

Rotation Synchronization via Deep Matrix Factorization

GK Tejus¹, Giacomo Zara², Paolo Rota², Andrea Fusiello³, Elisa Ricci² and Federica Arrigoni⁴

Abstract—In this paper we address the rotation synchronization problem, where the objective is to recover absolute rotations starting from pairwise ones, where the unknowns and the measures are represented as nodes and edges of a graph, respectively. This problem is an essential task for structure from motion and simultaneous localization and mapping. We focus on the formulation of synchronization via neural networks, which has only recently begun to be explored in the literature. Inspired by deep matrix completion, we express rotation synchronization in terms of matrix factorization with a deep neural network. Our formulation exhibits implicit regularization properties and, more importantly, is unsupervised, whereas previous deep approaches are supervised. Our experiments show that we achieve comparable accuracy to the closest competitors in most scenes, while working under weaker assumptions.

I. INTRODUCTION

Rotation synchronization arises in several applications in Robotics and Computer Vision. The most prominent ones are pose graph optimization [9], structure from motion (SfM) [32] and 3D point cloud registration [19]. Pose graph optimization – which is at the core of *Simultaneous Localization and Mapping* (SLAM) – estimates robot poses by a nonlinear optimization which iteratively refines the solution, starting from an initial guess. Rotation synchronization provides such a starting point. In much the same way, it produces an initial guess for *Bundle Adjustment*, the nonlinear optimization underlying SfM [39].

Rotation synchronization (a.k.a. *multiple rotation averaging* [22] or *rotation optimization* [39]) is an instance of the more general “synchronization” problem [3], where the task is to find elements of a group starting from pairwise ratios: we are concerned here with the Special Orthogonal Group $SO(3)$, whose elements are rotations in the 3D space. Many representations of rotations exist, including unit quaternions, Euler angles and angle-axis. In this paper we consider rotation matrices:

$$SO(3) = \{R \in \mathbb{R}^{3 \times 3} \text{ s.t. } RR^T = I \wedge \det(R) = 1\}. \quad (1)$$

In the *rotation synchronization* problem [37], the input is a set of relative/pairwise rotations R_{ij} for $(i, j) \in \mathcal{E}$ and the task is to compute a set of global/absolute rotations R_1, \dots, R_n such that the following equation is satisfied

$$R_{ij} = R_i R_j^T \quad \forall (i, j) \in \mathcal{E}. \quad (2)$$

The set \mathcal{E} is usually viewed as the edge set of a simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{1, 2, \dots, n\}$ (see Fig. 1).

It is well known that the solution to this problem is not unique: indeed, if R_i satisfies (2) for all $i = 1, \dots, n$ then also $S_i = R_i Q$ satisfies the same equation, for any fixed $Q \in SO(3)$. Indeed, when computing the product $S_i S_j^T = R_i Q Q^T R_j^T$, the inner term simplifies since Q is a rotation matrix satisfying $Q Q^T = I$. This phenomenon is referred to as “gauge ambiguity” and it is usually handled by fixing one rotation to the identity, e.g., $R_1 = I$.

The problem has a solution as soon as the graph is connected. The minimal case corresponds to a spanning tree, which has $n - 1$ edges. In this scenario, a solution can be found by sequential propagation: the root of the tree is fixed to the identity, and the remaining rotations are computed from $R_i = R_{ij} R_j$, which is equivalent to (2). Synchronization, however, has its strength in exploiting *redundancy* to achieve error compensation: indeed, in typical datasets the number of available relative rotations is much larger than the minimum required, i.e., $|\mathcal{E}| \gg n - 1$.

Practical rotation synchronization has to deal with some real-world nuisances, such as: relative rotations may be corrupted by noise with small variance, so that (2) is not satisfied exactly; a fraction of the relative rotations may be completely wrong (outliers); many of them are typically missing, i.e., the graph is not complete. Therefore, current research focuses on finding principled algorithms that address all of these challenges simultaneously.

Most methods (e.g., [28], [21], [12], [11], [16], [5]) follow a traditional perspective, where the problem is formalized as directly optimizing a suitable cost function that measures the discrepancy between the left and right sides in (2) over all




Fig. 1. Rotation synchronization problem. The task is to recover unknown absolute rotations (on the nodes) starting from measured relative rotations (on the edges).

