathbf{C}(\mathbf{q}_{g_n})$, and ${}^{I_0}_{g_0}\mathbf{R} = \mathbf{C}(\mathbf{q}_{g_0})$. This relation holds as rotational motion is uniform across all points of a rigid body. Sensor measurements corrected for bias introduce uncertainty due to zero-mean Gaussian white noise, as shown by $\widetilde{\mathbf{a}} - \mathbf{b}_a = \mathbf{n}_a$ from Eq. (1) and $\widetilde{\boldsymbol{\omega}} - \mathbf{b}_g = \mathbf{n}_g$ from Eq. (2), leading to accelerometer and gyroscope covariance matrices $\boldsymbol{\Sigma}_{\widetilde{\mathbf{a}}} = ({\sigma_a}^2/\Delta t)\cdot\mathbf{I}_{3\times 3}$ and $\boldsymbol{\Sigma}_{\widetilde{\boldsymbol{\omega}}} = ({\sigma_g}^2/\Delta t)\cdot\mathbf{I}_{3\times 3}$, respectively. Additional uncertainties from ${}^{I_0}\widetilde{\boldsymbol{\omega}}_k^2$ and ${}^{I_0}\boldsymbol{\alpha}_k$ introduced by ${}^{I_0}\widehat{\mathbf{a}}_k$ lead to covariance matrices $\boldsymbol{\Sigma}_{\widetilde{\boldsymbol{\omega}}^2} = 2\cdot\boldsymbol{\Sigma}^2_{\widetilde{\boldsymbol{\omega}}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\alpha}} = \sigma^2_{\boldsymbol{\alpha}}\cdot\mathbf{I}_{3\times 3}$, respectively. Consequently, the covariance matrix for $\mathbf{r}_a$ is $\boldsymbol{\Sigma}_{\mathbf{r}_a} = 2\cdot\boldsymbol{\Sigma}_{\widetilde{\mathbf{a}}} + 2\cdot\boldsymbol{\Sigma}^2_{\widetilde{\boldsymbol{\omega}}} + \boldsymbol{\Sigma}_{\boldsymbol{\alpha}}$ and that for $\mathbf{r}_g$ is $\mathbf{r}_g = 2\cdot\boldsymbol{\Sigma}_{\widetilde{\boldsymbol{\omega}}}$. Residuals $\mathbf{r}_{\mathbf{b}_a}$ and $\mathbf{r}_{\mathbf{b}_g}$ account for bias evolution:

$$ \mathbf{r}_{\mathbf{b}_a} = \mathbf{b}_{a_n,k+1} - \mathbf{b}_{a_n,k} $$

$$ \mathbf{r}_{\mathbf{b}_g} = \mathbf{b}_{g_n,k+1} - \mathbf{b}_{g_n,k}. $$

The covariance matrices for these residuals are $\boldsymbol{\Sigma}_{\mathbf{b}_a} = {\sigma_{\mathbf{b}_a}}^2\Delta t\cdot\mathbf{I}_{3\times 3}$ and $\boldsymbol{\Sigma}_{\mathbf{b}_g} = {\sigma_{\mathbf{b}_g}}^2\Delta t\cdot\mathbf{I}_{3\times 3}$, respectively.

F. Solution Approach

Upon identifying high-utility measurements, we formulate a nonlinear least-square problem using the chosen measurement subset. Beyond the previously mentioned residuals in Section IV-E, we incorporate additional residuals, $\mathbf{r}^*_{\mathbf{b}_a}$ and $\mathbf{r}^*_{\mathbf{b}_g}$, to bridge temporal gaps between adjacent selected segments, denoted as $\mathcal{D}^{info} = \{\mathcal{D}^m | m \in \mathcal{M}\}$ with $\mathcal{M} = \{\mathcal{M}_1, \ldots, \mathcal{M}_M\} \subseteq \{1, \ldots, L\}$:

$$ \mathbf{r}^*_{\mathbf{b}_a} = \mathbf{b}_{a_n,(\mathcal{M}_{m+1}-1)\cdot K+1} - \mathbf{b}_{a_n,\mathcal{M}_m\cdot K} $$

$$ \mathbf{r}^*_{\mathbf{b}_g} = \mathbf{b}_{g_n,(\mathcal{M}_{m+1}-1)\cdot K+1} - \mathbf{b}_{g_n,\mathcal{M}_m\cdot K}. $$

