A Probabilistic Approach to Pose Synchronization for Multi-Reference Alignment

Rob Romijnders[1], Gabriele Cesa[2], Christos Louizos[2], Kumar Pratik[2], Arash Behboodi[2]

[1]University of Amsterdam, QUvA lab; [2]Qualcomm AI research

Summary

We denoise unsynchronized MIMO signals by marginalizing over relative poses (random rotations). In contrast to previous work that estimates the relative poses from noisy observations, we denoise MIMO signals by 'marginalizing' over relative, unknown poses.

❶ **System:** The first MIMO denoising that explicitly incorporates relative rotations, addressing a limitation of existing methods.

❷ **Inference:** A direct and an iterative refinement algorithm for joint denoising and synchronization, grounded in a principled model.

❷ **Efficiency:** Inference operates on local subgraphs (triplets) and can, therefore, scale efficiently to large MIMO problems.

❸ **Results:** Theoretical insights and empirical evidence of the advantages of our approach in 5G-like MIMO scenarios.

Problem setting, model and method

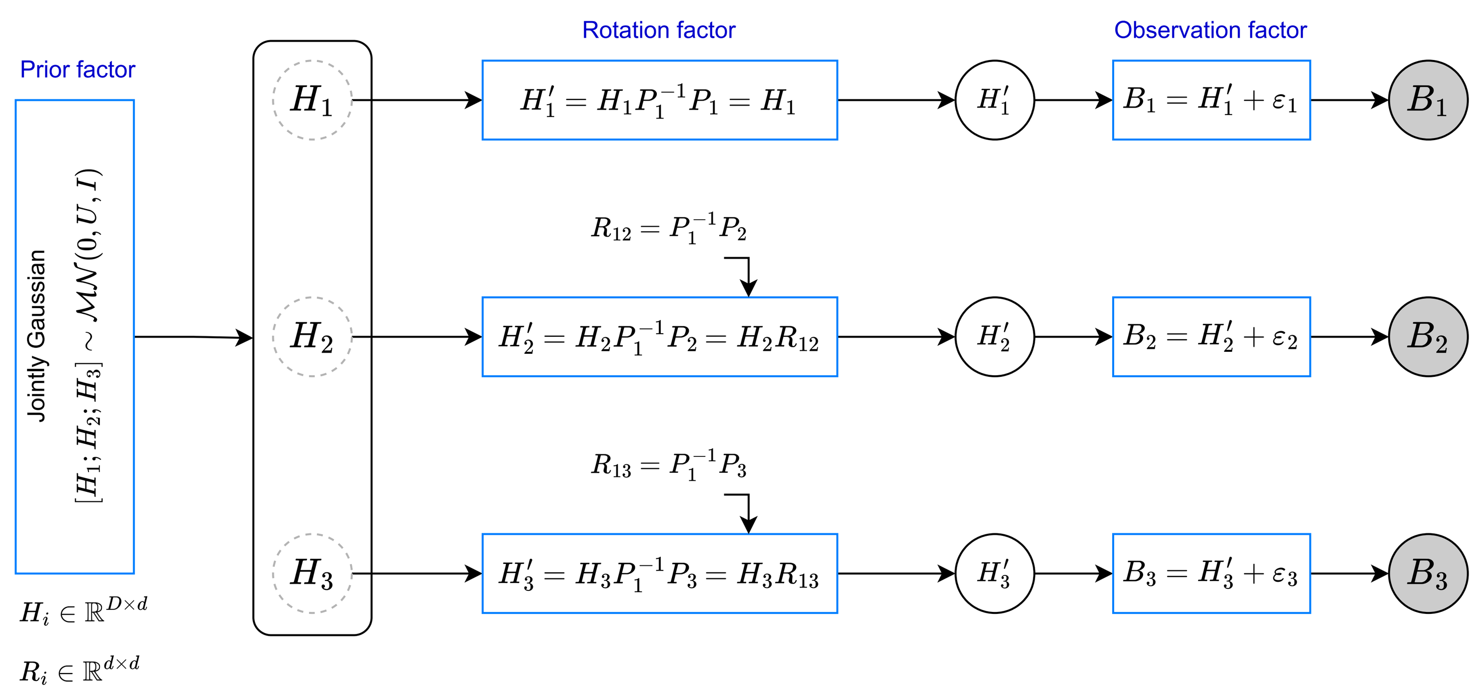

Structured Probabilistic Model to illustrate the joint distribution.