¹Indian Institute of Technology (ISM) Dhanbad, India

²University of Trento, Italy

³University of Udine, Italy

⁴Politecnico di Milano, Italy

the edges. A popular choice is the following objective¹

$$\min_{R_1, \dots, R_n \in SO(3)} \sum_{(i,j) \in \mathcal{E}} d(\hat{R}_{ij}, R_i R_j^T)^p \quad (3)$$

where d is a proper distance defined in $SO(3)$ and typically $p = 1$ or $p = 2$. See [22] for a review of common distances. As optimizing over rotations is a hard task [22], [42], many approaches typically relax the problem by ignoring rotation constraints in (3) and projecting the found solution onto $SO(3)$ at the end. Some examples include eigendecomposition and semidefinite programming [37], [1]. See Sec. II-A for more details and additional examples.

Recently, a few approaches have explored the usage of neural networks for rotation synchronization [34], [43], [25], based on the intuition that artificial neural networks are very good in learning patterns from noisy data. More precisely, given thousands of different problem instances with ground-truth solutions (i.e., graphs with input noisy relative rotations and associated output absolute rotations), the task is to learn a mapping that, at test time, regresses the sought parameters instead of optimizing them. In particular, a specific type of Graph Neural Networks are used in [34], [43], [25], namely Message Passing Neural Networks [17], by exploiting the fact that the rotation synchronization task exhibits a graph structure. This way of approaching the problem is still unexplored compared to the traditional perspective, thus motivating the need for further research, as done in our study. However, it should be noted that our method uses neural networks in a different way, as will be clarified later.

A. Contribution

Despite promising results were shown by [34], [43], [25], such methods suffer from a significant limitation: they are *supervised* approaches as they need a big dataset with known input-output relationship. However, such datasets are hardly available for rotation synchronization. In order to deal with limited dataset availability, the authors of [34], e.g., pre-train their network on synthetic graphs and fine-tune it on real graphs with a leave-one-out strategy.

To overcome this drawback, we propose a novel deep method for rotation synchronization which is *unsupervised*. In particular, our starting point is the observation that, since many relative rotations are missing, rotation synchronization can be viewed as a *matrix completion* problem [8], [4], namely the task of filling missing entries in a big low-rank matrix. In particular, we address the completion task via *deep matrix factorization* [2], where the sought matrix is written as the product of multiple terms. Our choice is motivated by the fact that gradient descent on a deep factorization has some interesting theoretical properties [2], namely it enhances an “implicit regularization” towards low-rank. In addition, as deep matrix factorization can be implemented as a neural network with *linear layers*, our approach offers interpretability [14]. Finally, in order to seek for robustness

¹To clearly distinguish between ground-truth and measured variables, hereafter we use the hat accent to denote the input measures, namely \hat{R}_{ij} denotes a noisy version of the relative rotation R_{ij} .

to outliers, we replace the ℓ_2 -norm typically used in general deep matrix factorization tasks with the ℓ_1 -norm, that best suits our rotation synchronization problem.

During the training phase our network minimizes a cost function via backpropagation: the intuition is to make this cost function be the one that solves a matrix completion problem on the instance presented at the input layer. Unlike traditional neural models, where to an input corresponds a predicted output, our solution regresses the complete matrix starting from a randomly initialised matrix of equal dimensions. During the training process, at each iteration the input matrix is equal to the transpose of the weight matrix of the first layer. This ploy enables the network to learn continuously, iteration after iteration, until convergence that is when the loss reaches a plateau. The network is unsupervised, in a sense that it only needs the known parameters of the incomplete target matrix. However the model solves a single problem instance for each training phase (inference is never carried out).

We evaluated our method on standard datasets used in the literature [41]. Our results show that we reach comparable accuracy to previous deep methods [34], [43], [25] in most cases. Considering that our approach is unsupervised whereas competing methods are supervised, this is an important achievement. Note also that outstanding accuracy is not needed as results are used as an initial guess for a subsequent optimization.

The paper is organized as follows. Section II reports previous work on rotation synchronization, Sec. III presents our method and Sec. IV shows experimental results. Finally, the conclusion is given in Sec. V.

II. RELATED WORK

In this section we review previous work on rotation synchronization, by considering separately methods based on explicit/analytical models and learned/neural ones. Our approach belongs to the latter category.