These additional residuals aim to maintain bias consistency across segments, with their covariance matrices $\boldsymbol{\Sigma}^*_{\mathbf{b}_a}$ and $\boldsymbol{\Sigma}^*_{\mathbf{b}_g}$ modeled after a random walk in discrete time as earlier discussed, resulting in $\boldsymbol{\Sigma}^*_{\mathbf{b}_a} = {\sigma_{\mathbf{b}_a}}^2\cdot(\mathcal{M}_{m+1} - \mathcal{M}_m)K\cdot\Delta t\cdot\mathbf{I}_{3\times 3}$ and $\boldsymbol{\Sigma}^*_{\mathbf{b}_g} = {\sigma_{\mathbf{b}_g}}^2\cdot(\mathcal{M}_{m+1} - \mathcal{M}_m)K\cdot\Delta t\cdot\mathbf{I}_{3\times 3}$.

These residuals allow for parameter estimation through solving a nonlinear least-squares problem

$$ \min\Bigg\{\sum_{l\in\mathcal{M}}\Bigg(\sum_{\substack{n\in\{1,\ldots,N\}\\ k\in\{1,\ldots,K\}}}\left(\|\mathbf{r}_a\|^2_{\boldsymbol{\Sigma}_{\mathbf{r}_a}} + \|\mathbf{r}_g\|^2_{\boldsymbol{\Sigma}_{\mathbf{r}_g}}\right) + \sum_{\substack{n\in\{0,\ldots,N\}\\ k\in\{1,\ldots,K-1\}}}\left(\|\mathbf{r}_{\mathbf{b}_a}\|^2_{\boldsymbol{\Sigma}_{\mathbf{r}_{\mathbf{b}_a}}} + \|\mathbf{r}_{\mathbf{b}_g}\|^2_{\boldsymbol{\Sigma}_{\mathbf{r}_{\mathbf{b}_g}}}\right)\Bigg) + \sum_{\substack{m\in\{1,\ldots,M-1\}\\ n\in\{0,\ldots,N\}}}\left(\|\mathbf{r}^*_{\mathbf{b}_a}\|^2_{\boldsymbol{\Sigma}^*_{\mathbf{r}_{\mathbf{b}_a}}} + \|\mathbf{r}^*_{\mathbf{b}_g}\|^2_{\boldsymbol{\Sigma}^*_{\mathbf{r}_{\mathbf{b}_g}}}\right)\Bigg\}, \tag{3} $$

where $\|\cdot\|^2_{\boldsymbol{\Sigma}}$ represents the Mahalanobis distance using covariance matrix $\boldsymbol{\Sigma}$.

V. Simulation Experiments

In this section, we assess multi-IMU extrinsic self-calibration using simulated trajectories from OpenVINS [21]. First, in Section V-B, we present an edge case highlighting the need for adopting the greedy algorithm for subset selection over M-largest selection, also noting the original greedy algorithm (Alg. 1)'s efficiency issue. A sensitivity analysis in Section V-C examines the impact of changes in parameter estimates on the measure of information—a function of the parameter estimates. This analysis supports the use of Alg. 2 as an alternative to Alg. 1. Section V-D compares the calibration performances of Alg. 1 (*Greedy (Original)*) and Alg. 2 (*Greedy (Init-Param)*) against the *baseline*, which uses the full set of measurements for calibration, showing Alg. 2 significantly reduces runtime—from minutes to roughly a quarter minute—without sacrificing accuracy.

A. Implementation Details

TABLE II: Reference IMU poses in simulation

index	position [cm]	orientation [deg]*
IMU0	[0, 0, 0]	[0, 0, 0]
IMU1	[20, 0, 0]	[180, 0, 0]
IMU2	[0, 20, 0]	[0, 180, 0]
IMU3	[0, 0, 20]	[0, 0, 180]

** in this table, orientation is given as XYZ Euler angles*

Table II shows the reference poses for four IMUs mounted on a rigid body. Accelerometer noise and bias instability were set at $\sigma_a = 2 \times 10^{-3}$ m/s$^2/\sqrt{\text{Hz}}$ and $\sigma_{\mathbf{b}_a} = 3 \times 10^{-3}$ m/s$^2 \cdot \sqrt{\text{Hz}}$, respectively, while for gyroscopes, these values were $\sigma_g = 1.6968 \times 10^{-4}$ rad/s$/\sqrt{\text{Hz}}$ and $\sigma_{\mathbf{b}_g} = 1.9393 \times 10^{-5}$ rad/s $\cdot \sqrt{\text{Hz}}$. Gyroscope misalignment was generated by rotating each IMU's reference orientation about a uniformly random axis by an angle from a zero-mean normal distribution with a standard deviation of 1°. IMU measurements were generated at 100 Hz.

