sensors

MDPI

Article

Fisher-Information-Based Cooperative Sensor Node Pre-Selection for UWB-Aided GNSS-Denied UAV Swarm Localization Under Heterogeneous Ranging Noise

Yanming Sun [1,2,*], Xiaoyan Du [1] and Pihong Gong [2]

[1] School of Transportation, Shandong University of Science and Technology, Qingdao 266590, China; 202482160015@sdust.edu.cn
[2] International Cooperation Center, National Development and Reform Commission, Beijing 100038, China; gphongwork@163.com
* Correspondence: skd994411@sdust.edu.cn or sun3311@126.com; Tel.: +86-19854216269

Highlights

What are the main findings?

- A Fisher-information-based generalized GDOP metric is developed to evaluate the local observability of UWB-aided cooperative ranging sensor networks under topology-coupled and heteroscedastic link-noise conditions.
- A resource-constrained greedy cooperative sensor node pre-selection heuristic is developed to jointly control the number of selected nodes, induced active ranging links, ranging-slot occupation, and general sensing-resource cost.

What are the implications of the main findings?

- The proposed method provides an interpretable information-theoretic criterion for cooperative sensor node pre-selection before cooperative localization or factor-graph optimization in GNSS-denied environments.
- The framework explicitly distinguishes the selected-node budget from the induced-link and ranging-slot budgets, while all mobile nodes remain in the cooperative localization state.

Abstract

Ultra-wideband (UWB) inter-node ranging provides relative-distance constraints for cooperative localization in GNSS-denied UAV swarms, but dense candidate networks can exceed the available ranging slots, communication bandwidth, computation, and energy. This paper proposes a Fisher-information-based cooperative sensor node pre-selection method under heterogeneous ranging noise. All mobile nodes remain in the localization state, while the selected nodes induce the active ranging-link set. Selected-node, induced-link, ranging-slot, and normalized general-resource budgets are represented separately. Using predicted geometry and estimated link-quality weights, a gauge-free normalized Fisher information matrix combines link geometry, link-quality-dependent weights, and topology-induced coupling. A trace-based generalized GDOP (G-GDOP) criterion is optimized by a two-stage greedy heuristic with recursive matrix updates. The experiments show that G-GDOP is a local observability and information-quality metric rather than a direct predictor of topology-level nonlinear recovery error. Within the same topology, normalized local RMSE increased from 0.698 [0.673, 0.752] in the Low G-GDOP group to 1.014 [0.999, 1.068] and 1.980 [1.806, 2.180] in the Medium and High groups. Increasing the selected-node budget from $K = 6$ to $K = 20$ reduced median RMSE from 0.550 to 0.148 m while increasing the median induced-link number from 109 to 214. Additional tests covered Gaussian and heterogeneous ranging noise, deterministic NLOS bias, online link-weight

Check for updates

Academic Editor: Andrey V. Savkin

Received: 23 June 2026
Revised: 12 August 2026
Accepted: 13 August 2026
Published: 14 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

errors, and predicted-position uncertainty. Direct Inversion and Woodbury Updating were numerically equivalent within a predefined tolerance in all 24 size–regime combinations, and a Woodbury runtime advantage was supported in 20 conditions. The proposed framework therefore provides an interpretable resource-aware pre-selection module without implying an unconditional real-time guarantee.

Keywords: UWB cooperative localization; cooperative sensor node pre-selection; Fisher information matrix; generalized GDOP; heterogeneous ranging noise; resource-constrained sensing; GNSS-denied navigation

1. Introduction

Reliable positioning and navigation are fundamental requirements for unmanned aerial vehicles (UAVs), mobile robots, and cooperative sensing systems operating in GNSS-denied or GNSS-degraded environments. In urban canyons, indoor–outdoor transition areas, cluttered low-altitude spaces, and post-disaster scenarios, GNSS signals are often affected by blockage, attenuation, multipath propagation, and non-line-of-sight (NLOS) effects, which may reduce positioning availability and degrade localization accuracy. To overcome the limitations of standalone GNSS positioning, various onboard and inter-node sensing sources, such as ultra-wideband (UWB) ranging, inertial measurement units (IMUs), visual-inertial odometry (VIO), simultaneous localization and mapping (SLAM), and range-based cooperative measurements, have been introduced into UAV and robotic localization systems [1–7]. Among these sensing technologies, UWB-based inter-node ranging is particularly attractive for cooperative localization because it can provide direct relative distance constraints with relatively low cost, low power consumption, and high temporal resolution. UWB waveform design and admissible emission spectra are constrained by regional regulations and pulse-shaping requirements, while the ranging error observed by a practical receiver depends strongly on propagation conditions, reflections, interference, and fading [8,9]. These physical effects motivate the heterogeneous random-error and residual bias stress tests used later in this paper, but the present work does not claim to reproduce a complete radio-frequency channel or transceiver model.

In a UAV swarm or multi-robot system equipped with UWB ranging sensors, the inter-node ranging links naturally form a cooperative ranging sensor network. Each UAV-borne sensor node can provide not only its own motion information, but also relative measurements to neighboring nodes. These relative measurements can improve the local observability of the network and reduce the accumulated positioning drift caused by GNSS outages or onboard sensor errors. Recent studies on UAV swarms and multi-robot systems have demonstrated the potential of formation-constrained cooperative localization, cluster-based localization, and multi-condition cooperative positioning in GNSS-denied environments [10–12]. However, in dense cooperative sensing scenarios, the number of potential ranging links is usually much larger than the number that can be activated simultaneously. Communication bandwidth, ranging time slots, onboard computation, and residual energy all restrict the scale of cooperative participation.

Therefore, a key issue in resource-constrained cooperative localization is not only how to estimate the positions of all mobile nodes from available measurements, but also which mobile nodes should be included in the selected node subset under limited resources. In the framework considered in this paper, individual ranging links are not optimized as independent decision variables. Instead, the active ranging-link set is induced by the selected node subset and the candidate ranging graph.

Geometric dilution of precision (GDOP) is a classical and interpretable metric for evaluating the influence of observation geometry on localization accuracy. It has been widely used in satellite navigation, wireless positioning, and geometric configuration analysis because of its clear physical meaning and low computational cost. Nevertheless, conventional GDOP is mainly developed for single-target localization under homoscedastic observation assumptions. This setting is different from cooperative ranging sensor networks, where multiple unknown nodes are estimated jointly and inter-node measurements introduce block-coupled information structures. Moreover, the quality of different UWB ranging links may vary significantly due to inter-node distance, antenna orientation, occlusion, multipath propagation, NLOS conditions, and environmental interference. As a result, the localization contribution of each link is jointly affected by link geometry, heterogeneous measurement noise, and cooperative network topology. A direct application of conventional GDOP is therefore insufficient to characterize the system-level information quality of topology-coupled cooperative localization networks [13–17].

The Fisher information matrix (FIM) and the Cramér–Rao lower bound (CRLB) provide a rigorous statistical framework for analyzing localization accuracy and theoretical error lower bounds. FIM- and equivalent-FIM-based methods can incorporate measurement geometry, noise statistics, and network topology into a unified information representation, and have been widely used in wireless sensor network localization, cooperative navigation, factor-graph-based localization, sensor selection, and link scheduling [14,15,18–23]. However, directly applying information-theoretic criteria to sensor node or link selection may introduce considerable computational burden, especially when the number of candidate nodes or links increases. Existing studies have investigated position-error-bound-based selection, game-theoretic resource allocation, topology optimization, and cooperative link selection strategies [19–23]. Although these methods have improved cooperative localization performance from different perspectives, an efficient and interpretable node selection criterion is still needed for UWB-aided cooperative localization, where link geometry, heterogeneous ranging noise, and multi-node topological coupling should be evaluated jointly while maintaining acceptable computational complexity.

The trace-inverse A-optimal objective, marginal-gain greedy selection, and Woodbury identity are established mathematical tools and are not claimed here as new primitives. The contribution of this study is their problem-specific integration into an anchor-free UWB cooperative sensor node pre-selection framework. The decision variable is the selected node subset; usable links are induced by that subset, heterogeneous online link-quality estimates enter the projected information matrix, and node, link, slot, and general resource limits are treated explicitly. Multiplication of the physical FIM by the positive reference factor σ_0^2 only changes the numerical scale. Likewise, division by the common observable dimension d_o does not change the minimizing subset when the total state and predicted geometry are fixed; it is introduced only to make reported objective values interpretable across network sizes. To provide a clearer comparison of existing studies on cooperative localization and node selection, representative works are summarized in Table 1.

Table 1. Comparison of representative studies on cooperative localization and resource-aware sensor node and measurement selection.

Research Category	Main Focus	Considered Factors	Selection	Remaining Gap
GNSS-denied and collaborative positioning studies [1–3]	Collaborative positioning and range-based localization under degraded GNSS conditions	Relative ranging, IMU information, and cooperative constraints	Limited	Mainly focus on localization feasibility and sensor fusion rather than resource-constrained node/link selection
Multi-sensor fusion and factor-graph-based localization [4–7]	Fusion of UWB, IMU, VIO, and SLAM measurements	Measurement uncertainty and inter-platform ranging constraints	Limited	Focus mainly on estimation and fusion accuracy rather than selecting informative cooperative measurements

Table 1. *Cont.*

Research Category	Main Focus	Considered Factors	Selection	Remaining Gap
UAV swarm and multi-robot cooperative localization [10–12]	Cooperative localization for UAV or multi-agent systems in GNSS-denied environments	Formation constraints, clustering, and multi-condition positioning	Partial	Node optimization under limited ranging and communication resources is not fully addressed
GDOP- and geometry-based localization analysis [16,17]	Geometric quality evaluation and configuration optimization	Mainly observation geometry	Partial	Difficult to characterize heterogeneous ranging noise and multi-node topological coupling
FIM/EFIM- and CRLB-based system analysis [14,15,18]	System-level accuracy analysis and theoretical error bounds	Geometry, noise statistics, and network topology	Partial	Direct information-theoretic evaluation may involve high computational cost in dense networks
Existing sensor node and link selection methods [19–23]	Cooperative node activation, link selection, or resource allocation	Partial consideration of accuracy, noise, and resource constraints	Explicit	Accuracy–complexity balance and interpretability remain challenging
Proposed method	Fisher-information-based cooperative sensor node pre-selection	Gauge-free geometry, heterogeneous ranging noise, induced active links, and resource constraints	Explicit	Dynamic scenarios, factor-graph integration, and hardware validation

As summarized in Table 1, existing studies have made important progress in GNSS-denied positioning, multi-sensor fusion, cooperative localization, topology-aware estimation, and sensor node or measurement-selection techniques. However, dense UWB cooperative ranging networks still lack an interpretable and computationally efficient pre-selection framework that can jointly evaluate link geometry, heterogeneous ranging reliability, and topology-induced multi-node coupling under limited ranging and communication resources. This motivates the development of a Fisher-information-based cooperative sensor node pre-selection framework for resource-constrained UWB cooperative localization.

The main contributions of this paper are summarized as follows:

1. A UWB cooperative ranging network model is established for resource-constrained cooperative sensor node pre-selection under heterogeneous ranging noise conditions. The proposed formulation distinguishes between the complete localization state and the selected node subset. All N mobile sensor nodes remain in the localization state, whereas only the selected nodes are scheduled as active UWB ranging initiators. The corresponding active ranging-link set is induced by the selected node subset and is not optimized independently.
2. A trace-based A-optimal criterion is specialized to the topology-coupled UWB pre-selection problem and reported through the G-GDOP notation. It combines UWB link-direction geometry, estimated heterogeneous ranging noise weights, and multi-node coupling in a gauge-free local information-quality measure; its normalization does not change the selected node subset for a fixed predicted network state.
3. A resource-constrained cooperative sensor node pre-selection model is formulated by separately representing the selected-node budget, induced active-link number, ranging-slot occupation, and general sensing-resource cost. A marginal-gain-based greedy heuristic with recursive information matrix updates is then developed to construct a feasible selected node subset without repeatedly inverting the complete candidate information matrix.
4. A comprehensive nine-experiment validation framework is established to clarify the applicability boundary and same-topology ranking behavior of G-GDOP, quantify node-, induced-link-, ranging-slot-, and general-resource trade-offs, and compare the proposed heuristic with the small-scale exhaustive optimum. The framework also evaluates Gaussian noise, heterogeneous link quality, deterministic NLOS bias, online weight-estimation errors, predicted-position uncertainty, and the tolerance-qualified numerical equivalence and runtime scaling of Direct Inversion and Woodbury Updat-

ing. Shared non-success and common-unobservable cases are retained in the formal status accounting rather than being silently discarded.

The remainder of this paper is organized as follows. Section 2 presents the UWB-aided cooperative sensor network model and constructs the generalized GDOP metric. Section 3 formulates the resource-constrained cooperative sensor node pre-selection problem and introduces the proposed greedy algorithm with recursive matrix updates. Section 4 reports the simulation settings and experimental results. Section 5 concludes the paper and outlines future research directions.

2. UWB-Aided Cooperative Sensor Network Modeling and Generalized GDOP Metric

2.1. Problem Definition of Cooperative Sensor Node Pre-Selection

Consider a GNSS-denied UAV swarm consisting of N mobile sensor nodes equipped with onboard motion sensors and UWB ranging modules. The candidate mobile sensor node set is defined as

$$\mathcal{C} = \{1, 2, \ldots, N\}.$$

An optional set of pre-surveyed reference nodes is denoted by $\mathcal{A}$. The candidate UWB ranging network is represented by the undirected graph

$$\mathcal{G}_c = (\mathcal{C} \cup \mathcal{A}, \mathcal{E}_c),$$

where $\mathcal{E}_c$ contains all candidate UWB ranging links that satisfy the communication-range and ranging-availability conditions. In dense cooperative ranging networks, activating all candidate sensor nodes and all available UWB ranging links may exceed the available ranging time slots, communication bandwidth, onboard computational capacity, or energy budget. Therefore, at each ranging-scheduling epoch, only a subset of cooperative sensor nodes is selected for active ranging. The selected node subset is denoted by

$$\mathcal{S} \subseteq \mathcal{C}, \qquad |\mathcal{S}| = K,$$

where K is the prescribed selected-node budget. This study optimizes the selected node subset $\mathcal{S}$ rather than independently selecting individual UWB ranging links. The selected nodes are scheduled as active UWB ranging initiators, whereas the remaining mobile nodes can still serve as ranging responders and remain part of the localization state. Once $\mathcal{S}$ is determined, the corresponding active ranging-link set is induced as

$$\begin{aligned} \mathcal{E}(\mathcal{S}) = \{(i,j) \in \mathcal{E}_c \mid i,j \in \mathcal{C},\ i \in \mathcal{S} \text{ or } j \in \mathcal{S}\} \\ \cup \{(i,a) \in \mathcal{E}_c \mid i \in \mathcal{S},\ a \in \mathcal{A}\}. \end{aligned} \tag{1}$$

Therefore, cooperative sensor node pre-selection controls the selected node subset and its induced ranging operations, rather than independently optimizing each individual UWB link. The main scenario considered in this study is anchor-free cooperative ranging, for which

$$\mathcal{A} = \varnothing.$$

In this case, the active ranging-link set reduces to

$$\mathcal{E}(\mathcal{S}) = \{(i,j) \in \mathcal{E}_c \mid i,j \in \mathcal{C},\ i \in \mathcal{S} \text{ or } j \in \mathcal{S}\}.$$

The formal problem and all reported experiments in this paper use the anchor-free case. The anchor-assisted expression is retained only to show how the same induced-link notation

could be extended when independently surveyed anchors are available; no known-position anchor is introduced into the reported GNSS-denied scenario.

For each candidate mobile-node ranging link $e = (u, v) \in \mathcal{E}_c$, define the selection-dependent link activation indicator as

$$\chi_e(\mathcal{S}) = \begin{cases} 1, & u \in \mathcal{S} \text{ or } v \in \mathcal{S}, \\ 0, & u \notin \mathcal{S} \text{ and } v \notin \mathcal{S}. \end{cases} \tag{2}$$

The induced active ranging-link set can equivalently be written as

$$\mathcal{E}(\mathcal{S}) = \{e \in \mathcal{E}_c \mid \chi_e(\mathcal{S}) = 1\}.$$

Accordingly, an inter-node ranging link becomes active when at least one of its two endpoints is selected as an active UWB ranging initiator. If both endpoints are unselected, the corresponding link remains inactive. Each candidate ranging link is included at most once in $\mathcal{E}(\mathcal{S})$, even when both of its endpoints belong to $\mathcal{S}$.

For any mobile sensor node $i \in \mathcal{C}$, its three-dimensional position is expressed as

$$\mathbf{p}_i = [x_i, y_i, z_i]^T \in \mathbb{R}^3.$$

The localization state contains the positions of all N mobile sensor nodes and is written as

$$\mathbf{x} = \left[\mathbf{p}_1^T, \mathbf{p}_2^T, \ldots, \mathbf{p}_N^T\right]^T \in \mathbb{R}^{3N}.$$

The selected node subset $\mathcal{S}$ determines the active ranging measurements but does not change the dimension of the localization state. Consequently, the information matrices corresponding to different candidate subsets have the same dimension, which enables consistent comparison across node budgets and recursive low-rank matrix updates.

A mobile node belonging to the selected node subset $\mathcal{S}$ is termed a selected nodeand is scheduled as an active UWB ranging initiator. Here, active UWB ranging initiator describes the ranging role of a selected node rather than a separate decision variable. Unselected mobile nodes remain unknown state variables and may serve as passive ranging responders through candidate UWB links. Thus, the selection decision controls the active ranging schedule and the induced measurement topology without changing the complete mobile-node state. Unless otherwise stated, the proposed method performs cooperative sensor node pre-selection only; individual UWB ranging links are not treated as independent binary decision variables.

In practical GNSS-denied operation, the information matrix used for cooperative sensor node pre-selection is constructed from predicted or estimated node positions, such as those propagated from inertial sensing, visual odometry, or the localization result obtained at the previous epoch. For mobile node i, the predicted position is denoted by

$$\widehat{\mathbf{p}}_i = [\widehat{x}_i, \widehat{y}_i, \widehat{z}_i]^T.$$

The predicted state of the complete mobile network is written as

$$\widehat{\mathbf{x}} = \left[\widehat{\mathbf{p}}_1^T, \widehat{\mathbf{p}}_2^T, \ldots, \widehat{\mathbf{p}}_N^T\right]^T.$$

The online cooperative sensor node pre-selection algorithm evaluates the ranging geometry at $\widehat{\mathbf{x}}$ rather than at the unknown true state $\mathbf{x}$. The true node positions are used only as an oracle reference and for performance evaluation in the simulation experiments. This distinc-

tion removes the unrealistic assumption that the online selector has access to ground-truth geometry. Section 4.1.3 therefore evaluates sensitivity to snapshot-wise Gaussian position perturbations while holding the candidate graph, link-quality information, measurement realization, and resource constraints fixed. The broader limitations of this controlled uncertainty model are discussed in Section 4.3.

Since an anchor-free range-only network is invariant to global translation and rotation, its Fisher information matrix is singular in the full $3N$-dimensional state space. The rigid-body gauge directions are therefore removed by projection onto their orthogonal complement. It should be emphasized that removing the global rigid-body gauge directions does not automatically guarantee local observability. Disconnected or non-rigid active-ranging topologies may introduce additional unobservable directions. Consequently, the projected information matrix must additionally satisfy a full-rank condition in the gauge-free subspace. The corresponding construction and numerical rank test are introduced in Section 2.3.

2.2. UWB Ranging Sensor Model and Link-Quality Weighting

Different UWB ranging links may exhibit heterogeneous random errors and systematic positive biases. For an active link $e = (i, j)$, the bias-aware ranging model is

$$z_e = h_e(\mathbf{x}) + b_e + n_e, \tag{3}$$

where b_e is the systematic bias and $n_e \sim \mathcal{N}(0, \sigma_e^2)$ is the zero-mean random error. For a mobile-node–mobile-node link, $h_e(\mathbf{x}) = \|\mathbf{p}_i - \mathbf{p}_j\|_2$, whereas for a mobile-node–anchor-node link, $h_e(\mathbf{x}) = \|\mathbf{p}_i - \mathbf{q}_a\|_2$.

The propagation state is modeled by $c_e \in \mathcal{L} = \{\mathrm{LOS}, \mathrm{wLOS}, \mathrm{NLOS}\}$, with $\sigma_e = \sigma_{c_e}$ and $b_e = \mu_{c_e}$. The class parameters satisfy $\sigma_{\mathrm{LOS}} < \sigma_{\mathrm{wLOS}} < \sigma_{\mathrm{NLOS}}$, $\mu_{\mathrm{LOS}} = \mu_{\mathrm{wLOS}} = 0$, and $\mu_{\mathrm{NLOS}} \geq 0$. The heteroscedastic random-noise experiment uses $\mu_{\mathrm{NLOS}} = 0$, whereas the NLOS-bias experiment uses $\mu_{\mathrm{NLOS}} = b_{\mathrm{NLOS}} > 0$.

In practical operation, c_e, σ_e^2, and b_e are not assumed known exactly. Link-quality features $\mathbf{f}_e$ are used to obtain $\widehat{\pi}_{e,c} = \Pr(c_e = c \mid \mathbf{f}_e)$ with $\sum_{c \in \mathcal{L}} \widehat{\pi}_{e,c} = 1$ and posterior-mean bias $\widehat{b}_e = \sum_{c \in \mathcal{L}} \widehat{\pi}_{e,c} \mu_c$. After subtracting this mean, moment matching gives

$$\widehat{\sigma}_{e,\mathrm{res}}^2 = \sum_{c \in \mathcal{L}} \widehat{\pi}_{e,c} \left[\sigma_c^2 + \left(\mu_c - \widehat{b}_e \right)^2 \right]. \tag{4}$$

A hard classification simply makes $\{\widehat{\pi}_{e,c}\}$ one-hot at $\widehat{c}_e$.

For a selected node subset $\mathcal{S}$, let $M_{\mathcal{S}} = |\mathcal{E}(\mathcal{S})|$. The actual measurement, estimated-bias, corrected-measurement, and residual-bias vectors are collected as

$$\begin{aligned} \mathbf{z}_{\mathcal{S}} &= \mathbf{h}_{\mathcal{S}}(\mathbf{x}) + \mathbf{b}_{\mathcal{S}} + \mathbf{n}_{\mathcal{S}}, & \widehat{\mathbf{b}}_{\mathcal{S}} &= [\widehat{b}_{e_1}, \ldots, \widehat{b}_{e_{M_{\mathcal{S}}}}]^T, \\ \widetilde{\mathbf{z}}_{\mathcal{S}} &= \mathbf{z}_{\mathcal{S}} - \widehat{\mathbf{b}}_{\mathcal{S}} = \mathbf{h}_{\mathcal{S}}(\mathbf{x}) + \widetilde{\mathbf{b}}_{\mathcal{S}} + \mathbf{n}_{\mathcal{S}}, & \widetilde{\mathbf{b}}_{\mathcal{S}} &= \mathbf{b}_{\mathcal{S}} - \widehat{\mathbf{b}}_{\mathcal{S}}. \end{aligned}$$

The online covariance is $\widehat{\mathbf{R}}_{\mathcal{S}} = \mathrm{diag}(\widehat{\sigma}_{e_1,\mathrm{res}}^2, \ldots, \widehat{\sigma}_{e_{M_{\mathcal{S}}},\mathrm{res}}^2)$. Linearization at $\widehat{\mathbf{x}}$ gives

$$\delta \widetilde{\mathbf{z}}_{\mathcal{S}} \approx \mathbf{H}_{\mathcal{S}} \delta \mathbf{x} + \widetilde{\mathbf{b}}_{\mathcal{S}} + \mathbf{n}_{\mathcal{S}}. \tag{5}$$

Accurate posterior-mean correction makes $\widetilde{\mathbf{b}}_{\mathcal{S}}$ small; without correction, $\widehat{\mathbf{b}}_{\mathcal{S}} = \mathbf{0}$ and the residual bias equals the injected physical bias.