A. Traditional Methods

Rotations have a rich structure that inspired a lot of research in synchronization methods over the years. Early attempts include methods based on linear least squares, where rotations are represented either as unit quaternions [18] or rotations matrices [28]. This approach was revisited in [37], [1] by proposing a compact way to store all rotations simultaneously, and solving the resulting problem via eigendecomposition or semidefinite programming. A semidefinite formulation is also adopted in [16], [15], [33] with a focus on global optimality. The authors of [21] address rotation synchronization in an iterative fashion, where each absolute rotation is iteratively updated with the L1 average of its neighbors, computed with the Weiszfeld algorithm. Another popular approach is Lie-group optimization [10], where the Lie-group structure of rotations is exploited and the problem is iteratively solved in the tangent space. Other methods include the Levenberg-Marquardt algorithm [12], low-rank and sparse decomposition [4], Riemannian optimization [5],

message passing [36] and maximum likelihood estimation [30]. Theoretical aspects are also investigated in [22], [6], [42], [40].

B. Methods Based on Neural Networks

NeuRoRa (Neural Robust Rotation Averaging) [34] is the first solution to rotation synchronization based on neural networks. In particular, the authors propose two networks named CleanNet and FineNet respectively: the former rectifies the input noisy relative rotations while also detecting outliers; the latter fine-tunes an initialization of absolute rotations, which in turn is derived from a spanning tree of the cleaned graph. Both architectures are based on Message Passing Neural Networks (MPNN) [17]. This approach is not end-to-end, for it entails a spanning tree initialization as an intermediate step between the two networks. Multi-source propagation (MSP) [43] constitutes the first end-to-end approach to rotation synchronization. The focus is on providing an initialization only: the main idea is to initialize absolute rotations according to multiple sources in a differentiable way, in contrast to the common practice of utilizing the spanning tree to initialize rotations according to a single reference node in a top-down manner. Moreover, the image context is also exploited to reduce the impact of outliers. At test time, results are improved with FineNet [34] combined with traditional optimization as in (3). PoGO-Net [25] addresses rotation synchronization with a MPNN architecture, similarly to [34], [43]. The main innovation is the introduction of a de-noising layer that addresses the outlier removal via an implicit edge-dropping scheme.

III. PROPOSED METHOD

In this section we derive our method: first, we define the notation we use; secondly, we show how to cast rotation synchronization in terms of deep matrix factorization; then, we report some implementation details; finally, we summarize the properties of our approach.

A. Notation

As done by other authors (e.g., [1], [16], [15]), we conveniently collect all the relative and absolute rotations in two block-matrices named Z and X respectively:

$$Z = \begin{bmatrix} I & R_{12} & \dots & R_{1n} \\ R_{21} & I & \dots & R_{2n} \\ \dots & \dots & \dots & \dots \\ R_{n1} & R_{n2} & \dots & I \end{bmatrix}, \quad X = \begin{bmatrix} R_1 \\ R_2 \\ \dots \\ R_n \end{bmatrix} \quad (4)$$

where I denotes the 3×3 identity matrix. Note that Z has dimension $3n \times 3n$ whereas X has dimension $3n \times 3$. It is evident from (2) that Z can be decomposed as $Z = XX^T$, thus it is symmetric, positive semidefinite and it has low-rank (in fact, it has rank 3). Similarly to Z , we define \widehat{Z} as the $3n \times 3n$ block-matrix containing the subset of *observed* pairwise rotations \widehat{R}_{ij} (when available) and

unspecified elsewhere. For example:

$$\widehat{Z} = \begin{bmatrix} I & \widehat{R}_{12} & \dots & ? \\ \widehat{R}_{21} & I & \dots & \widehat{R}_{2n} \\ \dots & \dots & \dots & \dots \\ ? & \widehat{R}_{n2} & \dots & I \end{bmatrix} \quad (5)$$

Observe that \widehat{Z} is a *partial* matrix as many relative rotations are missing in practice, i.e., the graph is not complete.