B. Limitations of Existing Methods for Self-Calibration with Informative Subset Selection

We assess multi-IMU extrinsic self-calibration using three measurement subsets: the full set ("baseline"), and subsets selected by M-largest selection and the greedy algorithm. We simulated a trajectory with sinusoidal oscillations in XYZ Euler angles every 20 seconds, maintaining consistent information for each interval:

- For $0 \leq t < 20$ seconds: $[\sin(1\pi t), 0, 0]$ [rad].
- For $20 \leq t < 40$ seconds: $[0, \sin(2\pi t), 0]$ [rad].
- For $40 \leq t < 60$ seconds: $[0, 0, \sin(3\pi t)]$ [rad].

For the greedy algorithm, the utility threshold λ was set to 0.5, following previous studies [14, 15]. For M-largest selection, we chose 15 segments, ensuring a similar subset size to the greedy algorithm. Initial parameter estimates $\mathbf{p}_{I_n}$, $\mathbf{q}_{I_n}$, and $\mathbf{q}_{g_n}$ were set to zero, with angular acceleration derived from numerical differentiation of IMU 0's gyroscope measurements and time-varying biases starting at zero.

Table III presents the calibration error and runtime for each process, and Fig. 1 illustrates the chosen segments. M-largest selection shows significantly higher calibration errors than both the baseline and greedy algorithm. While greedy algorithm selects new segments with utility (notably around 20s and 40s), the failure of M-largest selection in this regard likely results in poorer estimates. Despite lower errors, greedy algorithm's runtime is longer than M-largest selection and the baseline, highlighting the need for the efficient subset selection as introduced in Section III.

C. Sensitivity Analysis of Fisher Information Matrix

The underlying assumption of Alg. 2 is the Fisher information matrix ($\boldsymbol{\mathcal{I}} = \mathbf{J}^T|_{\boldsymbol{\theta}}\mathbf{J}|_{\boldsymbol{\theta}}$) — a function of the estimate for the parameters $\boldsymbol{\theta}$ — is insensitive to the parameter estimate. To investigate this, we examined the Spearman's correlation coefficient [22] between the distribution of information $-1/2 \cdot \log\left|\left[\boldsymbol{\mathcal{I}}^{-1}\right]_{\boldsymbol{\Theta}\boldsymbol{\Theta}}\right|$ at reference extrinsic parameters (*original distribution*) and at values altered by deviations in either position ($\delta\mathbf{p}$) or orientation ($\delta\mathbf{q}$) (*perturbed distribution*). The study utilized simulations across 48 different trajectories from five datasets provided by OpenVINS [21].

Fig. 2 shows changes in the correlation coefficient are smaller for positional deviations ($\delta\mathbf{p}$) but larger for orientational deviations ($\delta\mathbf{q}$), even with a few degrees. This indicates information sensitivity mainly lies with orientation guesses rather than position. Hence, ensuring accurate initial orientation guesses, as detailed in Section IV-D, is key for effective segment selection by Alg. 2.

D. Comparison of Self-Calibration with Greedy Algorithm against Benchmarks

We assess the self-calibration with subset selection in Alg. 2 ("Greedy (Init-Param)") against Alg. 1 ("Greedy (Original)") and a "Baseline" using all measurements. Our focus lies on runtime and calibration performance comparisons across 48 trajectories within five OpenVINS-supported datasets. For both the greedy methods, information evaluation occurred every second, using a $\lambda = 0.5$ threshold, with initial parameter guesses aligned with the earlier discussion.