For an unknown-node–unknown-node link, define $\mathbf{u}_{ij} = (\widehat{\mathbf{p}}_i - \widehat{\mathbf{p}}_j)/\|\widehat{\mathbf{p}}_i - \widehat{\mathbf{p}}_j\|_2$; the corresponding Jacobian row is $\mathbf{h}_{ij} = [\mathbf{0}, \ldots, \mathbf{u}_{ij}^T, \ldots, -\mathbf{u}_{ij}^T, \ldots, \mathbf{0}]$. For an unknown-node–anchor-node link, $\mathbf{u}_{ia} = (\widehat{\mathbf{p}}_i - \mathbf{q}_a)/\|\widehat{\mathbf{p}}_i - \mathbf{q}_a\|_2$, and only the position block of node i is nonzero.

Collecting all active rows gives $\mathbf{H}_{\mathcal{S}} \in \mathbb{R}^{M_{\mathcal{S}} \times 3N}$. With $\mathbf{R}_{\mathcal{S}} = \mathrm{diag}(\sigma_{e_1}^2, \ldots, \sigma_{e_{M_{\mathcal{S}}}}^2)$, the physical Fisher information matrix is

$$\mathbf{J}_{\mathrm{phys}}(\mathcal{S}) = \mathbf{H}_{\mathcal{S}}^T \mathbf{R}_{\mathcal{S}}^{-1} \mathbf{H}_{\mathcal{S}} + \mathbf{J}_{\mathrm{prior}}. \tag{6}$$

Here $\mathbf{J}_{\mathrm{prior}} \succeq \mathbf{0}$ contains only physically available prior information. In the anchor-free range-only analysis $\mathbf{J}_{\mathrm{prior}} = \mathbf{0}$, and $\widehat{\mathbf{x}}$ is only the linearization point. A reference variance σ_0^2 defines $\bar{\mathbf{J}}_{\mathrm{prior}} = \sigma_0^2 \mathbf{J}_{\mathrm{prior}}$.

For oracle benchmarking, $w_e^\star = \sigma_0^2/\sigma_e^2$ and $\mathbf{W}_{\mathcal{S}}^\star = \sigma_0^2 \mathbf{R}_{\mathcal{S}}^{-1}$, yielding

$$\bar{\mathbf{J}}_{\mathrm{phys}}^\star(\mathcal{S}) = \mathbf{H}_{\mathcal{S}}^T \mathbf{W}_{\mathcal{S}}^\star \mathbf{H}_{\mathcal{S}} + \bar{\mathbf{J}}_{\mathrm{prior}} = \sum_{e \in \mathcal{E}(\mathcal{S})} w_e^\star \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\mathrm{prior}}.$$

These oracle quantities are used only as ideal simulation references.

In practical online weighting, Equation (4) gives

$$\widehat{w}_e = \frac{\sigma_0^2}{\widehat{\sigma}_{e,\mathrm{res}}^2}, \tag{7}$$

with $\widehat{\mathbf{W}}_{\mathcal{S}} = \sigma_0^2 \widehat{\mathbf{R}}_{\mathcal{S}}^{-1}$ and

$$\widehat{\bar{\mathbf{J}}}_{\mathrm{phys}}(\mathcal{S}) = \mathbf{H}_{\mathcal{S}}^T \widehat{\mathbf{W}}_{\mathcal{S}} \mathbf{H}_{\mathcal{S}} + \bar{\mathbf{J}}_{\mathrm{prior}}. \tag{8}$$

Equivalently, $\widehat{\bar{\mathbf{J}}}_{\mathrm{phys}}(\mathcal{S}) = \sum_{e \in \mathcal{E}(\mathcal{S})} \widehat{w}_e \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\mathrm{prior}}$.

Subsequent theoretical and algorithmic sections use $w_e \equiv \widehat{w}_e$, $\mathbf{W}_{\mathcal{S}} \equiv \widehat{\mathbf{W}}_{\mathcal{S}}$, and $\bar{\mathbf{J}}_{\mathrm{phys}}(\mathcal{S}) \equiv \widehat{\bar{\mathbf{J}}}_{\mathrm{phys}}(\mathcal{S})$. Thus $\bar{\mathbf{J}}_{\mathrm{phys}}(\mathcal{S}) = \sum_{e \in \mathcal{E}(\mathcal{S})} w_e \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\mathrm{prior}}$. Using Equation (2), the same model is $\mathbf{J}_{\mathrm{phys}}(\mathcal{S}) = \mathbf{J}_{\mathrm{prior}} + \sum_{e \in \mathcal{E}_c} \chi_e(\mathcal{S}) \widehat{w}_e \mathbf{h}_e^T \mathbf{h}_e$. This separates fixed per-link information from subset-dependent link activation.

Unless $(\cdot)^\star$ is shown, all weights and normalized information matrices are estimated online quantities. The covariance-based Fisher matrix does not explicitly represent deterministic residual NLOS bias, and $\mathbf{R}_{\mathcal{S}}$ and $\widehat{\mathbf{R}}_{\mathcal{S}}$ serve different roles: the former describes physical random errors and oracle benchmarks, whereas the latter supplies practical online weights. Therefore, when $\widehat{\mathbf{R}}_{\mathcal{S}} \neq \mathbf{R}_{\mathcal{S}}$, the inverse estimated information matrix is not generally the actual covariance of the mismatched WLS estimator; the corresponding first-order covariance and bias propagation are analyzed in Section First-Order Error Propagation Under Covariance and Bias Mismatch.

2.3. Gauge-Free Subspace Treatment for Anchor-Free Ranging

In an anchor-free range-only network, common translation and rotation leave all inter-node distances unchanged, so the full $3N$-dimensional Jacobian and Fisher matrix contain rigid-body gauge directions and are singular. Center the predicted positions using $\widetilde{\mathbf{p}}_i = \widehat{\mathbf{p}}_i - N^{-1} \sum_{j=1}^N \widehat{\mathbf{p}}_j$ and, for $\mathbf{p} = [p_x, p_y, p_z]^T$, define

$$[\mathbf{p}]_\times = \begin{bmatrix} 0 & -p_z & p_y \\ p_z & 0 & -p_x \\ -p_y & p_x & 0 \end{bmatrix}.$$

The translational and rotational gauge directions are then collected in

$$\mathbf{G} = \begin{bmatrix} \mathbf{1}_N \otimes \mathbf{I}_3 & \begin{bmatrix} -[\widetilde{\mathbf{p}}_1]_\times \\ -[\widetilde{\mathbf{p}}_2]_\times \\ \vdots \\ -[\widetilde{\mathbf{p}}_N]_\times \end{bmatrix} \end{bmatrix} \in \mathbb{R}^{3N\times 6}. \tag{9}$$

Let $\mathbf{G} = \mathbf{U}_g\mathbf{\Sigma}_g\mathbf{V}_g^T$ have singular values $\gamma_1 \geq \gamma_2 \geq \cdots \geq 0$. The numerical gauge rank and gauge-free dimension are

$$d_g = \#\{\ell : \gamma_\ell > \tau_g\gamma_1\}, \qquad \tau_g = 10^{-10}, \tag{10}$$

and

$$d_o = 3N - d_g. \tag{11}$$

Thus $d_g = 6$ and $d_o = 3N - 6$ for a nondegenerate three-dimensional configuration, while degenerate configurations use the numerical rank directly. The gauge-free basis $\mathbf{Q}_o = \mathrm{null}(\mathbf{G}^T) \in \mathbb{R}^{3N\times d_o}$ satisfies $\mathbf{Q}_o^T\mathbf{Q}_o = \mathbf{I}_{d_o}$ and $\mathbf{Q}_o^T\mathbf{G} = \mathbf{0}$. This basis removes only the rigid gauge; it does not itself guarantee complete local observability.

For the anchor-free range-only model, $\mathbf{H}_\mathcal{S}\mathbf{G} = \mathbf{0}$ and, when $\mathbf{J}_{\text{prior}} = \mathbf{0}$, $\mathbf{J}_{\text{phys}}(\mathcal{S})\mathbf{G} = \mathbf{0}$. Projection onto the gauge-free coordinates gives

$$\mathbf{J}_o(\mathcal{S}) = \mathbf{Q}_o^T\mathbf{J}_{\text{phys}}(\mathcal{S})\mathbf{Q}_o \in \mathbb{R}^{d_o\times d_o}. \tag{12}$$

Disconnected, insufficiently constrained, or non-rigid topologies can still be rank deficient after this projection, so

$$r_o(\mathcal{S}) = \#\{k : \lambda_k(\mathbf{J}_o(\mathcal{S})) > \tau_J\lambda_{\max}(\mathbf{J}_o(\mathcal{S}))\}, \qquad \tau_J = 10^{-8},$$

with $r_o(\mathcal{S}) = 0$ when $\lambda_{\max}(\mathbf{J}_o(\mathcal{S})) = 0$. Local observability modulo the rigid-body gauge therefore requires

$$r_o(\mathcal{S}) = d_o. \tag{13}$$

Equivalently, the symmetrized projected information matrix must be positive definite under the adopted relative eigenvalue threshold. Infeasible subsets are rejected rather than made invertible by artificial regularization. For fixed predicted geometry, $\mathbf{G}$, $\mathbf{Q}_o$, d_g, and d_o are computed once and kept unchanged across candidate subsets.

First-Order Error Propagation Under Covariance and Bias Mismatch

The covariance-based Fisher matrix describes local random-error information, whereas deterministic residual NLOS bias must be propagated separately. Under covariance mismatch, the inverse of the estimated information matrix is also not generally the actual estimation-error covariance. Define

$$\mathbf{H}_{o,\mathcal{S}} = \mathbf{H}_\mathcal{S}\mathbf{Q}_o, \qquad \mathbf{R}_\mathcal{S} = \mathrm{diag}(\sigma_{e_1}^2, \ldots, \sigma_{e_{M_\mathcal{S}}}^2),$$
$$\widehat{\mathbf{R}}_\mathcal{S} = \mathrm{diag}(\widehat{\sigma}_{e_1,\text{res}}^2, \ldots, \widehat{\sigma}_{e_{M_\mathcal{S}},\text{res}}^2), \quad \mathbf{J}_{o,\text{prior}} = \mathbf{Q}_o^T\mathbf{J}_{\text{prior}}\mathbf{Q}_o.$$

The oracle and practical projected matrices and their normalized forms are

$$\mathbf{J}_{o,\text{true}}^\star(\mathcal{S}) = \mathbf{H}_{o,\mathcal{S}}^T\mathbf{R}_\mathcal{S}^{-1}\mathbf{H}_{o,\mathcal{S}} + \mathbf{J}_{o,\text{prior}},$$
$$\widehat{\mathbf{J}}_{o,\text{WLS}}(\mathcal{S}) = \mathbf{H}_{o,\mathcal{S}}^T\widehat{\mathbf{R}}_\mathcal{S}^{-1}\mathbf{H}_{o,\mathcal{S}} + \mathbf{J}_{o,\text{prior}},$$
$$\overline{\mathbf{J}}_o^\star(\mathcal{S}) = \sigma_0^2\mathbf{J}_{o,\text{true}}^\star(\mathcal{S}), \qquad \widehat{\overline{\mathbf{J}}}_o(\mathcal{S}) = \sigma_0^2\widehat{\mathbf{J}}_{o,\text{WLS}}(\mathcal{S}).$$

Unless $(\cdot)^\star$ is shown, the subsequent pre-selection algorithm uses the normalized estimated quantity $\widehat{\widetilde{\mathbf{J}}}_o(\mathcal{S})$.

For a feasible topology, define $\widehat{\mathbf{A}}_\mathcal{S} = \widehat{\mathbf{J}}_{o,\mathrm{WLS}}(\mathcal{S})^{-1}\mathbf{H}_{o,\mathcal{S}}^T\widehat{\mathbf{R}}_\mathcal{S}^{-1}$ and $\widetilde{\mathbf{b}}_\mathcal{S} = \mathbf{b}_\mathcal{S} - \widehat{\mathbf{b}}_\mathcal{S}$. For the anchor-free case with $\mathbf{J}_{\mathrm{prior}} = \mathbf{0}$, the first-order error and bias are

$$\delta\widehat{\mathbf{y}}_\mathcal{S} \approx \widehat{\mathbf{A}}_\mathcal{S}(\mathbf{n}_\mathcal{S} + \widetilde{\mathbf{b}}_\mathcal{S}), \qquad \delta\widehat{\mathbf{y}}_b(\mathcal{S}) = \widehat{\mathbf{A}}_\mathcal{S}\widetilde{\mathbf{b}}_\mathcal{S}.$$

Since the physical random error has covariance $\mathbf{R}_\mathcal{S}$, the actual first-order random covariance is the sandwich expression

$$\mathbf{C}_{\mathrm{rand}}(\mathcal{S}) = \widehat{\mathbf{A}}_\mathcal{S}\mathbf{R}_\mathcal{S}\widehat{\mathbf{A}}_\mathcal{S}^T. \tag{14}$$

The dimension-normalized local MSE surrogate combines random and bias contributions,

$$\Psi_{\mathrm{local}}(\mathcal{S}) = \frac{1}{d_o}\Big\{\mathrm{tr}[\mathbf{C}_{\mathrm{rand}}(\mathcal{S})] + \|\delta\widehat{\mathbf{y}}_b(\mathcal{S})\|_2^2\Big\},$$

with RMSE surrogate

$$\mathrm{RMSE}_{\mathrm{local}}^{\mathrm{sur}}(\mathcal{S}) = \sqrt{\Psi_{\mathrm{local}}(\mathcal{S})}. \tag{15}$$

The oracle gain is $\mathbf{A}_\mathcal{S}^\star = \mathbf{J}_{o,\mathrm{true}}^\star(\mathcal{S})^{-1}\mathbf{H}_{o,\mathcal{S}}^T\mathbf{R}_\mathcal{S}^{-1}$. When $\widehat{\mathbf{R}}_\mathcal{S} = \mathbf{R}_\mathcal{S}$, $\widetilde{\mathbf{b}}_\mathcal{S} = \mathbf{0}$, and $\mathbf{J}_{\mathrm{prior}} = \mathbf{0}$, $\mathbf{C}_{\mathrm{rand}}(\mathcal{S}) = \mathbf{J}_{o,\mathrm{true}}^\star(\mathcal{S})^{-1}$; otherwise $\mathbf{C}_{\mathrm{rand}}(\mathcal{S}) \neq \widehat{\mathbf{J}}_{o,\mathrm{WLS}}(\mathcal{S})^{-1}$ in general. Thus the inverse online information matrix must not be interpreted as the actual covariance or an oracle CRLB under model mismatch. The surrogate in Equation (15) remains a secondary diagnostic, whereas the primary NLOS-bias metric is the rigid-aligned nonlinear localization RMSE obtained directly from the biased UWB measurements.

2.4. Conventional GDOP Metric and Its Limitations

The geometric dilution of precision (GDOP) is a commonly used geometric performance metric in single-target localization. Under the assumption of independent and identically distributed measurement noise, the observation model can be expressed as

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2\mathbf{I}).$$

Accordingly, the conventional GDOP can be written as

$$\mathrm{GDOP} = \sqrt{\mathrm{tr}\left[(\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}\right]}.$$

In this case, the root mean square error (RMSE) of the estimator can be approximately expressed as

$$\mathrm{RMSE} \approx \sigma_0 \cdot \mathrm{GDOP}.$$

The essence of conventional GDOP is to measure the amplification effect of observation geometry on measurement errors. It is a dimensionless metric with clear physical interpretation and low computational complexity. Although conventional GDOP has been widely used in single-target localization, satellite navigation, wireless positioning, and base-station layout problems, it has the following limitations when directly applied to cooperative ranging sensor networks.

First, conventional GDOP is based on the single-target localization assumption and is insufficient to characterize the error propagation mechanism caused by multi-node state coupling. In cooperative sensor network localization, the variables to be estimated consist of the positions of multiple mobile sensor nodes, and the state dimension is usually $3N_{\mathrm{u}}$. UWB inter-node ranging observations introduce significant block-coupled structures

into the error covariance matrix. Therefore, single-target GDOP is unable to reflect the system-level information quality of multi-node joint estimation.

Second, conventional GDOP cannot effectively handle heteroscedastic measurement noise. The conventional GDOP adopts the form of $\mathbf{H}^{\mathrm{T}}\mathbf{H}$, which implicitly assumes that all observations follow the same noise variance model. In practical UAV-borne cooperative sensing systems, the accuracy of UWB ranging is affected by inter-node distance, occlusion, line-of-sight conditions, non-line-of-sight effects, and link quality. As a result, different ranging observations usually have different noise levels, and conventional GDOP cannot accurately reflect the differences in information contributions among different UWB links.

Third, conventional GDOP cannot fully characterize the influence of cooperative ranging sensor network topology on system-level localization performance. In cooperative localization systems, the information matrix is affected not only by local geometric relationships, but also by graph connectivity, information propagation paths, and the selection of cooperative sensor nodes or ranging links. Conventional GDOP mainly describes local geometric quality and is therefore insufficient to characterize the local observability and error-bound level of topology-coupled cooperative localization networks.

Figure 1 contrasts the conventional single-target GDOP setting with the proposed G-GDOP formulation. Conventional GDOP mainly describes the geometric quality of known reference-to-target observations. In the proposed anchor-free cooperative model, all mobile nodes remain unknown state variables, whereas only nodes in the selected node subset are scheduled as active UWB ranging initiators. The selected node subset induces the ranging topology, and G-GDOP evaluates its estimated local information quality after link-quality weighting and removal of the rigid-body gauge directions.

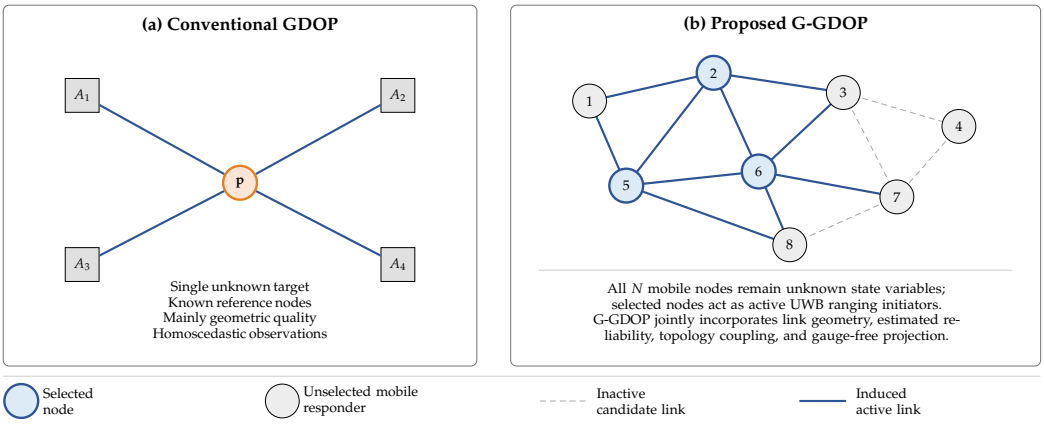


Figure 1. Conceptual comparison between conventional single-target GDOP and the proposed G-GDOP for anchor-free, topology-coupled multi-node UWB cooperative localization.

Generalized GDOP can be regarded as an extension of conventional GDOP from single-target homoscedastic localization to multi-node and heteroscedastic cooperative ranging sensor networks. At the application level, it extends the use of GDOP from geometric quality evaluation to system-level information assessment, cooperative sensor node pre-selection, and topology-aware cooperative localization. Therefore, it is necessary to construct a system-level generalized GDOP metric under the Fisher information matrix framework, so as to uniformly reflect the combined effects of link geometry, network topology, and ranging noise weights on the local information quality of cooperative localization systems.

2.5. Definition and Property Analysis of the Generalized GDOP Metric

In this paper, G-GDOP is defined within the FIM/CRLB framework for UWB-aided cooperative ranging sensor networks. It retains the error-amplification interpretation of conventional GDOP while incorporating link-quality-dependent noise weighting and multi-node topological coupling, thereby allowing local observability differences among cooperative ranging topologies to be compared.

2.5.1. A-Optimal G-GDOP in the Gauge-Free Subspace

The trace-based A-optimal criterion is evaluated from the normalized Fisher information matrix projected onto the gauge-free subspace. Topologies that remain unobservable after removal of the rigid-body gauge are excluded by defining

$$\Phi_{\mathrm{A}}(\mathcal{S}) = \begin{cases} \mathrm{tr}\left[\mathbf{J}_o(\mathcal{S})^{-1}\right], & r_o(\mathcal{S}) = d_o, \\ +\infty, & r_o(\mathcal{S}) < d_o, \end{cases} \tag{16}$$

with the corresponding generalized GDOP

$$\text{G-GDOP}(\mathcal{S}) = \sqrt{\Phi_{\mathrm{A}}(\mathcal{S})}. \tag{17}$$

For the practical pre-selection method, $\mathbf{J}_o(\mathcal{S})$ is constructed from predicted geometry and estimated link-error covariance, so $\Phi_{\mathrm{A}}(\mathcal{S})$ is an online information surrogate for ranking feasible subsets. Its inverse-trace value coincides with the normalized oracle CRLB only when the estimated and true covariances are identical; under covariance mismatch, the actual first-order estimator covariance is instead given by Equation (14).

For comparisons across network sizes, define

$$\Phi_{\mathrm{A,norm}}(\mathcal{S}) = \frac{1}{d_o}\Phi_{\mathrm{A}}(\mathcal{S}), \tag{18}$$

and

$$\text{G-GDOP}_{\mathrm{norm}}(\mathcal{S}) = \sqrt{\Phi_{\mathrm{A,norm}}(\mathcal{S})}. \tag{19}$$

For fixed network size and predicted geometry, d_o is common to all feasible selected node subsets, and therefore

$$\arg\min_{\mathcal{S}} \Phi_{\mathrm{A}}(\mathcal{S}) = \arg\min_{\mathcal{S}} \Phi_{\mathrm{A,norm}}(\mathcal{S}). \tag{20}$$

The normalized form is used for cross-size reporting, whereas the unnormalized form is retained when discussing the connection with conventional GDOP.

2.5.2. Physical Interpretation of the Generalized GDOP Metric

For a feasible topology induced by the selected node subset with $r_o(\mathcal{S}) = d_o$, $\bar{\mathbf{J}}_o^{\star}(\mathcal{S})^{-1}$ is the normalized local oracle CRLB under the true random-error covariance. By contrast, the practical criterion uses $\hat{\bar{\mathbf{J}}}_o(\mathcal{S})$, constructed from estimated link-error statistics, and its inverse trace is therefore an estimated local information-quality surrogate. A smaller practical G-GDOP indicates a more favorable estimated information structure under the adopted online geometry and weighting model, but does not by itself guarantee smaller nonlinear localization error under covariance mismatch, residual NLOS bias, or poor initialization.

The practical and oracle projected information models differ only in their link weights:

$$\mathbf{J}_o(\mathcal{S}) = \mathbf{Q}_o^T \left[\sum_{e \in \mathcal{E}(\mathcal{S})} \widehat{w}_e \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\text{prior}} \right] \mathbf{Q}_o, \quad \widehat{w}_e = \frac{\sigma_0^2}{\widehat{\sigma}_{e,\text{res}}^2},$$

$$\mathbf{J}_o^\star(\mathcal{S}) = \mathbf{Q}_o^T \left[\sum_{e \in \mathcal{E}(\mathcal{S})} w_e^\star \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\text{prior}} \right] \mathbf{Q}_o, \quad w_e^\star = \frac{\sigma_0^2}{\sigma_e^2}.$$

Thus G-GDOP jointly reflects link geometry, heterogeneous measurement reliability, and the position at which each link couples into the multi-node state. The link direction determines the spatial information direction, the variance determines its reliability, and the induced topology determines how that information couples the network state. This unified information representation is the main advantage of G-GDOP for heteroscedastic multi-node cooperative ranging.

2.5.3. Property Analysis of the Generalized GDOP Criterion

The suitability of the proposed criterion for cooperative sensor node pre-selection is examined through monotonicity, spectral bounds, and sensitivity to ranging topology and link reliability in the gauge-free subspace.