B. Formulation

Recall that \widehat{Z} represents the input to rotation synchronization whereas X represents the desired output. Our objective here is to recover Z from \widehat{Z} , or, in other terms, at the same time *denoising* relative rotations and *completing* missing ones. Indeed, once Z is known, the sought absolute rotations X can be straightforwardly recovered from the eigendecomposition of Z [1]. In other words, rotation synchronization can be viewed as a *matrix completion* problem [8], [4].

In general, matrix completion deals with the recovery of missing values of a matrix from a subset of its entries. In this case, the task is to find W such that

$$W \odot \Omega \approx \widehat{Z} \odot \Omega \quad (6)$$

where \odot denotes the entry-wise product and Ω denotes the corresponding binary mask having zeros in correspondence of missing entries². Ω is also known as the *sampling set* in matrix completion literature [7]. Problem (6) means that we are looking for a *full* matrix W that is as close as possible to the input partial matrix \widehat{Z} *on the set of available entries*. Such a task is clearly ill-posed, as there are infinitely-many solutions. In order to make the problem tractable, additional assumptions are typically introduced. A popular choice is the *low-rank* assumption [7], [8], meaning that W is supposed to belong to a low-dimensional subspace. Note that such condition is satisfied in rotation synchronization, as observed in Sec. III-A.

Matrix completion is a highly researched topic comprising several solvers (e.g., [24], [27], [26], [13]) that work well for a variety of applications, the most prominent one being recommender systems (see the survey [31] for additional references). Most of the solutions typically enforce the low-rank assumption via *explicit* constraints such as the nuclear norm regularization [23]. Another possibility is to adopt *implicit* regularization, as done in [20], [2]. In particular, in [2] the unknown matrix W is expressed as the product of several factors:

$$W = W_d \cdot W_{d-1} \dots W_2 \cdot W_1 \quad (7)$$

and the following problem is addressed

$$\min_{W_1, \dots, W_d} \frac{1}{|\Omega|} \|(W_d \cdot W_{d-1} \dots W_2 \cdot W_1) \odot \Omega - \widehat{Z} \odot \Omega\|_F^2 \quad (8)$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $|\Omega|$ denotes the cardinality of the sampling set, i.e., the number of

²Observe that Ω has a block structure in the rotation synchronization problem: each missing rotation/pair results in a 3×3 zero block.




Fig. 2. Architecture, input and output of the network. Our solution is based on deep matrix factorization, implemented via linear layers (see Sec. III).

available entries. Equation (7) is also called *deep matrix factorization* and the number of factors d is called *depth*. It was observed in [2] that in practice, even when the rank is unconstrained, performing gradient descent on a matrix factorization tends to produce low-rank solutions, i.e., the gradient descent algorithm has *implicit regularization* properties. Moreover, increasing depth leads to a more accurate completion when the number of missing entries is high. In particular, the authors of [2] noted that “adding depth to the matrix factorization leads to more significant gaps between the singular values of the product matrix, i.e., to a stronger implicit bias towards low rank”.

In this paper we adopt the deep matrix factorization formulation from [2] and we apply it – for the first time – to the rotation synchronization problem. Such a framework is suitable for our task, as it involves a low-rank matrix with missing data, as already observed. However, a practical method has to deal also with wrong measures, i.e., outliers. To this end, we substitute the ℓ_2 -norm typically used for deep matrix factorization – see (8) – with the ℓ_1 -loss, motivated by the fact that the latter is more robust to outliers than the former. Thus we consider the following optimization problem:

$$\min_{W_1, \dots, W_d} \frac{1}{|\Omega|} \|(W_d \cdot W_{d-1} \dots W_2 \cdot W_1) \odot \Omega - \widehat{Z} \odot \Omega\|_{1,1}, \quad (9)$$

where $\|\cdot\|_{1,1}$ is the entry-wise ℓ_1 norm (in the same way as Frobenius is the entry-wise ℓ_2 norm). Following [2], we cast the above problem as training a *linear neural network* with d layers. See the next section for more details. Please note that the matrix completion problem is *not* solved in the forward phase, as one might expect. In fact, it is the backward phase that is equivalent to minimizing the cost function of matrix completion with factorization, as will be clarified later. Solving Problem (9) will produce an optimal $W = W_d \cdot W_{d-1} \dots W_2 \cdot W_1$, which represents a *denoised and complete* version of the input \widehat{Z} , or, equivalently, an approximation of the ground-truth Z . Then, in order to recover the sought absolute rotations, we apply eigendecomposition to the optimal W , as explained in [1].