The receiver observes the effective channels $\boldsymbol{H}_i' = \boldsymbol{H}_i \boldsymbol{P}_i$ via some pilot messages at certain time-frequency positions (demodulation reference signal). This signal is used to estimate the full channel via linear predictors. An unknown rotation is modelled as orthogonal matrices $\boldsymbol{P} \in O(d) \subset \mathbb{R}^{d \times d}$. The joint probability distribution is:

$$p(\boldsymbol{B}, \boldsymbol{P}, \boldsymbol{H}, \boldsymbol{H}') = p(\boldsymbol{B}|\boldsymbol{H}')\ p(\boldsymbol{H}'|\boldsymbol{H}, \boldsymbol{P})\ p(\boldsymbol{P})\ p(\boldsymbol{H})$$

Both prior and observation model **are rotation invariant** and any attempt at inference thus requires symmetry breaking. Therefore, we work with relative poses $\boldsymbol{R} = \{\boldsymbol{R}_{i,j} := \boldsymbol{P}_i^{-1} \boldsymbol{P}_j\}_{i,j}$.

Inference: Estimating the channel $\boldsymbol{H}'$ can be done by writing the posterior distribution of $\boldsymbol{H}'$ given the data $\boldsymbol{B}$:

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

The *iterative refinement* is loosely inspired by Expectation-Maximization. However, there is only a point estimate, $q(\boldsymbol{R}) = \delta_{\boldsymbol{R}}$ in $\log p(\boldsymbol{B}) \geq \max_{q(\boldsymbol{R})} \mathbb{E}_{\boldsymbol{R} \sim q(\boldsymbol{R})}[\log p(\boldsymbol{B}, \boldsymbol{R}, \boldsymbol{H}_t') - \log q(\boldsymbol{R})]$.

Matrices $\boldsymbol{U}_a, \boldsymbol{U}_b, \cdots$ are submatrices of the covariance matrix $\boldsymbol{U}$ of $p(\boldsymbol{H})$. The relative rotations can be estimated with:

$$\begin{cases} \boldsymbol{R}_{12} = \arg\max_{\boldsymbol{R}_{12}} \langle \boldsymbol{R}_{21}, \boldsymbol{H}_{2,t}'^T \boldsymbol{U}_a^T \boldsymbol{H}_{1,t}' \rangle_F + \langle \boldsymbol{R}_{31}, \boldsymbol{H}_{3,t}'^T U_b^T \boldsymbol{H}_{1,t}' \rangle_F + \langle \boldsymbol{R}_{21}, \boldsymbol{H}_{2,t}'^T \boldsymbol{U}_c \boldsymbol{H}_{3,t}' \boldsymbol{R}_{31} \rangle_F, \\ \boldsymbol{R}_{13} = \arg\max_{\boldsymbol{R}_{13}} \langle \boldsymbol{R}_{21}, \boldsymbol{H}_{2,t}'^T U_a^T \boldsymbol{H}_{1,t}' \rangle_F + \langle \boldsymbol{R}_{31}, \boldsymbol{H}_{3,t}'^T \boldsymbol{U}_b^T \boldsymbol{H}_{1,t}' \rangle_F + \langle \boldsymbol{R}_{31}, \boldsymbol{H}_{3,t}'^T \boldsymbol{U}_c^T \boldsymbol{H}_{2,t}' \boldsymbol{R}_{21} \rangle_F. \end{cases}$$

