
A Probabilistic Approach to Pose Synchronization for Multi-Reference Alignment with Applications to MIMO Wireless Communication Systems

Rob Romijnders
University of Amsterdam

Gabriele Cesa
Qualcomm AI Research

Christos Louizos
Qualcomm AI Research

Kumar Pratik
Qualcomm AI Research

Arash Behboodi
Qualcomm AI Research

Abstract

From molecular imaging to wireless communications, the ability to align and reconstruct signals from multiple misaligned observations is crucial for system performance. We study the problem of multi-reference alignment (MRA), which arises in many real world problems, such as cryo-EM, computer vision and, in particular, wireless communication systems. Using a probabilistic approach to model MRA, we find a new algorithm that uses relative poses as nuisance variables to marginalize out – thereby removing the global symmetries of the problem and allowing for more direct solutions and improved convergence. The decentralization of this approach enables significant computational savings by avoiding the cubic scaling of centralized methods through cycle consistency. Both proposed algorithms achieve lower reconstruction error across experimental settings.

1 Introduction

Multi-reference alignment (MRA) is a fundamental problem in signal processing and statistical inference that arises when reconstructing a signal from multiple noisy observations, and each observation has been transformed by an unknown group action such as rotation, translation, or other geometric transformation. This problem appears in diverse applications across scientific domains, from cryo-electron microscopy (Cryo-EM) [11, 9] and cryo-electron tomography (Cryo-ET) for molecular structure determination [17, 16, 14], to computer vision applications like Structure-from-Motion (SfM) and Simultaneous Localization and Mapping (SLAM), and even narrow-band 5G NR [13].

A core challenge in MRA lies in the fact that unknown transformations of the signal create an ambiguity. The signal could be recovered if the poses were known, and the unknown poses could be recovered when the signal was known. Moreover, the original signal can only be recovered up to a global transformation and, likewise, the observations’ misalignments (i.e., their absolute poses) are only defined up to a joint transformation. In other words, any global transformation applied to all observations’ poses yields an equally valid solution. Traditional approaches to MRA typically treat the absolute poses of the observations as nuisance variables to be estimated or marginalized over in an intermediate step before signal reconstruction. This includes various probabilistic approaches based on variants of the expectation-maximization algorithm [14] or other message-passing algorithms [12]. However, these methods can face challenges due to the inherent global symmetry of the problem.

Indeed, any disjoint estimate of the absolute poses necessarily breaks the global symmetry of the problem, which can lead to optimization difficulties in expectation-maximization or maximization-maximization algorithms. The fundamental issue is that global poses are not uniquely defined.

Any global transformation of the poses leads to another equally valid assignment. This can create convergence issues or trap optimization algorithms in local minima, e.g., due to subsets of observations initially converging to different global alignments.

In this work, we propose a novel probabilistic approach to MRA that addresses these challenges by reformulating the problem in terms of relative poses rather than absolute ones. Our key insight is that relative poses, in contrast to absolute ones, are uniquely defined and do not suffer from the global symmetries that can cause optimization difficulties. However, not all possible assignments of relative poses correspond to a set of absolute ones. As we prove in Appendix E, relative poses that are cycle-consistent are in one-to-one correspondence with equivalence classes of absolute pose assignments. Hence, to ensure that the relative poses represent valid configurations, we incorporate a cycle consistency constraint into our probabilistic model, which enforces that walks along any cycle in the relative-poses graph lead to identity transformations.

As a concrete example of the multi-reference alignment problem, we consider the problem of channel estimation (CE) in narrow-band 5G NR wireless communication systems [1, 13]. In the multi-input multi-output (MIMO) setting, the 5G NR specification allows channels to be precoded individually. Still, signals in neighboring PRG blocks can be correlated (Orthogonal Frequency Modulation and Demodulation). Current receivers do not utilize this correlation, resulting in suboptimal performance. We are the first to combine the group-theoretic approaches from multi-reference alignment with the application of 5G decoding using statistical methods, demonstrating how our relative pose framework can be applied to real-world communication systems.

Finally, we propose different algorithms to efficiently perform inference on this graph and evaluate them on multiple synthetic datasets, showing significant improvements over traditional approaches that rely on absolute pose estimation. In summary, this paper makes the following key contributions:

- We present the first MIMO decoding framework that explicitly incorporates relative rotations, addressing a fundamental limitation of existing methods.
- We develop both a direct method and an iterative refinement algorithm for joint denoising and synchronization, grounded in a principled probabilistic model, also known as a structured probabilistic graphical model.
- We provide theoretical insights and empirical evidence of the advantages of our approach in 5G-like MIMO scenarios.

The remainder of this paper is organized as follows: Section 2 formulates the problem, introduces our probabilistic graphical model, and reviews related work on MIMO synchronization and pose estimation. Section 3.2 details the proposed algorithms. Section 4 presents experimental results, and Section 5 concludes the paper with directions for future research.

2 Problem formulation

Background on Narrow-band 5G NR MIMO system The 5G NR communication system uses an orthogonal-frequency division multiplexing (OFDM) scheme to communicate messages along orthogonal subcarrier signals at each time step. As a result, the communication channel is divided into a time-frequency grid, with each cell (or resource element, RE) carrying a single symbol. In a multiple-input-multiple-output (MIMO) scheme, both the transmitter and receiver have multiple antennas, which enables the transmission of multiple symbols on the same time-frequency location. Formally, assuming d transmitting and receiving antennas, the transmission at a frequency f and time step t of a symbol vector $x_{f,t} \in \mathbb{C}^d$ is

$$y_{f,t} = \mathbf{H}_{f,t} \mathbf{P}_{f,t} x_{f,t} + \varepsilon \quad (1)$$

where $\mathbf{H}_{f,t} \in \mathbb{C}^{d \times d}$ is the MIMO channel matrix, $y \in \mathbb{C}^d$ the received signal, $\varepsilon \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2 \mathbf{I})$ is Gaussian noise and $\mathbf{P} \in U(n)$ is a unitary matrix which defines the *precoding*. The transmitter precodes the transmitted messages to help minimize interference and improve SNR, e.g., by aligning the message along the singular vectors of the channel matrix. For practical reasons, in narrow-band 5G NR, the precoding matrix is shared by multiple neighboring resource elements within the same physical resource block group (PRG). The challenge for the receiver is estimating the effective channel matrices (i.e., after precoding) from a few noisy observations obtained via some pilot messages. We consider a multi-reference alignment problem inspired by this narrow-band 5G

NR channel estimation problem. Our problem formulation allows us to abstract away from many technical details of the communication protocol, while preserving its relevant properties and the underlying synchronization challenges. For more details, we refer to Chen et al. [5], Pratik et al. [13] for a more precise description of the real channel estimation problem in narrow-band 5G NR.

Abstract formulation For ease of modelling, we model the channel matrix as real numbers and the precoding matrices as real orthogonal matrices $\mathbf{P} \in O(d) \subset \mathbb{R}^{d \times d}$. We denote the channel in the j -th PRG block by $\mathbf{H}_i \in \mathbb{R}^{D \times d}$, where the dimension D indexes all time-frequency positions within the block and the output antennas. In a realistic 5G MIMO scenario, the receiver observes the effective channels $\mathbf{H}'_i = \mathbf{H}_i \mathbf{P}_i$ via some pilot messages only at certain time-frequency positions (demodulation reference signal, DMRS). This signal is then used to estimate the full channel via linear predictors. We assume the receiver observes the full channel¹ corrupted by real Gaussian noise, i.e. $\mathbf{B}_i = \mathbf{H}_i \mathbf{P}_i + \varepsilon_i \in \mathbb{R}^{D \times d}$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2 \mathbf{I})$.

Taken all random variables together, \mathbf{H} , \mathbf{P} , and \mathbf{B} , the joint probability distribution is:

$$p(\mathbf{B}, \mathbf{P}, \mathbf{H}) = p(\mathbf{B}|\mathbf{H}') p(\mathbf{H}'|\mathbf{H}, \mathbf{P}) p(\mathbf{P}) p(\mathbf{H}) \quad (2)$$

where $p(\mathbf{B}|\mathbf{H}') = \mathcal{N}(\mathbf{B}|\mathbf{H}', \sigma^2 \mathbf{I})$, $p(\mathbf{H}'|\mathbf{H}, \mathbf{P}) = \delta_{\mathbf{H}' - \mathbf{H}\mathbf{P}}$ is a deterministic factor applying the precoding to the channel. We assume the prior $p(\mathbf{P})$ on the rotation precoding matrices is uniform over the $O(d)$ group. Finally, $p(\mathbf{H})$ specifies a rotation invariant² prior via a matrix normal distribution over all underlying channel matrices $\mathcal{MN}(0, \mathbf{U}, \mathbf{I})$ with \mathbf{I} the identity matrix and matrix \mathbf{U} modelling the row-covariance. Note the rotation invariance: if $\mathbf{P} \in O(d)$ and $\mathbf{X} \sim \mathcal{MN}(0, \mathbf{U}, \mathbf{I})$ then $\mathbf{X}\mathbf{P} \sim \mathcal{MN}(0, \mathbf{U}, \mathbf{P}^T \mathbf{P} = \mathbf{I})$. We also refer to this prior as ‘spatial correlation,’ which is shorthand for the correlation that occurs along the time and frequency axes among the PRG blocks due to Rayleigh scattering [18, 5]. This spatial correlation enables the estimation of the channel in a PRG block \mathbf{H}'_i from its noisy observations \mathbf{B}_i . However, the unknown precoding matrices prevent the receiver from leveraging the correlation between neighboring blocks. Our key insight is that by modeling the relative rotations $\mathbf{R}_{ij} = \mathbf{P}_i^{-1} \mathbf{P}_j$ between adjacent blocks in the time-frequency grid, we can leverage the spatial and temporal correlations to achieve more robust channel estimation.