C. Implementation Details

We follow the deep matrix completion approach from [2]. More precisely, the network is implemented in PyTorch as

a plain sequence of linear layers, whose input and output space is fixed across the architecture and the hidden layers, matching the size of the matrix associated with the problem (see Fig. 2). This allows the network to address the optimization problem outlined in (9). In particular, the weight matrices associated to the network are initialized according to a random Gaussian distribution. Given the absence of any activation function, the composition of functions carried out with the forward pass across the network can be described by the product reported in (7). At training time, the prediction of the network is compared with the observed entries in the incomplete matrix, and the parameters are updated by back-propagation according to a standard ℓ_1 loss. At every iteration, the input matrix is obtained by transposing the weights of the first layer, reducing the input noise variability while the training reaches convergence. At test time, the reconstructed matrix can be compared with the complete ground truth for evaluation. An RTX 2080 Nvidia GPU has been used to carry out our experiments. To optimize the network and its parameters, the SGD optimizer is used with a learning rate of 0.3 and a momentum of 0.9.

D. Properties

We summarize here the most relevant properties of our method, henceforth named DMF-SYNCH from “deep matrix factorization for synchronization”.

Unsupervised: Note that the matrix completion problem is solved for a single matrix at a time³: given relative rotations from a single synchronization problem, we initialize the network with random weights and use backpropagation to minimize the loss function (9). In this context, we may view each entry in the incomplete input matrix as a data point: observed entries constitute our training set, whereas the average reconstruction error over the unobserved entries is the test error, quantifying generalization. In this sense, we may view our method as an *unsupervised* approach. This is opposed to what is done by previous deep methods for rotation synchronization [34], [43], [25], which can be regarded as *supervised* approaches: they assume as input a dataset with multiple rotation synchronization problems whose solution is known, i.e., the absolute rotations are known for each problem instance comprising a set of relative rotations; hence the model is trained on such dataset and then tested on novel scenes. Following the terminology used in [29] for the related problem of projective factorization, we can say that our approach follows a *single scene* setup, whereas previous deep methods [34], [43], [25] perform *learning from multiple scenes*.

Interpretability: Our method is based on deep matrix factorization [2], which in turn is implemented by a neural network with linear layers. As noted in [14], the main motivation of deep matrix factorization is to combine both interpretability, as in classical matrix factorization (of which it is an extension), and the extraction of multiple hierarchical features, as allowed by multilayer architectures. In general, linearity has an interpretability power.

³This is customary in deep matrix completion literature [20], [2].




Fig. 3. Mean angular errors [degrees] of our approach versus fraction of missing edges. For each configuration, different choices of depth are analyzed.

Implicit Regularization: Our method is inspired by [2] which proposes an over parametrization of the unknown matrix via deep matrix factorization: this induces a form of implicit regularization when applying gradient-based optimization, namely a bias towards models of low complexity. Although this property is only proved for ℓ_2 -norm in [2], in Sec. IV we experimentally confirm this finding also for the ℓ_1 -norm, which is used by our formulation (see (9)). Note also that over-parameterization has been recently shown to be beneficial for 3D vision tasks like camera calibration [35].

IV. EXPERIMENTS

We report experiments on both synthetic data, where we validate some properties of DMF-SYNCH, and real data, where we compare our approach to the state of the art.

A. Synthetic Data

In order to analyze the regularization properties of our approach, we consider synthetic datasets generated as follows. First, n ground truth absolute rotations are obtained by random sampling Euler angles. The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is drawn from the Erdős-Rényi model with parameters (n, p) , i.e., given a vertex set $\mathcal{V} = \{1, 2, \dots, n\}$ each edge (i, j) is in the set \mathcal{E} with probability $p \in [0, 1]$, independently of all other edges. Only connected graphs are considered and a fraction of the pairwise rotations – equal to 0.4 in our experiments – was highly corrupted, thereby representing outliers. The remaining pairwise rotations were corrupted by small multiplicative noise $\widehat{R}_{ij} = R_{ij}N_{ij}$ where, with reference to the angle-axis representation, $N_{ij} \in \text{SO}(3)$ has random axis and angle following a Gaussian distribution with zero mean and standard deviation equal to 5° , thus representing a small perturbation of the identity matrix. In order to evaluate the results, we follow standard practices: first, an optimal rotation is computed [22] such that the solution gets “aligned” with the ground truth rotations, hence becoming comparable; then, for each estimated absolute rotation, the distance (e.g., angular) to the ground-truth one can be calculated and a single statistic (e.g., the mean or median over all rotations) is reported.