The signal denoising can be solved in $O(D^3)$ time:

$$\begin{bmatrix} \boldsymbol{H}_{1,t+1}' \\ \boldsymbol{H}_{2,t+1}' \hat{\boldsymbol{R}}_{21} \\ \boldsymbol{H}_{3,t+1}' \hat{\boldsymbol{R}}_{31} \end{bmatrix} = \left(\sigma_\epsilon^2 \boldsymbol{U}^{-1} + \boldsymbol{I}_{3D} \right)^{-1} \begin{bmatrix} \boldsymbol{H}_{1,t}' \\ \boldsymbol{H}_{2,t}' \hat{\boldsymbol{R}}_{21} \\ \boldsymbol{H}_{3,t}' \hat{\boldsymbol{R}}_{31} \end{bmatrix}.$$

Above refinement of signal $\boldsymbol{H}_{i,t+1}'$ by marginalizing over the rotations $\boldsymbol{R}_{ij}$ usually converges in three to five steps (see Figure on the right).

Experimental results

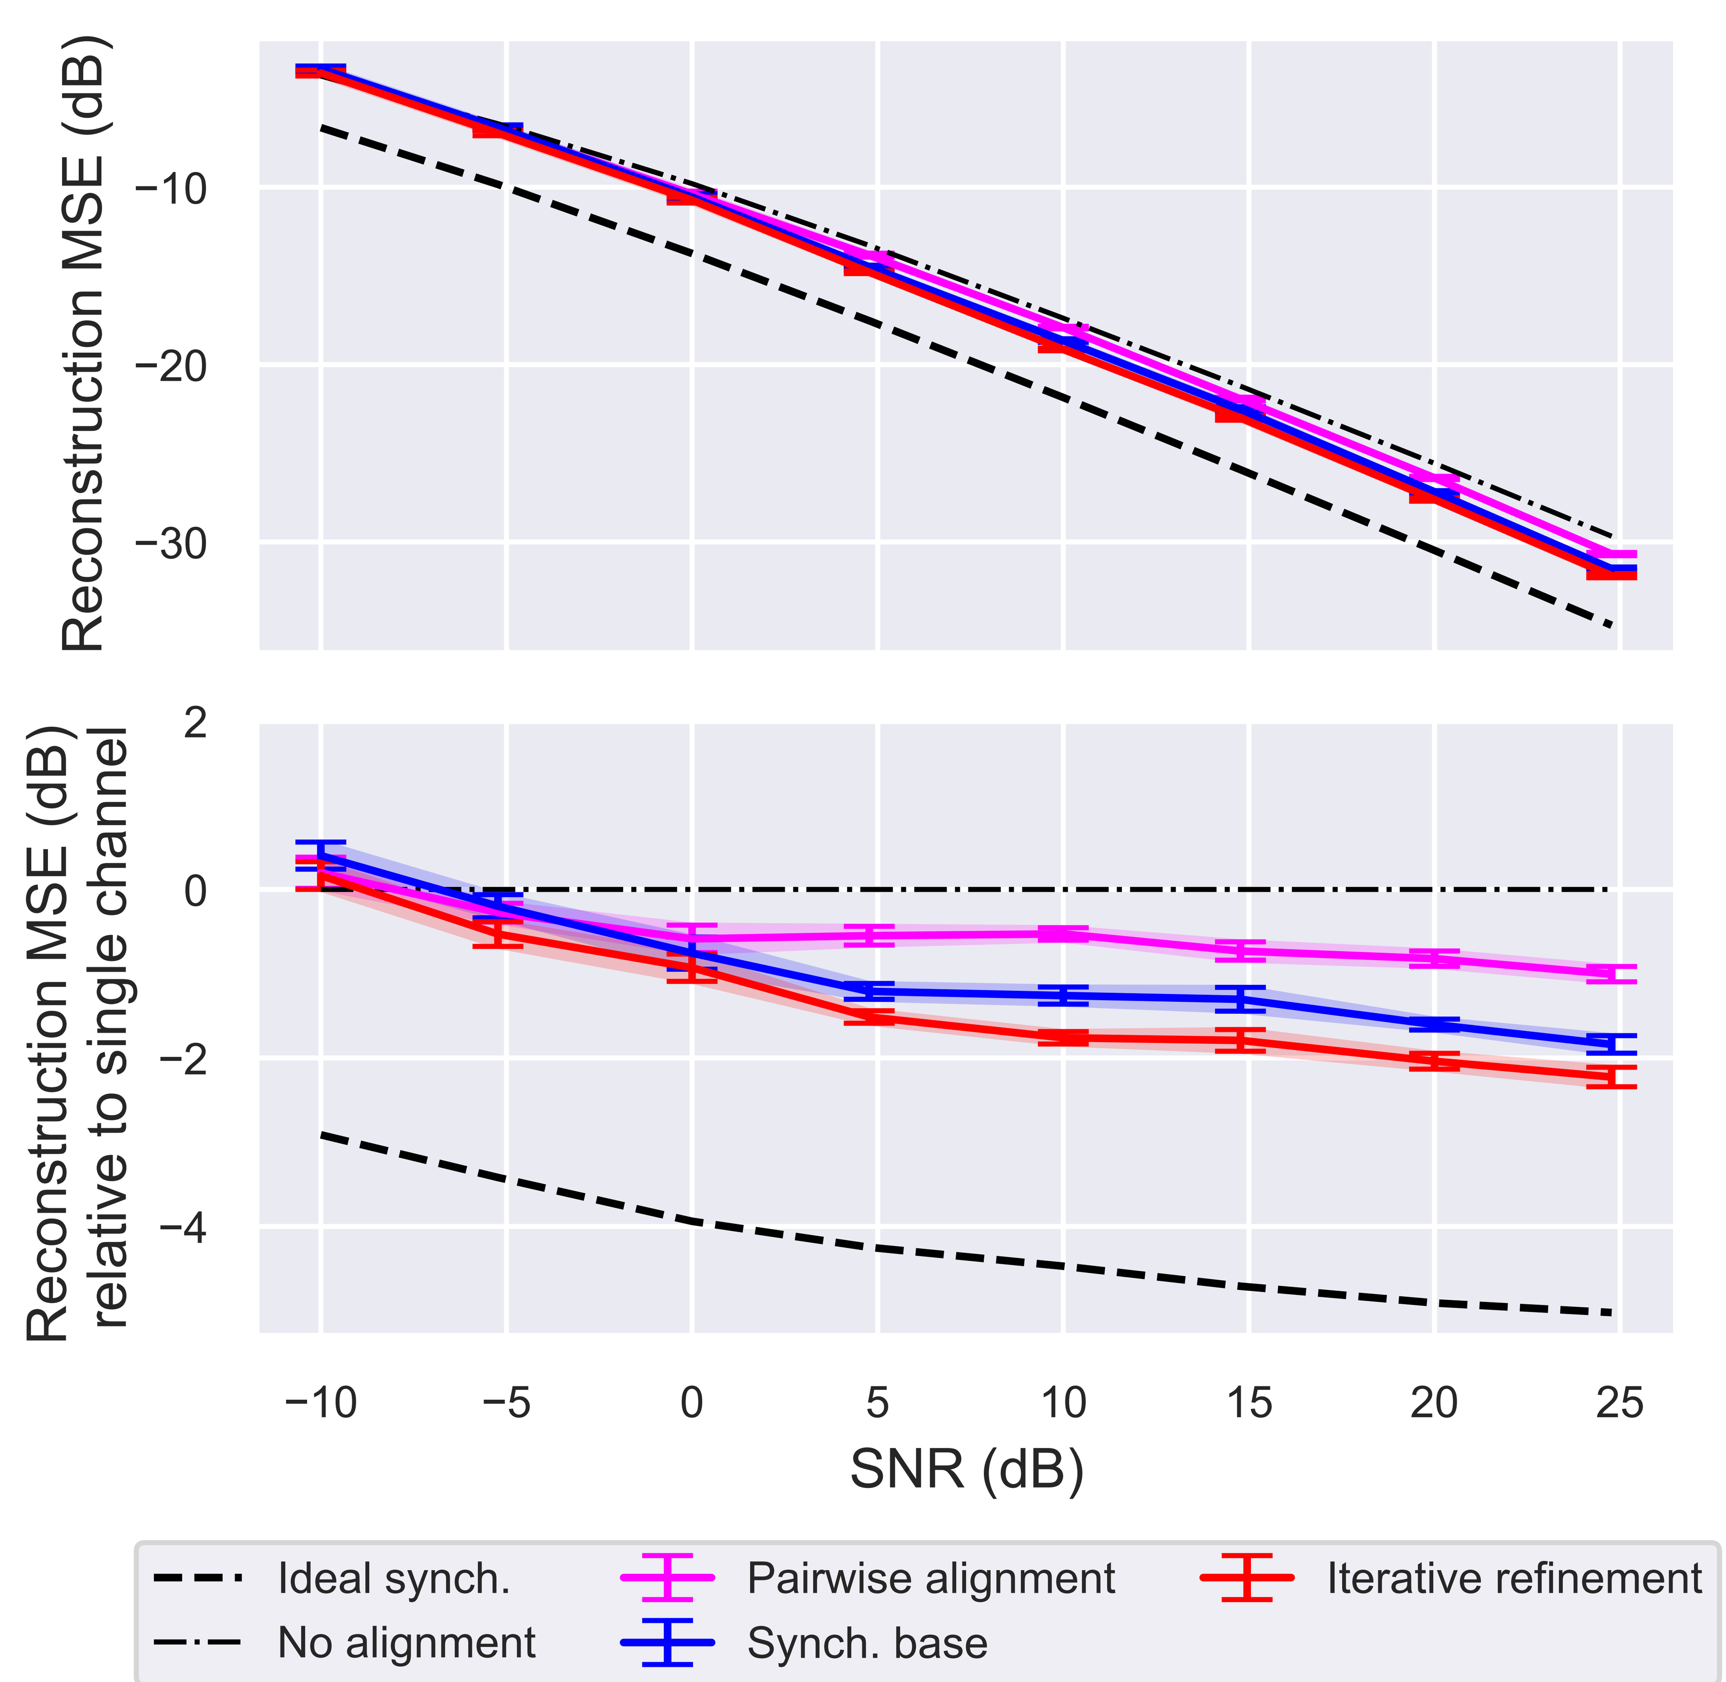

Iterative refinement achieves lower reconstruction error than signal denoising only. Standard error among 25 random seeds. The "synchronization base" approach already achieves lower reconstruction error than prior art for reasonable signal-to-noise ratios.