Absolute to relative poses Multi-reference alignment problems typically present a global symmetry, which implies that absolute poses and the underlying signal are not uniquely defined. Indeed, the set of channels $\mathbf{H} = \{\mathbf{H}_i\}_i$ and precoding matrices $\mathbf{P} = \{\mathbf{P}_i\}_i$ generate the same observations $\mathbf{B} = \{\mathbf{B}_i\}_i$ as the channels $\tilde{\mathbf{H}} = \{\mathbf{H}_i \mathbf{Q}\}_i$ and the precoding $\tilde{\mathbf{P}} = \{\mathbf{Q}^{-1} \mathbf{P}_i\}_i$ for any transformation $\mathbf{Q} \in O(d)$. Any attempt to estimate these quantities requires some form of symmetry breaking. However, this problem vanishes when taking a local viewpoint and, therefore, considering relative poses; see Appendix E.1 For this reason, we develop our method to recover relative poses $\mathbf{R} = \{\mathbf{R}_{i,j} := \mathbf{P}_i^{-1} \mathbf{P}_j\}_{i,j}$, which are uniquely defined and suitable for an efficient distributed and local computational scheme.

Relation to prior work in synchronization: synchronization over compact groups has been approached with approximate message passing methods [12] or semi-definite programming [3]. Other works studied CryoEM [14] and stochastic block models [2]. The fundamental difference in our work, though, is that we aim to estimate the signal itself. [12] assume noisy observation of the relative poses and strive to find an optimal representation of the absolute poses. In our case, we have noisy observations of the underlying signal at each location. The goal is to efficiently and effectively marginalize over the relative poses to obtain a better estimate of the channel. Although our approach utilizes similar estimations of the relative poses, it does not synchronize the absolute poses. Instead, it can be said informally that we ‘directly synchronize’ the underlying signal. [13] first describes the MRA of narrow-band 5G NR channel estimation, but adopted complex deep learning architectures. **Relation to wider related work** Synchronization problems have been studied in a more abstract context for their group-theoretic properties [7, 6]. Some works have studied pairwise relative pose estimation by candidate elimination [19], but we are the first structured statistical model to marginalize over the unknown poses to reconstruct the underlying signal effectively and efficiently.

¹Note that pilot DMRS signals provide linear measurements of the channel, so our simplification reduces to the use of a diagonal full rank covariance matrix between \mathbf{B}_i and \mathbf{H}'_i , rather than a more generic one.

²This is a reasonable assumption since it is equivalent to assuming zero mean and that the correlation depends only on the frequency and time position and is the same among all the antennas, since the precoding transformation only acts on the antennas’ dimension of the channels matrices.

3 Method

3.1 Direct estimation

A direct approach is to estimate the relative rotations \mathbf{R}_{ij} between adjacent nodes \mathbf{H} in the time-frequency grid. This is done by maximizing the likelihood of the data $\log p(\mathbf{B})$. Observations $\mathbf{B} \in \mathbb{R}^{D \times d}$ is the matrix of D times a d dimensional observation. However, the direct approach has two disadvantages: estimating in pairs does not enable global synchronization, and direct estimates can be suboptimal as rotation and observation noise are confounded, where estimating one changes the other. We address both properties in the next section.

Estimating the channel \mathbf{H}' can be done by writing the posterior given the data \mathbf{B} directly:

$$\arg \max_{\mathbf{H}'} p(\mathbf{H}' | \mathbf{B}) = \int_{\mathbf{R}} p(\mathbf{H}' | \mathbf{B}, \mathbf{R}) p(\mathbf{R} | \mathbf{B}) d\mathbf{R} \quad (3)$$

The posterior $p(\mathbf{R} | \mathbf{B})$ can be heuristically approximated with a delta distribution – meaning that we can take a point estimate of the rotation. Assuming only two blocks $\{\mathbf{H}_1, \mathbf{H}_2\}$, if the covariance matrix of the prior on the correlated channel is \mathbf{U} , then one can find the relative rotation \mathbf{R}_{12} between the two blocks from the following projection: $\arg \max_{\mathbf{R}_{12}} \langle \mathbf{R}_{21}, \mathbf{H}_2'^T \mathbf{U}_a^T \mathbf{H}_1' \rangle_F$.

Likewise, for a triplet of three nodes $\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3$ and relative poses \mathbf{R}_{12} and \mathbf{R}_{13} (note that $\mathbf{R}_{23} := \mathbf{R}_{12}^T \mathbf{R}_{13}$ and that $\mathbf{R}_{ij} = \mathbf{R}_{ji}^T$), one can iterate the two updates below in Equation 4 until convergence. Matrices $\mathbf{U}_a, \mathbf{U}_b, \dots$ are submatrices of the covariance matrix \mathbf{U} , as described with Equation 30 in the Appendix.

$$\begin{cases} \mathbf{R}_{13} = \arg \max \langle \mathbf{R}_{21}, \mathbf{B}_2^T \mathbf{U}_a^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T \mathbf{U}_b^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T \mathbf{U}_c^T \mathbf{B}_2 \mathbf{R}_{21} \rangle_F \\ \mathbf{R}_{12} = \arg \max \langle \mathbf{R}_{21}, \mathbf{B}_2^T \mathbf{U}_a^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T \mathbf{U}_b^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{21}, \mathbf{B}_2^T \mathbf{U}_c^T \mathbf{B}_3 \mathbf{R}_{31} \rangle_F. \end{cases} \quad (4)$$

The derivation is outlined in Appendix C. Crucially, the above solution arises from using properties of the Matrix Normal distribution. Ending up with the above Frobenius inner-products is a result of the Orthogonal Procrustes problem [8, 15], which is a mean-square solution when optimizing over the set of unitary matrices. The projection is a closed-form maximizer of the Procrustes problem. Therefore, each step of the iteration is non-decreasing in the objective value, and we find that this usually converges in three to five steps.

Secondly, obtaining the channel estimate from the observations and rotation matrices $\arg \max p(\mathbf{H}' | \mathbf{B}, \mathbf{R})$ is a linear problem, which can be solved in $O(D^3)$ time:

$$\begin{bmatrix} \mathbf{H}_1' \\ \mathbf{H}_2' \hat{\mathbf{R}}_{21} \\ \mathbf{H}_3' \hat{\mathbf{R}}_{31} \end{bmatrix} = \left(\sigma_\epsilon^2 \mathbf{U}^{-1} + \mathbf{I}_{3D} \right)^{-1} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \hat{\mathbf{R}}_{21} \\ \mathbf{B}_3 \hat{\mathbf{R}}_{31} \end{bmatrix} \quad (5)$$

The above solutions are worked out for a ‘triplet’ of nodes, meaning that we consider three nodes and derive the estimation equations. This could be worked out for an arbitrary number of nodes. However, this has two downsides: (a) the computation scales cubically in the number of measurements and thus cubically in the number of nodes. Hence, scaling this approach on the wide time-frequency lattices of up to hundreds blocks per time slot is impractical, especially under the resource and power constraints of modern 5G telecom modems [5]. The second downside (b) to including more nodes is that synchronization becomes ‘harder.’ After all, only noisy nodes are available, and harmonizing among more noisy observations pertaining to random rotations has more numerical issues.

Therefore, we suggest decentralizing the estimation into triplets. A triplet of nodes is the smallest subgraph in which a local synchronization can imply a global synchronization over the entire grid, see Appendix E. Meanwhile, a triplet of nodes is most economical with the local ‘spatial correlation’ that exists among the channels. Due to Rayleigh scattering [18], only nearby PRG blocks will have correlation, and denoising among local subgraphs most economically uses the compute resources.