Table IV presents the chosen measurement subset ratio and runtime for each calibration process, with Fig. 3 showing the absolute extrinsic parameter errors for all IMUs. The runtime breakdown and memory usage are presented in Tables V and VI, respectively. Greedy (Init-Param) consistently reduced runtime, with significant improvements in the `kaist` dataset, while Greedy (Original) showed less or negative improvement. Calibration accuracy remained consistent both for greedy methods and the baseline, achieving sub-centimeter and sub-degree precision, except in the `kaist` dataset. Greedy methods particularly outperformed the baseline in the `kaist`, which was collected in a driving car and primarily consists of low-utility measurements, by excluding these low-utility measurements to improve problem observability. The subset selection addresses the challenges of unobservable or weakly observable calibration scenarios that can arise in self-calibration with full measurements.

TABLE III: ERROR IN ESTIMATED EXTRINSIC PARAMETERS AND RUNTIME FOR BASELINE, M-LARGEST SELECTION, AND GREEDY ALGORITHM

Baseline				M-largest Selection				Greedy Algorithm			
$\mathbf{p}$ [cm]	$\mathbf{q}$ [deg]	$\mathbf{q}_g$ [deg]	t [s]	$\mathbf{p}$ [cm]	$\mathbf{q}$ [deg]	$\mathbf{q}_g$ [deg]	t [s]	$\mathbf{p}$ [cm]	$\mathbf{q}$ [deg]	$\mathbf{q}_g$ [deg]	t [s]
0.0151	0.3685	2.1002	7.58	19.8605	1.4150	62.3996	1.91	0.0229	0.3090	0.8152	8.04

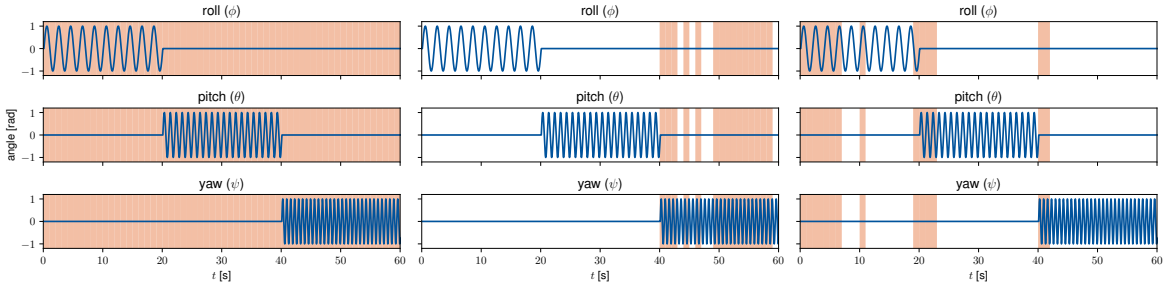

Fig. 1: The intervals of subset within a roll-pitch-yaw profile, which is used by the calibration process for the baseline (left), M-largest selection (middle), and greedy algorithm (right), are highlighted. The background color highlights the segments selected and used in each calibration process.

TABLE IV: RATIO OF SELECTED SEGMENTS AND RUNTIME FOR BASELINE, ORIGINAL GREEDY, AND GREEDY WITH UTILITY EVALUATION AT INITIAL CALIBRATION PARAMETERS OVER EACH DATASET

Dataset	# of traj.	Length [m]	Duration [s]	Ratio of selected segments [%]			Runtime [s]		
				Baseline	Greedy (Original)	Greedy (Init-Param)	Baseline	Greedy (Original)	Greedy (Init-Param)
uzh_fpv	15	210.02 ± 144.69	39.03 ± 17.68	100.00 ± 0.00	59.30 ± 17.37	66.15 ± 21.01	5.98 ± 3.46	7.12 ± 4.44	**3.41 ± 0.94**
euroc_mav	11	81.22 ± 23.44	122.58 ± 27.58	100.00 ± 0.00	26.06 ± 8.43	24.88 ± 7.41	18.21 ± 4.78	23.16 ± 6.22	**5.72 ± 1.16**
tum_vi	6	115.25 ± 36.98	135.10 ± 12.52	100.00 ± 0.00	25.94 ± 7.03	26.33 ± 5.11	19.82 ± 1.99	32.45 ± 10.06	**5.72 ± 1.06**
kaist_vio	11	37.10 ± 16.54	179.44 ± 30.00	100.00 ± 0.00	17.76 ± 9.74	17.66 ± 6.39	25.28 ± 11.88	30.91 ± 19.29	**5.32 ± 1.79**
kaist	5	9618.87 ± 1958.50	1533.96 ± 656.45	100.00 ± 0.00	3.87 ± 2.26	3.44 ± 2.11	379.42 ± 135.80	460.81 ± 211.07	**20.08 ± 6.60**