Monotonicity

For fixed predicted geometry, estimated link weights, and gauge-free basis $\mathbf{Q}_o$, let $\mathcal{S}_1 \subseteq \mathcal{S}_2$. The endpoint-activation rule in Equation (2) implies the following complete ordering chain:

$$\begin{aligned}
&\chi_e(\mathcal{S}_1) \le \chi_e(\mathcal{S}_2), \quad \forall e \in \mathcal{E}_c, && \mathcal{E}(\mathcal{S}_1) \subseteq \mathcal{E}(\mathcal{S}_2), \\
&\Delta\mathbf{J}_{\text{phys}} = \sum_{e \in \mathcal{E}(\mathcal{S}_2) \setminus \mathcal{E}(\mathcal{S}_1)} \widehat{w}_e \mathbf{h}_e^T \mathbf{h}_e \succeq \mathbf{0}, && \Delta\mathbf{J}_o = \mathbf{Q}_o^T \Delta\mathbf{J}_{\text{phys}} \mathbf{Q}_o \succeq \mathbf{0}, \\
&\bar{\mathbf{J}}_{o,2} = \bar{\mathbf{J}}_{o,1} + \Delta\bar{\mathbf{J}}_o \succeq \bar{\mathbf{J}}_{o,1}, && \bar{\mathbf{J}}_{o,2}^{-1} \preceq \bar{\mathbf{J}}_{o,1}^{-1}, \\
&\text{tr}(\bar{\mathbf{J}}_{o,2}^{-1}) \le \text{tr}(\bar{\mathbf{J}}_{o,1}^{-1}), && \text{G-GDOP}_2 \le \text{G-GDOP}_1, \\
& && \text{G-GDOP}_{\text{norm},2} \le \text{G-GDOP}_{\text{norm},1}.
\end{aligned}$$

The inverse inequalities require both projected information matrices to be positive definite. Because $d_o = 3N - d_g$ is fixed for the same network size and linearization geometry, adding valid ranging information cannot increase either A-optimal G-GDOP measure under these fixed conditions. The result does not compare different predicted geometries, because changing geometry requires recomputation of $\mathbf{G}$, d_g, d_o, $\mathbf{Q}_o$, and the ranging Jacobians.

Spectral Bounds

Let $0 < \lambda_{\min}(\bar{\mathbf{J}}_o) \le \lambda_{\max}(\bar{\mathbf{J}}_o)$ and $d_o = 3N - d_g$; for the nondegenerate three-dimensional configurations used in the main experiments, $d_g = 6$ and $d_o = 3N - 6$. The trace and both G-GDOP forms satisfy the compact set of bounds

$$\begin{aligned}
\text{tr}(\bar{\mathbf{J}}_o^{-1}) &= \sum_{k=1}^{d_o} \frac{1}{\lambda_k(\bar{\mathbf{J}}_o)}, \\
\frac{d_o}{\lambda_{\max}(\bar{\mathbf{J}}_o)} &\le \text{tr}(\bar{\mathbf{J}}_o^{-1}) \le \frac{d_o}{\lambda_{\min}(\bar{\mathbf{J}}_o)}, \\
\sqrt{\frac{d_o}{\lambda_{\max}(\bar{\mathbf{J}}_o)}} &\le \text{G-GDOP} \le \sqrt{\frac{d_o}{\lambda_{\min}(\bar{\mathbf{J}}_o)}}, \\
\frac{1}{\sqrt{\lambda_{\max}(\bar{\mathbf{J}}_o)}} &\le \text{G-GDOP}_{\text{norm}} \le \frac{1}{\sqrt{\lambda_{\min}(\bar{\mathbf{J}}_o)}}.
\end{aligned}$$

Hence all gauge-free eigen-directions influence G-GDOP, with particular sensitivity to weakly observed directions. As $\lambda_{\min}(\bar{\mathbf{J}}_o) \to 0$, conditioning deteriorates and the upper local error-amplification bound increases.

Sensitivity to Topology, Geometry, and Link Reliability

For a selected node subset $\mathcal{S}$,

$$\bar{\mathbf{J}}_o(\mathcal{S}) = \mathbf{Q}_o^T \left[\sum_{e \in \mathcal{E}(\mathcal{S})} w_e \mathbf{h}_e^T \mathbf{h}_e + \bar{\mathbf{J}}_{\text{prior}} \right] \mathbf{Q}_o. \tag{21}$$

In the anchor-free range-only experiments $\bar{\mathbf{J}}_{\text{prior}} = \mathbf{0}$. Equation (21) therefore exposes three coupled factors directly: $\mathcal{E}(\mathcal{S})$ determines the active topology, $\mathbf{h}_e$ determines geometric direction and state coupling, and w_e represents estimated UWB measurement reliability. Changing the selected node subset, an active link, its reliability weight, or the predicted geometry consequently changes $\bar{\mathbf{J}}_o(\mathcal{S})$ and therefore G-GDOP.

In summary, G-GDOP is a trace-based A-optimal information measure evaluated in the gauge-free subspace and is finite only when the induced ranging topology is locally observable modulo the global rigid-body gauge. It extends conventional GDOP to topology-coupled, heteroscedastic cooperative ranging networks, while its validity remains local to the adopted predicted geometry and linearized measurement model. Dense networks may still exceed communication, sensing, and computational budgets; the next section therefore formulates the resource-constrained pre-selection problem and develops a marginal-gain-based greedy heuristic.

3. G-GDOP-Based Cooperative Sensor Node Pre-Selection

3.1. G-GDOP-Based Cooperative Sensor Node Pre-Selection Problem

In dense UWB-aided cooperative ranging sensor networks, bandwidth, ranging slots, energy, and onboard computation may prevent all mobile nodes from actively initiating ranging exchanges within one scheduling epoch. Let $\mathcal{C} = \{1, 2, \ldots, N\}$ denote the complete mobile-node set and $\mathcal{S} \subseteq \mathcal{C}$ the selected node subset. Under the fixed selected-node-budget case, $|\mathcal{S}| = K$, while all N nodes remain in the localization state and unselected nodes may still serve as ranging responders.

As discussed in Section 2.5.1, the selection criterion is G-GDOP$(\mathcal{S}) = \sqrt{\Phi_{\text{A}}(\mathcal{S})}$, where for a feasible topology $r_o(\mathcal{S}) = d_o$, $\Phi_{\text{A}}(\mathcal{S}) = \text{tr}[\mathbf{J}_o(\mathcal{S})^{-1}]$. The basic fixed-budget problem is therefore

$$\begin{aligned} \underset{\mathcal{S} \subseteq \mathcal{C}}{\text{minimize}} \quad & \text{tr}\left[\mathbf{J}_o(\mathcal{S})^{-1}\right] \\ \text{subject to} \quad & |\mathcal{S}| = K, \\ & r_o(\mathcal{S}) = d_o. \end{aligned} \tag{22}$$

Because the selected node subset induces the active ranging-link set, a candidate node does not have a subset-independent information increment. Define the fixed projected prior and per-link contributions as

$$\mathbf{J}_{o,\text{base}} = \mathbf{Q}_o^T \mathbf{J}_{\text{prior}} \mathbf{Q}_o, \qquad \mathbf{B}_e = \widehat{w}_e \mathbf{Q}_o^T \mathbf{h}_e^T \mathbf{h}_e \mathbf{Q}_o \succeq \mathbf{0}.$$

All practical per-link matrices use estimated online weights; oracle weights are reserved for explicitly labeled benchmark and diagnostic evaluations. Using the activation indicator $\chi_e(\mathcal{S})$ gives

$$\mathbf{J}_o(\mathcal{S}) = \mathbf{J}_{o,\text{base}} + \sum_{e \in \mathcal{E}_c} \chi_e(\mathcal{S}) \mathbf{B}_e. \tag{23}$$

Equivalently, $\mathbf{J}_o(\mathcal{S}) = \mathbf{J}_{o,\text{base}} + \sum_{e\in\mathcal{E}(\mathcal{S})} \mathbf{B}_e$. For a current subset $\mathcal{S}_t$ and candidate $i \notin \mathcal{S}_t$, define

$$\Delta\mathcal{E}_i(\mathcal{S}_t) = \mathcal{E}(\mathcal{S}_t \cup \{i\}) \setminus \mathcal{E}(\mathcal{S}_t), \qquad \Delta\mathbf{J}_{o,i}(\mathcal{S}_t) = \sum_{e\in\Delta\mathcal{E}_i(\mathcal{S}_t)} \mathbf{B}_e.$$

The candidate update is consequently $\mathbf{J}_o(\mathcal{S}_t \cup \{i\}) = \mathbf{J}_o(\mathcal{S}_t) + \Delta\mathbf{J}_{o,i}(\mathcal{S}_t)$, which makes its selection dependence explicit.

Induced active-link number. The selected-node budget controls selected nodes, not the number of ranging links generated by their induced topology. Define

$$M(\mathcal{S}) = |\mathcal{E}(\mathcal{S})| = \sum_{e\in\mathcal{E}_c} \chi_e(\mathcal{S}). \tag{24}$$

For the complete candidate set, $M(\mathcal{C}) = |\mathcal{E}_c|$, whereas the number of links newly activated by candidate i is $\Delta M_i(\mathcal{S}_t) = M(\mathcal{S}_t \cup \{i\}) - M(\mathcal{S}_t) = |\Delta\mathcal{E}_i(\mathcal{S}_t)|$. Thus the incremental link number equals the iteration-dependent update width before possible linear dependence among projected link vectors is considered.

Ranging-slot occupation. Let $s_e \in \mathbb{N}_+$ denote the scheduling slots required by candidate link e. The total slot occupation is

$$T_{\text{slot}}(\mathcal{S}) = \sum_{e\in\mathcal{E}_c} \chi_e(\mathcal{S}) s_e = \sum_{e\in\mathcal{E}(\mathcal{S})} s_e. \tag{25}$$

Adding candidate i requires $\Delta T_{\text{slot},i}(\mathcal{S}_t) = \sum_{e\in\Delta\mathcal{E}_i(\mathcal{S}_t)} s_e$. When $s_e = 1$ for every candidate link, $T_{\text{slot}}(\mathcal{S}) = M(\mathcal{S})$; the general definition also accommodates repeated ranging, handshakes, or multi-slot protocols.

3.1.1. General Sensing-Resource Cost

Let $c_i^{\text{node}} \geq 0$ denote the normalized cost of selecting node i as an active UWB ranging initiator and $c_e^{\text{link}} \geq 0$ the normalized communication, ranging, energy, or processing cost associated with active link e. The total cost is

$$C_{\text{res}}(\mathcal{S}) = \sum_{i\in\mathcal{S}} c_i^{\text{node}} + \sum_{e\in\mathcal{E}_c} \chi_e(\mathcal{S}) c_e^{\text{link}}. \tag{26}$$

Equivalently, $C_{\text{res}}(\mathcal{S}) = \sum_{i\in\mathcal{S}} c_i^{\text{node}} + \sum_{e\in\mathcal{E}(\mathcal{S})} c_e^{\text{link}}$, and the incremental cost of candidate i is $\Delta C_{\text{res},i}(\mathcal{S}_t) = c_i^{\text{node}} + \sum_{e\in\Delta\mathcal{E}_i(\mathcal{S}_t)} c_e^{\text{link}}$. The node and link coefficients must use a common normalized resource unit. If hardware-specific energy or communication measurements are unavailable, $|\mathcal{S}|$, $M(\mathcal{S})$, and $T_{\text{slot}}(\mathcal{S})$ are reported as the primary physical resource indicators rather than interpreting C_{res} as exact energy.

For the anchor-free range-only observability analysis, $\mathbf{J}_{\text{prior}} = \mathbf{0}$ and $\mathbf{J}_{o,\text{base}} = \mathbf{0}$; inverse-based A-optimal candidate evaluation is therefore performed only after a full-rank observable seed has been constructed. Beyond the selected-node budget and observability requirement, communication availability and residual energy can impose additional feasibility conditions.

3.1.2. Communication-Availability Constraint

A candidate link belongs to $\mathcal{E}_c$ only when its inter-node distance satisfies

$$d_{ij} \leq d_{\max}, \qquad \forall (i,j) \in \mathcal{E}_c. \tag{27}$$

Hence $\mathcal{E}(\mathcal{S}) \subseteq \mathcal{E}_c$ automatically, and this range condition is enforced during candidate-graph construction rather than repeated as a selection-dependent constraint in Equation (29).

3.1.3. Energy-Feasibility Constraint

A selected node must satisfy

$$E_i \geq E_{\min}, \qquad \forall i \in \mathcal{S}, \tag{28}$$

where E_i is the residual energy of node i and $E_{\min}$ the minimum threshold for initiating ranging. This is an eligibility condition on an individual node, whereas $C_{\text{res}}(\mathcal{S})$ is an aggregate per-epoch cost of the selected node subset and its induced topology.

3.1.4. Local-Observability Constraint

The projected information matrix must reach the complete gauge-free dimension, $r_o(\mathcal{S}) = d_o$, which is equivalent under the adopted numerical threshold to requiring the symmetrized $\mathbf{J}_o(\mathcal{S})$ to be positive definite. Combining the resource and observability requirements gives

$$\begin{aligned} \underset{\mathcal{S}\subseteq\mathcal{C}}{\text{minimize}} \quad & \Phi_{\text{A}}(\mathcal{S}) = \text{tr}\left[\mathbf{J}_o(\mathcal{S})^{-1}\right] \\ \text{subject to} \quad & |\mathcal{S}| \leq K, \\ & M(\mathcal{S}) \leq B_{\text{link}}, \\ & T_{\text{slot}}(\mathcal{S}) \leq B_{\text{slot}}, \\ & C_{\text{res}}(\mathcal{S}) \leq B_{\text{res}}, \\ & E_i \geq E_{\min}, \qquad \forall i \in \mathcal{S}, \\ & r_o(\mathcal{S}) = d_o. \end{aligned} \tag{29}$$

Here K, B_{link}, B_{slot}, and B_{res} are respectively the node, induced-link, ranging-slot, and general-resource budgets. The candidate graph has already been formed under

$$\mathcal{G}_c = (\mathcal{C}, \mathcal{E}_c), \qquad \mathcal{E}(\mathcal{S}) \subseteq \mathcal{E}_c, \qquad d_{ij} \leq d_{\max} \quad \forall (i,j) \in \mathcal{E}_c,$$

so these are graph properties rather than additional binary selection constraints. Fixed-K experiments use sufficiently loose link, slot, and general-resource budgets and impose $|\mathcal{S}| = K$; the complete resource-constrained formulation uses $|\mathcal{S}| \leq K$ because another budget may become active first.

3.1.5. Information–Resource Trade-Off Indicators

Let $\mathcal{S}_{\text{all}} = \mathcal{C}$ denote the All-Node Reference. If its induced topology is locally observable, define the quality-retention, resource-usage, and saving ratios by

$$\begin{aligned} Q_{\text{A}}(\mathcal{S}) &= \frac{\Phi_{\text{A}}(\mathcal{S}_{\text{all}})}{\Phi_{\text{A}}(\mathcal{S})}, \quad 0 < Q_{\text{A}}(\mathcal{S}) \leq 1, \\ \rho_{\text{link}}(\mathcal{S}) &= \frac{M(\mathcal{S})}{|\mathcal{E}_c|}, \qquad \rho_{\text{slot}}(\mathcal{S}) = \frac{T_{\text{slot}}(\mathcal{S})}{T_{\text{slot}}(\mathcal{S}_{\text{all}})}, \\ \rho_{\text{res}}(\mathcal{S}) &= \frac{C_{\text{res}}(\mathcal{S})}{C_{\text{res}}(\mathcal{S}_{\text{all}})}, \qquad s_q(\mathcal{S}) = 1 - \rho_q(\mathcal{S}), \quad q \in \{\text{link}, \text{slot}, \text{res}\}. \end{aligned}$$

A value of Q_{A} closer to one indicates greater retention of All-Node A-optimal information quality. A feasible solution $\mathcal{S}_a$ dominates $\mathcal{S}_b$ in the information–resource plane when

$$\Phi_{\text{A}}(\mathcal{S}_a) \leq \Phi_{\text{A}}(\mathcal{S}_b), \qquad C_{\text{res}}(\mathcal{S}_a) \leq C_{\text{res}}(\mathcal{S}_b),$$

with at least one strict inequality; the nondominated solutions form an empirical information–resource Pareto frontier.

The resulting optimization remains combinatorial: for fixed cardinality there are $\binom{N}{K}$ possible subsets. This growth makes exhaustive search suitable only for small-scale optimality verification and impractical for online dense-network pre-selection, motivating the marginal-gain greedy heuristic introduced next.

Figure 2 summarizes the resource-constrained formulation. The candidate graph, predicted geometry, and estimated link-quality statistics form the inputs, and the selected node subset is the decision variable. Active ranging links are induced by the endpoint-activation rule. Feasible selected node subsets are evaluated through the gauge-free projected information matrix under the common node, link, slot, energy, resource-cost, and local-observability constraints.

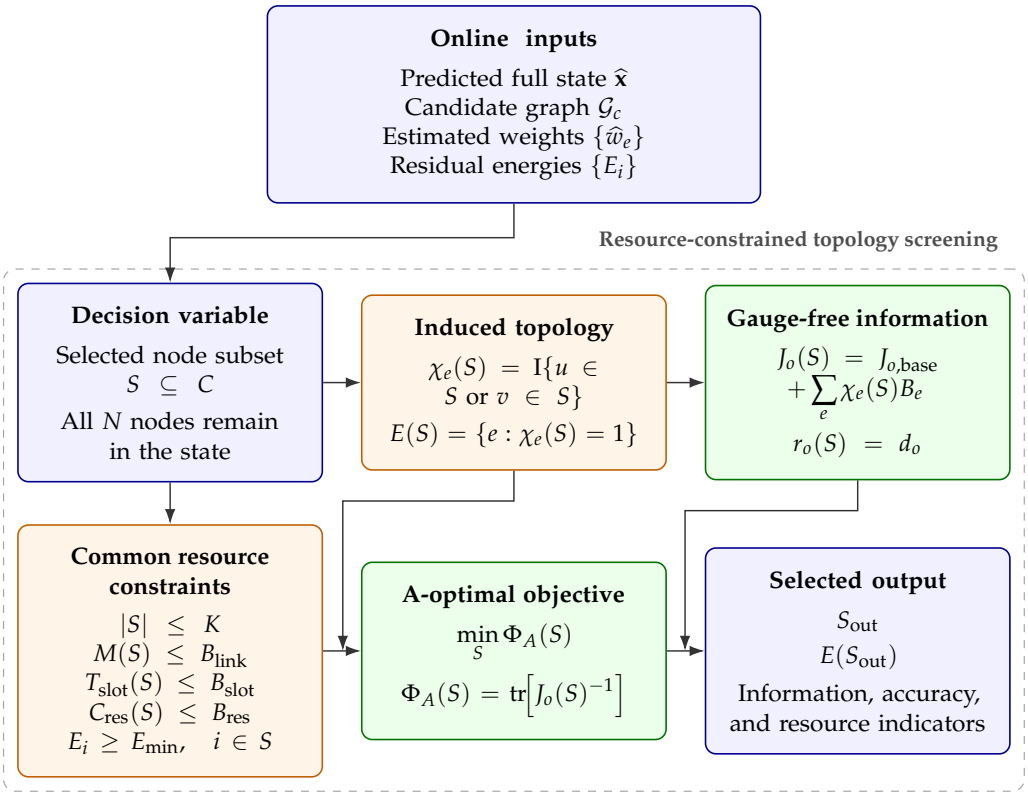


Figure 2. Resource-constrained cooperative sensor node pre-selection problem with induced ranging links and gauge-free information evaluation. Arrows indicate the dependency flow from the online inputs and selection decision to the induced topology, information matrix, feasibility constraints, objective, and final selected-node output.

3.2. Greedy Cooperative Sensor Node Pre-Selection Heuristic Based on G-GDOP Marginal Gain

The method consists of resource-feasible observable-seed construction followed by post-seed A-optimal selection. Let $\mathcal{S}_t \subseteq \mathcal{C}$ denote the selected node subset after iteration t. Both stages use

$$\mathcal{F}_t = \Big\{ i \in \mathcal{C} \setminus \mathcal{S}_t \;\Big|\; E_i \ge E_{\min},\; M(\mathcal{S}_t) + \Delta M_i(\mathcal{S}_t) \le B_{\mathrm{link}},\; T_{\mathrm{slot}}(\mathcal{S}_t) + \Delta T_{\mathrm{slot},i}(\mathcal{S}_t) \le B_{\mathrm{slot}},\; C_{\mathrm{res}}(\mathcal{S}_t) + \Delta C_{\mathrm{res},i}(\mathcal{S}_t) \le B_{\mathrm{res}} \Big\}.$$

Algorithm 1 gives the stage-specific evaluation, stopping, and failure logic.

Algorithm 1: Greedy Cooperative Sensor Node Pre-Selection Heuristic Based on G-GDOP

Input: C, $G_C = (C, \mathcal{E}_C)$, K, B_{link}, B_{slot}, B_{res}, $E_{\min}$, $\{E_i\}$, $\{s_e\}$, $\{c_i^{\text{node}}\}$, $\{c_e^{\text{link}}\}$, $\hat{\mathbf{x}}$, $\{\hat{w}_e\}$, τ_J, and ϵ_{tie}; use $+\infty$ to disable an upper budget.

Output: *status*, S_{out}, $\mathcal{E}(S_{\text{out}})$, M, T_{slot}, and C_{res}.