3.2 Synchronization and denoising via iterative refinement

From the direct solution, we suggest an iterative refinement method that is loosely inspired by the Expectation Maximization algorithm. The direct solution alternates between two maximizations,




Figure 1: Iterative refinement achieves the lowest reconstruction error compared to signal denoising only. This is a synthetic 6-by-6 grid of 12 measurements each. Error bars indicate standard error among 25 random seeds. Additionally, the ‘synchronization base’ approach, by itself, achieves lower reconstruction error than prior art for reasonable signal-to-noise ratios.

over $p(\mathbf{H}'|\mathbf{B}, \mathbf{R})$ and over $p(\mathbf{R}|\mathbf{B})$. The latter maximization assumes that channel observations \mathbf{B} are made and the relative poses can be estimated from them. By extension, \mathbf{H}' will serve as a ‘denoised version’ of the channel. As such, we will report the Mean-Square error between \mathbf{H}' and \mathbf{B} . Then, as \mathbf{H}' is a denoised version of \mathbf{B} , we can refine the rotation estimate from $p(\mathbf{R}|\mathbf{H}')$ instead. Continuing, the relative poses can be found from $p(\mathbf{H}'|\mathbf{B}, \mathbf{R})$.

Loosely, this is inspired by the Expectation-Maximization algorithm [4, 10]. Analogous to the M-step, we estimate the parameters of the channel using an estimate for the rotation random variable. Analogous to the E-step, we find an estimate for the relative poses. In this case, however, we take only a point estimate, $q(\mathbf{R}) = \delta_{\mathbf{R}}$ in $\log p(\mathbf{B}) \geq \max_{q(\mathbf{R})} \mathbb{E}_{\mathbf{R} \sim q(\mathbf{R})} [\log p(\mathbf{B}, \mathbf{R}, \mathbf{H}'_t) - \log q(\mathbf{R})]$.

The estimation for the relative poses can be done using similar Frobenius projections, starting from $\arg \max_{\mathbf{R}} p(\mathbf{R}|\mathbf{H}')$. This time, however, the estimates for the channel are used:

$$\begin{cases} \mathbf{R}_{12} = \arg \max_{\mathbf{R}_{12}} \langle \mathbf{R}_{21}, \mathbf{H}'_{2,t} \mathbf{U}_a^T \mathbf{H}'_{1,t} \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_{3,t} \mathbf{U}_b^T \mathbf{H}'_{1,t} \rangle_F + \langle \mathbf{R}_{21}, \mathbf{H}'_{2,t} \mathbf{U}_c \mathbf{H}'_{3,t} \mathbf{R}_{31} \rangle_F \\ \mathbf{R}_{13} = \arg \max_{\mathbf{R}_{13}} \langle \mathbf{R}_{21}, \mathbf{H}'_{2,t} \mathbf{U}_a^T \mathbf{H}'_{1,t} \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_{3,t} \mathbf{U}_b^T \mathbf{H}'_{1,t} \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_{3,t} \mathbf{U}_c^T \mathbf{H}'_{2,t} \mathbf{R}_{21} \rangle_F \end{cases} \quad (6)$$

The channel estimation is derived in Appendix D and results in an estimator similar to Equation 5:

$$\begin{bmatrix} \mathbf{H}'_{1,t+1} \\ \mathbf{H}'_{2,t+1} \hat{\mathbf{R}}_{21} \\ \mathbf{H}'_{3,t+1} \hat{\mathbf{R}}_{31} \end{bmatrix} = \left(\sigma_c^2 \mathbf{U}^{-1} + \mathbf{I}_{3D} \right)^{-1} \begin{bmatrix} \mathbf{H}'_{1,t} \\ \mathbf{H}'_{2,t} \hat{\mathbf{R}}_{21} \\ \mathbf{H}'_{3,t} \hat{\mathbf{R}}_{31} \end{bmatrix} \quad (7)$$

Together, Equations 6 and 7 can be repeated to refine the channel estimate and the rotation estimate. We typically find that this stabilizes after three or four iterations, which is illustrated in the experimental results section.

4 Experimental results

Our experimental evaluation demonstrates the advantages of our proposed approach. We compare three methods. Starting with the ‘pairwise alignment’ that operates on pairs of nodes in the time-frequency grid. It serves as the baseline, and, while computationally efficient, it suffers from the fundamental limitation that local synchronization in pairs does not guarantee global consistency across the entire network. Addressing this limitation, our ‘synchronization base’ solution operates on triplets of nodes, providing the crucial advantage that each local synchronization within a triplet inherently implies global synchronization (Appendix E. This triplet-based approach not only improves synchronization accuracy but also maintains computational tractability with $O(D^3)$ complexity.

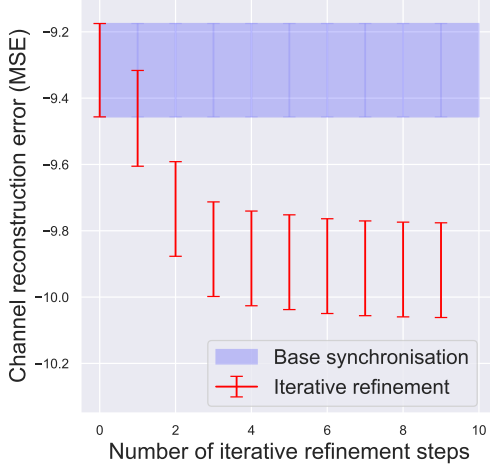


Figure 2: The reconstruction error for varying number of iterative refinement steps. Most improvement over the direct approach (blue shaded region) is achieved in two to five steps.




Figure 3: Numerical results along varying length scales, which determine the correlation properties. Our synchronization improves the channel reconstruction across these settings.

Building upon the triplet foundation, we then demonstrate the benefits of an ‘Iterative Refinement’ algorithm, which operates on the same triplet structure but leverages the denoised channel estimates \mathbf{H}' rather than the noisy observations \mathbf{B} . This final experiment demonstrates how our method enables more effective denoising strategies, resulting in superior overall performance in 5G-like scenarios. Each plot includes both a ‘single channel’ and an ‘ideal line.’ These are theoretical results for either not using synchronization at all or assuming perfect synchronization (explained in Appendix F).

For the experiments, we sample from the ground-truth graphical model and use each method to reconstruct the signal. The problem consists of $6 \times 6 = 36$ grids of PRG blocks, with $D = 3 \times 4 = 12$ measurements made in each block. This means that each block is 3 by 4 measurement cells. The spatial correlation has a length scale of five measurements and follows a squared-exponential decay, which is a simplified model for the Rayleigh-fading scatter [18]. Different values of the length scale indicate varying amounts of scatter in the environment and a comparison for multiple values of the length scale is shown in Figure 3.

Figure 1 shows the main experimental result. It compares all three methods on a 5G MIMO problem. In general, for larger Signal-to-Noise (SNR) ratios, the channel reconstruction achieves lower decibels, and this is reflected in the ‘ideal synchronization’ and ‘single channel ideal’ comparison lines. The right figure redraws each line compared to the ‘single channel’ MSE and indicates the relative improvement. Beyond 0 dB SNR, all three MRA methods exhibit improved performance, with synchronization achieving a lower MSE than pairwise alignment, and Iterative Refinement achieving the lowest MSE across experimental settings.

5 Conclusion

Our work presents a significant advancement in multi-reference alignment (MRA) by introducing a novel structured probabilistic framework that addresses the fundamental challenges of MIMO channel denoising. By reformulating the MRA problem in terms of relative poses within a principled statistical model, we eliminate the global redundancies that hinder optimization. We use cycle consistency and arrive at a computationally efficient method that can scale to realistic system sizes.

The iterative refinement algorithm consistently outperforms direct estimation methods and methods that use only pairwise alignment. Across experimental settings, our approach that is decentralized and refines relative pose and channel estimates achieves better reconstruction error consistently. The success of our approach demonstrates the promise of iterative refinement techniques for MIMO synchronization. It opens up new directions for developing more sophisticated denoising strategies that can leverage the global synchronization achieved through our triplet-based cycle consistency.