TABLE V: RUNTIME BREAKDOWN FOR BASELINE, ORIGINAL GREEDY, AND GREEDY WITH UTILITY EVALUATION AT INITIAL CALIBRATION PARAMETERS OVER EACH DATASET

Dataset	Baseline			Greedy (Original)			Greedy (Init-Param)		
	Evaluate	Calibrate	Total	Evaluate	Calibrate	Total	Evaluate	Calibrate	Total
uzh_fpv	-	5.94 ± 3.43	5.98 ± 3.46	2.52 ± 2.05	4.50 ± 2.39	7.12 ± 4.44	0.15 ± 0.07	3.19 ± 0.86	3.41 ± 0.94
euroc_mav	-	18.09 ± 4.76	18.21 ± 4.78	13.84 ± 4.11	8.91 ± 2.94	23.16 ± 6.22	0.46 ± 0.13	5.12 ± 1.19	5.72 ± 1.16
tum_vi	-	19.69 ± 1.98	19.82 ± 1.99	18.50 ± 6.46	13.48 ± 3.89	32.45 ± 10.06	0.54 ± 0.06	5.04 ± 1.06	5.72 ± 1.06
kaist_vio	-	25.13 ± 11.82	25.28 ± 11.88	21.11 ± 14.18	9.27 ± 5.43	30.91 ± 19.29	0.58 ± 0.26	4.56 ± 1.51	5.32 ± 1.79
kaist	-	377.85 ± 135.22	379.42 ± 135.80	371.62 ± 172.61	81.47 ± 41.68	460.81 ± 211.07	4.61 ± 1.98	13.46 ± 5.33	20.08 ± 6.60

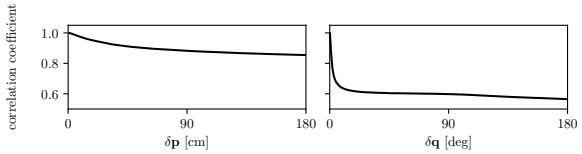

Fig. 2: The correlation coefficient of the information as a function of the difference ($\delta\mathbf{p}$, $\delta\mathbf{q}$) in guessed and reference extrinsic parameters in the multi-IMU system. A coefficient of 1 indicates perfect correlation with the true value, while 0 denotes independence.

TABLE VI: MEMORY FOOTPRINT (RAM USAGE IN MB) FOR BASELINE, ORIGINAL GREEDY, AND GREEDY WITH UTILITY EVALUATION AT INITIAL CALIBRATION PARAMETERS OVER EACH DATASET

Dataset	Baseline	Greedy (Original)	Greedy (Init-Param)
uzh_fpv	92.69 ± 40.62	55.38 ± 14.88	65.13 ± 32.34
euroc_mav	235.46 ± 50.42	82.70 ± 14.43	69.10 ± 33.19
tum_vi	264.21 ± 24.86	92.54 ± 17.72	100.73 ± 38.73
kaist_vio	274.96 ± 124.70	80.69 ± 45.38	85.89 ± 47.05
kaist	1678.32 ± 350.95	141.08 ± 11.22	135.75 ± 8.31

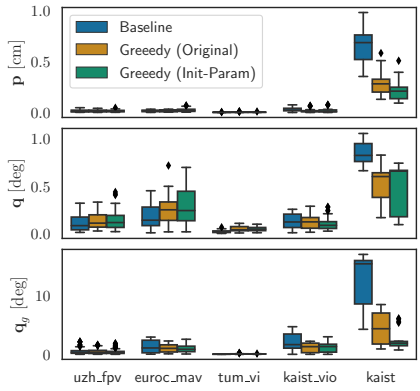

Fig. 3: Absolute error in estimated extrinsic parameters for baseline, original greedy, and greedy with utility evaluation at initial calibration parameters over each dataset in simulation.

VI. Hardware Experiments

This section evaluates self-calibration with measurement subsets using real-world data from a dataset featured in our prior work [13]. As in simulations, a sensitivity analysis in Section VI-B examines how changes in parameter estimates affect the measure of information, supporting the use of Alg. 2 as an alternative to Alg. 1. Section VI-C compares the calibration performances of Alg. 1 and Alg. 2 with a full-measurement baseline, highlighting Alg. 2's reduction in runtime—from approximately one minute to just a second—without loosing calibration accuracy.