Results are given in Fig. 3, which shows the mean angular errors versus fraction of missing edges in the graph for two different scenarios (100 nodes and 200 nodes). Multiple values for the depth in the factorization are considered, in order to study implicit regularization for increasing depth. Indeed, it is claimed in [2] that – with reference to (8) – increasing depth leads to a more accurate solution to matrix completion, especially when the amount of missing data is high. From Fig. 3 we can appreciate that the same holds also for the ℓ_1 -norm used by our approach – see (9). In particular, we can observe that using depth 2 results in the worst errors. Hence we have experimentally verified that also the ℓ_1 -norm promotes low-rank solutions when combined with deep matrix factorization and gradient descent, or, in other terms, DMF-SYNCH has implicit regularization properties. Hereafter we will use depth 5 in our experiments.

Finally, it is worth observing that, when using more nodes, it results in a simpler problem. The reason lies in the fact that, for a fixed percentage of edges in the graph, then the effective number of edges does not scale linearly but quadratically with the number of nodes. For example, if the amount of missing data is fixed to 80%, then the number of edges is $0.2 \cdot n(n-1)/2$, meaning that there are $990 \approx 9n$ edges with 100 nodes and $3980 \approx 20n$ edges with 200 nodes. In the latter case there is significantly more redundancy compared to the minimal case (i.e., a spanning tree with $n-1$ edges), resulting in an easier problem.

B. Real Data

We consider the IDSfM benchmark [41], which has been widely adopted in the rotation synchronization literature. For each scene, both the graph and the associated relative rotations are available, which represents the input to our approach. As done in the literature, the output of Bundler [38] is regarded as the ground-truth solution. In order to make a fair evaluation, we focus on methods based on neural networks, namely NeuRoRa [34], MSP [43] and PoGO-Net [25]. We refer the reader to such papers for a detailed comparison between deep methods and traditional approaches. Observe that DMF-SYNCH is *unsupervised* whereas competing methods are supervised, as already mentioned in Sec. III-

TABLE I

MEDIAN ANGULAR ERRORS [DEGREES] FOR SEVERAL METHODS ON THE 1DSfM DATASET [41]. FOR EACH SCENE, THE FOLLOWING STATISTICS ARE REPORTED: THE NUMBER OF CAMERAS; THE PERCENTAGE OF EDGES IN THE GRAPH; THE HARDNESS MEASURE (THE LOWER THE HARDER) DEFINED IN [42]. CASES WHERE OUR METHOD (DMF-SYNCH) IS COMPARABLE TO COMPETING APPROACHES (I.E., THE ERROR DIFFERS BY AT MOST 0.5 DEGREES) ARE HIGHLIGHTED. RESULTS OF L1-IRLS ARE COPIED FROM [43] AND THE OTHERS ARE TAKEN FROM THE RESPECTIVE PAPERS.

Dataset				Median Error					
Name	# cam.	% Edges	Hardness	DMF-SYNCH	NeuRoRa [34]	MSP [43]	PoGO-Net [25]	L1-IRLS [11]	
Alamo	577	50	0.0017	1.2	1.2	1.1	0.9	1.1	
Ellis Island	227	50	0.0088	0.8	0.6	0.8	0.4	0.5	
Gendarmenmarkt	677	18	0.0021	10.5	2.9	3.7	-	7.7	
Madrid Metropolis	341	31	0.0029	2.3	1.1	1.1	1.0	1.2	
Montreal Notre Dame	450	47	0.0034	0.6	0.6	0.5	0.4	0.5	
NYC Library	332	29	0.0030	1.8	1.1	1.1	0.9	1.3	
Piazza del Popolo	338	40	0.0027	1.0	0.7	0.8	0.8	0.8	
Roman Forum	1084	11	0.0009	1.8	1.3	1.2	0.7	1.5	
Tower of London	472	19	0.0021	2.7	1.4	1.3	0.4	2.4	
Union Square	789	6	0.0006	4.4	2.0	1.9	1.3	3.9	
Vienna Cathedral	836	25	0.0010	1.6	1.5	1.1	1.4	1.2	
Yorkminster	437	27	0.0022	1.7	0.9	0.9	0.7	1.6	