References

- [1] 3GPP. Third generation partnership project (3gpp). *Technical Specification (TS)*, 38(10), 2020.
- [2] Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal Machine Learning Research*, 2017.
- [3] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Amelia Perry, Jonathan Weed, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *arXiv preprint arXiv:1712.10163*, 2017.
- [4] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006.
- [5] Wanshi Chen, Peter Gaal, Juan Montojo, and Haris Zisimopoulos. *Fundamentals of 5G communications: connectivity for enhanced mobile broadband and beyond*. McGraw-Hill Education, 2021.
- [6] Yifeng Fan, Tingran Gao, and Zhizhen Jane Zhao. Unsupervised co-learning on g -manifolds across irreducible representations. *Advances in neural information processing systems*, 32, 2019.
- [7] Tingran Gao, Jacek Brodzki, and Sayan Mukherjee. The geometry of synchronization problems and learning group actions. *Discrete & Computational Geometry*, 2021.
- [8] John C Gower and Garnt B Dijksterhuis. *Procrustes problems*. Oxford university press, 2004.
- [9] R. Henderson, J. M. Baldwin, T. A. Ceska, F. Zemlin, E. Beckmann, and K. H. Downing. Model for the structure of bacteriorhodopsin based on high-resolution electron cryo-microscopy. *Journal of Molecular Biology*, June 1990.
- [10] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [11] Eva Nogales. The development of cryo-EM into a mainstream structural biology technique. *Nature methods*, 2016. Publisher: Nature Publishing Group.
- [12] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Message-passing algorithms for synchronization problems over compact groups. *Communications on Pure and Applied Mathematics*, 2018.
- [13] Kumar Pratik, Pouriya Sadeghi, Gabriele Cesa, Sanaz Barghi, Joseph B Soriaga, Yuanning Yu, Supratik Bhattacharjee, and Arash Behboodi. Requestnet: A foundational learning model for channel estimation. *arXiv 2508.08790*, 2025.
- [14] Sjors HW Scheres. Relion: implementation of a bayesian approach to cryo-em structure determination. *Journal of structural biology*, 2012.
- [15] Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 1966.
- [16] Amit Singer. Mathematics for cryo-electron microscopy. In *Proceedings of the International Congress of Mathematicians*. World Scientific, 2018.
- [17] Amit Singer and Fred J Sigworth. Computational methods for single-particle electron cryomicroscopy. *Annual Review of Biomedical Data Science*, 2020.
- [18] Bernard Sklar. Rayleigh fading channels in mobile digital communication systems. i. characterization. *IEEE Communications magazine*, 2002.
- [19] Yifan Sun and Qixing Huang. Pose synchronization under multiple pair-wise relative poses. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

Acknowledgements

This work is financially supported by Qualcomm Technologies Inc., the University of Amsterdam and the allowance Top consortia for Knowledge and Innovation (TKIs) from the Netherlands Ministry of Economic Affairs and Climate Policy. Qualcomm AI research is an initiative of Qualcomm Technologies, Inc. and/or its subsidiaries. Correspondence may go to romijndersrob@gmail.com.

A Notation

An overview of the notation in this paper. Additionally, Figure 4 presents a factor graph to illustrate how the random variables relate to each other. That factor graph is for three PRG blocks.

Notation

D, d	Number of measurements and number of dimensions
H, W	Number of rows and columns in the MIMO grid
$N = H \cdot W$	The number of blocks in the MIMO grid
$P_1, P_2, P_3, \dots, P_J$	Rotation matrix of size $\mathbb{R}^{d \times d}$
$p(P_i)$	Prior distribution of P_i , uniform over the rotation angles, $P_i \in SO(d)$
$R_{ij} = P_i^{-1} P_j$	Relative rotation matrix of size $\mathbb{R}^{d \times d}$
$H_1, H_2, H_3, \dots, H_J$	Hidden matrices of size $\mathbb{R}^{D \times d}$
$H = [H_1; H_2; H_3]$	Stacked hidden matrices of size $\mathbb{R}^{3D \times d}$
$H'_i = H_i P_i$	H-accent is the channel matrix that has undergone an unknown rotation, size $\mathbb{R}^{D \times d}$
$B_1, B_2, B_3, \dots, B_J$	Observed matrices of size $\mathbb{R}^{D \times d}$
$B = [B_1; B_2; B_3]$	Stacked observed matrices of size $\mathbb{R}^{3D \times d}$
$\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$	Observation noise, shape clear from context
$B_i = H'_i + \varepsilon_i$	Observation model with iid noise
$\mathcal{MN}(\mathbf{0}, U, V)$	Matrix-normal distribution with mean 0, row covariance U , column covariance V
$p(H) = \mathcal{MN}(\mathbf{0}, U, I)$	Matrix normal distribution with mean 0, row covariance U

A.1 Relative poses and channel prior

We model the channel with a Matrix-Normal distribution. The Matrix Normal has the following properties: if $X \sim \mathcal{MN}(M, U, V)$ then $\text{vec}(X) = \mathcal{N}(\text{vec}(M), V \otimes U)$. Here, $\text{vec}(X)$ is the Frobenius vectorization of a matrix X and \otimes is the Kronecker product.

Let's denote the prior as $p(\cdot)$, then the following equivalence holds for the $\mathcal{MN}(0, U, I)$ matrix normal:

$$p\left(\begin{bmatrix} H_1 \\ H_2 \\ H_3 \\ \vdots \end{bmatrix}\right) = p\left(\begin{bmatrix} H_1 P_1^{-1} \\ H_2 P_1^{-1} \\ H_3 P_1^{-1} \\ \vdots \end{bmatrix}\right) \quad (8)$$

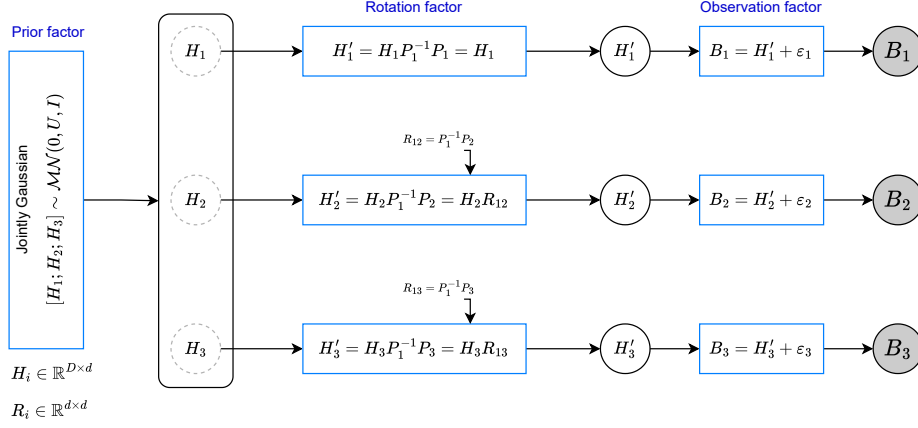


Figure 4: Structured Probabilistic Graphical Model to illustrate the joint distribution in Equation 2

This uses the rotated matrices:

$$\begin{bmatrix} H_1' \\ H_2' \\ H_3' \\ \vdots \end{bmatrix} = \begin{bmatrix} H_1 P_1^{-1} P_1 \\ H_2 P_1^{-1} P_2 \\ H_3 P_1^{-1} P_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} H_1 \\ H_2 R_{12} \\ H_3 R_{13} \\ \vdots \end{bmatrix} \quad (9)$$

B Estimate channel using a rotation estimate

We first derive the estimation of the channel if the rotations were known. We present the case for reconstructing four matrix-valued signals in the presence of relative rotation. The dots, \dots , will indicate generality towards sets of more matrices. The variable J indicates the number of blocks considered in the denoising. For this example, $J = 4$.

The derivation starts from $\log p(\mathbf{B}, \hat{\mathbf{R}}, \mathbf{H})$. This assumes that one has an estimate of the Rotation matrices, \mathbf{R} , and the goal is to find the channel matrices \mathbf{H} with the largest likelihood.

$$\arg \max_{\mathbf{H}} \log p(\mathbf{B}, \hat{\mathbf{R}}, \mathbf{H}) \quad (10)$$

$$\begin{aligned} &= \arg \max_{\mathbf{H}} \log \int_{\mathbf{H}_1} \cdots \int_{\mathbf{H}_4} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) p(\mathbf{H}_1' | \mathbf{H}_1) \prod_{j \in [2, 3, 4]} p(\mathbf{H}_j' | \mathbf{H}_j, \hat{\mathbf{R}}_{1j}) \\ &\quad \cdot p(\mathbf{B}_1 | \mathbf{H}_1') \left(\prod_{j \in [2, 3, 4]} p(\mathbf{B}_j | \mathbf{H}_j') \right) d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 d\mathbf{H}_4 \end{aligned} \quad (11)$$

$$\begin{aligned} &= \arg \max_{\mathbf{H}} \log \int_{\mathbf{H}_1} \cdots \int_{\mathbf{H}_4} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) \delta_{\mathbf{H}_1 = \mathbf{H}_1'} \prod_{j \in [2, 3, 4]} \delta_{\mathbf{H}_j = \mathbf{H}_j' \hat{\mathbf{R}}_{1j}} \\ &\quad \cdot \left(\prod_{j \in [1, \dots, 4]} p(\mathbf{B}_j | \mathbf{H}_j') \right) d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 d\mathbf{H}_4 \end{aligned} \quad (12)$$

We integrate over the delta distributions and introduce a change of variables $H_2' = H_2 \hat{R}_{12}$ and $H_3' = H_3 \hat{R}_{13}$ and $H_4' = H_4 \hat{R}_{14}$.

$$\arg \max_{\mathbf{H}} \log p(\mathbf{H}_1', \mathbf{H}_2' \hat{\mathbf{R}}_{21}, \mathbf{H}_3' \hat{\mathbf{R}}_{31}, \mathbf{H}_4' \hat{\mathbf{R}}_{41}) \prod_{j \in [1, \dots, 4]} p(\mathbf{B}_j | \mathbf{H}_j') \quad (13)$$

$$= \arg \max_{\mathbf{H}} \log p(\mathbf{H}_1', \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) p(\mathbf{B}_1 | \mathbf{H}_1') \prod_{j \in [2, 3, 4]} p(\mathbf{B}_j | \mathbf{H}_j \hat{\mathbf{R}}_{1j}) \quad (14)$$