A. Implementation Details

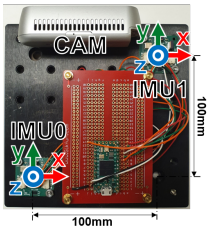

Fig. 4: Sensor rig used for data collection.

Fig. 4 shows the sensor rig used for data collection, featuring two IMUs (IMU 0 and IMU 1) and a camera (CAM), with only the IMU data utilized for evaluation. The reference relative position $\mathbf{p}_{I_1}$ and orientation $\mathbf{q}_{I_1}$ between the IMUs are $[-10, -10, 0]$ [cm] and $[0, 0, 0]$ [deg] in XYZ Euler angles, with assumed zero gyroscope misalignment; $\mathbf{q}_{g_0}$ and $\mathbf{q}_{g_1}$ are also $[0, 0, 0]$ [deg]. Allan variance analysis provided white noise and bias instability of the accelerometers $\sigma_a = 3.3 \times 10^{-3}$ m/s$^2/\sqrt{\text{Hz}}$ and $\sigma_{\mathbf{b}_a} = 2.4 \times 10^{-4}$ m/s$^2 \cdot \sqrt{\text{Hz}}$ and of the gyroscopes $\sigma_g = 8.3 \times 10^{-5}$ rad/s$/\sqrt{\text{Hz}}$ and $\sigma_{\mathbf{b}_g} = 3.7 \times 10^{-6}$ rad/s $\cdot \sqrt{\text{Hz}}$.

Acknowledging that comparing estimated calibration parameters to reference values might not entirely reflect accuracy, as noted in our previous work, we also measure reprojection error between the IMUs, with a lower error indicating more precise estimation.

B. Sensitivity Analysis of Fisher Information Matrix

As in the sensitivity analysis conducted in simulations (in Section V-C), we assessed how deviations from true calibration parameters affect the measured information using the hardware data. In Fig.5, we examined the Spearman's correlation coefficient [22] between the information distribution, $-1/2 \cdot \log\left|\left[\boldsymbol{\mathcal{I}}^{-1}\right]_{\boldsymbol{\Theta\Theta}}\right|$, evaluated at reference extrinsic parameters and at deviations in either relative position ($\delta\mathbf{p}$) or orientation ($\delta\mathbf{q}$).

Similar to that in simulations, sensitivity to orientation deviations ($\delta\mathbf{q}$) was evident, showing a greater decline than for position deviations ($\delta\mathbf{p}$). This highlights the importance of good initial guesses for relative orientation, as advocated in Section IV-D, for sucessful measurement subset selection.

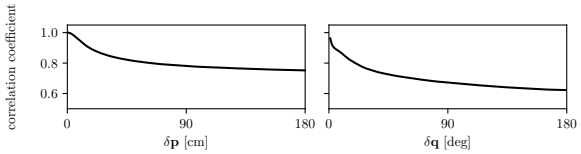

Fig. 5: The correlation coefficient of the information as a function of the difference ($\delta\mathbf{p}$, $\delta\mathbf{q}$) in guessed and reference extrinsic parameters in the two-IMU system on hardware. A coefficient of 1 indicates perfect correlation with the true value, while 0 denotes independence.

TABLE VII: Estimated extrinsics, segment selection ratio, and runtime for baseline, original greedy, and greedy with utility evaluation at initial calibration parameters on hardware data

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**

TABLE VIII: IMU measurement reprojection errors for baseline, original greedy, and greedy with utility evaluation at initial calibration parameters on hardware data

	Dataset	Baseline	Greedy (Original)	Greedy (Init-Param)
accelerometer ([m/s^2])	baseline	0.8160 ± 0.4947	**0.6015 ± 0.4947**	**0.4749 ± 0.4760**
	blurry	1.0981 ± 0.5470	0.5069 ± 0.5470	0.5041 ± 0.4257
	ill-lit	1.6424 ± 0.3951	**0.4697 ± 0.3951**	**0.8968 ± 0.3911**
gyroscope ([rad/s])	baseline	0.0942 ± 0.0139	0.0942 ± 0.0139	0.0942 ± 0.0141
	blurry	0.0996 ± 0.0209	0.0996 ± 0.0219	0.0996 ± 0.0219
	ill-lit	0.0926 ± 0.0101	0.0925 ± 0.0103	0.0930 ± 0.0107