D. Hence DMF-SYNCH works under weaker assumptions than its competitors, for it does not require training on a big dataset with known absolute rotations. Moreover, it is worth observing that DMF-SYNCH, NeuRoRa [34] and PoGO-Net [25] assume as input relative rotations and associated graph only, whereas MSP [43] exploits additional information (namely raw images), as summarized in Tab. II. As a reference, we also include in the comparison a traditional approach, namely L1-IRLS [11].

Results are shown in Tab. I, which reports the median angular error for all the analyzed methods. Some statistics about each scene are also reported, namely the number of cameras (i.e., the number of nodes in the graph representation), the percentage of edges (with respect to a complete graph) and also the hardness measure (the lower the harder) defined in [42]. In general, synchronization problems with few edges are harder to solve, since there is low redundancy, as reflected in the hardness measure. This is also reflected in the errors achieved by all methods, that get slightly higher for the most difficult problems. Table I shows that our approach exhibits good performances, with an angular error lower than two degrees in almost all of the scenes, thus showing the viability of deep matrix factorization for rotation synchronization.

Our method is on-par with L1-IRLS [11]. Moreover, DMF-SYNCH reaches comparable accuracy to previous deep methods in many cases: the fact that our approach is not the best is not surprising since we are tackling a more difficult task, as clarified in Tab. II. Note also that the property of being unsupervised is relevant in practice, given the lack of datasets for training (as confirmed by previous work [34]).

V. CONCLUSION

In this paper we considered the rotation synchronization problem and we showed the feasibility of a novel approach

TABLE II

DMF-SYNCH IS UNSUPERVISED WHEREAS COMPETING METHODS ARE SUPERVISED. ONLY MSP [43] USES INFORMATION COMING FROM RAW IMAGES. OUR METHOD CONSIDERS WEAKER ASSUMPTIONS.

	supervised	raw images
DMF-SYNCH	✗	✗
NeuRoRa [34]	✓	✗
MSP [43]	✓	✓
PoGO-Net [25]	✓	✗

based on deep matrix factorization. Our method follows a different perspective than previous deep techniques: here neural networks are not used for learning from data but as a powerful tool to minimize an objective function. In other terms, our approach is unsupervised and optimizes the network for a single scene at a time. Hence it does not need to acquire the ground truth, which is hard to obtain in real scenarios.

Although performances are only comparable to competing supervised approaches without outperforming them, this is a first attempt to investigate role of unsupervised neural networks towards rotation averaging, opening the way to new developments and application scenarios.

ACKNOWLEDGEMENTS

This work was supported by the EU Horizon 2020 Research and Innovation Programme under the project SPRING (No. 871245).

REFERENCES

- [1] M. Arie-Nachimson, S. Z. Kovalsky, I. Kemelmacher-Shlizerman, A. Singer, and R. Basri. Global motion estimation from point matches. *Proceedings of the Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2012.