Next, we write down the actual likelihoods that are Matrix Normal distribution by \mathcal{MN} and align the

$$\text{writing of said distributions } \begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \mathbf{H}_3 \\ \mathbf{H}_4 \end{bmatrix} \sim \mathcal{MN}(0, \mathbf{U}, \mathbf{I}).$$

The operation $\text{vec}(\cdot)$ indicates the ‘‘Frobenius vectorization’’ of matrix \mathbf{X} , which concatenates all columns from left to right in one long column vector. This means that we can conveniently write down the Matrix Normal distribution. If \mathbf{X} has a matrix normal distribution $\mathbf{X} \sim \mathcal{MN}(0, \mathbf{U}, \mathbf{I})$, then $\text{vec}(\mathbf{X})$ has Normal distribution $\text{vec}(\mathbf{X}) \sim \mathcal{N}(0, \mathbf{I} \otimes \mathbf{U})$ where \otimes indicates the Kronecker product.

$$\begin{aligned} \arg \max_{\mathbf{H}} \log \mathcal{N}(\text{vec}(\mathbf{H}'_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) | 0, \mathbf{I}_d \otimes \mathbf{U}) + \log \mathcal{N}(\mathbf{H}'_1 | \mathbf{B}_1, \sigma_\epsilon^2 \mathbf{I}_{dD}) \\ + \sum_{j \in [2,3,4]} \log \mathcal{N}(\mathbf{H}_j | \mathbf{B}_j \hat{\mathbf{R}}_{j1}, \sigma_\epsilon^2 \mathbf{I}_{dD}) \end{aligned} \quad (15)$$

$$\begin{aligned} = \arg \max_{\mathbf{H}} \log \mathcal{N}(\text{vec}(\mathbf{H}'_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) | 0, \mathbf{I}_d \otimes \mathbf{U}) \\ + \log \mathcal{N}(\text{vec}(\mathbf{H}'_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) | \text{vec} \left(\begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \hat{\mathbf{R}}_{21} \\ \mathbf{B}_3 \hat{\mathbf{R}}_{31} \\ \mathbf{B}_4 \hat{\mathbf{R}}_{41} \end{bmatrix} \right), \mathbf{I}_d \otimes \sigma_\epsilon^2 \mathbf{I}_{4D}) \end{aligned} \quad (16)$$

$$= \arg \max_{\mathbf{H}} \log \mathcal{N}(\text{vec}(\mathbf{H}'_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4) | \boldsymbol{\Sigma}_H \text{vec}(\sigma_\epsilon^{-2} \mathbf{B}_1, \sigma_\epsilon^{-2} \mathbf{B}_2 \hat{\mathbf{R}}_{21}, \sigma_\epsilon^{-2} \mathbf{B}_3 \hat{\mathbf{R}}_{31}, \sigma_\epsilon^{-2} \mathbf{B}_4 \hat{\mathbf{R}}_{41}), \boldsymbol{\Sigma}_H) \quad (17)$$

We use shorthand for the denoising correlation matrix:

$$\begin{aligned} \boldsymbol{\Sigma}_H &= ((\mathbf{I}_d \otimes \mathbf{U})^{-1} + (\mathbf{I}_d \otimes \sigma_\epsilon^2 \mathbf{I}_{JD})^{-1})^{-1} \\ &= ((\mathbf{I}_d \otimes \mathbf{U})^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{JDd})^{-1} \\ &= \left(\begin{bmatrix} (\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{JD})^{-1} & 0 & 0 & \dots \\ 0 & (\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{JD})^{-1} & 0 & \dots \\ 0 & 0 & (\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{JD})^{-1} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \right)^{-1} \\ &= \mathbf{I}_d \otimes (\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{4D})^{-1}. \end{aligned} \quad (18)$$

The argmax of a Gaussian likelihood is the mean parameter of the said Gaussian distribution. Therefore, we find the mean parameter of \mathbf{H}' from:

$$\begin{bmatrix} \mathbf{H}'_{1,t+1} \\ \mathbf{H}'_{2,t+1} \hat{\mathbf{R}}_{21} \\ \mathbf{H}'_{3,t+1} \hat{\mathbf{R}}_{31} \\ \mathbf{H}'_{4,t+1} \hat{\mathbf{R}}_{41} \end{bmatrix} = (\mathbf{I}_d \otimes ((\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{4D})^{-1})) \begin{bmatrix} \sigma_\epsilon^{-2} \mathbf{B}_1 \\ \sigma_\epsilon^{-2} \mathbf{B}_2 \hat{\mathbf{R}}_{21} \\ \sigma_\epsilon^{-2} \mathbf{B}_3 \hat{\mathbf{R}}_{31} \\ \sigma_\epsilon^{-2} \mathbf{B}_4 \hat{\mathbf{R}}_{41} \end{bmatrix}. \quad (19)$$

The general equation for general J would be:

$$\begin{bmatrix} \mathbf{H}'_{1,t+1} \\ \mathbf{H}'_{2,t+1} \hat{\mathbf{R}}_{21} \\ \mathbf{H}'_{3,t+1} \hat{\mathbf{R}}_{31} \\ \mathbf{H}'_{4,t+1} \hat{\mathbf{R}}_{41} \\ \vdots \end{bmatrix} = (\mathbf{I}_d \otimes ((\mathbf{U}^{-1} + \sigma_\epsilon^{-2} \mathbf{I}_{JD})^{-1})) \begin{bmatrix} \sigma_\epsilon^{-2} \mathbf{B}_1 \\ \sigma_\epsilon^{-2} \mathbf{B}_2 \hat{\mathbf{R}}_{21} \\ \sigma_\epsilon^{-2} \mathbf{B}_3 \hat{\mathbf{R}}_{31} \\ \sigma_\epsilon^{-2} \mathbf{B}_4 \hat{\mathbf{R}}_{41} \\ \vdots \end{bmatrix}. \quad (20)$$

C Directly estimate rotations

For the direct approach, we estimate the rotation matrices within each subgraph. This step occurs even before iterative refinement and provides the first estimate of the rotation matrices based on the noisy observations.

C.1 MLE, three obs, directly from B

The comments prefixed with ‘...’ in blue are comments to explain the derivation steps.

$$\begin{aligned} & \hat{\mathbf{R}}_{12}, \hat{\mathbf{R}}_{13} \\ &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log p(\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3 | \mathbf{R}_{12}, \mathbf{R}_{13}) \end{aligned} \quad (21)$$

... Introduce and marginalize over \mathbf{H}_1 , \mathbf{H}_2 , and \mathbf{H}_3

$$\begin{aligned} &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}_1} \int_{\mathbf{H}_2} \int_{\mathbf{H}_3} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3) \int_{\mathbf{H}'_1} \int_{\mathbf{H}'_2} \int_{\mathbf{H}'_3} p(\mathbf{B}_1 | \mathbf{H}'_1) p(\mathbf{B}_2 | \mathbf{H}'_2) p(\mathbf{B}_3 | \mathbf{H}'_3) \\ & \quad p(\mathbf{H}'_1 | \mathbf{H}_1) p(\mathbf{H}'_2 | \mathbf{H}_2, \mathbf{R}_{12}) p(\mathbf{H}'_3 | \mathbf{H}_3, \mathbf{R}_{13}) d\mathbf{H}'_1 d\mathbf{H}'_2 d\mathbf{H}'_3 d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 \end{aligned} \quad (22)$$

$$\begin{aligned} &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}_1} \int_{\mathbf{H}_2} \int_{\mathbf{H}_3} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3) \int_{\mathbf{H}'_1} \int_{\mathbf{H}'_2} \int_{\mathbf{H}'_3} p(\mathbf{B}_1 | \mathbf{H}'_1) p(\mathbf{B}_2 | \mathbf{H}'_2) p(\mathbf{B}_3 | \mathbf{H}'_3) \\ & \quad \delta_{\mathbf{H}_1 = \mathbf{H}'_1} \delta_{\mathbf{H}_2 = \mathbf{H}'_2 \mathbf{R}_{21}} \delta_{\mathbf{H}_3 = \mathbf{H}'_3 \mathbf{R}_{31}} d\mathbf{H}'_1 d\mathbf{H}'_2 d\mathbf{H}'_3 d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 \end{aligned} \quad (23)$$

... Integrate against delta distributions

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}'_1} \int_{\mathbf{H}'_2} \int_{\mathbf{H}'_3} p(\mathbf{H}'_1, \mathbf{H}'_2 \mathbf{R}_{21}, \mathbf{H}'_3 \mathbf{R}_{31}) p(\mathbf{B}_1 | \mathbf{H}'_1) p(\mathbf{B}_2 | \mathbf{H}'_2) p(\mathbf{B}_3 | \mathbf{H}'_3) d\mathbf{H}'_1 d\mathbf{H}'_2 d\mathbf{H}'_3 \quad (24)$$