C. Comparison of Self-Calibration with Efficient Greedy Algorithm against Benchmarks

Table VII shows estimated calibration parameters $\mathbf{p}_{I_1}$, $\mathbf{q}_{I_1}$, $\mathbf{q}_{g_0}$, and $\mathbf{q}_{g_1}$ along with the chosen measurement subset ratio and runtime for each calibration process. Table VIII shows the IMU reprojection error across trajectories for the three methods. Significant differences between the greedy methods, determined through t-tests, are highlighted in bold. Initial parameter guesses were as previously discussed in simulations. The results indicate the greedy algorithms use less than 3% of total measurements to achieve calibration results that closely align with reference values. Greedy (Init-Param) significantly reduces runtime from over a minute to approximately a second, whereas Greedy (Original) achieves a modest reduction, cutting the runtime by about fifteen seconds from the baseline.

VII. Conclusion

This paper presents a multi-IMU extrinsic calibration by selecting high-utility measurement subsets. We hypothesize

that, in our system, utility—a function of the parameter estimates—is largely insensitive to the specific choice of parameters, allowing evaluation at an initial guess and reducing the need for frequent recalibrations. This significantly decreases computation time in both simulations and experiments, making it suitable for resource-limited platforms.

We acknowledge that the hypothesis regarding the insensitivity in parameter choice for subset selection could have been examined more thoroughly. Future research could explore this in greater depth, for example by comparing the selected segments using both the original and the modified Greedy approaches.

Appendix I
Proof of Time Complexity of Original Greedy Algorithm

In this appendix, we derive the time complexity of the original greedy algorithm (Alg. 1).

We will use the following symbols:

- L means the number of total segments.
- M means the number of finally chosen segments (i.e., $M \leq L$).
- K means the number of timesteps (i.e., data) in each segment.
- p means the number of parameters being estimated.

We assume $K > p$ in general.

1) Line 5 in Alg. 1 iterates L times.
2) Line 6 conducts nonlinear least-square on $m \times K$ data, where m is the number of segments in $\mathcal{D}^{info} \cup \mathcal{D}^{new}$ ($\mathcal{O}(mK \cdot p^2)$).
3) Line 7 evaluates Jacobians $\mathbf{J}(\mathcal{D}^{info})_{\hat{\boldsymbol{\theta}}^+}$ (size: $(m-1)K \times p$) and $\mathbf{J}(\mathcal{D}^{new})_{\hat{\boldsymbol{\theta}}^+}$ (size: $K \times p$) and computes the Fisher information matrix ($\mathcal{O}(mKp^2)$).
4) Lines 8-9 calculate the Fisher information matrix's inverse and determinant (both $\mathcal{O}(p^3)$).

Operations on lines 6 and 7 within the loop are computational bottlenecks. The sum of m in the loop can reach L^2 in the worst case and L in the best case, leading to Alg. 1's time complexity expressed as:

$$T(n) = \mathcal{O}(\sum_m mKp^2) = \begin{cases} \mathcal{O}(LKp^2) & \text{in the best} \\ \mathcal{O}(L^2Kp^2) & \text{in the worst} \end{cases}$$

Appendix II
Proof of Time Complexity of Greedy Algorithm Using Only Initial Calibration Parameters

In this appendix, we derive the time complexity of the greedy algorithm with utility evaluation at initial calibration parameters (Alg. 2). We use the same symbols as introduced in Appendix I.

1) Line 4 in Alg. 2 iterates L times.
2) Line 5 evaluates Jacobian $\mathbf{J}(\mathcal{D}^{new})|_{\boldsymbol{\theta}^0}$ (size: $K \times p$) and computes the Fisher information matrix ($\mathcal{O}(Kp^2)$).
3) Lines 6-7 calculates the Fisher information matrix's inverse and determinant (both $\mathcal{O}(p^3)$).
4) After the loop, line 12 conducts nonlinear least-square on $M \times K$ data ($\mathcal{O}(MK \cdot p^2)$).