Compute Q_o, d_g, d_o, and the projected weighted link-information vectors $\{a_e\}_{e\in\mathcal{E}_C}$.
if preprocessing fails **then return** `SELECTION_NUMERICAL_FAILURE`.
$S \leftarrow \varnothing$; $\mathcal{E} \leftarrow \varnothing$; $M \leftarrow 0$; $T_{\text{slot}} \leftarrow 0$; $C_{\text{res}} \leftarrow 0$; $J_o \leftarrow 0$.
function STATS(i, S).
  $\Delta\mathcal{E}_i \leftarrow \mathcal{E}(S \cup \{i\}) \setminus \mathcal{E}(S)$.
  $\Delta M_i \leftarrow |\Delta\mathcal{E}_i|$; $\Delta T_i \leftarrow \sum_{e\in\Delta\mathcal{E}_i} s_e$.
  $\Delta C_i \leftarrow c_i^{\text{node}} + \sum_{e\in\Delta\mathcal{E}_i} c_e^{\text{link}}$.
  $\Delta J_{o,i} \leftarrow \sum_{e\in\Delta\mathcal{E}_i} a_e a_e^{\mathsf{T}}$.
  return $(\Delta\mathcal{E}_i, \Delta M_i, \Delta T_i, \Delta C_i, \Delta J_{o,i})$.
end function.
function FEASIBLE$(i, \Delta M_i, \Delta T_i, \Delta C_i)$.
  return $(E_i \geq E_{\min}) \wedge (|S| + 1 \leq K) \wedge (M + \Delta M_i \leq B_{\text{link}}) \wedge (T_{\text{slot}} + \Delta T_i \leq B_{\text{slot}}) \wedge (C_{\text{res}} + \Delta C_i \leq B_{\text{res}})$.
end function.
Stage I: observable-seed construction
while $\text{rank}_{\tau_J}(J_o) < d_o$ **do**.
  if $|S| \geq K$ **then return** `UNOBSERVABLE`.
  $C_{\text{rem}} \leftarrow C \setminus S$.
  if $C_{\text{rem}} = \varnothing$ **then return** `UNOBSERVABLE`.
  $\mathcal{F}_{\text{seed}} \leftarrow \varnothing$; $i^\star \leftarrow \varnothing$; $\kappa^\star \leftarrow \varnothing$; $r_{\text{cur}} \leftarrow \text{rank}_{\tau_J}(J_o)$.
  for each $i \in C_{\text{rem}}$ in increasing node-ID order **do**.
    $(\Delta\mathcal{E}_i, \Delta M_i, \Delta T_i, \Delta C_i, \Delta J_{o,i}) \leftarrow$ STATS(i, S).
    if $\neg$FEASIBLE$(i, \Delta M_i, \Delta T_i, \Delta C_i)$ **then continue**.
    $\mathcal{F}_{\text{seed}} \leftarrow \mathcal{F}_{\text{seed}} \cup \{i\}$; $J_o^{+i} \leftarrow J_o + \Delta J_{o,i}$.
    $r_i \leftarrow \text{rank}_{\tau_J}(J_o^{+i})$; compute the retained $\lambda_{\min}^{+i}$ of J_o^{+i}.
    $\kappa_i \leftarrow (r_i - r_{\text{cur}}, \lambda_{\min}^{+i}, -\Delta T_i, -\Delta C_i, -i)$.
    if $i^\star = \varnothing$ **or** $\kappa_i \succ_{\text{lex}} \kappa^\star$ **then** $i^\star \leftarrow i$ and $\kappa^\star \leftarrow \kappa_i$.
  end for.
  if $\mathcal{F}_{\text{seed}} = \varnothing$ **then return** `RESOURCE_INFEASIBLE`.
  $(\Delta\mathcal{E}_{i^\star}, \Delta M_{i^\star}, \Delta T_{i^\star}, \Delta C_{i^\star}, \Delta J_{o,i^\star}) \leftarrow$ STATS$(i^\star, S)$; $S \leftarrow S \cup \{i^\star\}$; $\mathcal{E} \leftarrow \mathcal{E} \cup \Delta\mathcal{E}_{i^\star}$.
  $M \leftarrow M + \Delta M_{i^\star}$; $T_{\text{slot}} \leftarrow T_{\text{slot}} + \Delta T_{i^\star}$; $C_{\text{res}} \leftarrow C_{\text{res}} + \Delta C_{i^\star}$; $J_o \leftarrow J_o + \Delta J_{o,i^\star}$.
end while.
$S_{\text{seed}} \leftarrow S$; $k_{\text{seed}} \leftarrow |S_{\text{seed}}|$.
if J_o^{-1} cannot be evaluated **then return** `SELECTION_NUMERICAL_FAILURE`.
$P \leftarrow J_o^{-1}$.
Stage II: post-seed A-optimal refinement
while $|S| < K$ **do**.
  $i^\star \leftarrow \varnothing$; $g^\star \leftarrow -\infty$; $\Delta T^\star \leftarrow +\infty$; $\Delta C^\star \leftarrow +\infty$; $P^\star \leftarrow P$.
  for each $i \in C \setminus S$ in increasing node-ID order **do**.
    $(\Delta\mathcal{E}_i, \Delta M_i, \Delta T_i, \Delta C_i, \Delta J_{o,i}) \leftarrow$ STATS(i, S).
    if $\neg$FEASIBLE$(i, \Delta M_i, \Delta T_i, \Delta C_i)$ **then continue**.
    if $\Delta\mathcal{E}_i = \varnothing$ **then** $P^{(i)} \leftarrow P$; **else** compute $P^{(i)}$ using Equation (37).
    $g_i \leftarrow \text{tr}(P) - \text{tr}(P^{(i)})$.
    if $i^\star = \varnothing$ **or** $g_i > g^\star + \epsilon_{\text{tie}}$ **then** store $(i, g_i, \Delta T_i, \Delta C_i, P^{(i)})$ as the current candidate.
    else if $|g_i - g^\star| \leq \epsilon_{\text{tie}}$ **and** $\Delta T_i < \Delta T^\star - \epsilon_{\text{tie}}$ **then** store $(i, g_i, \Delta T_i, \Delta C_i, P^{(i)})$.
    else if $|g_i - g^\star| \leq \epsilon_{\text{tie}}$, $|\Delta T_i - \Delta T^\star| \leq \epsilon_{\text{tie}}$, and $\Delta C_i < \Delta C^\star - \epsilon_{\text{tie}}$ **then** store $(i, g_i, \Delta T_i, \Delta C_i, P^{(i)})$.
    else if $|g_i - g^\star| \leq \epsilon_{\text{tie}}$, $|\Delta T_i - \Delta T^\star| \leq \epsilon_{\text{tie}}$, $|\Delta C_i - \Delta C^\star| \leq \epsilon_{\text{tie}}$, and $i < i^\star$ **then** store $(i, g_i, \Delta T_i, \Delta C_i, P^{(i)})$.
  end for.
  if $i^\star = \varnothing$ **then break**.
  $(\Delta\mathcal{E}_{i^\star}, \Delta M_{i^\star}, \Delta T_{i^\star}, \Delta C_{i^\star}, \Delta J_{o,i^\star}) \leftarrow$ STATS$(i^\star, S)$.
  $S \leftarrow S \cup \{i^\star\}$; $\mathcal{E} \leftarrow \mathcal{E} \cup \Delta\mathcal{E}_{i^\star}$.
  $M \leftarrow M + \Delta M_{i^\star}$; $T_{\text{slot}} \leftarrow T_{\text{slot}} + \Delta T_{i^\star}$; $C_{\text{res}} \leftarrow C_{\text{res}} + \Delta C_{i^\star}$.
  $J_o \leftarrow J_o + \Delta J_{o,i^\star}$; $P \leftarrow P^\star$.
end while.
if exact-K selection is required **and** $|S| < K$ **then return** `RESOURCE_INFEASIBLE`.
if $\text{rank}_{\tau_J}(J_o) < d_o$ **then return** `UNOBSERVABLE`.
Evaluate the final A-optimal metric and resource statistics.
if the final metric evaluation fails **then return** `SELECTION_NUMERICAL_FAILURE`.
$S_{\text{out}} \leftarrow S$.
return `SUCCESS`, S_{out}, $\mathcal{E}(S_{\text{out}})$, M, T_{slot}, C_{res}.

After a full-rank seed exists, the marginal gain for $i \in \mathcal{F}_t$ is

$$\Delta(i \mid \mathcal{S}_t) = \Phi_{\text{A}}(\mathcal{S}_t) - \Phi_{\text{A}}(\mathcal{S}_t \cup \{i\}). \tag{30}$$

Its dependence on $\mathcal{S}_t$ reflects the fact that some links incident to i may already have been activated. The greedy rule is

$$i^\star = \arg\max_{i \in \mathcal{F}_t} \Delta(i \mid \mathcal{S}_t), \tag{31}$$

with numerically tied gains resolved first by smaller $\Delta T_{\mathrm{slot},i}(\mathcal{S}_t)$ and then by smaller $\Delta C_{\mathrm{res},i}(\mathcal{S}_t)$. The subset update is simply $\mathcal{S}_{t+1} = \mathcal{S}_t \cup \{i^\star\}$.

Because an empty anchor-free network is rank deficient, Stage I first constructs $\mathcal{S}_{\mathrm{seed}}$ satisfying

$$r_o(\mathcal{S}_{\mathrm{seed}}) = d_o. \tag{32}$$

Writing $k_{\mathrm{seed}} = |\mathcal{S}_{\mathrm{seed}}|$, if $k_{\mathrm{seed}} = K$ the node budget is already exhausted and the returned selected node subset is determined entirely by the Stage I lexicographic observable-seed ordering. Otherwise, Stage II starts from $\mathcal{S}_{k_{\mathrm{seed}}} = \mathcal{S}_{\mathrm{seed}}$ and $\mathbf{P}_{k_{\mathrm{seed}}} = \bar{\mathbf{J}}_o(\mathcal{S}_{\mathrm{seed}})^{-1}$. The detailed ordering, resource checks, and failure returns remain those in Algorithm 1.

3.2.1. Set-Dependent Induced-Link Information Increment

For $i \in \mathcal{F}_t$, the newly activated link set is

$$\Delta\mathcal{E}_i(\mathcal{S}_t) = \mathcal{E}(\mathcal{S}_t \cup \{i\}) \setminus \mathcal{E}(\mathcal{S}_t). \tag{33}$$

This set difference prevents double counting under endpoint activation. The associated projected increment is

$$\Delta\mathbf{J}_{o,i}(\mathcal{S}_t) = \sum_{e \in \Delta\mathcal{E}_i(\mathcal{S}_t)} \mathbf{B}_e, \tag{34}$$

equivalently $\Delta\mathbf{J}_{o,i}(\mathcal{S}_t) = \sum_{e \in \Delta\mathcal{E}_i(\mathcal{S}_t)} \widehat{w}_e \mathbf{Q}_o^T \mathbf{h}_e^T \mathbf{h}_e \mathbf{Q}_o$. Hence $\mathbf{J}_o(\mathcal{S}_t \cup \{i\}) = \mathbf{J}_o(\mathcal{S}_t) + \Delta\mathbf{J}_{o,i}(\mathcal{S}_t)$. Both the increment and marginal gain must therefore be reconstructed at every iteration. If $\Delta\mathcal{E}_i(\mathcal{S}_t) = \varnothing$, then $\Delta\mathbf{J}_{o,i}(\mathcal{S}_t) = \mathbf{0}$ and $\Delta(i \mid \mathcal{S}_t) = 0$. After selecting $i^\star$, $\mathcal{E}_{t+1} = \mathcal{E}_t \cup \Delta\mathcal{E}_{i^\star}(\mathcal{S}_t)$ with $\mathcal{E}_t = \mathcal{E}(\mathcal{S}_t)$.

3.2.2. Low-Rank Recursive Inverse Update

Candidate evaluation uses the estimated weight $\widehat{w}_e$ from Equation (7). For each newly activated link define $\mathbf{a}_e = \sqrt{w_e}\mathbf{Q}_o^T \mathbf{h}_e^T \in \mathbb{R}^{d_o}$, and collect the $r_{i,t} = |\Delta\mathcal{E}_i(\mathcal{S}_t)|$ vectors in $\mathbf{U}_i(\mathcal{S}_t) = [\mathbf{a}_{e_1}, \ldots, \mathbf{a}_{e_{r_{i,t}}}] \in \mathbb{R}^{d_o \times r_{i,t}}$. The low-rank information increment is

$$\Delta\bar{\mathbf{J}}_{o,i}(\mathcal{S}_t) = \mathbf{U}_i(\mathcal{S}_t)\mathbf{U}_i(\mathcal{S}_t)^T. \tag{35}$$

With current inverse

$$\mathbf{P}_t = \bar{\mathbf{J}}_o(\mathcal{S}_t)^{-1}, \tag{36}$$

the candidate inverse is obtained from

$$\begin{aligned} \mathbf{P}_t^{(i)} = \mathbf{P}_t &- \mathbf{P}_t\mathbf{U}_i(\mathcal{S}_t) \\ &\times \left[\mathbf{I}_{r_{i,t}} + \mathbf{U}_i(\mathcal{S}_t)^T\mathbf{P}_t\mathbf{U}_i(\mathcal{S}_t)\right]^{-1} \\ &\times \mathbf{U}_i(\mathcal{S}_t)^T\mathbf{P}_t. \end{aligned} \tag{37}$$

After $i^\star$ is selected, $\bar{\mathbf{J}}_o(\mathcal{S}_{t+1}) = \bar{\mathbf{J}}_o(\mathcal{S}_t) + \Delta\bar{\mathbf{J}}_{o,i^\star}(\mathcal{S}_t)$ and

$$\mathbf{P}_{t+1} = \mathbf{P}_t^{(i^\star)}. \tag{38}$$

All candidate matrices remain in the same $d_o = 3N - d_g$ gauge-free subspace for fixed predicted geometry, so the Woodbury update is dimensionally consistent after observable-

seed construction. For the nondegenerate three-dimensional experiments, $d_g = 6$ and $d_o = 3N - 6$.

Algorithm 1 provides the executable selection procedure, Figure 3 summarizes the high-level two-stage workflow, and Figure 4 gives a concise structural illustration of the corresponding induced-topology evolution.

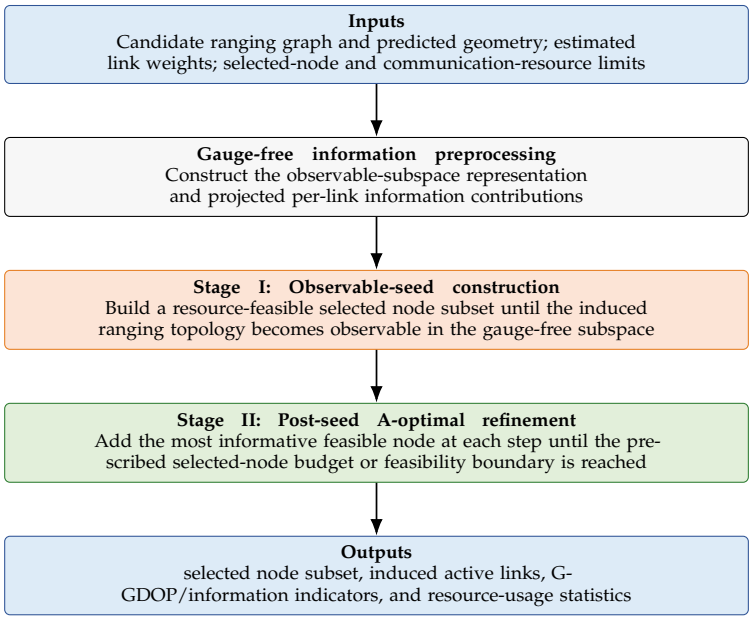


Figure 3. High-level architecture of the proposed two-stage cooperative sensor node pre-selection framework. Algorithm 1 provides the executable candidate-screening, update, tie-breaking, and failure logic, whereas Figure 4 illustrates the corresponding induced-topology evolution.

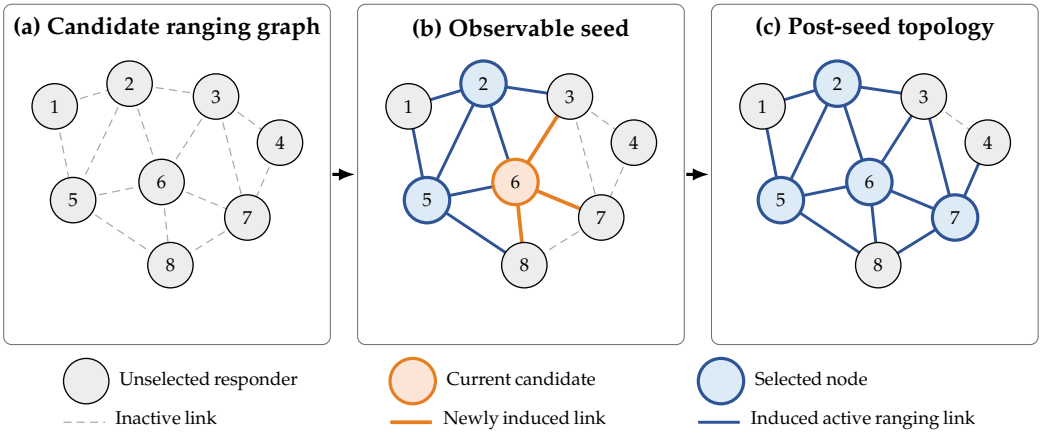


Figure 4. Illustration of induced-topology evolution from the candidate ranging graph to an observable seed and subsequent post-seed refinement. Colors distinguish selected nodes, unselected mobile nodes, candidate links, induced active ranging links, and the links newly induced by the highlighted candidate; the detailed decision logic is given in Algorithm 1. The numbers denote mobile-node IDs.

The proposed procedure is treated as a computationally efficient greedy heuristic for engineering implementation. For the general optimization problem

$$\min_{\substack{\mathcal{S}\subseteq\mathcal{C}\\ |\mathcal{S}|=K}} \operatorname{tr}\left[\bar{\mathbf{J}}_o(\mathcal{S})^{-1}\right].$$

the inverse-trace A-optimal objective is not generally submodular without additional structural assumptions. In particular, the information contribution of a candidate node

may depend on the currently selected node subset and the corresponding induced ranging topology. Therefore, no general constant-factor approximation guarantee is claimed for the proposed method. Instead, the solution quality of the proposed greedy heuristic is evaluated empirically through repeated exhaustive-search comparisons in small-scale networks and paired Monte Carlo experiments under different operating conditions.

For the following complexity analysis, the candidate graph $\mathcal{G}_c = (\mathcal{C}, \mathcal{E}_c)$ is represented using adjacency lists together with a Boolean active-link mask. The projected weighted information vector $\mathbf{a}_e \in \mathbb{R}^{d_o}$ is precomputed for each candidate ranging link $e \in \mathcal{E}_c$.

Candidate nodes are evaluated sequentially in the complexity bounds reported below. Parallel candidate evaluation is possible within one greedy iteration, but its additional synchronization and memory costs are not included in the serial bounds.

3.3. Algorithm Complexity and Scalability Analysis

The complexity depends on the candidate-node count, graph density, gauge-free dimension, update width, and resource-screening cost. The main quantities are $N = |\mathcal{C}|$, $M_c = |\mathcal{E}_c|$, $d_o = 3N - d_g$, $k_{\text{seed}} = |\mathcal{S}_{\text{seed}}|$, and $K_{\text{out}} = |\mathcal{S}_{\text{out}}| \le K$. Under a fixed-$K$ experiment with loose resource budgets, $K_{\text{out}} = K$; otherwise a resource constraint may cause $K_{\text{out}} < K$. At iteration t, $\mathcal{F}_t \subseteq \mathcal{C} \setminus \mathcal{S}_t$ is the resource-feasible candidate set, $n_t = |\mathcal{F}_t|$ is the number requiring method-specific evaluation, and $r_{i,t} = |\Delta\mathcal{E}_i(\mathcal{S}_t)|$ is the number of newly activated links for candidate i. Define $r_{\max} = \max_{t,i} r_{i,t}$. Since $r_{i,t}$ is the number of columns of $\mathbf{U}_i(\mathcal{S}_t)$, the algebraic rank of an information increment is at most $r_{i,t}$.

3.3.1. Preprocessing Complexity

Preprocessing consists of candidate-graph construction, gauge-free basis construction, and precomputation of projected per-link vectors. Direct pairwise graph construction costs $\mathcal{O}(N^2)$, although this term is replaced by the corresponding graph-construction cost when the graph is supplied externally or obtained through spatial indexing. With $\mathbf{G} \in \mathbb{R}^{3N \times d_g}$, $\mathbf{Q}_o \in \mathbb{R}^{3N \times d_o}$, and $d_g \le 6$, constructing the explicit orthogonal complement requires approximately $\mathcal{O}(Nd_o d_g)$ operations. Each projected vector

$$\mathbf{a}_e = \sqrt{\widehat{w}_e}\, \mathbf{Q}_o^T \mathbf{h}_e^T \in \mathbb{R}^{d_o}$$

can be formed in $\mathcal{O}(d_o)$ time by exploiting Jacobian block sparsity, so all M_c vectors require $\mathcal{O}(M_c d_o)$. The total direct-graph preprocessing cost is therefore

$$\mathcal{C}_{\text{pre}} = \mathcal{O}\Big(N^2 + Nd_o d_g + M_c d_o\Big). \tag{39}$$

When the candidate graph is already available, the $\mathcal{O}(N^2)$ term is omitted.

3.3.2. Complexity of Exhaustive Search

For fixed cardinality $|\mathcal{S}| = K$, exhaustive search evaluates $\mathcal{N}_{\text{exh}} = \binom{N}{K}$ subsets; for the general $|\mathcal{S}| \le K$ case, up to $\mathcal{N}_{\text{exh}}^{(\le K)} = \sum_{k=0}^{K} \binom{N}{k}$ subsets may require consideration before resource pruning. For one subset, induced-link and resource evaluation costs at most $\mathcal{O}(M_c)$, accumulation of rank-one projected information contributions costs $\mathcal{O}(M(\mathcal{S})d_o^2)$, and dense numerical-rank plus inverse-trace evaluation costs $\mathcal{O}(d_o^3)$. A fixed-cardinality upper-order expression is therefore

$$\mathcal{C}_{\text{exh}} = \mathcal{O}\left[\binom{N}{K}\left(M_c + M_c d_o^2 + d_o^3\right)\right]. \tag{40}$$

Resource constraints reduce the feasible set in practice but do not remove worst-case combinatorial growth, so exhaustive search is used only for repeated small-scale optimality-gap verification.

3.3.3. Resource-Feasibility Screening Cost

Let $\deg_c(i)$ denote the candidate-graph degree of node i. For an unselected candidate, newly activated links and the increments ΔM_i, $\Delta T_{\mathrm{slot},i}$, and $\Delta C_{\mathrm{res},i}$ are obtained by scanning its incident candidate links, giving $\mathcal{O}(1+\deg_c(i))$ screening cost. At iteration t,

$$\begin{aligned}\mathcal{C}_{\mathrm{screen},t} &= \mathcal{O}\left[(N-t)+\sum_{i\in\mathcal{C}\setminus\mathcal{S}_t}\deg_c(i)\right]\\ &= \mathcal{O}(N+M_c),\end{aligned}$$

because the degree sum is at most $2M_c$. Over the full procedure,

$$\mathcal{C}_{\mathrm{screen}} = \mathcal{O}[K_{\mathrm{out}}(N+M_c)].$$

This dependence exposes the role of graph density: $M_c = \mathcal{O}(N)$ for a sparse graph but can reach $\mathcal{O}(N^2)$ in a dense graph.

3.3.4. Complexity of Observable-Seed Construction

Before the projected matrix reaches full gauge-free rank, the inverse-trace objective is unavailable. Stage I therefore compares the numerical rank increase and minimum retained positive eigenvalue produced by each resource-feasible candidate. For candidate i at seed iteration t,

$$\mathbf{J}_o(\mathcal{S}_t\cup\{i\}) = \mathbf{J}_o(\mathcal{S}_t) + \mathbf{U}_i(\mathcal{S}_t)\mathbf{U}_i(\mathcal{S}_t)^T.$$

Explicit formation costs $\mathcal{O}(d_o^2 r_{i,t})$, while a dense symmetric eigendecomposition for the rank and minimum positive eigenvalue costs $\mathcal{O}(d_o^3)$ per candidate. Hence

$$\mathcal{C}_{\mathrm{seed}} = \mathcal{O}\left[\sum_{t=0}^{k_{\mathrm{seed}}-1}\sum_{i\in\mathcal{F}_t}\left(d_o^2 r_{i,t}+d_o^3\right)\right],$$

and using $n_t \le N-t$ and $r_{i,t}\le r_{\max}$ gives

$$\mathcal{C}_{\mathrm{seed}} = \mathcal{O}\left[k_{\mathrm{seed}}N\left(d_o^2 r_{\max}+d_o^3\right)\right]. \tag{41}$$

This seed cost is part of the complete method; the recursive A-optimal complexity below applies only after a full-rank seed exists. A rank-revealing incremental factorization could reduce the seed-stage cost but is not assumed in the reported bound.

3.3.5. Post-Seed Greedy Evaluation with Direct Inversion

After seed construction, a direct-inversion implementation evaluates each feasible candidate by forming and inverting its complete projected information matrix. For candidate i at iteration t, forming the low-rank increment costs $\mathcal{O}(d_o^2 r_{i,t})$ and directly inverting the resulting $d_o\times d_o$ matrix costs $\mathcal{O}(d_o^3)$. Consequently,

$$\mathcal{C}_{\mathrm{direct}} = \mathcal{O}\left[\sum_{t=k_{\mathrm{seed}}}^{K_{\mathrm{out}}-1}\sum_{i\in\mathcal{F}_t}\left(d_o^2 r_{i,t}+d_o^3\right)\right],$$

with upper bound

$$\mathcal{C}_{\text{direct}} = \mathcal{O}\Big[(K_{\text{out}} - k_{\text{seed}})N\Big(d_o^2 r_{\max} + d_o^3\Big)\Big]. \tag{42}$$

3.3.6. Post-Seed Greedy Evaluation with Woodbury Updates

For candidate i at iteration t, $\mathbf{U}_i(\mathcal{S}_t) \in \mathbb{R}^{d_o \times r_{i,t}}$ contains the projected vectors of newly activated links and $\mathbf{P}_t = \mathbf{J}_o(\mathcal{S}_t)^{-1}$. One Woodbury candidate evaluation requires $\mathcal{O}(d_o^2 r_{i,t})$ for $\mathbf{P}_t\mathbf{U}_i$, $\mathcal{O}(d_o r_{i,t}^2)$ for the inner product, $\mathcal{O}(r_{i,t}^3)$ for inversion of the inner matrix, and the same leading-order multiplication terms for the trace correction and inverse update. Therefore,

$$\mathcal{C}_{i,t}^{\text{WB}} = \mathcal{O}\Big(d_o^2 r_{i,t} + d_o r_{i,t}^2 + r_{i,t}^3\Big). \tag{43}$$

The exact post-seed recursive cost is

$$\mathcal{C}_{\text{A}} = \mathcal{O}\left[\sum_{t=k_{\text{seed}}}^{K_{\text{out}}-1} \sum_{i \in \mathcal{F}_t} \Big(d_o^2 r_{i,t} + d_o r_{i,t}^2 + r_{i,t}^3\Big)\right],$$

and the bounds $n_t \le N - t$ and $r_{i,t} \le r_{\max}$ yield

$$\begin{aligned}\mathcal{C}_{\text{A}} = \mathcal{O}\Big[&(K_{\text{out}} - k_{\text{seed}})N \\ &\times \Big(d_o^2 r_{\max} + d_o r_{\max}^2 + r_{\max}^3\Big)\Big].\end{aligned} \tag{44}$$

For a sparse bounded-degree candidate graph, $r_{\max}$ is effectively constant and $M_c = \mathcal{O}(N)$, reducing the post-seed stage approximately to $\mathcal{O}[(K_{\text{out}} - k_{\text{seed}})Nd_o^2]$. This simplification does not apply automatically in a dense network, where $r_{\max}$ can grow with N.