... Make Matrix Normal explicit

$$\begin{aligned} &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}'_1} \int_{\mathbf{H}'_2} \int_{\mathbf{H}'_3} \mathcal{MN} \left(\begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix}; 0, \mathbf{U}, \mathbf{I}_d \right) \mathcal{MN}(\mathbf{H}'_1 | \mathbf{B}_1, \sigma^2 \mathbf{I}_D, \mathbf{I}_d) \\ & \quad \mathcal{MN}(\mathbf{H}'_2 \mathbf{R}_{21}; \mathbf{B}_2 \mathbf{R}_{21}, \sigma^2 \mathbf{I}_D, \mathbf{I}_d) \mathcal{MN}(\mathbf{H}'_3 \mathbf{R}_{31}; \mathbf{B}_3 \mathbf{R}_{31}, \sigma^2 \mathbf{I}_D, \mathbf{I}_d) d\mathbf{H}'_1 d\mathbf{H}'_2 d\mathbf{H}'_3 \end{aligned} \quad (25)$$

... Fuse Matrix normals

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}'_1} \int_{\mathbf{H}'_2} \int_{\mathbf{H}'_3} \mathcal{MN} \left(\begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix}; 0, \mathbf{U}, \mathbf{I}_d \right) \mathcal{MN} \left(\begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix} \middle| \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix}, \sigma^2 \mathbf{I}_{3D}, \mathbf{I}_d \right) d\mathbf{H}'_1 d\mathbf{H}'_2 d\mathbf{H}'_3 \quad (26)$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \mathcal{MN} \left(\begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix}; 0, \mathbf{U} + \sigma^2 \mathbf{I}, \mathbf{I}_d \right) \quad (27)$$

The simplification from two to one Matrix Normal uses some tricks: the multiplication of two normal distributions yields a normal distribution, and the integral will not change under rotation. Therefore, we can use standard properties of the Matrix Normal product. The argmax over the matrix normal can be obtained by writing out the Trace and removing scalar constants:

$$\begin{aligned} & \hat{\mathbf{R}}_{12}, \hat{\mathbf{R}}_{13} \\ &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \frac{-1}{2} \text{Tr} \left[\begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix}^T (\mathbf{U} + \sigma^2 \mathbf{I})^{-1} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix} \right] \\ &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \frac{-1}{2} \text{Tr} \left[\begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \mathbf{R}_{21} \\ \mathbf{B}_3 \mathbf{R}_{31} \end{bmatrix}^T (\mathbf{U} + \sigma^2 \mathbf{I})^{-1} \right] \end{aligned} \quad (28)$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} - \text{Tr} \left[\begin{bmatrix} \mathbf{B}_1 \mathbf{B}_1^T & \mathbf{B}_1 \mathbf{R}_{12} \mathbf{B}_2^T & \mathbf{B}_1 \mathbf{R}_{13} \mathbf{B}_3^T \\ \mathbf{B}_2 \mathbf{R}_{21} \mathbf{B}_1^T & \mathbf{B}_2 \mathbf{B}_2^T & \mathbf{B}_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{B}_3^T \\ \mathbf{B}_3 \mathbf{R}_{31} \mathbf{B}_1^T & \mathbf{B}_3 \mathbf{R}_{31} \mathbf{R}_{12} \mathbf{B}_2^T & \mathbf{B}_3 \mathbf{B}_3^T \end{bmatrix} (\mathbf{U} + \sigma^2 \mathbf{I})^{-1} \right] \quad (29)$$

Equation 29 considers the matrix as composed of submatrices. Since the \mathbf{U} matrix is always composed of $J \cdot D$ by $J \cdot D$ elements, we can always split it in J by J submatrices of size $D \times D$:

$$-(U + \sigma^2 I)^{-1} = \begin{bmatrix} U_1 & U_a & U_b \\ U_1^T & U_2 & U_c \\ U_b^T & U_c^T & U_3 \end{bmatrix} \quad (30)$$

$$\hat{\mathbf{R}}_{12}, \hat{\mathbf{R}}_{13} = \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \text{Tr} \left[\begin{bmatrix} \mathbf{B}_1 \mathbf{B}_1^T & \mathbf{B}_1 \mathbf{R}_{12} \mathbf{B}_2^T & \mathbf{B}_1 \mathbf{R}_{13} \mathbf{B}_3^T \\ \mathbf{B}_2 \mathbf{R}_{21} \mathbf{B}_1^T & \mathbf{B}_2 \mathbf{B}_2^T & \mathbf{B}_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{B}_3^T \\ \mathbf{B}_3 \mathbf{R}_{31} \mathbf{B}_1^T & \mathbf{B}_3 \mathbf{R}_{31} \mathbf{R}_{12} \mathbf{B}_2^T & \mathbf{B}_3 \mathbf{B}_3^T \end{bmatrix} \begin{bmatrix} U_1 & U_a & U_b \\ U_1^T & U_2 & U_c \\ U_b^T & U_c^T & U_3 \end{bmatrix} \right] \quad (31)$$

...Split to individual traces

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} 2 \text{Tr} \left[\mathbf{B}_1 \mathbf{R}_{12} \mathbf{B}_2^T U_a^T \right] + 2 \text{Tr} \left[\mathbf{B}_1 \mathbf{R}_{13} \mathbf{B}_3^T U_b^T \right] + 2 \text{Tr} \left[\mathbf{B}_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{B}_3^T U_c^T \right] \quad (32)$$

The argument in Equation 32 is linear in both rotation matrices. In each argument, there is a closed-form expression for the optimum over the set of orthogonal matrices. Therefore, we suggest to iteratively determine this optimum. With each closed form solution, each step is guaranteed to non-decrease the objective value and we find that this usually converges to machine precision in less than eight steps.

$$\text{iterate} \begin{cases} \mathbf{R}_{13} = \arg \max \langle \mathbf{R}_{21}, \mathbf{B}_2^T U_a^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T U_b^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T U_c^T \mathbf{B}_2 \mathbf{R}_{21} \rangle_F \\ \mathbf{R}_{12} = \arg \max \langle \mathbf{R}_{21}, \mathbf{B}_2^T U_a^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{B}_3^T U_b^T \mathbf{B}_1 \rangle_F + \langle \mathbf{R}_{21}, \mathbf{B}_2^T U_c \mathbf{B}_3 \mathbf{R}_{31} \rangle_F \end{cases} \quad (33)$$

The closed-form solution is known as the Orthogonal Procrustes problem [8, 15]. The projection is achieved by removing the singular values from the Singular-Value Decomposition:

$$\arg \max_{\mathbf{R}} \langle \mathbf{R}, \mathbf{X} \rangle_F = \mathbf{A} \mathbf{L} \mathbf{V}^T \quad (34)$$

from Singular Value Decomposition $\mathbf{X} = \mathbf{A} \mathbf{\Sigma} \mathbf{V}^T$ and with $\mathbf{L} = \mathbf{I}$ if $\det(\mathbf{X}) > 0$ or $\mathbf{L} = \text{diag}(1, 1, 1, \dots, -1)$ otherwise (assuming the smallest singular-value of \mathbf{X} is at the bottom of $\mathbf{\Sigma}$).

D Iterative refinement 3

Deriving the iterative refinement looks much like Equation 21. However, it crucially starts from optimizing the rotation matrices given the channel estimates \mathbf{H} , instead of the noisy estimates \mathbf{B} .

The comments prefixed with ‘...’ are comments to explain the derivation steps.

$$\begin{aligned} & \hat{\mathbf{R}}_{12}, \hat{\mathbf{R}}_{13} \\ &= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log p(\mathbf{R}_{12}, \mathbf{R}_{13} | \mathbf{H}') \end{aligned} \quad (35)$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log p(\mathbf{H}'_1, \mathbf{H}'_2, \mathbf{H}'_3 | \mathbf{R}_{12}, \mathbf{R}_{13}) \quad (36)$$

... Introduce and marginalize over \mathbf{H}_1 , \mathbf{H}_2 , and \mathbf{H}_3

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}_1} \int_{\mathbf{H}_2} \int_{\mathbf{H}_3} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3) p(\mathbf{H}'_1 | \mathbf{H}_1) p(\mathbf{H}'_2 | \mathbf{H}_2, \mathbf{R}_{12}) p(\mathbf{H}'_3 | \mathbf{H}_3, \mathbf{R}_{13}) d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 \quad (37)$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log \int_{\mathbf{H}_1} \int_{\mathbf{H}_2} \int_{\mathbf{H}_3} p(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3) \delta_{\mathbf{H}_1 = \mathbf{H}'_1} \delta_{\mathbf{H}_2 = \mathbf{H}'_2 \mathbf{R}_{21}} \delta_{\mathbf{H}_3 = \mathbf{H}'_3 \mathbf{R}_{31}} d\mathbf{H}_1 d\mathbf{H}_2 d\mathbf{H}_3 \quad (38)$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \log p(\mathbf{H}'_1, \mathbf{H}'_2 \mathbf{R}_{21}, \mathbf{H}'_3 \mathbf{R}_{31}) \quad (39)$$