The operation on line 5 within the loop is a computational bottleneck, leading to Alg. 2's time complexity expressed as:

$$T(n) = \mathcal{O}(L \times Kp^2 + (MK)p^2) = \mathcal{O}(LKp^2).$$

References

[1] J.-O. Nilsson and I. Skog, "Inertial sensor arrays—a literature review," in *Proc. IEEE Eur. Navigation Conf*, 2016.

[2] I. Skog, J.-O. Nilsson, P. Händel, and A. Nehorai, "Inertial sensor arrays, maximum likelihood, and cramér–rao bound," *IEEE Trans. Signal Process.*, vol. 64, no. 16, pp. 4218–4227, 2016.

[3] K. Parsa, J. Angeles, and A. K. Misra, "Estimation of the flexural states of a macro-micro manipulator using point-acceleration data," *IEEE Trans. Robot.*, vol. 21, no. 4, pp. 565–573, 2005.

[4] I. B. Wijayasinghe, M. N. Saadatzi, S. Abubakar, and D. O. Popa, "A study on optimal placement of accelerometers for pose estimation of a robot arm," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 1444–1451.

[5] M. Zhang, X. Xu, Y. Chen, and M. Li, "A lightweight and accurate localization algorithm using multiple inertial measurement units," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1508–1515, 2020.

[6] K. Eckenhoff, P. Geneva, and G. Huang, "Mimc-vins: A versatile and resilient multi-imu multi-camera visual-inertial navigation system," *IEEE Trans. Robot.*, vol. 37, no. 5, pp. 1360–1380, 2021.

[7] S. Y. Cho and C. G. Park, "A calibration technique for a redundant imu containing low-grade inertial sensors," *ETRI J.*, vol. 27, no. 4, pp. 418–426, 2005.

[8] P. Schopp, L. Klingbeil, C. Peters, and Y. Manoli, "Design, geometry evaluation, and calibration of a gyroscope-free inertial measurement unit," *Sensors Actuators A: Phys.*, vol. 162, no. 2, pp. 379–387, 2010.

[9] K. He, J. Han, and Y. Shao, "A novel redundant inertial measurement unit and calibration algorithm," in *Proc. IEEE Int. Conf. Optoelectron. Microelectronics*, 2013, pp. 18–23.

[10] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 4304–4311.

[11] P. Schopp, H. Graf, W. Burgard, and Y. Manoli, "Self-calibration of accelerometer arrays," *IEEE Trans. Instrum. Meas.*, vol. 65, no. 8, pp. 1913–1925, 2016.

[12] D. Kim, S. Shin, and I. S. Kweon, "On-line initialization and extrinsic calibration of an inertial navigation system with a relative preintegration method on manifold," *IEEE Trans. Autom. Sci. Eng.*, vol. 15, no. 3, pp. 1272–1285, 2017.

[13] 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.

[14] J. Maye, P. Furgale, and R. Siegwart, "Self-supervised calibration for robotic systems," in *Proc. IEEE Intell. Veh. Symp.*, 2013, pp. 473–480.

[15] 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.

[16] C. R. Rao, "Minimum variance and the estimation of several parameters," in *Math. Proc. Camb. Phil. Soc.*, vol. 43, no. 2. Cambridge University Press, 1947, pp. 280–283.

[17] T. Schneider, M. Li, C. Cadena, J. Nieto, and R. Siegwart, "Observability-aware self-calibration of visual and inertial sensors for ego-motion estimation," *IEEE Sensors J.*, vol. 19, no. 10, pp. 3846–3860, 2019.

[18] J. Lv, X. Zuo, K. Hu, J. Xu, G. Huang, and Y. Liu, "Observability-aware intrinsic and extrinsic calibration of lidar-imu systems," *IEEE Trans. Robot.*, 2022.

[19] R. Zamir, "A proof of the fisher information inequality via a data processing argument," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 1246–1250, 1998.

[20] Z. Yang and S. Shen, "Monocular visual–inertial state estimation with online initialization and camera–imu extrinsic calibration," *IEEE Trans. Autom. Sci. Eng.*, vol. 14, no. 1, pp. 39–51, 2016.

[21] P. Geneva, K. Eckenhoff, W. Lee, Y. Yang, and G. Huang, "Openvins: A research platform for visual-inertial estimation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 4666–4672.

[22] J. H. Zar, "Spearman rank correlation: overview," *Wiley StatsRef: Statistics Reference Online*, 2014.