3.3.7. Total Time Complexity of the Proposed Procedure

Combining preprocessing, resource screening, seed construction, and post-seed recursive evaluation gives $\mathcal{C}_{\text{total}} = \mathcal{C}_{\text{pre}} + \mathcal{C}_{\text{screen}} + \mathcal{C}_{\text{seed}} + \mathcal{C}_{\text{A}}$, or explicitly

$$\begin{aligned}\mathcal{C}_{\text{total}} = \mathcal{O}\Big[&N^2 + Nd_o d_g + M_c d_o \\ &+ K_{\text{out}}(N + M_c) \\ &+ \sum_{t=0}^{k_{\text{seed}}-1} \sum_{i \in \mathcal{F}_t} \Big(d_o^2 r_{i,t} + d_o^3\Big) \\ &+ \sum_{t=k_{\text{seed}}}^{K_{\text{out}}-1} \sum_{i \in \mathcal{F}_t} \Big(d_o^2 r_{i,t} + d_o r_{i,t}^2 + r_{i,t}^3\Big)\Big].\end{aligned} \tag{45}$$

Equation (45) shows that the method avoids combinatorial subset enumeration but still depends on candidate-graph density, resource screening, observable-seed size, and iteration-dependent update widths. The compact $\mathcal{O}(KNd_o^2)$ expression is therefore only an approximate post-seed description under a sparse bounded-degree graph, not the unconditional complexity of the complete algorithm.

3.3.8. Space Complexity

For serial evaluation, adjacency lists and active-link masks require $\mathcal{O}(N + M_c)$ storage, the explicit gauge-free basis requires $\mathcal{O}(Nd_o)$, projected per-link vectors require $\mathcal{O}(M_c d_o)$, and the current projected information matrix plus inverse require $\mathcal{O}(d_o^2)$. The largest

temporary Woodbury matrices require $\mathcal{O}(d_o r_{\max} + r_{\max}^2)$. Hence the total serial space complexity is

$$\begin{aligned} \mathcal{C}_{\text{space}} = \mathcal{O}\big(&N + M_c + N d_o + M_c d_o \\ &+ d_o^2 + d_o r_{\max} + r_{\max}^2\big). \end{aligned} \tag{46}$$

A dense $\mathbf{B}_e \in \mathbb{R}^{d_o \times d_o}$ is not stored for every candidate link because that would require $\mathcal{O}(M_c d_o^2)$ memory; only $\mathbf{a}_e$ is retained and its rank-one contribution is formed when needed. Likewise, no fixed candidate-node increment is stored because $\Delta\mathcal{E}_i(\mathcal{S}_t)$ and $\mathbf{U}_i(\mathcal{S}_t)$ depend on the current selected node subset. Parallel evaluation within an iteration would require additional memory proportional to the number of concurrent candidates; the bound in Equation (46) assumes sequential candidate evaluation.

As summarized in Table 2, exhaustive search retains combinatorial dependence on the candidate-node number. The proposed heuristic replaces subset enumeration with sequential candidate evaluation and uses low-rank recursive updates after a full-rank observable seed has been constructed.

Table 2. Complexity comparison of exhaustive and greedy cooperative sensor node pre-selection procedures.

Procedure	Time Complexity	Main Dependence
Preprocessing	$\mathcal{O}(N^2 + N d_o d_g + M_c d_o)$	Graph, gauge basis, and per-link vectors
Fixed-cardinality exhaustive search	$\mathcal{O}\left[\binom{N}{K}(M_c + M_c d_o^2 + d_o^3)\right]$	Combinatorial subset enumeration
Resource-feasibility screening	$\mathcal{O}[K_{\text{out}}(N + M_c)]$	Candidate graph and resource budgets
Observable-seed construction	$\mathcal{O}\left[\sum_{t<k_{\text{seed}}} \sum_{i\in\mathcal{F}_t} (d_o^2 r_{i,t} + d_o^3)\right]$	Rank and eigenvalue tests
Post-seed direct-inversion greedy	$\mathcal{O}\left[\sum_{t\geq k_{\text{seed}}} \sum_{i\in\mathcal{F}_t} (d_o^2 r_{i,t} + d_o^3)\right]$	Full candidate inversion
Post-seed recursive greedy	$\mathcal{O}\left[\sum_{t\geq k_{\text{seed}}} \sum_{i\in\mathcal{F}_t} \left(d_o^2 r_{i,t} + d_o r_{i,t}^2 + r_{i,t}^3\right)\right]$	Woodbury update width
Sparse bounded-degree special case	$\mathcal{O}[(K_{\text{out}} - k_{\text{seed}}) N d_o^2]$	Post-seed stage only

The resulting procedure is polynomial in the number of candidate evaluations, but its practical runtime is not determined by N alone. Candidate-graph density, the observable-seed stage, resource screening, and the distribution of $r_{i,t}$ also affect the computational cost. Consequently, no unconditional real-time claim is made solely from the asymptotic expression. Empirical runtime and memory measurements are reported in the scalability experiment.

Candidate evaluations within the same iteration are conditionally independent given the current selected node subset and may be parallelized. However, the experiments in this paper use a serial implementation so that all reported running times correspond to the complexity model derived above.

4. Simulation Validation

4.1. Experimental Settings and Evaluation Metrics

To evaluate the proposed G-GDOP-based UWB cooperative sensor node pre-selection method, simulation experiments are conducted in a controllable three-dimensional environment. The experiments evaluate the applicability boundary of G-GDOP, information–accuracy–resource trade-offs under selected-node and ranging-slot budgets, repeated small-scale comparisons with exhaustive enumeration, robustness to Gaussian ranging noise, heterogeneous ranging noise, and NLOS bias, sensitivity to predicted positions and link-quality estimates, and computational scalability.

4.1.1. Simulation Scenario and General Parameter Settings

A three-dimensional UWB cooperative localization simulation environment is constructed in this study. The simulation region is set as 1000 m × 1000 m × 300 m. Simulated UAV-borne sensor nodes are uniformly and randomly distributed in the three-dimensional

space. When the distance between two nodes is smaller than the communication/ranging radius R_c, an effective UWB ranging link is assumed to be established between them. Unless otherwise specified, the default total number of nodes is set to $N = 30$, and the default number of selected nodes is set to $K = 10$. Unless otherwise specified, the experiments use the anchor-free range-only formulation: all N mobile nodes remain in the localization state, K nodes are scheduled as active UWB ranging initiators, and the corresponding active links are induced by that selection. Each Monte Carlo trial represents one static network snapshot; closed-loop UAV motion and communication dynamics are outside the present simulation scope. The G-GDOP, A-optimal, and D-optimal criteria use the gauge-free projected Fisher information matrix and full-rank observability test defined in Section 2.3. For localization-error evaluation, the estimated configuration of all N mobile nodes is rigidly aligned with the corresponding ground-truth configuration before the RMSE is calculated. Hence, the reported RMSE measures the recovery accuracy of the complete mobile-node configuration under the ranging topology induced by the selected nodes, rather than the error of only the selected nodes. Table 3 lists the general simulation parameters used in this study.

Table 3. General simulation parameter settings.

Parameter	Description	Default Value
Simulation region	Three-dimensional cooperative localization space	1000 m × 1000 m × 300 m
N	Total number of UAV-borne mobile sensor nodes	30
K	Number of selected nodes	10
R_c	Communication/ranging radius	450 m
σ	Standard deviation of ranging noise	0.1 m
Node distribution	Spatial distribution of UAV nodes	Three-dimensional random distribution
Localization scenario	Cooperative localization mode	Anchor-free full-state inter-node ranging localization
RMSE calculation	Error calculation method	Calculated for all N mobile nodes after rigid alignment
Information metric calculation	Treatment of rigid-body gauge directions	Gauge-free projection with a numerical full-rank test
Gauge-rank threshold	Relative singular-value threshold for $\mathbf{G}$	$\tau_g = 10^{-10}$
Projected-FIM rank threshold	Relative eigenvalue threshold for $\mathbf{J}_o$	$\tau_J = 10^{-8}$
Active-link budget	Maximum number of induced active links	Experiment dependent
Ranging-slot model	Slot occupation of one active ranging link	$s_e = 1$ by default
Ranging-slot budget	Maximum total ranging-slot occupation	Experiment dependent
Selected-node cost	Normalized cost of one selected node	$c_i^{\text{node}} = 1$ in the illustrative cost model
Active-link cost	Normalized cost of one induced ranging link	$c_e^{\text{link}} = 1$ in the illustrative cost model
General resource budget	Maximum normalized resource cost	Experiment dependent
Monte Carlo trials	Number of paired trials per operating condition	$R_{\text{MC}} = 100$
Bootstrap repetitions	Number of bootstrap resamples	$B_{\text{boot}} = 10{,}000$
Runtime repetitions	Repeated executions after one warm-up run	10
Statistical summary	Primary summary for continuous metrics	Median and pointwise 95% bootstrap CI
Binary-rate interval	Confidence interval for feasibility and success rates	95% Wilson CI

The UWB ranging-noise parameters are chosen to represent decimeter-level UWB ranging accuracy and the degradation of ranging precision under LOS/NLOS conditions. The simulation region, node scale, and communication/ranging radius are used to construct a controllable three-dimensional cooperative ranging network. Their effects will be further analyzed in the subsequent node-budget experiment, noise robustness experiment, and scalability experiment.

4.1.2. UWB Ranging Error and Estimated Link-Quality Settings

For each measurable UWB link $e = (i, j)$, the simulated observation is generated according to

$$z_e = \left\|\mathbf{p}_i - \mathbf{p}_j\right\|_2 + b_e + n_e, \qquad n_e \sim \mathcal{N}\left(0, \sigma_e^2\right). \tag{47}$$

Three ranging-error regimes are considered. First, in the homoscedastic Gaussian-noise experiments, all active links use

$$\sigma_e = 0.10 \text{ m}, \qquad b_e = 0.$$

Second, in the heteroscedastic random-noise experiment, the links are divided into LOS, wLOS, and NLOS/degraded classes with proportions of 50%, 30%, and 20%, respectively. Their random-error standard deviations are set to 0.05 m, 0.10 m, and 0.30 m, while all systematic bias terms are set to zero. Third, in the NLOS-bias experiment, the same link classes and random standard deviations are retained, and a positive bias is introduced only into the NLOS/degraded links:

$$b_e = \begin{cases} 0, & c_e \in \{\text{LOS}, \text{wLOS}\}, \\ b_{\text{NLOS}}, & c_e = \text{NLOS}, \end{cases}$$

where

$$b_{\text{NLOS}} \in \{0.30, 0.50, 0.70, 1.00\} \text{ m}.$$

Two link-quality information settings are considered. In the oracle setting, the true class and true random-error variance are provided to the pre-selection method. This setting is used only as an ideal upper-reference benchmark for evaluating the weighting mechanism.

In the practical estimated-quality setting, the link class is imperfectly identified. Let $\eta \in [0, 1]$ denote the adjacent-class misclassification probability. The class-confusion matrix is defined as

$$\mathbf{\Gamma}(\eta) = \begin{bmatrix} 1-\eta & \eta & 0 \\ \eta/2 & 1-\eta & \eta/2 \\ 0 & \eta & 1-\eta \end{bmatrix}, \tag{48}$$

where the rows correspond to the true classes $\{\text{LOS}, \text{wLOS}, \text{NLOS}\}$ and the columns correspond to the estimated classes.

In addition to class misclassification, a relative standard-deviation estimation error is introduced as

$$\widehat{\sigma}_e = \max\left\{\sigma_{\widehat{c}_e}(1 + \epsilon_e), \sigma_{\min}\right\}, \qquad \epsilon_e \sim \mathcal{U}(-\rho, \rho), \tag{49}$$

where ρ controls the variance-estimation uncertainty and $\sigma_{\min} > 0$ prevents a nonphysical nonpositive standard deviation. The confusion matrix and variance perturbation form a controlled sensitivity model rather than a trained LOS/NLOS classifier or a complete non-Gaussian UWB channel model. Thus, the experiment tests the effect of imperfect quality information within the prescribed model, while validation with classifier outputs

and measured UWB error distributions remains outside the scope of the present study. The estimated weight supplied to the practical pre-selection method is

$$\widehat{w}_e = \frac{\sigma_0^2}{\widehat{\sigma}_e^2}.$$

All competing methods in a paired Monte Carlo trial use the same node topology, true link classes, random ranging noise realization, NLOS bias realization, and estimated-quality realization. The physical UWB measurements are generated using the true random-error covariance $\mathbf{R}_{\mathcal{S}}$. The practical weighted pre-selection method and the practical weighted nonlinear localization solver use $\widehat{\mathbf{R}}_{\mathcal{S}}$. The true covariance is supplied to a method only when that method is explicitly labeled as an oracle benchmark. In the unmodeled-NLOS-bias stress test, the injected positive bias magnitude is not included in the selection weights and is not supplied to the nonlinear localization solver. Thus, $\widehat{\mathbf{b}}_{\mathcal{S}} = \mathbf{0}$ and the experiment evaluates sensitivity to a residual systematic error rather than an oracle bias-aware estimator. Therefore, performance differences within each trial are caused by the pre-selection strategy rather than by independently generated random conditions.

4.1.3. Controlled Predicted-Position Uncertainty Experiment

A controlled predicted-position uncertainty experiment is introduced to quantify how errors in the geometry available to the online pre-selection algorithm affect the information objective, selected node subset, observability assessment, and downstream nonlinear localization. For each of the $R_{\mathrm{MC}} = 100$ paired Monte Carlo trials, the true node positions $\mathbf{p}_i$ are generated once. A trial-specific standardized position-error realization $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3)$ is also generated once and is shared across all tested uncertainty levels. The predicted position used by the online selector is constructed as

$$\widehat{\mathbf{p}}_i(\sigma_p) = \mathbf{p}_i + \sigma_p \mathbf{z}_i, \tag{50}$$

where σ_p denotes the per-coordinate standard deviation of the prediction error. The tested levels are

$$\sigma_p \in \{0,\ 0.5,\ 1,\ 2,\ 5,\ 10,\ 20,\ 40\}\ \mathrm{m}.$$

The default prediction-error setting used in the preceding experiments corresponds to $\sigma_p = 2$ m per coordinate.

To isolate predicted-geometry uncertainty from topology-discovery and link-quality-estimation errors, the true node configuration, candidate ranging graph, link-quality classes, ranging noise realization, resource parameters, localization initialization, and standardized position-error directions are held fixed across all σ_p levels within the same trial. The candidate graph is constructed once from the true node configuration using $R_{\mathrm{c}} = 550$ m and is not regenerated from the perturbed predicted positions. The experiment uses $N = 30$ nodes, a selected-node budget of $K = 10$, the heterogeneous link classes in Table 4, zero link-quality estimation error ($\eta = \rho = 0$), and no injected systematic NLOS bias. Thus, the controlled variable is the geometry supplied to the pre-selection stage.

Table 4. True UWB link-quality classes and ranging-error parameters.

Link Class	Proportion	Random-Error Standard Deviation	Systematic Bias
LOS	50%	0.05 m	0
wLOS	30%	0.10 m	0
NLOS/degraded	20%	0.30 m	0 or 0.30–1.00 m

At each uncertainty level, the practical Weighted G-GDOP method constructs the gauge-free basis, ranging Jacobian, observability test, and A-optimal objective using $\widehat{\mathbf{p}}_i(\sigma_p)$. A true-geometry reference is obtained by running the same Weighted G-GDOP greedy procedure with $\widehat{\mathbf{p}}_i = \mathbf{p}_i$, while keeping the candidate graph, link weights, resource constraints, and all other trial-level quantities unchanged. This reference is therefore a geometry-knowledge reference for the same greedy method rather than a globally optimal subset.

For the selected node subset obtained using predicted geometry, $\mathcal{S}_{\text{pred}}$, the absolute relative mismatch between the online objective and the objective evaluated for the same subset under true geometry is

$$\delta_{\Phi} = \frac{\left|\widehat{\Phi}_{\text{A}}^{\text{pred}}(\mathcal{S}_{\text{pred}}) - \widehat{\Phi}_{\text{A}}^{\text{true-geom}}(\mathcal{S}_{\text{pred}})\right|}{\widehat{\Phi}_{\text{A}}^{\text{true-geom}}(\mathcal{S}_{\text{pred}})}. \tag{51}$$

Both quantities in Equation (51) use the same estimated link weights; only the geometry used to construct the projected information matrix is changed. Subset agreement with the true-geometry reference selected node subset $\mathcal{S}_{\text{true}}$ is quantified using

$$J_{\text{sub}} = \frac{\left|\mathcal{S}_{\text{pred}} \cap \mathcal{S}_{\text{true}}\right|}{\left|\mathcal{S}_{\text{pred}} \cup \mathcal{S}_{\text{true}}\right|}. \tag{52}$$

The information-quality loss caused by geometry-dependent subset selection is additionally measured under true geometry as

$$\Delta_{\Phi}^{\text{sel}} = \frac{\widehat{\Phi}_{\text{A}}^{\text{true-geom}}(\mathcal{S}_{\text{pred}}) - \widehat{\Phi}_{\text{A}}^{\text{true-geom}}(\mathcal{S}_{\text{true}})}{\widehat{\Phi}_{\text{A}}^{\text{true-geom}}(\mathcal{S}_{\text{true}})}. \tag{53}$$

The same selected node subset obtained using predicted geometry is also checked for observability under both predicted and true geometry using the common projected-FIM rank threshold τ_J. Nonlinear localization is then performed with the common downstream solver, and the paired RMSE difference is defined as

$$\Delta\text{RMSE} = \text{RMSE}_{\text{pred}} - \text{RMSE}_{\text{true-geom}}. \tag{54}$$

Observability-status changes, additional selection or localization failures relative to the true-geometry-reference-success trials, and the contributing sample counts are retained explicitly. Continuous metrics are summarized by the median and pointwise 95% percentile bootstrap confidence interval, whereas binary rates use the common 95% Wilson interval defined in the paired Monte Carlo protocol.

4.1.4. Comparison Methods and Ablation Versions

To comprehensively evaluate the performance of the proposed method, the comparison set includes Random Selection, Nearest Neighbor, Max–Min Distance, D-Optimal, Feasible Traditional GDOP, Unweighted G-GDOP, Weighted G-GDOP, and the All-Node Reference. Meanwhile, to verify the effects of topology-coupled modeling and heteroscedastic link-noise weighting in the proposed model, two ablation versions, namely Unweighted G-GDOP and Weighted G-GDOP, are introduced for comparison. Table 5 summarizes the comparison methods and G-GDOP variants used in the subsequent experiments.

Table 5. Comparison methods and G-GDOP variants.

Method	Description
Random Selection	Randomly constructs one size-K selected node subset in each paired Monte Carlo trial using a pre-generated random seed and the same common resource-feasibility screening. An unobservable or localization-failed result is retained as a failure and is not redrawn.
Nearest Neighbor	A distance-priority local greedy node selection strategy.
Max–Min Distance	A geometric dispersion method that maximizes the minimum distance among the selected nodes.
D-Optimal	A node selection method that maximizes the determinant of the information matrix.
Feasible Traditional GDOP	A resource-feasible node selection method based on minimization of the traditional GDOP metric.
Unweighted G-GDOP	A G-GDOP-based method that considers topology-coupled multi-node information but does not use heteroscedastic link-noise weighting.
Weighted G-GDOP	The proposed G-GDOP-based method that jointly considers topology-coupled information and heteroscedastic link-noise weighting.
All-Node Reference	Activates all candidate nodes and all available induced ranging links. It is used as a resource-unconstrained reference rather than as an accuracy upper bound.

For a fair comparison, all-node pre-selection methods except the All-Node Reference are evaluated under the same node budget K. In the heteroscedastic link-noise ablation experiment, the comparison between Feasible Traditional GDOP and Unweighted G-GDOP is used to analyze the effect of topology-induced multi-node coupling. The comparison between Unweighted G-GDOP and Weighted G-GDOP is used to analyze the contribution of heteroscedastic link-noise weighting.

In the resource-constrained experiments, all competing sequential selection methods use the same candidate graph, selected-node budget, induced-link budget, ranging-slot budget, residual-energy threshold, and general resource-cost budget. A candidate that would violate any common budget is removed from the feasible candidate set before the method-specific selection score is evaluated.

For exhaustive search, only subsets satisfying all common resource and local observability constraints are enumerated. The All-Node Reference is reported as a resource-unconstrained reference and is not included as a feasible competitor when its resource consumption exceeds the prescribed budget.

4.1.5. Nonlinear Localization Backend Used for RMSE Evaluation

The nonlinear localization backend is used only to evaluate the ranging topology returned by a pre-selection method; it does not provide ground-truth geometry, an objective value, or an initialization to the selector. Within each paired trial, every competing method uses the same localization backend and the same trial-level initialization protocol.

A classical multidimensional-scaling (MDS) configuration is first constructed from the finite symmetric pairwise-range matrix supplied to the localization subproblem. Let $\mathbf{D}$ denote this range matrix and $\mathbf{H}_c = \mathbf{I}_N - \mathbf{1}_N\mathbf{1}_N^T/N$. The double-centered Gram matrix is

$$\mathbf{B} = -\frac{1}{2}\mathbf{H}_c(\mathbf{D} \odot \mathbf{D})\mathbf{H}_c.$$

After numerical symmetrization, the three largest nonnegative eigenvalues and their eigenvectors form the three-dimensional MDS initialization; a rank-deficient result is padded with zero coordinates. If this initialization is invalid, a centered random initialization with scale determined by the median positive range is used. Ground-truth coordinates are not used to initialize MDS or the iterative estimator. The diagnostic backend does not perform method-specific oracle completion of missing measurements: a localization subproblem that cannot provide the required finite input is retained as a localization failure.

The subsequent weighted Gauss–Newton stage minimizes

$$\min_{\mathbf{x}} \sum_{e \in \mathcal{E}(\mathcal{S})} \frac{\left[z_e - \widehat{b}_e - h_e(\mathbf{x})\right]^2}{\widehat{\sigma}_e^2},$$

where $\widehat{b}_e = 0$ in the unmodeled-bias stress test. At iteration q, the damped normal equations are

$$\left(\mathbf{H}_q^T\widehat{\mathbf{W}}\mathbf{H}_q + \lambda\mathbf{I}\right)\delta_q = \mathbf{H}_q^T\widehat{\mathbf{W}}\mathbf{r}_q, \qquad \lambda = 10^{-4}.$$

The maximum iteration number is 50. Iteration terminates when $\|\delta_q\|_2 < 10^{-5}$ or when the change in residual norm is below 10^{-6}. A normal system with reciprocal condition number below 10^{-12}, an empty usable Jacobian, or a nonfinite update is treated as a solver failure rather than being repaired using true coordinates.

Because range-only localization is identifiable only up to a rigid transformation, the final estimate is aligned to the ground-truth configuration solely for RMSE evaluation by an SVD-based orthogonal Procrustes transform with translation and rotation but no scaling; reflections are rejected. A trial is classified as localization-success only when the backend returns a finite $N \times 3$ estimate and the rigid-aligned RMSE is finite. Failed trials are retained in status accounting and are not regenerated. This common backend is a diagnostic evaluation tool, not a claim that MDS–Gauss–Newton is the only or preferred downstream estimator; factor-graph and incremental-smoothing validation remains future work.

4.1.6. Evaluation Metrics and Statistical Settings

To evaluate the performance of different cooperative sensor node pre-selection methods from multiple perspectives, this study adopts localization accuracy, the A-optimal metric, the G-GDOP metric, running time, and the proportion of NLOS/degraded links as evaluation metrics.

Localization accuracy.