We'll write down the matrix-normal form of the channel prior: $\begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \mathbf{H}_3 \end{bmatrix} \sim \mathcal{MN}(0, \mathbf{U}, \mathbf{I})$:

$$\begin{aligned}
& \hat{\mathbf{R}}_{12}, \hat{\mathbf{R}}_{13} \\
&= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \frac{-1}{2} \text{Tr} \left[\begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix}^T \mathbf{U}^{-1} \begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix} \right] \\
&= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \frac{-1}{2} \text{Tr} \left[\begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix} \begin{bmatrix} \mathbf{H}'_1 \\ \mathbf{H}'_2 \mathbf{R}_{21} \\ \mathbf{H}'_3 \mathbf{R}_{31} \end{bmatrix}^T \mathbf{U}^{-1} \right] \tag{40}
\end{aligned}$$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} -\text{Tr} \left[\begin{bmatrix} \mathbf{H}'_1 \mathbf{H}'_1{}^T & \mathbf{H}'_1 \mathbf{R}_{12} \mathbf{H}'_2{}^T & \mathbf{H}'_1 \mathbf{R}_{13} \mathbf{H}'_3{}^T \\ \mathbf{H}'_2 \mathbf{R}_{21} \mathbf{H}'_1{}^T & \mathbf{H}'_2 \mathbf{H}'_2{}^T & \mathbf{H}'_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{H}'_3{}^T \\ \mathbf{H}'_3 \mathbf{R}_{31} \mathbf{H}'_1{}^T & \mathbf{H}'_3 \mathbf{R}_{31} \mathbf{R}_{12} \mathbf{H}'_2{}^T & \mathbf{H}'_3 \mathbf{H}'_3{}^T \end{bmatrix} \mathbf{U}^{-1} \right] \tag{41}$$

...Consider $-\mathbf{U}^{-1}$ a block matrix $\begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_a & \mathbf{U}_b \\ \mathbf{U}_a^T & \mathbf{U}_2 & \mathbf{U}_c \\ \mathbf{U}_b^T & \mathbf{U}_c^T & \mathbf{U}_3 \end{bmatrix}$ with $\mathbf{U} \in \mathbb{R}^{3D \times 3D}$ and $\mathbf{U}_a \in \mathbb{R}^{D \times D}$

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} \text{Tr} \left[\begin{bmatrix} \mathbf{H}'_1 \mathbf{H}'_1{}^T & \mathbf{H}'_1 \mathbf{R}_{12} \mathbf{H}'_2{}^T & \mathbf{H}'_1 \mathbf{R}_{13} \mathbf{H}'_3{}^T \\ \mathbf{H}'_2 \mathbf{R}_{21} \mathbf{H}'_1{}^T & \mathbf{H}'_2 \mathbf{H}'_2{}^T & \mathbf{H}'_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{H}'_3{}^T \\ \mathbf{H}'_3 \mathbf{R}_{31} \mathbf{H}'_1{}^T & \mathbf{H}'_3 \mathbf{R}_{31} \mathbf{R}_{12} \mathbf{H}'_2{}^T & \mathbf{H}'_3 \mathbf{H}'_3{}^T \end{bmatrix} \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_a & \mathbf{U}_b \\ \mathbf{U}_a^T & \mathbf{U}_2 & \mathbf{U}_c \\ \mathbf{U}_b^T & \mathbf{U}_c^T & \mathbf{U}_3 \end{bmatrix} \right] \tag{42}$$

...Split to individual traces

$$= \arg \max_{\mathbf{R}_{12}, \mathbf{R}_{13}} 2 \text{Tr} \left[\mathbf{H}'_1 \mathbf{R}_{12} \mathbf{H}'_2{}^T \mathbf{U}_a^T \right] + 2 \text{Tr} \left[\mathbf{H}'_1 \mathbf{R}_{13} \mathbf{H}'_3{}^T \mathbf{U}_b^T \right] + 2 \text{Tr} \left[\mathbf{H}'_2 \mathbf{R}_{21} \mathbf{R}_{13} \mathbf{H}'_3{}^T \mathbf{U}_c^T \right] \tag{43}$$

The optimization argument in Equation 43 is linear in both rotation matrices. In each argument, there is a closed-form expression for the optimum over the set of orthogonal matrices. Therefore, we suggest to iteratively determine this optimum. With each closed form solution, each step is guaranteed to non-decrease the objective value and we find that this usually converges to machine precision in less than eight steps.

iterate until convergence

$$\begin{cases} \mathbf{R}_{12} = \arg \max_{\mathbf{R}_{12}} \langle \mathbf{R}_{21}, \mathbf{H}'_2{}^T \mathbf{U}_a^T \mathbf{H}'_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_3{}^T \mathbf{U}_b^T \mathbf{H}'_1 \rangle_F + \langle \mathbf{R}_{21}, \mathbf{H}'_2{}^T \mathbf{U}_c \mathbf{H}'_3 \mathbf{R}_{31} \rangle_F \\ \mathbf{R}_{13} = \arg \max_{\mathbf{R}_{13}} \langle \mathbf{R}_{21}, \mathbf{H}'_2{}^T \mathbf{U}_a^T \mathbf{H}'_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_3{}^T \mathbf{U}_b^T \mathbf{H}'_1 \rangle_F + \langle \mathbf{R}_{31}, \mathbf{H}'_3{}^T \mathbf{U}_c^T \mathbf{H}'_2 \mathbf{R}_{21} \rangle_F \end{cases}$$

The closed-form solution to the latter argmax is the same as the Procrustes problem in Section C. Each step is a closed-form maximizer, so the iteration is guaranteed to converge.

E Local to Global synchronization

This section includes the theoretical proofs behind our main arguments for reducing the absolute pose estimation problem to a relative pose synchronization problem. In particular, we will show

1. that the set of absolute poses and the set of cycle-consistent relative poses are in one-to-one correspondence (up to a global transformation of the absolute poses) and that
2. if all 3-cycles (triplets) in a surface mesh (e.g., the physical resource blocks in the frequency-time lattice in the 5G NR channel) are cycle-consistent, then the whole mesh is cycle consistent.

First, we define the notion of cycle consistency precisely. Consider a **connected graph** $G = (V, E)$ with a set V of N nodes and a set E of edges. Then, an assignment of relative poses $R = \{R_{ij}\}_{(i,j) \in E}$ is cycle consistent along a k -cycle $\gamma = [\gamma_1, \gamma_2, \dots, \gamma_k, \gamma_1]$ (with $\gamma_i \in V$ and $(\gamma_i, \gamma_{i+1}) \in E$) if

$$R_{\gamma_1, \gamma_2} R_{\gamma_2, \gamma_3} \cdots R_{\gamma_k, \gamma_1} = I \tag{44}$$

An assignment of relative poses $\{R_{ij}\}_{(i,j) \in E}$ is cycle consistent if it is cycle-consistent over all cycles in the graph G .

E.1 Absolute to relative poses

We need to prove two directions. First, any set of absolute poses $P = \{P_i\}_{i=1}^N$ is associated with one unique set of cycle-consistent relative poses $R = \{R_{ij}\}_{1 \leq i, j \leq N}$. This true by definition, since $R_{ij} := P_i^{-1}P_j$ and for any cycle $\gamma = [\gamma_1, \gamma_2, \dots, \gamma_k, \gamma_1]$:

$$R_{\gamma_1, \gamma_2} R_{\gamma_2, \gamma_3} \cdots R_{\gamma_k, \gamma_1} = P_{\gamma_1}^{-1} P_{\gamma_2} P_{\gamma_2}^{-1} P_{\gamma_3} \cdots P_{\gamma_k} P_{\gamma_1} = P_{\gamma_1}^{-1} P_{\gamma_1} = I \quad (45)$$

Second, we show that any set of cycle-consistent relative poses corresponds to an *equivalence class* of absolute poses (defined up to a global transformation). First, we note that cycle-consistency implies that, for any pair of nodes i, j , walking along any path connecting them generates the same transformation, i.e., the relative pose R_{ij} is well defined (even if edge $(i, j) \notin E$). For any pair of paths α, β of length a and b connecting i to j , i.e. $\alpha_1 = \beta_1 = i$ and $\alpha_a = \beta_b = j$, concatenating α and β reversed form a cycle $\gamma = [\alpha_1 = i, \alpha_2, \alpha_3, \dots, \alpha_a = j = \beta_b, \dots, \beta_3, \beta_2, \beta_1 = i]$. Since the relative poses R are cycle-consistent, it follows that

$$R_{\alpha_1=i, \alpha_2} R_{\alpha_2, \alpha_3} \cdots R_{\alpha_{a-1}, \alpha_a=j} R_{\beta_b=j, \beta_{b-1}} \cdots R_{\beta_2, \beta_1=i} = I \quad (46)$$

$$R_{\alpha_1=i, \alpha_2} R_{\alpha_2, \alpha_3} \cdots R_{\alpha_{a-1}, \alpha_a=j} = R_{\beta_2, \beta_1=i}^{-1} \cdots R_{\beta_b=j, \beta_{b-1}}^{-1} \quad (47)$$

$$R_{\alpha_1=i, \alpha_2} R_{\alpha_2, \alpha_3} \cdots R_{\alpha_{a-1}, \alpha_a=j} = R_{\beta_1=i, \beta_2} \cdots R_{\beta_{b-1}, \beta_b=j} =: R_{i,j} \quad (48)$$

Hence, one can construct a set of absolute poses P by choosing an arbitrary pose for the first node P_0 and then defining $P_i := P_0 R_{0,i}$. Note that any choice of P_0 is valid, but leads to a different set of absolute poses. In particular, any two such equivalent sets of absolute poses P and P' differ precisely by a global transformation, i.e., $\exists T : \forall i, P'_i = T P_i$. Moreover, picking a different starting node j rather than 0 results in the equivalent solution one would have obtained by instead picking $P_0 = T R_{j,0}$.