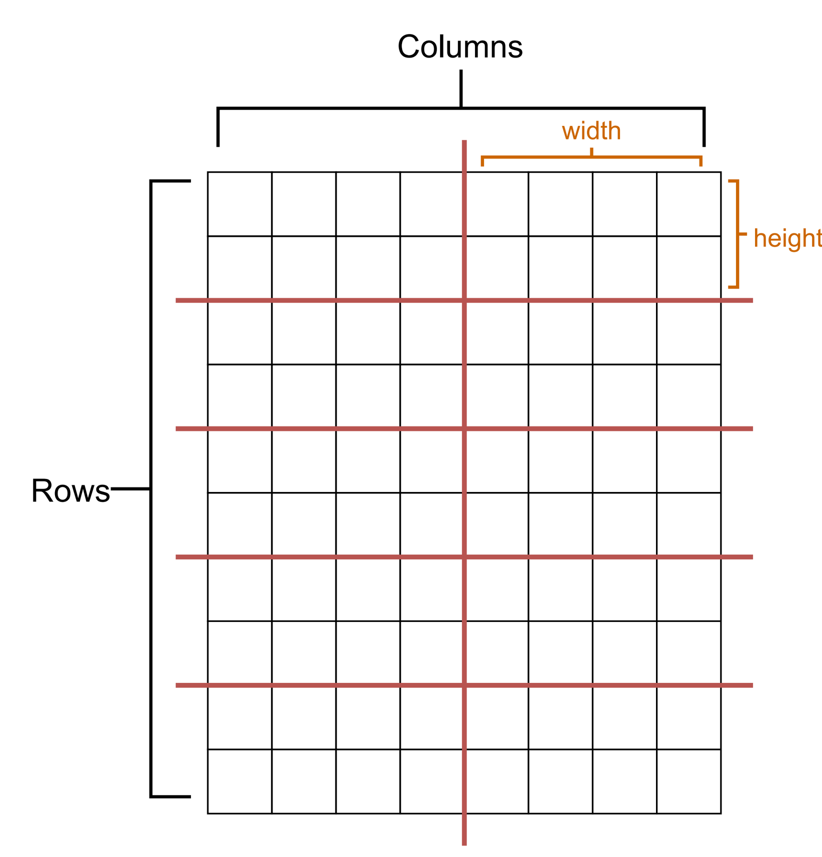

Overview of grid structure in the experiments. This is a model for OFDM time-frequency blocks in 5G MIMO.

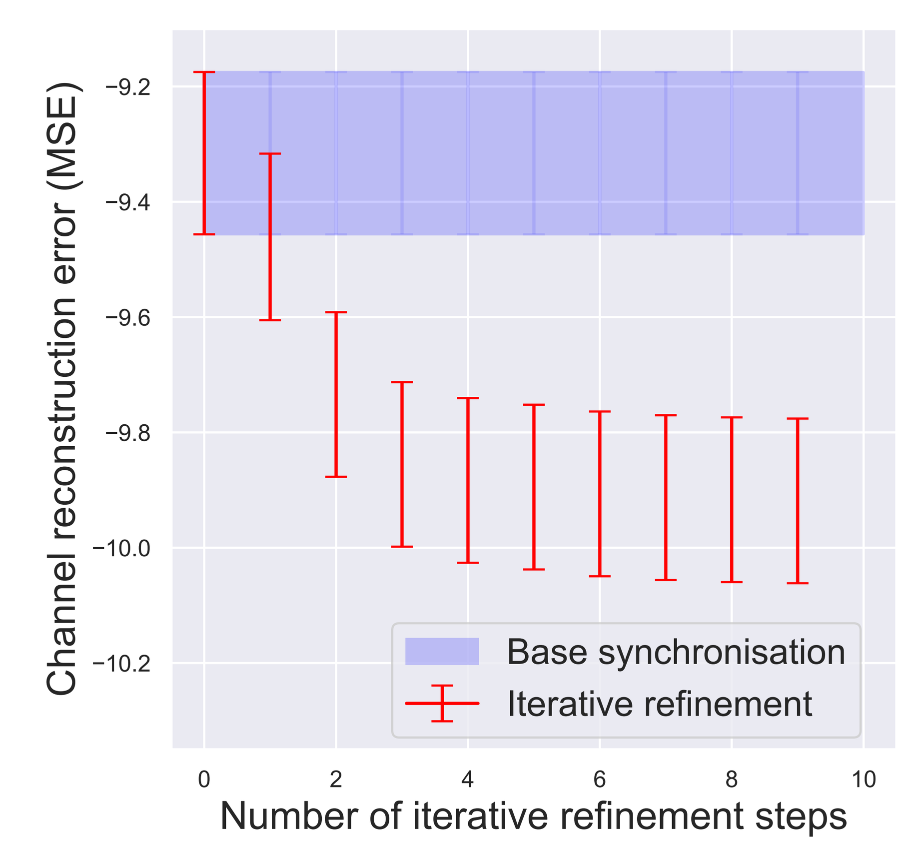

The error for varying number of iterative refinement steps. Most improvement is achieved in the first few steps.

Conclusion and future outlook

A significant advancement in multi-reference alignment (MRA) is made by introducing a structured probabilistic framework. The algorithm is computationally efficient, due to a triplet-based setting, and thus can scale to realistic grid sizes. The iterative refinement algorithm consistently outperforms direct estimation methods. It opens up new directions for developing more sophisticated denoising strategies that can leverage global synchronization.

References:

- Pratik et al. (2025) "Requestnet: A foundational learning model for channel estimation";
- Gower and Dijksterhuis. (2004). "Procrustes problems";
- Singer (2018). "Mathematics for cryo-electron microscopy";
- Perry et al. (2018). "Message-passing algorithms for synchronization problems over compact groups."