Let $\widehat{\mathbf{p}}_i^{\text{ali}}$ denote the estimated position of mobile node i after rigid alignment with the ground-truth configuration, and let $\mathbf{p}_i$ denote its true position. The full-state localization RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left\|\widehat{\mathbf{p}}_i^{\text{ali}} - \mathbf{p}_i\right\|_2^2}. \tag{55}$$

The RMSE is calculated over all N mobile sensor nodes, including both the selected nodes and the unselected ranging responders. Therefore, changing the selected-node budget K changes the measurement topology but does not change the set of states included in the RMSE calculation.

Online A-optimal selection metric.

For a feasible topology induced by the selected node subset, the practical online A-optimal selection metric is

$$\widehat{\Phi}_{\mathrm{A}}(\mathcal{S}) = \frac{1}{d_o} \operatorname{tr}\left[\widehat{\bar{\mathbf{J}}}_o(\mathcal{S})^{-1}\right]. \tag{56}$$

This metric is constructed from predicted coordinates and estimated link-error variances and is used by the practical pre-selection algorithm.

Oracle true-covariance A-optimal metric.

For diagnostic evaluation, the selected node subset is also assessed under the true random-error covariance:

$$\Phi_{\mathrm{A,true}}(\mathcal{S}) = \frac{1}{d_o} \operatorname{tr}\left[\bar{\mathbf{J}}_o^{\star}(\mathcal{S})^{-1}\right]. \tag{57}$$

This metric is unavailable to the practical online algorithm and is reported only to quantify the effect of covariance-estimation errors.

First-Order Local RMSE Surrogate Under Model Mismatch

The first-order local RMSE surrogate is calculated using Equation (15). It incorporates the true physical random-error covariance, the covariance used by the practical estimator, and the residual bias vector. This metric is used only as a secondary diagnostic and is not substituted for the actual nonlinear localization RMSE.

Generalized GDOP metric.

For a feasible topology induced by the selected node subset, the generalized GDOP metric is defined as

$$\text{G-GDOP}(\mathcal{S}) = \sqrt{\operatorname{tr}[\mathbf{J}_o(\mathcal{S})^{-1}]}. \tag{58}$$

This metric describes the local error amplification level in the gauge-free coordinates under the joint effects of link geometry, network topology, and estimated link-quality weighting.

Induced active-link number.

The number of induced active ranging links is evaluated as

$$M(\mathcal{S}) = |\mathcal{E}(\mathcal{S})|.$$

This metric measures the number of UWB ranging operations activated by the selected node subset.

Ranging-slot occupation.

The total ranging-slot occupation is evaluated as

$$T_{\mathrm{slot}}(\mathcal{S}) = \sum_{e \in \mathcal{E}(\mathcal{S})} s_e.$$

In the default unit-slot experiment, $s_e = 1$, so that $T_{\mathrm{slot}}(\mathcal{S}) = M(\mathcal{S})$. A nonuniform-slot sensitivity experiment may assign different s_e values to different ranging protocols or link-quality classes.

General resource cost.

The normalized general resource cost is evaluated using Equation (26). Because this quantity is an engineering cost proxy rather than a hardware-calibrated energy measurement, the raw selected-node number, induced-link number, and ranging-slot occupation are reported together with it.

A-Optimal Quality Retention and Resource Saving

The dimensionless A-optimal quality-retention ratio $Q_{\mathrm{A}}(\mathcal{S})$ and the resource-usage ratios ρ_{link}, ρ_{slot}, and ρ_{res} are used to characterize the information–resource trade-off relative to the All-Node Reference.

Selection-Stage Running Time

The reported running time includes only the cooperative sensor node pre-selection procedure. UWB measurement generation, nonlinear localization, rigid alignment, result serialization, and figure generation are excluded.

For the proposed method, the following runtime components are recorded separately:

1. candidate-graph and per-link-vector preprocessing time;
2. resource-feasibility screening time;
3. observable-seed construction time;
4. post-seed A-optimal candidate-evaluation time; and
5. total pre-selection time.

Each timing instance is executed once for numerical-library warm-up and is then repeated ten times using exactly the same input data. The median of the ten repeated measurements is retained as the trial-level running time. Across paired Monte Carlo trials, the reported point estimate is the median of the trial-level running times, and the uncertainty band is the corresponding pointwise 95% percentile bootstrap confidence interval.

Graph and Update-Structure Indicators

To interpret the observed runtime, the following structural quantities are recorded for every trial:

$$M_c = |\mathcal{E}_c|,$$

$$\rho_c = \frac{2M_c}{N(N-1)},$$

$$N_{\mathrm{eval}} = \sum_{t=0}^{K_{\mathrm{out}}-1} |\mathcal{F}_t|,$$

$$\bar{r} = \frac{\sum_t \sum_{i\in\mathcal{F}_t} r_{i,t}}{N_{\mathrm{eval}}},$$

and

$$r_{\max} = \max_{t,i} r_{i,t}.$$

The observable-seed size k_{seed} and the returned selected-node count K_{out} are also recorded. These quantities allow the empirical running time to be interpreted together with the theoretical complexity terms.

Peak Memory Consumption

Peak memory consumption is recorded during the pre-selection stage. The measurement excludes the memory used by nonlinear localization, plotting, and result files. The reported value is used to evaluate the practical influence of explicitly storing $\mathbf{Q}_o$, the per-link projected vectors, and the current inverse information matrix.

Proportion of NLOS/Degraded Links

The proportion of NLOS/degraded links is defined as

$$\rho_{\mathrm{NLOS}} = \frac{N_{\mathrm{NLOS}}^{\mathrm{sel}}}{N_{\mathrm{link}}^{\mathrm{sel}}}, \tag{59}$$

where $N_{\mathrm{NLOS}}^{\mathrm{sel}}$ denotes the number of NLOS/degraded links in the subnetwork induced by the selected node subset, and $N_{\mathrm{link}}^{\mathrm{sel}}$ denotes the total number of effective ranging links in the subnetwork induced by the selected node subset. This metric is used to evaluate whether a method can reduce the inclusion of low-quality ranging links in the cooperative localization subnetwork.

4.1.7. Paired Monte Carlo Protocol, Failure Handling, and Statistical Reporting

Unless otherwise stated, each operating condition is evaluated using $R_{\mathrm{MC}} = 100$ paired Monte Carlo trials. For the r-th trial, the true node configuration, predicted node configuration, candidate ranging graph, true link-quality classes, estimated link-quality realization, standardized ranging noise realization, NLOS-bias realization, resource parameters, and nonlinear localization initialization are generated only once and are shared by all competing methods.

The true node coordinates are used only to generate the physical UWB measurements and to calculate the rigid-aligned full-state localization RMSE. All practical pre-selection methods construct their ranging Jacobians and online information matrices using the same predicted node coordinates. Similarly, the true link-error covariance is used to generate the physical ranging errors, whereas a practical weighted method uses the same estimated link-error covariance available in that trial. The true covariance is supplied only to methods explicitly labeled as oracle benchmarks.

Common random numbers are also used across parameter levels whenever the corresponding physical interpretation permits. In the Gaussian noise sweep, a standardized realization $\xi_e^{(r)}$ is generated once and scaled by the specified noise standard deviation. In the NLOS-bias sweep, the node topology, link classes, standardized random errors, estimated link-quality realization, and localization initialization are held fixed while only the injected bias magnitude is changed. In the node-, link-, and ranging-slot-budget sweeps, the candidate topology and measurement realization are kept unchanged within each trial.

No method-specific trial regeneration is permitted. If a method cannot construct a selected node subset satisfying the common resource constraints, the result is recorded as resource-infeasible. If the returned projected information matrix does not satisfy

$$r_o(\mathcal{S}) = d_o,$$

the result is recorded as unobservable. An unobservable matrix is not made artificially invertible through numerical regularization. If the nonlinear localization solver does not converge to a finite solution, or if the rigid-aligned RMSE is non-finite, the result is recorded as a localization failure. Failed trials are not replaced by newly generated topologies, noise realizations, or initial values.

The resource-feasibility rate, observability rate, selection-success rate, and nonlinear localization success rate are calculated using all R_{MC} common trials as the denominator. These binary rates are reported together with 95% Wilson confidence intervals.

For a continuous metric y, the reported point estimate for method m is the median over the trials in which that metric is defined:

$$\widetilde{y}_m = \text{median}_{r \in \mathcal{V}_m} y_m^{(r)}.$$

Here, $\mathcal{V}_m$ is the set of trials in which the metric is defined. Information matrix metrics are summarized over observable trials, whereas nonlinear localization RMSE is summarized over localization-success trials. The number of contributing trials is reported with every continuous summary.

Uncertainty intervals for continuous metrics are calculated using $B_{\text{boot}} = 10{,}000$ percentile-bootstrap repetitions. Monte Carlo trial indices are resampled with replacement. Within one bootstrap repetition, the same resampled trial indices are applied jointly to all competing methods so that the within-trial pairing is preserved. The 2.5th and 97.5th percentiles of the bootstrap distribution form the pointwise 95% confidence interval.

For a pairwise comparison between method m and reference method b, only trials in which both methods provide a valid value are used. The paired median difference is defined as

$$\widetilde{\Delta}_{m,b} = \text{median}_{r \in \mathcal{V}_m \cap \mathcal{V}_b} \left(y_m^{(r)} - y_b^{(r)} \right).$$

Its 95% confidence interval is obtained by jointly resampling the paired trial indices. The number of common valid trials n_{pair} is reported for every pairwise comparison.

Unless otherwise stated, curves and point estimates show medians, and shaded regions or error bars show pointwise 95% bootstrap confidence intervals. Interquartile ranges may be shown only as supplementary distributional summaries and are not used as the primary uncertainty intervals.

4.1.8. Computational Environment

Since the experiments in this study include running-time comparisons, all formal experiments are conducted under a unified computational environment to ensure the interpretability of the computational-efficiency evaluation. The Exp01–Exp09 selection, evaluation, and reporting pipelines use the validated `ggdop_core_step9` backend under Python 3.13.12 on a 64-bit Windows 11 computer equipped with an AMD Ryzen 7 4800U processor (Advanced Micro Devices, Inc., Santa Clara, CA, USA) and 16 GB DDR4-3200 memory. The reported running time is the median pre-selection time obtained from ten repeated executions for each network topology after one numerical-library warm-up run. All methods are executed serially under the same process and numerical-thread configuration. Measurement generation, nonlinear localization, rigid alignment, plotting, and file serialization are excluded from the reported pre-selection time.

4.2. Simulation Results

4.2.1. Applicability Boundary and Local Diagnostic Consistency

Table 6 summarizes the applicability boundary and the local diagnostic behavior of G-GDOP. In the global diagnostic set, the median G-GDOP was 24.95 [19.97, 28.34], and the topology-level median RMSE was 0.419 [0.326, 0.482] m. The local analysis yielded median values of 0.1316 [0.0896, 0.1737] for the CRLB proxy and 0.2723 [0.1976, 0.3419] for the empirical local MSE. The clearest ordering appeared in the normalized local RMSE, which increased from 0.698 [0.673, 0.752] in the Low group to 1.014 [0.999, 1.068] in the Medium group and 1.980 [1.806, 2.180] in the High group. This ordered separation supports the use of G-GDOP as a ranking indicator within the tested domain, without implying a coefficient-based global association. The same-topology ordering is visualized in Figure 5.

Table 6. Applicability-boundary and local-diagnostic summaries for G-GDOP.

Global Applicability Diagnostics			
Metric	**Valid Samples**	**Median [95% CI]**	**Unit**
G-GDOP	59	24.95 [19.97, 28.34]	–
Topology-level median RMSE	59	0.419 [0.326, 0.482]	m
Local Diagnostic Consistency			
Metric	**Valid Samples**	**Median [95% CI]**	**Unit**
Local CRLB proxy	59	0.1316 [0.0896, 0.1737]	m^2
Empirical local MSE	59	0.2723 [0.1976, 0.3419]	m^2
Within-Topology Ranking			
G-GDOP Group	**Valid Samples**	**Normalized Local RMSE [95% CI]**	
Low	20	0.698 [0.673, 0.752]	
Medium	20	1.014 [0.999, 1.068]	
High	20	1.980 [1.806, 2.180]	

Notes: Lower values indicate better information quality or localization accuracy. Shared non-success cases remain in the formal status accounting. The local CRLB proxy is $\sigma_0^2 \Phi_A$ and therefore has unit m^2. Empirical local MSE is the median squared rigid-aligned RMSE over successful diagnostic repetitions and also has unit m^2. Normalized local RMSE is dimensionless. The table reports ordered groups only and does not report Pearson, Spearman, or coefficient-of-determination statistics.

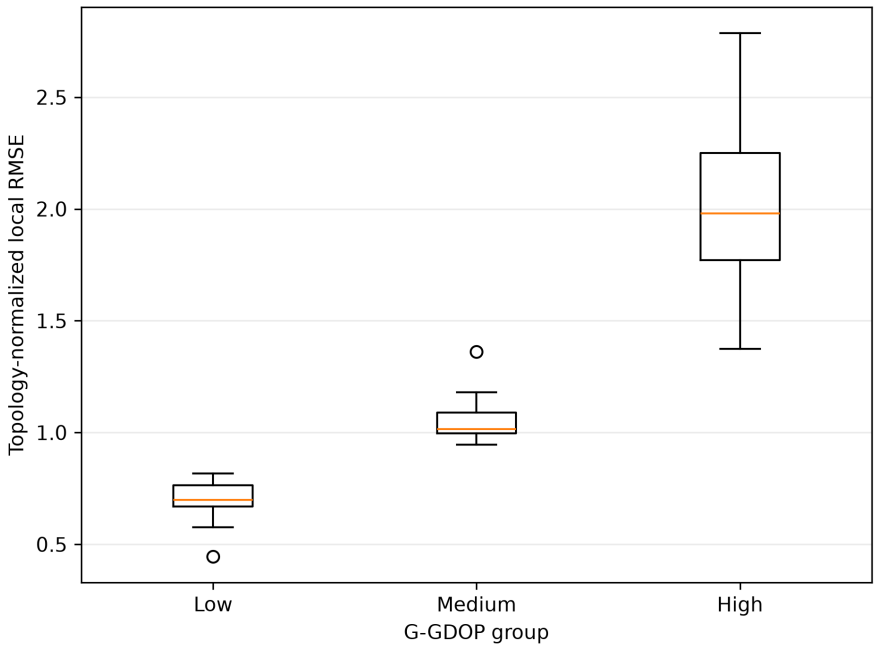


Figure 5. Same-topology ranking validation of G-GDOP. Topology-normalized local RMSE is shown for the Low, Medium, and High G-GDOP groups.

4.2.2. Information–Accuracy–Resource Trade-Offs

The node- and ranging-slot-budget results in Table 7 reveal the expected information–accuracy–resource trade-off. For Weighted G-GDOP, increasing K from 6 to 20 raised the median number of active links from 109 to 214 and the median runtime from 0.0578 to 0.0963 s, while reducing the median normalized objective from 7.994 to 0.777 and the median RMSE from 0.550 to 0.148 m. Relaxing the ranging-slot budget from $\beta_{\text{slot}} = 0.45$ to 1.0 produced a corresponding change: the active-link count increased from 112 to 214, whereas the objective and RMSE decreased from 6.694 to 0.777 and from 0.435 to 0.148 m, respectively. At $K = 20$, Weighted G-GDOP activated 214 links, compared with 229 for the All-Node Reference, and achieved median objective and RMSE values of 0.777 and 0.148 m; the corresponding all-node values were 0.745 and 0.143 m. Shared non-success

cases remain included in the formal status accounting and are not attributed to individual methods. Figures 6 and 7 show the corresponding runtime growth and RMSE reduction as K increases.

Table 7. Information–accuracy–resource trade-offs under node and ranging-slot budgets.

K	Method	Links	Norm. objective	RMSE (m)	Time (s)
Node-Budget Trade-Off					
6	Weighted G-GDOP	109 [105.0, 115.5]	7.994 [5.636, 25.360]	0.550 [0.418, 0.814]	0.0578 [0.0571, 0.0584]
8	Weighted G-GDOP	128 [124.0, 131.0]	4.239 [3.412, 5.399]	0.321 [0.298, 0.368]	0.0687 [0.0678, 0.0693]
10	Weighted G-GDOP	145 [142.0, 149.0]	1.914 [1.729, 2.144]	0.233 [0.203, 0.249]	0.0750 [0.0742, 0.0764]
12	Weighted G-GDOP	164 [157.0, 166.0]	1.293 [1.154, 1.416]	0.187 [0.167, 0.201]	0.0807 [0.0798, 0.0819]
15	Weighted G-GDOP	186 [178.0, 188.0]	0.961 [0.877, 1.005]	0.160 [0.149, 0.170]	0.0877 [0.0866, 0.0887]
20	Weighted G-GDOP	214 [206.0, 218.0]	0.777 [0.723, 0.819]	0.148 [0.131, 0.159]	0.0963 [0.0953, 0.0975]
20	All-Node Reference	229 [224.0, 237.0]	0.745 [0.685, 0.783]	0.143 [0.131, 0.156]	0.0034 [0.0034, 0.0035]
Ranging-Slot-Budget Trade-Off					
0.35	Weighted G-GDOP	96 [96.0, 96.0]	99.347 [99.347, 99.347]	10.454 [10.454, 10.454]	0.0491 [0.0491, 0.0491]
0.45	Weighted G-GDOP	112 [107.0, 115.0]	6.694 [4.727, 11.572]	0.435 [0.363, 0.565]	0.0594 [0.0583, 0.0605]
0.55	Weighted G-GDOP	127 [125.0, 131.0]	3.529 [3.036, 5.098]	0.321 [0.283, 0.360]	0.0688 [0.0675, 0.0699]
0.70	Weighted G-GDOP	161 [156.0, 165.0]	1.312 [1.184, 1.451]	0.187 [0.171, 0.202]	0.0802 [0.0789, 0.0818]
1.00	Weighted G-GDOP	214 [206.0, 218.0]	0.777 [0.723, 0.820]	0.148 [0.131, 0.159]	0.0969 [0.0955, 0.0978]
1.00	All-Node Reference	229 [222.0, 237.0]	0.745 [0.685, 0.783]	0.143 [0.131, 0.156]	0.0035 [0.0034, 0.0035]

Notes: The All-Node Reference is included only at the matching full-budget endpoint. Shared structural non-success cases are not attributed to individual methods. The node-budget and ranging-slot-budget dimensions are kept in separate panels and are not collapsed into one sweep.

The monotonic statement above applies to the reported medians of Weighted G-GDOP, not to every comparison method or every trial. In Figure 7, the Nearest Neighbor curve can increase over part of the budget range. Adding selected nodes according to distance alone changes the induced topology but does not guarantee improved geometric diversity or link reliability, and the final nonlinear RMSE also depends on initialization and convergence. Accordingly, the manuscript does not claim that RMSE must decrease monotonically with K for all selection rules, nor that Weighted G-GDOP is the lowest-RMSE method under every tested budget.

4.2.3. Small-Scale Comparison with Exhaustive Enumeration

Table 8 provides an absolute small-scale comparison between Weighted G-GDOP and the Exhaustive Optimum. For $K = 3$, 4, 5, and 6, the respective median normalized objective pairs (Weighted G-GDOP versus Exhaustive Optimum) were 8.179 versus 3.901, 2.210 versus 1.702, 1.375 versus 1.250, and 1.060 versus 1.021. The absolute objective values therefore became progressively closer over the tested range of K. Meanwhile, the runtime of Weighted G-GDOP increased only from 0.0050 to 0.0074 s, whereas exhaustive enumeration increased from 0.0308 to 0.1279 s. These results are used solely as an absolute benchmark and do not establish a general approximation guarantee. Figure 8 visualizes the convergence of the tested objective values over the small-scale range of K.

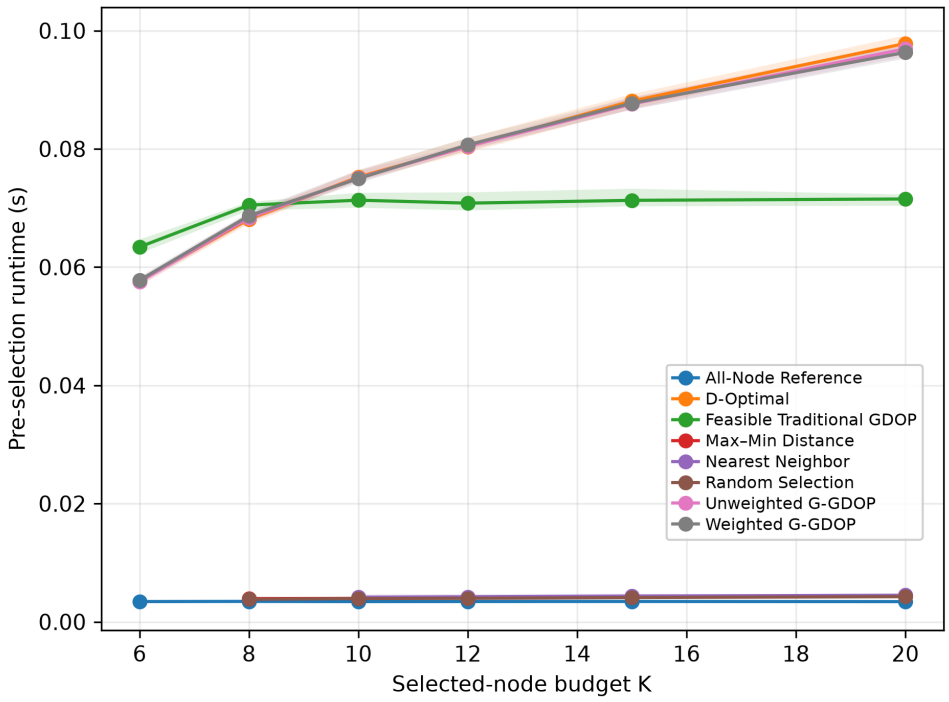


Figure 6. Pre-selection runtime versus the selected-node budget K in the node-budget experiment for the tested selection methods and the All-Node Reference.

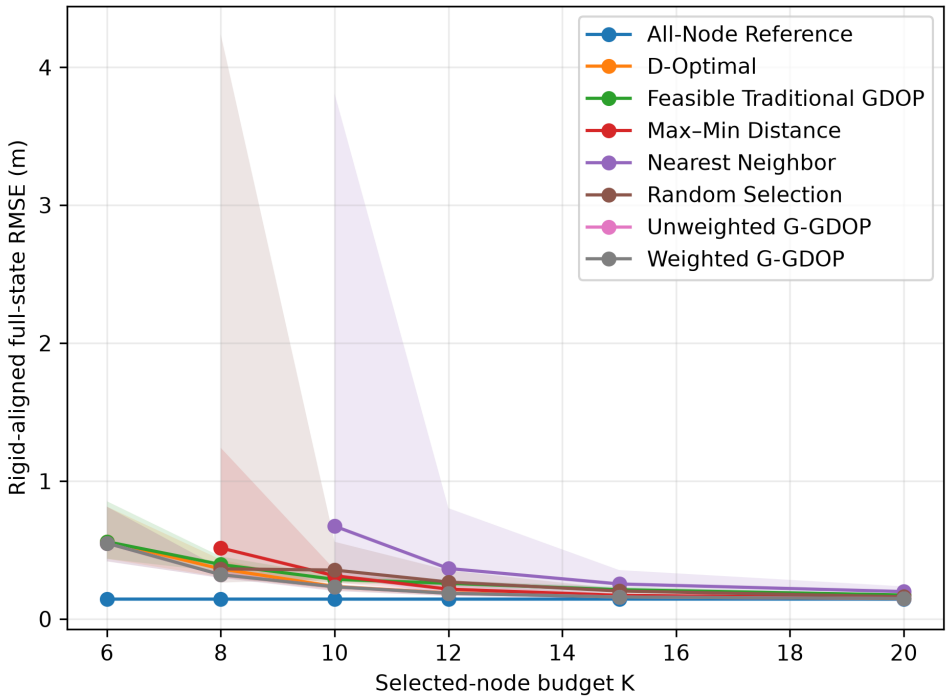


Figure 7. Rigid-aligned full-state RMSE versus the selected-node budget K in the node-budget experiment.

Table 8. Small-scale comparison of Weighted G-GDOP with exhaustive enumeration.