Finally, any of these equivalent solutions generates the initial set of relative poses; hence, we have proved the bi-directional correspondence.

E.2 Sufficiency of 3-cycles

Consider a graph $G = (V, E)$ representing the edges of a triangulated **planar** mesh, i.e. each node in V can be embedded in \mathbb{R}^2 with no edges crossing. Given a certain assignment of relative poses R on G , if all 3-cycles in G (i.e. all triangles in the mesh) are cycle-consistent, then any cycle γ in G is also cycle-consistent, i.e.:

$$R_\gamma = R_{\gamma_1, \gamma_2} R_{\gamma_2, \gamma_3} \cdots R_{\gamma_k, \gamma_1} = I \quad (49)$$

To prove this, note that a cycle γ is the boundary of an enclosed area $\mathcal{U}_\gamma \subset \mathbb{R}^2$ which includes a subset $V_\gamma \subset V$ of nodes and a subset $E_\gamma \subset E$ of edges (including those in γ). We prove that $R_\gamma = I$ by *induction* on the size of E_γ .

Base case $|E_\gamma| = 3$: this implies the cycle γ contains a single triangle and, therefore, it's already cycle-consistent by assumption, i.e. $R_\gamma = I$.

Inductive case $|E_\gamma| > 3$: consider any edge $(i, j) \in \gamma$. There exists a node $k \in V_\gamma$ such that (i, k) and $(j, k) \in E_\gamma$. Then, (i, j, k) for a cycle-consistent triangle. Now we consider three cases.

- If $(j, k) \in \gamma$, we can define a new cycle γ' by replacing both edges (i, j) and (j, k) with the edge (i, k) . Note that $R_\gamma = R_{\gamma'}$ since $R_{ij} R_{jk} = R_{ik}$, but $E_{\gamma'} = E_\gamma \cup \{(i, j), (j, k)\}$ and, therefore, $E_{\gamma'} = E_\gamma - 2$.
- If $(i, k) \in \gamma$, one can apply a similar construction.
- Otherwise, we can define a new cycle γ' by replacing the edge (i, j) with the two edge (i, k) and (k, j) . Note that $R_\gamma = R_{\gamma'}$ since $R_{ij} = R_{ik} R_{jk}^{-1}$, but $E_{\gamma'} = E_\gamma \cup \{(i, j)\}$ and, therefore, $E_{\gamma'} = E_\gamma - 1$.

Thus, the relative poses R are cycle-consistent on γ .

More generally, the set of all sufficient cycles for a generic graph is the *cycle basis* for that graph. When considering meshes over more general surfaces, the result proved above is not necessarily true, i.e. even if all triangles are cycle consistent, there might exist cycles which are not. These cycles can be precisely identified in the first homology group of the surface. For instance, a 2-torus has two additional cycles, but a similar proof works for a spherical surface.

F Performance under ideal synchronization

For many of the results, we are seeking a comparison for the empirical results. Without the precoding, the MIMO task would amount to reconstructing a signal under structured Gaussian noise. As such, we can characterize the Mean-Squared Error under ‘ideal synchronization’ in closed form. This Appendix section aims to characterize the MSE value so that we can plot it in the diagrams.

The ‘ideal synchronization’ MSE is instructive to have, but it is not a strict lower bound, however. Due to finite sample sizes, some reconstructions could achieve lower MSE. Conversely, due to finite sample sizes, the performance might not be attainable. Therefore, it shouldn’t be considered a ‘target performance’ either. The latter is due to the interaction between rotation and signal noise: two signals could be noised such that the optimal rotation on a finite sample under some probabilistic model would not match the true rotation.

In general, the case where rotations are known could be modeled as the reconstruction of a correlated Gaussian distribution. Consider the following system. Matrix \mathbf{X} is drawn from a correlated Gaussian distribution. After additive i.i.d. Gaussian noise on each of the elements in the matrix, a noisy matrix \mathbf{Y} is observed. Then, we aim to characterize the reconstruction error $\mathbb{E}[\|\mathbf{X} - \hat{\mathbf{X}}(\mathbf{Y})\|^2]$.

$$\mathbf{X} \sim \mathcal{N}(0, \Sigma_X); \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}); \quad \mathbf{Y} = \mathbf{X} + \sigma\epsilon \quad (50)$$

We are interested in the expected MSE of estimating \mathbf{X} from \mathbf{Y} , as a function of Σ_X and σ .

$$\mathcal{L}(\Sigma_X, \sigma) = \mathbb{E}[\|\mathbf{X} - \hat{\mathbf{X}}(\mathbf{Y})\|^2] = \mathbb{E}[\mathbf{X}^T \mathbf{X} + \hat{\mathbf{X}}(\mathbf{Y})^T \hat{\mathbf{X}}(\mathbf{Y}) - 2\mathbf{X}^T \hat{\mathbf{X}}(\mathbf{Y})] \quad (51)$$

The covariance matrix Σ_X is of size ND by ND elements, which is the covariance between all measurements in all blocks. The optimal estimator is the MMSE estimator $\hat{\mathbf{X}}(\mathbf{Y}) = \mathbb{E}[\mathbf{X}|\mathbf{Y}]$. This follows from the conditional distribution: $p(\mathbf{X}|\mathbf{Y}) \propto p(\mathbf{Y}|\mathbf{X})p(\mathbf{X}) \propto \mathcal{N}(\mu_c, \Sigma_c)$.

This follows from estimating \mathbf{X} with the MMSE estimator:

$$\Sigma_c = \left(\Sigma_X^{-1} + \frac{1}{\sigma^2} \mathbf{I} \right)^{-1} = \sigma^2 (\sigma^2 \Sigma_X^{-1} + \mathbf{I})^{-1} \quad (52)$$

$$\mu_c = \Sigma_c \left(\frac{1}{\sigma^2} \mathbf{Y} \right) \quad (53)$$

$$\hat{\mathbf{X}}(\mathbf{Y}) = \left(\Sigma_X^{-1} + \frac{1}{\sigma^2} \mathbf{I} \right)^{-1} \left(\frac{1}{\sigma^2} \mathbf{Y} \right) = (\sigma^2 \Sigma_X^{-1} + \mathbf{I})^{-1} \mathbf{Y} \quad (54)$$

For the derivation, we will need the second moment matrices: $\mathbb{E}[\mathbf{X}\mathbf{X}^T] = \Sigma_X$, $\mathbb{E}[\mathbf{Y}\mathbf{Y}^T] = \Sigma_X + \sigma^2 \mathbf{I}$, $\mathbb{E}[\mathbf{Y}\mathbf{X}^T] = \Sigma_X$.

Linear MMSE The covariance of the linear MMSE estimator is given by $\Sigma_X - \Sigma_{XY} \Sigma_Y^{-1} \Sigma_{YX}$. This follows from searching over the space of linear estimators: $\hat{\mathbf{X}}(\mathbf{Y}) = \mathbf{W}\mathbf{Y} + b$. The minimal error is achieved when $\mathbf{W} = \Sigma_{XY} \Sigma_Y^{-1}$. Therefore, the MMSE is

$$\hat{\mathcal{L}}(\Sigma_X, \sigma) = \Sigma_X - \Sigma_X (\Sigma_X + \sigma^2 \mathbf{I})^{-1} \Sigma_X \quad (55)$$

Equation 55 can also be used to characterize the performance of the per-block MSE. That is, the performance of considering any rotations at all and treating each PRG block in isolation. In that case, $\Sigma_X \in \mathbb{R}^{D \times D}$ is the correlated matrix of all measurements within one cell, instead of $\Sigma_X \in \mathbb{R}^{JD \times JD}$.