K	Method	Normalized Objective	Time (s)
3	Weighted G-GDOP	8.179 [6.520, 9.809]	0.0050 [0.0050, 0.0050]
3	Exhaustive optimum	3.901 [3.547, 4.195]	0.0308 [0.0307, 0.0309]

Table 8. *Cont.*

K	Method	Normalized Objective	Time (s)
4	Weighted G-GDOP	2.210 [2.073, 2.442]	0.0059 [0.0059, 0.0059]
4	Exhaustive optimum	1.702 [1.590, 1.835]	0.0684 [0.0684, 0.0685]
5	Weighted G-GDOP	1.375 [1.291, 1.460]	0.0067 [0.0067, 0.0068]
5	Exhaustive optimum	1.250 [1.154, 1.344]	0.1094 [0.1093, 0.1095]
6	Weighted G-GDOP	1.060 [0.996, 1.149]	0.0074 [0.0074, 0.0075]
6	Exhaustive optimum	1.021 [0.933, 1.085]	0.1279 [0.1277, 0.1280]

Notes: Only absolute objective and runtime statistics are reported. The table does not state a general approximation ratio, optimality gap guarantee, or performance bound.

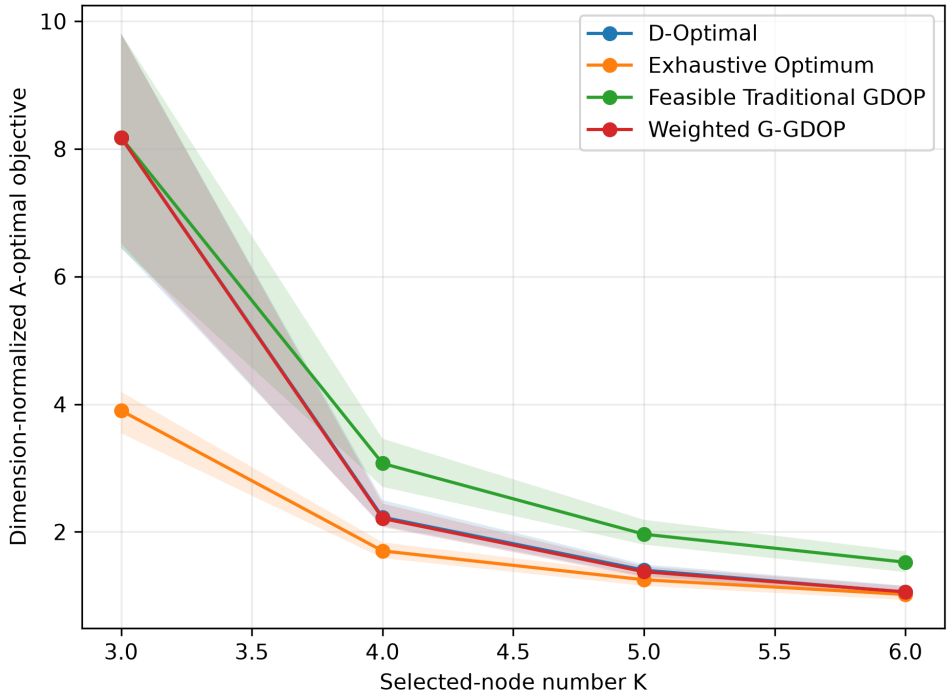


Figure 8. Dimension-normalized A-optimal objective versus the selected node number K in the small-scale exhaustive-enumeration benchmark.

The comparatively large difference at $K = 3$ was examined separately by tracing the same 100 paired Monte Carlo trials used in the formal comparison. In all 100 trials, the observable-seed size was $k_{\text{seed}} = 3$ and the final selected-node count was also $k_{\text{out}} = 3$; hence, the post-seed A-optimal Stage II performed zero node-addition iterations. The Stage I seed was identical to the exhaustive-optimal selected node subset in only 7 of 100 trials, and the median paired true-objective gap relative to exhaustive enumeration was 79.0%.

The first divergence from the exhaustive-optimal selected node subset was attributed to the retained minimum-eigenvalue criterion in 83 trials and to the deterministic node-ID tie-break in 10 trials; the remaining seven trials did not diverge. Across the 300 Stage I seed decisions, numerical rank increase reduced the feasible candidate set in 12 decisions, retained minimum eigenvalue further reduced the surviving candidate set in all 300 decisions, and node ID was finally decisive in 31. The incremental induced-link number, ranging-slot occupation, and general resource cost did not discriminate among feasible candidates in any of these 300 decisions. In this $N = 12$, $K = 3$ setting, every final selected node subset activated 30 links, occupied 30 ranging slots, and had a general resource cost of 33. As a diagnostic check, when the observable seed was selected directly by the final

A-optimal objective instead of the Stage I lexicographic ordering, the exhaustive optimum was recovered in all 100 trials.

4.2.4. Robustness, Ablation, and Sensitivity Analyses

Tables 9–12 consolidate the Gaussian-noise, heterogeneous-noise, NLOS-bias, and online weight-estimation analyses. **Gaussian ranging noise (Panel A).** As σ increased from 0.01 to 1.0, the median RMSE of Weighted G-GDOP increased from 0.0223 to 2.228 m. Its median objective remained 0.01754, and its median runtime remained approximately 0.0756 s, because the candidate geometry and selection rule were unchanged across the injected Gaussian-noise levels. At $\sigma = 1.0$, Weighted G-GDOP achieved a median RMSE of 2.228 m, compared with 2.256 m for D-Optimal, 2.976 m for Feasible Traditional GDOP, and 5.896 m for Nearest Neighbor. The resource-unconstrained All-Node Reference recorded a median RMSE of 1.262 m at this noise level; it is reported as a reference and is not treated as a universal accuracy bound. The complete RMSE trend is shown in Figure 9.

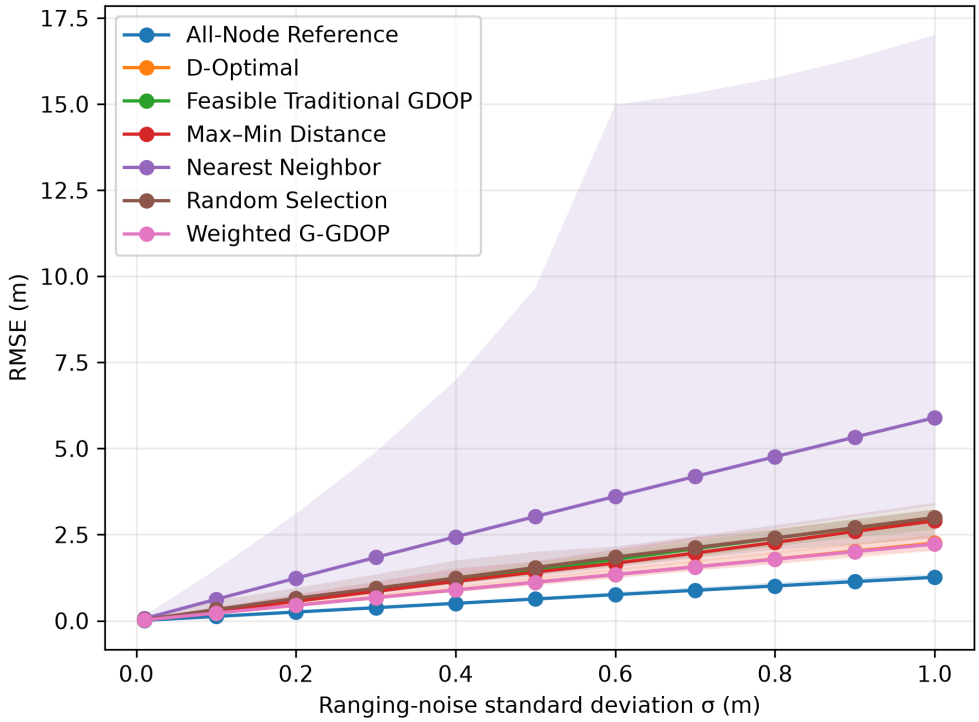


Figure 9. Rigid-aligned RMSE under increasing Gaussian ranging noise standard deviation σ for the tested methods.

Table 9. Gaussian ranging noise robustness.

σ (m)	Method	Normalized Objective	RMSE (m)	Time (s)
0.01	All-Node Reference	0.00672 [0.00606, 0.00756]	0.0127 [0.0124, 0.0135]	0.0035 [0.0034, 0.0035]
0.01	D-Optimal	0.02109 [0.01786, 0.02298]	0.0225 [0.0207, 0.0251]	0.0752 [0.0745, 0.0763]
0.01	Feasible Traditional GDOP	3.42817 [2.74044, 3.98743]	0.0295 [0.0265, 0.0332]	0.0713 [0.0706, 0.0726]
0.01	Max–Min Distance	0.03256 [0.02311, 0.05777]	0.0274 [0.0248, 0.0390]	0.0040 [0.0039, 0.0041]
0.01	Nearest Neighbor	0.31088 [0.03210, 2.01513]	0.0627 [0.0332, 0.1441]	0.0043 [0.0040, 0.0044]
0.01	Random Selection	0.04212 [0.02986, 0.09134]	0.0326 [0.0279, 0.0586]	0.0041 [0.0038, 0.0041]
0.01	Weighted G-GDOP	0.01754 [0.01554, 0.02241]	0.0223 [0.0205, 0.0242]	0.0756 [0.0752, 0.0767]
0.50	All-Node Reference	0.00672 [0.00606, 0.00756]	0.6321 [0.6199, 0.6752]	0.0035 [0.0034, 0.0035]
0.50	D-Optimal	0.02109 [0.01789, 0.02298]	1.1204 [1.0351, 1.2555]	0.0752 [0.0745, 0.0763]
0.50	Feasible Traditional GDOP	3.42817 [2.73387, 3.98743]	1.4839 [1.3511, 1.6752]	0.0713 [0.0706, 0.0726]
0.50	Max–Min Distance	0.03256 [0.02311, 0.05777]	1.4142 [1.2119, 1.7212]	0.0040 [0.0039, 0.0041]
0.50	Nearest Neighbor	0.31088 [0.03210, 2.01513]	3.0273 [1.6688, 9.6621]	0.0043 [0.0040, 0.0044]

Table 9. *Cont.*

σ (m)	Method	Normalized Objective	RMSE (m)	Time (s)
0.50	Random Selection	0.04212 [0.02986, 0.09134]	1.5451 [1.3018, 2.0021]	0.0041 [0.0038, 0.0041]
0.50	Weighted G-GDOP	0.01754 [0.01554, 0.02241]	1.1195 [1.0210, 1.2102]	0.0756 [0.0752, 0.0767]
1.00	All-Node Reference	0.00672 [0.00614, 0.00756]	1.2618 [1.2375, 1.3591]	0.0035 [0.0034, 0.0035]
1.00	D-Optimal	0.02109 [0.01786, 0.02410]	2.2556 [2.0280, 2.4969]	0.0752 [0.0745, 0.0764]
1.00	Feasible Traditional GDOP	3.42817 [2.73387, 4.05261]	2.9763 [2.6400, 3.2315]	0.0713 [0.0706, 0.0726]
1.00	Max–Min Distance	0.03252 [0.02311, 0.04825]	2.9006 [2.4124, 3.2275]	0.0040 [0.0039, 0.0041]
1.00	Nearest Neighbor	0.31088 [0.03210, 2.01513]	5.8956 [3.3559, 17.0095]	0.0043 [0.0040, 0.0044]
1.00	Random Selection	0.04212 [0.02986, 0.09134]	2.9996 [2.3935, 3.4209]	0.0041 [0.0038, 0.0041]
1.00	Weighted G-GDOP	0.01754 [0.01559, 0.02096]	2.2275 [2.0464, 2.4035]	0.0756 [0.0750, 0.0767]

Notes: The table retains representative lower, intermediate, and upper noise levels; the complete σ grid remains available in the formal results and figure.

Table 10. Heterogeneous-noise ablation.

Method	G-GDOP	Local RMSE Surrogate (m)	Normalized Objective	Empirical RMSE (m)	Time (s)
D-Optimal	12.634 [11.903, 15.229]	0.1379 [0.1299, 0.1662]	1.901 [1.682, 2.761]	0.2172 [0.2015, 0.2503]	0.1536 [0.1478, 0.1551]
Feasible Traditional GDOP	17.123 [16.017, 21.543]	0.1868 [0.1748, 0.2351]	3.264 [2.841, 3.927]	0.2848 [0.2550, 0.3269]	0.1481 [0.1435, 0.1513]
Oracle-Weighted G-GDOP	11.710 [10.406, 12.497]	0.1278 [0.1132, 0.1364]	1.632 [1.289, 1.859]	0.1999 [0.1844, 0.2230]	0.1528 [0.1473, 0.1551]
Unweighted G-GDOP	11.789 [10.515, 14.031]	0.1286 [0.1152, 0.1531]	1.696 [1.479, 2.017]	0.2047 [0.1885, 0.2300]	0.1526 [0.1493, 0.1544]

Notes: The first-order local RMSE surrogate is the square root of the corresponding local MSE surrogate and has unit m. It is a secondary local diagnostic and is not substituted for the empirical nonlinear localization RMSE.

Table 11. NLOS-bias sensitivity.

NLOS Bias (m)	Method	Local RMSE Surrogate (m)	Empirical RMSE (m)
0.0	D-Optimal	0.1672 [0.1500, 0.1816]	0.2430 [0.2197, 0.2670]
0.0	Weighted G-GDOP	0.1486 [0.1364, 0.1609]	0.2197 [0.2071, 0.2407]
0.0	Feasible Traditional GDOP	0.2399 [0.2144, 0.2703]	0.3373 [0.2999, 0.3756]
0.0	Oracle-Weighted G-GDOP	0.1361 [0.1230, 0.1486]	0.2078 [0.1878, 0.2278]
0.0	Unweighted G-GDOP	0.1521 [0.1440, 0.1651]	0.2542 [0.2302, 0.2697]
0.5	D-Optimal	0.2429 [0.2119, 0.2677]	0.3536 [0.3108, 0.4291]
0.5	Weighted G-GDOP	0.2231 [0.1999, 0.2400]	0.3622 [0.3160, 0.3993]
0.5	Feasible Traditional GDOP	0.3572 [0.3081, 0.4116]	0.5508 [0.4915, 0.6794]
0.5	Oracle-Weighted G-GDOP	0.1871 [0.1594, 0.2005]	0.2661 [0.2362, 0.3205]
0.5	Unweighted G-GDOP	0.2269 [0.2091, 0.2452]	0.3696 [0.3405, 0.4183]
1.0	D-Optimal	0.3767 [0.3280, 0.4338]	0.6081 [0.4981, 0.7154]
1.0	Weighted G-GDOP	0.3475 [0.3172, 0.3949]	0.5946 [0.5278, 0.6449]
1.0	Feasible Traditional GDOP	0.5703 [0.4820, 0.6306]	0.9528 [0.8430, 1.0891]
1.0	Oracle-Weighted G-GDOP	0.2672 [0.2367, 0.3073]	0.4033 [0.3579, 0.5093]
1.0	Unweighted G-GDOP	0.3761 [0.3296, 0.3943]	0.6148 [0.5366, 0.7010]

Notes: The local RMSE surrogate incorporates the true random-error covariance, estimator covariance, and residual bias. It has unit m and remains distinct from the rigid-aligned nonlinear empirical RMSE.

Table 12. Online link-weight error sensitivity.

η	ρ	Method/Reference	Pairs	Relative True-Objective Difference	True-Objective Difference	Subset-Jaccard Difference
0.0	0.0	Weighted G-GDOP	97	0.0000 [0.0000, 0.0000]	0.0000 [0.0000, 0.0000]	0.0000 [0.0000, 0.0000]
0.2	0.2	Weighted G-GDOP	97	0.0330 [0.0000, 0.0892]	0.0408 [0.0000, 0.0860]	−0.3333 [−0.3333, −0.1818]
0.3	0.3	Weighted G-GDOP	97	0.0621 [0.0019, 0.0967]	0.0735 [0.0067, 0.1041]	−0.3333 [−0.3333, −0.3333]
–	–	Unweighted G-GDOP (full-grid median)	97	0.1400	0.2351	−0.3333

Notes: All differences are calculated as the tested method minus Oracle-Weighted G-GDOP. Negative subset Jaccard differences indicate lower subset overlap. The unweighted row is an explicitly derived full-grid median, not an additional trial condition.

Heterogeneous-noise comparison (Panel B). Oracle-Weighted G-GDOP attained the lowest median G-GDOP (11.710), local RMSE surrogate (0.1278), normalized objective (1.632), and empirical RMSE (0.1999 m). Unweighted G-GDOP was the closest comparator, with median objective and RMSE values of 1.696 and 0.2047 m, whereas Feasible Traditional GDOP yielded 3.264 and 0.2848 m. Median runtimes varied only from 0.1481 to 0.1536 s across the panel, indicating that the observed differences primarily reflect selection quality. Figure 10 compares the corresponding local RMSE surrogates.

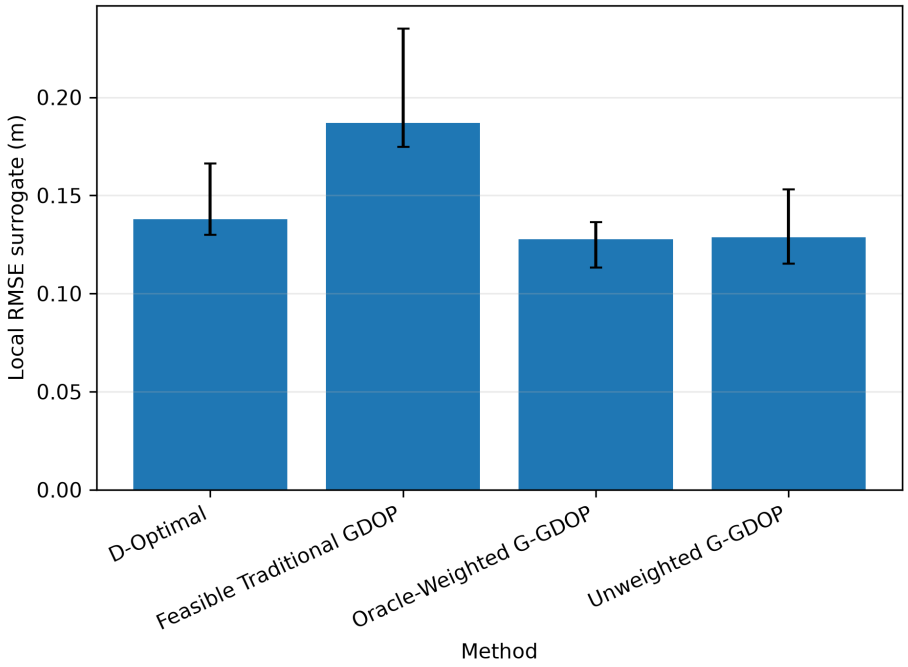


Figure 10. Local RMSE surrogate in the heterogeneous-noise ablation for D-Optimal, Feasible Traditional GDOP, Oracle-Weighted G-GDOP, and Unweighted G-GDOP.

NLOS-bias sensitivity (Panel C). Both the local surrogate and empirical RMSE deteriorated as the positive NLOS bias increased. The median RMSE of Oracle-Weighted G-GDOP rose from 0.2078 to 0.4033 m between bias levels of 0 and 1.0 m, while Weighted G-GDOP increased from 0.2197 to 0.5946 m. At the largest bias, Weighted G-GDOP retained a local surrogate of 0.3475, compared with 0.3767 for D-Optimal, 0.3761 for Unweighted G-GDOP, and 0.5703 for Feasible Traditional GDOP. Because the NLOS-link ratios and runtimes changed little with bias, the degradation is associated with bias exposure rather than a

change in the nominal selection budget. Figure 11 shows the first-order local-surrogate degradation with increasing bias.

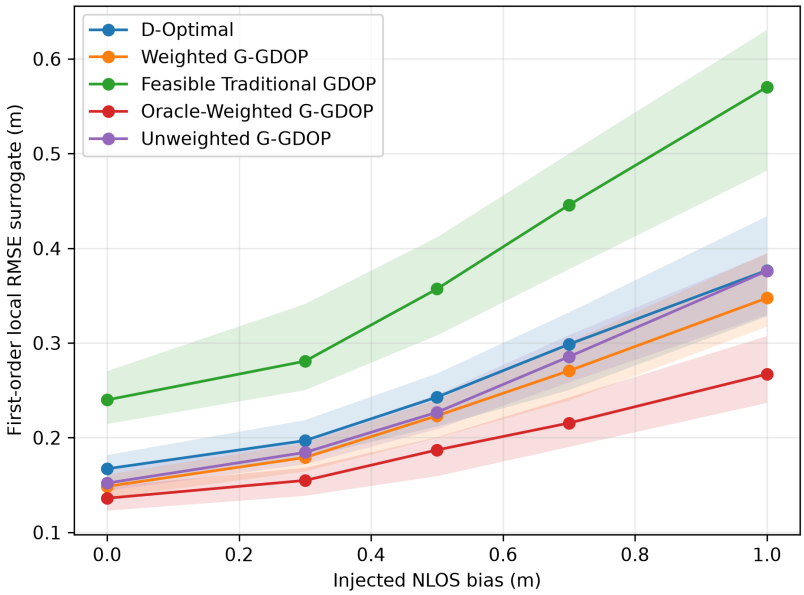


Figure 11. First-order local RMSE surrogate versus the injected NLOS bias for the tested methods.

Online weight-estimation errors (Panel D). Each η–ρ condition included 97 paired trials. At $\eta = \rho = 0$, the median relative true-objective difference, true-objective difference, and subset Jaccard difference between Weighted G-GDOP and the oracle were all zero. At the largest tested perturbation, $\eta = \rho = 0.3$, the corresponding median differences were 0.0621, 0.0735, and -0.3333. Across the complete η–ρ grid, Unweighted G-GDOP retained a median relative true-objective difference of 0.1400. Weighted G-GDOP therefore remained closer to the oracle objective throughout the tested perturbation grid, although subset agreement decreased as the weight errors increased. Figure 12 shows the RMSE response along the $\rho = 0$ slice.

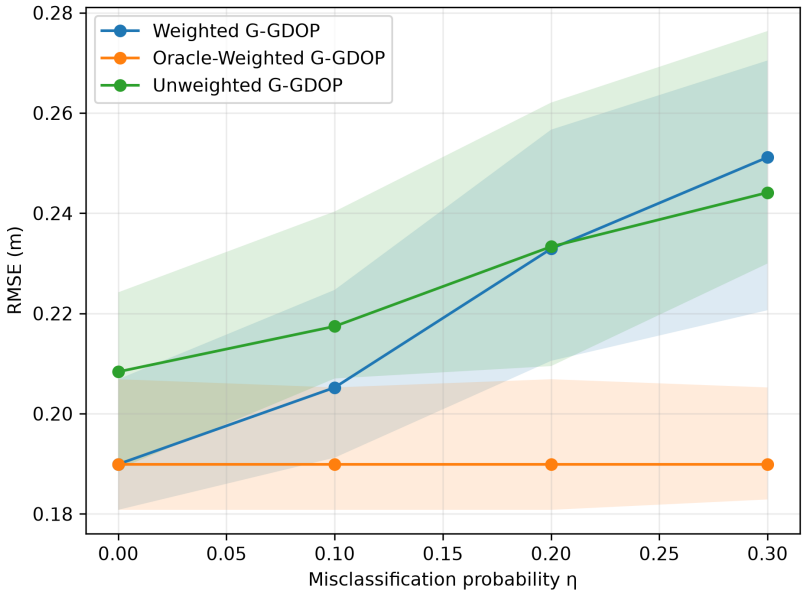


Figure 12. RMSE sensitivity to the link-quality misclassification probability η at $\rho = 0$ for Weighted G-GDOP, Oracle-Weighted G-GDOP, and Unweighted G-GDOP.

4.2.5. Numerical Equivalence and Computational Scalability

Across all 24 tested network-size–graph-regime combinations, Direct Inversion and Woodbury Updating were numerically equivalent under the stated `np.isclose` tolerance (`rtol` = 1×10^{-8} and `atol` = 1×10^{-8}). Across the 589 common-success paired groups, the maximum absolute objective difference was 1.376×10^{-8}, confirming that the Woodbury update preserves the objective values to numerical precision. For runtime, Woodbury Updating was faster in 20 of the 24 tested size–regime combinations; the largest median speedup was 1.947 at $N = 100$ in the fixed-radius regime. Complete condition-by-condition numerical-equivalence and runtime results are provided in Supplementary Tables S1 and S2, respectively.

Figure 13 visualizes the runtime speedup relative to Direct Inversion as network size increases in the dense-graph regime.

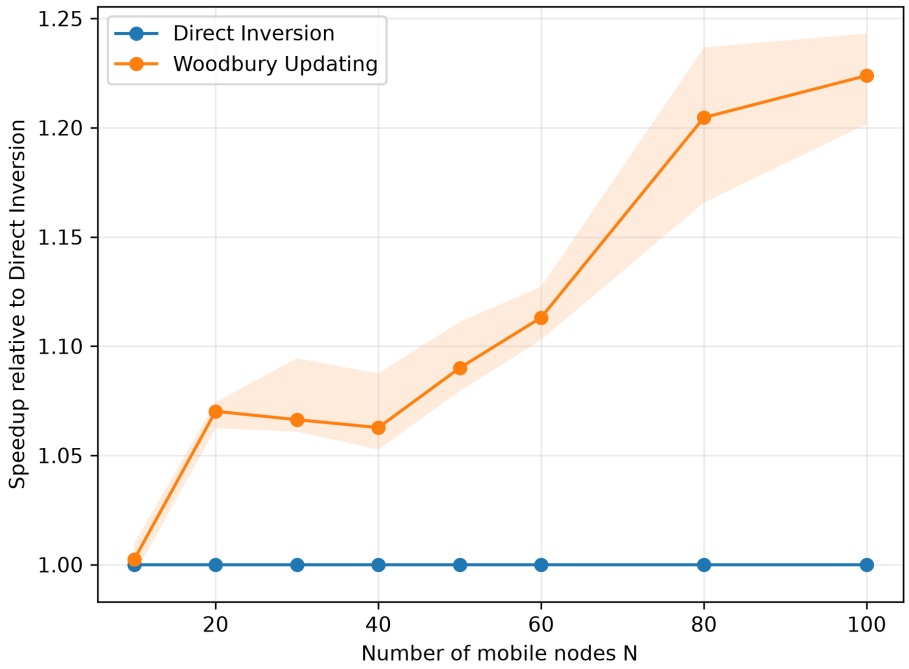


Figure 13. Runtime speedup relative to Direct Inversion versus network size N in the dense graph regime. A speedup above one favors Woodbury Updating.

4.2.6. Sensitivity to Predicted-Position Uncertainty

The controlled position uncertainty experiment shows that errors in the predicted geometry first affect the online information objective and subset choice, whereas the downstream nonlinear localization error changes more gradually. At the default setting of $\sigma_p = 2$ m per coordinate, the median realized 3-D position RMSE was 3.446 m. The median absolute predicted-versus-true geometry objective mismatch δ_Φ was 0.840% with a 95% bootstrap confidence interval of [0.626%, 1.123%], while the median subset Jaccard index was 1.000 [1.000, 1.000]. The median true-geometry selection gap $\Delta_\Phi^{\mathrm{sel}}$ and the paired nonlinear localization-RMSE difference were both zero at this default uncertainty level.

As σ_p increased, the median objective mismatch rose to 2.41%, 5.47%, 12.00%, and 22.93% at $\sigma_p = 5$, 10, 20, and 40 m, respectively, while the corresponding median subset Jaccard indices decreased to 0.818, 0.667, 0.538, and 0.538. The median true-geometry selection gap remained zero through $\sigma_p = 10$ m, then increased to 4.91% [0, 9.49%] at 20 m and 16.63% [10.53%, 21.04%] at 40 m. The paired nonlinear-RMSE difference remained comparatively small, with medians of 6.26 mm [0, 11.96 mm] and 8.77 mm [0, 20.48 mm] at the two largest uncertainty levels. Figure 14 summarizes these paired effects.

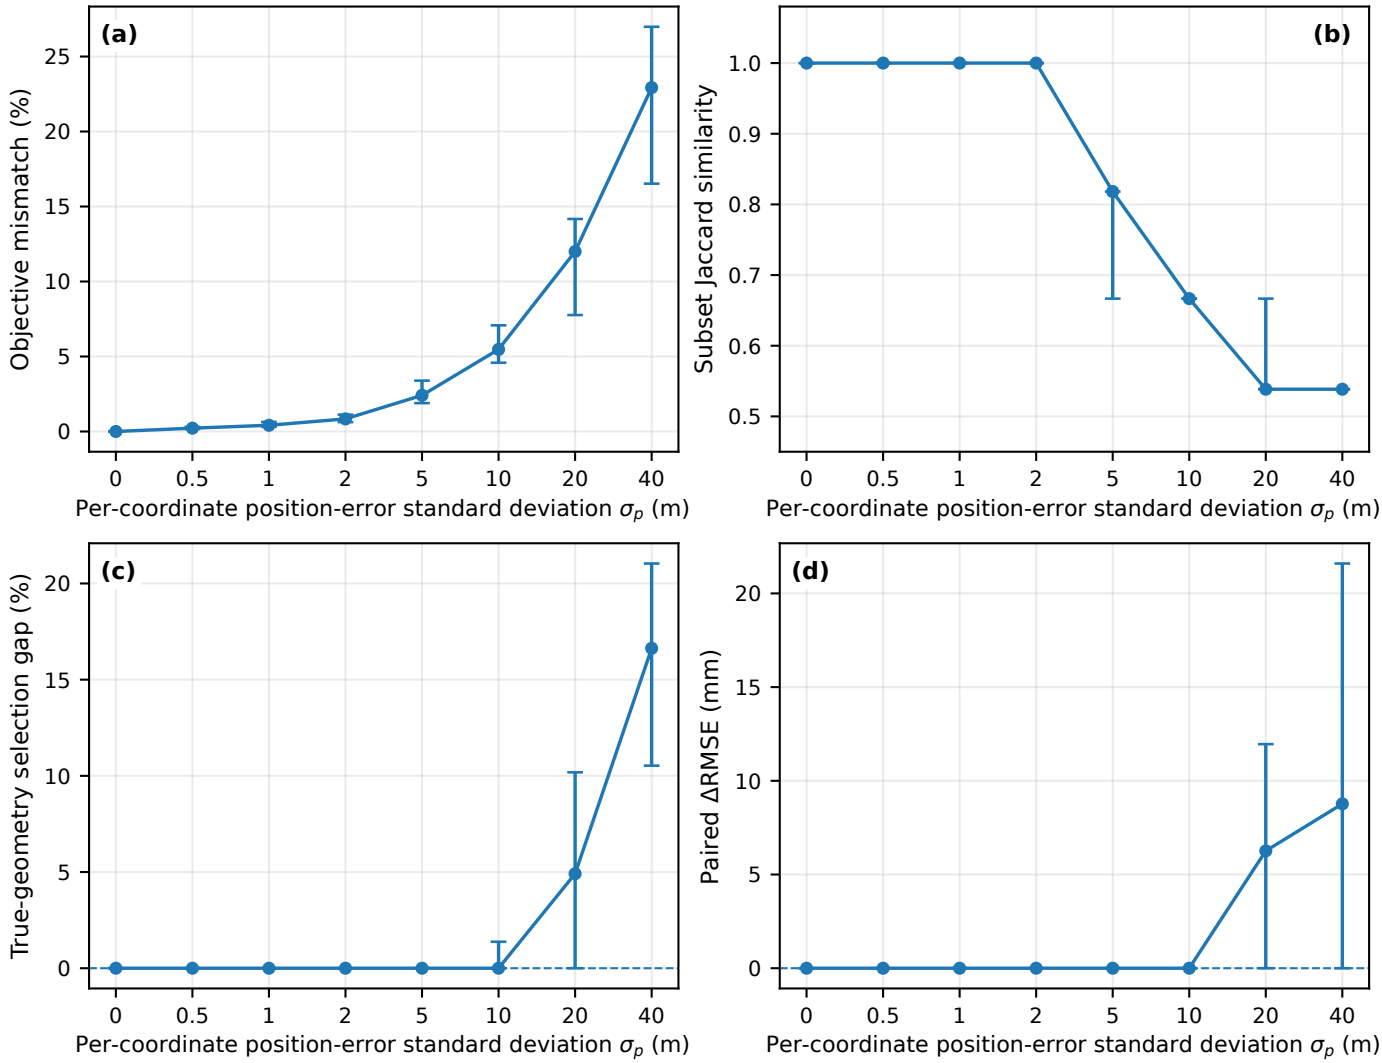


Figure 14. Sensitivity of cooperative sensor node pre-selection to predicted-position uncertainty. The per-coordinate prediction-error standard deviation σ_p is varied using paired Gaussian error realizations while other trial-level quantities are held fixed. (**a**) Absolute relative mismatch of the A-optimal objective between predicted and true geometry for the selected node subset obtained using predicted geometry. (**b**) Jaccard similarity between selected node subsets obtained under predicted and true geometry. (**c**) Signed true-geometry objective gap relative to the true-geometry greedy reference. (**d**) Paired difference in rigid-aligned nonlinear localization RMSE. Markers show medians and error bars show pointwise 95% percentile bootstrap confidence intervals over the 97 true-geometry-reference-success trials.

Three of the 100 trial groups were unobservable under the true-geometry reference and remained common-unobservable cases across all tested uncertainty levels. Among the 97 true-geometry-reference-success trials, the additional geometry-induced selection-failure rate, localization-failure rate, and predicted-versus-true observability-status-mismatch rate were each 0/97 = 0% at every tested σ_p (95% Wilson CI: 0–3.81%).

4.3. Discussion and Practical Implications

The validation results define a clear applicability boundary for the proposed criterion. Across different topology realizations, G-GDOP is not treated as a direct predictor of the final error of a nonlinear localization pipeline because that error also depends on initialization, convergence, rigid alignment, and the realized noise and bias. Within a fixed

topology and local linearization regime, however, the ordered normalized local RMSE values of 0.698, 1.014, and 1.980 for the Low, Medium, and High groups demonstrate that G-GDOP provides a meaningful information-quality ranking. It should therefore be interpreted as a local observability and structural-screening metric rather than as an absolute global-error predictor.

The resource experiments further show why the selected node subset and the induced active-link set must be distinguished. Raising K from 6 to 20 increased the median induced-link number from 109 to 214 and the median pre-selection time from 0.0578 to 0.0963 s, while reducing the median normalized objective from 7.994 to 0.777 and the median RMSE from 0.550 to 0.148 m. The ranging-slot sweep produced the same information–accuracy–resource pattern. At the full tested budget, Weighted G-GDOP approached the All-Node Reference while activating fewer links. The small-scale exhaustive comparison also showed that the absolute objective values became progressively closer as K increased, but this empirical result is not presented as a general approximation guarantee.

The $K = 3$ result further identifies a boundary behavior of the two-stage heuristic. At this smallest tested budget, the observable seed consumed the entire node budget in every paired trial, so the large objective difference cannot be attributed to the post-seed A-optimal Stage II; that stage was never entered for an additional node selection iteration. Instead, the final selected node subset of size three was fixed entirely by the observability-first Stage I ordering. The diagnostic attribution showed that retained minimum eigenvalue was the dominant source of the first divergence from the exhaustive optimum, whereas deterministic node-ID tie-breaking explained a smaller number of cases. Induced-link, ranging-slot, and general resource increments did not explain the $K = 3$ difference in this particular dense small-scale configuration because they were equal across the feasible candidate choices. Thus, $K = 3$ should be viewed as an observable-seed boundary case rather than as a direct measure of the typical approximation behavior of the post-seed A-optimal greedy stage.

The robustness analyses identify two distinct sources of uncertainty: link-quality errors and predicted-geometry errors. Under heterogeneous ranging noise, link-dependent weighting improved the information objective and empirical RMSE relative to the tested baselines, whereas positive NLOS bias degraded both the local surrogate and nonlinear localization accuracy, confirming that covariance weighting alone does not remove deterministic bias. Online link-class and variance errors also increased the true-objective difference and reduced subset agreement relative to the oracle, although Weighted G-GDOP remained closer to the oracle objective than Unweighted G-GDOP over the tested perturbation grid.

The predicted-position experiment further shows that geometry uncertainty affects the information-based pre-selection stage before it produces a comparable change in downstream nonlinear localization accuracy. At the default $\sigma_p = 2$ m per coordinate, the median predicted-versus-true objective mismatch was only 0.840%, the median subset Jaccard index remained 1.000, and both the median true-geometry selection gap and paired nonlinear-RMSE difference were zero. With increasing uncertainty, subset agreement decreased and the objective mismatch increased; at $\sigma_p = 40$ m, the median objective mismatch reached 22.93%, the median Jaccard index was 0.538, and the median true-geometry selection gap reached 16.63%. The corresponding median paired nonlinear-RMSE difference remained 8.77 mm. No predicted-versus-true observability-status mismatch was observed, and no additional geometry-induced failure occurred among the 97 true-geometry-reference-success trials. These results support the use of predicted geometry under the tested default and moderate uncertainty levels, but they do not imply invariance to large position errors.

The implementation results likewise separate numerical correctness from runtime benefit. Direct Inversion and Woodbury Updating were numerically equivalent within the predefined tolerance in all 24 tested network-size and graph-regime combinations, while a Woodbury runtime advantage was supported in 20 conditions. The benefit was weaker in several small or sparse cases and reached its largest median speedup of 1.947 at $N = 100$ in the fixed-radius regime. The recursive update should therefore be interpreted as a numerically consistent acceleration mechanism whose practical benefit depends on network size, graph structure, and update width rather than as an unconditional real-time guarantee.

From an engineering perspective, the proposed framework is a snapshot-level cooperative sensor node pre-selection layer intended to precede a downstream nonlinear or factor-graph estimator. The present experiments do not validate a specific GTSAM, iSAM2, or sliding-window implementation, and the normalized resource cost is not a hardware-calibrated energy measurement. Moreover, the controlled geometry experiment deliberately holds the candidate graph fixed and uses snapshot-wise Gaussian position perturbations; it therefore does not represent topology-discovery errors, temporally correlated state-prediction drift, latency, packet loss, or closed-loop dynamic state propagation. Link-quality uncertainty is represented by controlled class-confusion and variance-error models rather than by a measured classifier or a fully non-Gaussian UWB error model. The validation is also Monte Carlo based and does not include a public measured UWB dataset, hardware-in-the-loop testing, or UAV flight experiments. Accordingly, the practical claims are restricted to the information quality, subset selection, observability behavior, nonlinear RMSE, resource use, and computational behavior observed under the reported simulation conditions. Future work should investigate jointly uncertain geometry and topology, temporally correlated prediction errors, robust bias and outlier treatment, factor-graph integration, temporal scheduling, and validation using UWB-equipped UAV platforms or hardware-in-the-loop systems.

5. Conclusions

This paper developed a Fisher-information-based cooperative sensor node pre-selection framework for UWB-aided GNSS-denied UAV swarm localization. All mobile nodes remain in the localization state, while nodes in the resource-limited selected node subset are scheduled as active UWB ranging initiators, thereby inducing the active ranging-link set. The formulation separately represents the selected-node budget, induced-link number, ranging-slot occupation, residual-energy eligibility, and a normalized general resource cost.

A gauge-free projected Fisher information matrix was used to remove the rigid-body ambiguity of anchor-free range-only localization, and a trace-based G-GDOP/A-optimal criterion was constructed from link geometry, heterogeneous ranging reliability, and cooperative topology. A two-stage greedy heuristic first builds an observable seed and then uses marginal information gain with recursive Woodbury updates. The complexity analysis explicitly separates preprocessing, resource-feasibility screening, observable-seed construction, and post-seed candidate evaluation.

The nine formal experiments establish both the usefulness and the applicability boundary of the proposed method. G-GDOP did not serve as a direct predictor of global nonlinear recovery error, but it produced clear same-topology local-error ordering. Increasing node or slot resources improved the information objective and localization RMSE at the cost of more induced links and computation. Small-scale exhaustive enumeration provided an absolute optimality benchmark without implying a general approximation ratio. Gaussian-noise, heterogeneous-noise, NLOS-bias, online weight-error, and predicted-position-uncertainty

experiments showed that link-dependent weighting is beneficial when its reliability estimates are informative, while deterministic bias, weight-estimation mismatch, and large predicted-position errors degrade information quality, localization performance, or subset agreement.

Direct Inversion and Woodbury Updating remained numerically consistent, while Woodbury provided a runtime advantage in most tested conditions. Its benefit nevertheless depended on network size and graph regime; the framework should therefore be viewed as an interpretable structural-screening module rather than an unconditional real-time or bias-robustness guarantee.

The present evidence remains limited to snapshot-level simulations with predicted-state and parameterized link-quality inputs. Future work will address dynamic and jointly uncertain topology/geometry, robust bias and outlier handling, factor-graph integration, hardware-calibrated resource models, and UWB-equipped UAV or hardware-in-the-loop validation.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/s26165164/s1, Table S1: Numerical equivalence of Direct Inversion and Woodbury Updating across network sizes and graph regimes; Table S2: Runtime comparison of Direct Inversion and Woodbury Updating across all tested network sizes and graph regimes.

Author Contributions: Conceptualization: Y.S., X.D. and P.G.; methodology: Y.S. and X.D.; software: Y.S.; validation: Y.S., X.D. and P.G.; formal analysis: Y.S. and X.D.; investigation: Y.S.; resources: X.D. and P.G.; data curation: Y.S.; writing—original draft preparation: Y.S.; writing—review and editing: X.D. and P.G.; visualization: Y.S.; supervision: X.D. and P.G.; project administration: X.D. and P.G.; funding acquisition: X.D. and P.G. All authors have read and agreed to the published version of the manuscript.

Funding: This study was funded by the Natural Science Foundation of Shandong Province, grant number ZR2021MG021, and the Youth Innovation Technology Project of Higher School in Shandong Province, grant number 2021RW030.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

UAV	Unmanned Aerial Vehicle
GNSS	Global Navigation Satellite System
UWB	Ultra-Wideband
IMU	Inertial Measurement Unit
LOS	Line-of-Sight
wLOS	Weak Line-of-Sight
NLOS	Non-Line-of-Sight
VIO	Visual-Inertial Odometry
SLAM	Simultaneous Localization and Mapping

GDOP	Geometric Dilution of Precision
G-GDOP	Generalized Geometric Dilution of Precision
FIM	Fisher Information Matrix
EFIM	Equivalent Fisher Information Matrix
CRLB	Cramér–Rao Lower Bound
RMSE	Root Mean Square Error
MDS	Multidimensional Scaling

References

1. Pascacio, P.; Casteleyn, S.; Torres-Sospedra, J.; Lohan, E.-S.; Nurmi, J. Collaborative Indoor Positioning Systems: A Systematic Review. *Sensors* **2021**, *21*, 1002. [CrossRef] [PubMed]
2. Gabela, J.; Retscher, G.; Goel, S.; Perakis, H.; Masiero, A.; Toth, C.; Gikas, V.; Kealy, A.; Koppanyi, Z.; Blaszczak-Bak, W.; et al. Experimental Evaluation of a UWB-Based Cooperative Positioning System for Pedestrians in GNSS-Denied Environment. *Sensors* **2019**, *19*, 5274. [CrossRef] [PubMed]
3. Han, Y.; Wei, C.; Li, R.; Wang, J.; Yu, H. A Novel Cooperative Localization Method Based on IMU and UWB. *Sensors* **2020**, *20*, 467. [CrossRef] [PubMed]
4. Wang, R.; Xu, C.; Li, R.; Duan, S.; Zhang, X. Cooperative Localization and Mapping Based on UWB/IMU Fusion Using Factor Graphs. *IEEE Sens. J.* **2024**, *24*, 21931–21940. [CrossRef]
5. Cao, Y.; Beltrame, G. VIR-SLAM: Visual, Inertial, and Ranging SLAM for Single and Multi-Robot Systems. *Auton. Robot.* **2021**, *45*, 905–917. [CrossRef]
6. Xu, W.; Chen, Y.; Liu, S.; Nie, A.; Chen, R. Multi-Robot Cooperative Simultaneous Localization and Mapping Algorithm Based on Sub-Graph Partitioning. *Sensors* **2025**, *25*, 2953. [CrossRef] [PubMed]
7. Wang, C.; Li, B.; Duan, Y.; Sui, X.; Shi, Z.; Gao, S.; Zhang, Z.; Chen, J. An Improved Robust ESKF Fusion Positioning Method with a Novel UWB-VIO Initialization. *Sensors* **2026**, *26*, 1804. [CrossRef] [PubMed]
8. Nowakowski, M.; Idzkowski, A. Ultra-Wideband Signal Transmission According to European Regulations and Typical Pulses. In Proceedings of the 2020 International Conference Mechatronic Systems and Materials (MSM), Bialystok, Poland, 1–3 July 2020; pp. 1–4.
9. Efremova, E.V.; Kuzmin, L.V. Wireless Ranging by Evaluating Received Signal Strength of UWB Chaotic Radio Pulses: Effects of Signal Propagation Conditions. *Technologies* **2024**, *12*, 141. [CrossRef]
10. Li, Q.; Wang, P.; Li, X.; Zhang, J.; Luo, Y.; Yu, W.; Cheng, H. Formation-Constrained Cooperative Localization for UAV Swarms in GNSS-Denied Environments. *Sensors* **2026**, *26*, 1984. [CrossRef] [PubMed]
11. Zhang, S.; Zhang, H.; Zhan, Y.; Wei, X.; Liu, Y. Cluster Networking and Cooperative Localization Based on Biogeography Optimization and Improved Super-Multidimensional Scaling for Multi-Unmanned Aerial Vehicles. *Sensors* **2025**, *25*, 2887. [CrossRef] [PubMed]
12. Liu, H.; Jiang, W.; Long, Q.; Xia, Q.; Chen, X. A High-Precision Cooperative Localization Method for UAVs Based on Multi-Condition Constraints. *Sensors* **2026**, *26*, 1641. [CrossRef] [PubMed]
13. Patwari, N.; Ash, J.N.; Kyperountas, S.; Hero, A.O.; Moses, R.L.; Correal, N.S. Locating the Nodes: Cooperative Localization in Wireless Sensor Networks. *IEEE Signal Process. Mag.* **2005**, *22*, 54–69. [CrossRef]
14. Shen, Y.; Wymeersch, H.; Win, M.Z. Fundamental Limits of Wideband Localization—Part II: Cooperative Networks. *IEEE Trans. Inf. Theory* **2010**, *56*, 4981–5000. [CrossRef]
15. Win, M.Z.; Dai, W.; Shen, Y. A Theoretical Foundation of Network Localization and Navigation. *Proc. IEEE* **2018**, *106*, 1136–1165. [CrossRef]
16. Quan, Q. The Tightest Geometry for Lowest GDOP in Range-Based 2-D Wireless Location Systems. In Proceedings of the World Congress on Computer Science and Information Engineering, Los Angeles, CA, USA, 31 March–2 April 2009.
17. Zhou, R.; Sun, H.; Li, H.; Li, W. TDOA and Track Optimization of UAV Swarm Based on D-Optimality. *J. Syst. Eng. Electron.* **2020**, *31*, 1140–1151. [CrossRef]
18. Dellaert, F.; Kaess, M. Factor Graphs for Robot Perception. *Found. Trends Robot.* **2017**, *6*, 1–139. [CrossRef]
19. Bel, A.; Vicario, J.L.; Seco-Granados, G. Localization Algorithm with On-Line Path Loss Estimation and Node Selection. *Sensors* **2011**, *11*, 6905–6925. [CrossRef] [PubMed]
20. Tang, C.; Dou, L. An Improved Game Theory-Based Cooperative Localization Algorithm for Eliminating the Conflicting Information of Multi-Sensors. *Sensors* **2020**, *20*, 5579. [CrossRef] [PubMed]
21. Yin, T.; Zou, D.; Zhang, R.; Jing, D.; Li, Y.; Lu, X. Evolutionary Overlapping Coalitional Game-Based Link Selection for Distributed Cooperative Localization in Mobile Networks. *Wirel. Commun. Mob. Comput.* **2022**, *2022*, 6588197. [CrossRef]

22. Qiu, F.; Zhang, W. Efficient Cooperative Localization Method with Node Selection Based on Position Error Bound. *Signal Process.* **2023**, *209*, 109037. [CrossRef]
23. Chen, G.; Cheng, L.; Zeng, Q.; Shen, F.; Zhang, Y.-D. Willingness Allocation-Assisted Cooperative Localization Algorithm Based on Competitive Game for Resource-Constrained Environment. *IEEE Trans. Green Commun. Netw.* **2025**, *9*, 498–512. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.