- [2] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [3] Federica Arrigoni and Andrea Fusiello. Synchronization problems in computer vision with closed-form solutions. *International Journal of Computer Vision*, 128:26–52, 2020.
- [4] Federica Arrigoni, Beatrice Rossi, Pasqualina Fragneto, and Andrea Fusiello. Robust synchronization in $SO(3)$ and $SE(3)$ via low-rank and sparse matrix decomposition. *Computer Vision and Image Understanding*, 174:95–113, 2018.
- [5] Tolga Birdal, Michael Arbel, Umüt Simsekli, and Leonidas J. Guibas. Synchronizing probability measures on rotations via optimal transport. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1566–1576, 2020.
- [6] Nicolas Boumal, Amit Singer, P. A. Absil, and Vincent D. Blondel. Cramer-Rao bounds for synchronization of rotations. *Information and Inference: A Journal of the IMA*, 3(1):1 – 39, 2014.
- [7] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [8] Emmanuel J. Candès and Terence Tao. The power of convex relaxation: near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [9] Luca Carlone, Roberto Tron, Kostas Daniilidis, and Frank Dellaert. Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2015.
- [10] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proceedings of the International Conference on Computer Vision*, 2013.
- [11] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [12] David Crandall, Andrew Owens, Noah Snavely, and Daniel P. Huttenlocher. Discrete-continuous optimization for large-scale structure from motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3001–3008, 2011.
- [13] Mark A. Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- [14] Pierre De Handschutter, Nicolas Gillis, and Xavier Siebert. A survey on deep matrix factorizations. *Computer Science Review*, 42, 2021.
- [15] Frank Dellaert, David M. Rosen, Jing Wu, Robert Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing $so(p)^n$. In *Computer Vision – ECCV 2020*, pages 292–308. Springer International Publishing, 2020.
- [16] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 127–135, 2018.
- [17] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the International Conference on Machine Learning*, page 1263–1272, 2017.
- [18] V. M. Govindu. Combining two-view constraints for motion estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2001.
- [19] Venu M. Govindu and A. Pooja. On averaging multiview relations for 3D scan registration. *IEEE Transactions on Image Processing*, 23(3):1289–1302, 2014.
- [20] Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [21] R. Hartley, K. Aftab, and J. Trunpf. L1 rotation averaging using the Weiszfeld algorithm. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3041–3048, 2011.
- [22] R. I. Hartley, J. Trunpf, Y. Dai, and H. Li. Rotation averaging. *International Journal of Computer Vision*, 2013.
- [23] Yao Hu, Debing Zhang, Jieping Ye, Xuelong Li, and Xiaofei He. Fast and accurate matrix completion via truncated nuclear norm regularization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2117–2130, 2013.
- [24] Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Trans. Inf. Theor.*, 2009.
- [25] Xinyi Li and Haibin Ling. Pogo-net: Pose graph optimization with graph neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5895–5905, October 2021.
- [26] Zhouchen Lin, Minming Chen, and Yi Ma. The augmented lagrange multiplier method for exact recovery of corrupted Low-Rank matrices. eprint arXiv:1009.5055, 2010.
- [27] Shiqian Ma, Donald Goldfarb, and Lifeng Chen. Fixed point and Bregman iterative methods for matrix rank minimization. *Mathematical Programming*, 128(1-2):321–353, 2011.
- [28] D. Martinec and T. Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
- [29] Dror Moran, Hodaya Koslowsky, Yoni Kasten, Haggai Maron, Meirav Galun, and Ronen Basri. Deep permutation equivariant structure from motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5976–5986, 2021.
- [30] Gabriel Moreira, Manuel Marques, and João Paulo Costeira. Rotation averaging in a split second: A primal-dual method and a closed-form for cycle graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5452–5460, 2021.
- [31] Luong Trung Nguyen, Junhan Kim, and Byonghyo Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.
- [32] Onur Ozyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *Acta Numerica*, 26:305 – 364, 2017.
- [33] Alvaro Parra, Shin-Fang Chng, Tat-Jun Chin, Anders Eriksson, and Ian Reid. Rotation coordinate descent for fast globally optimal rotation averaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4298–4307, 2021.
- [34] Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 137–154. Springer International Publishing, 2020.
- [35] Thomas Schops, Viktor Larsson, Marc Pollefeys, and Torsten Sattler. Why having 10,000 parameters in your camera model is better than twelve. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [36] Yunpeng Shi and Gilad Lerman. Message passing least squares framework and its application to rotation synchronization. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8796–8806. PMLR, 2020.
- [37] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20 – 36, 2011.
- [38] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: exploring photo collections in 3D. In *SIGGRAPH: International Conference on Computer Graphics and Interactive Techniques*, pages 835–846, 2006.
- [39] Roberto Tron, Xiaowei Zhou, and Kostas Daniilidis. A survey on rotation optimization in structure from motion. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2016.
- [40] K. Wilson and D. Bindel. On the distribution of minima in intrinsic-metric rotation averaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6030–6038, 2020.
- [41] K. Wilson and N. Snavely. Robust global translations with 1DSfM. In *Proceedings of the European Conference on Computer Vision*, pages 61–75, 2014.
- [42] Kyle Wilson, David Bindel, and Noah Snavely. When is rotations averaging hard? In *Proceedings of the European Conference on Computer Vision*, pages 255 – 270, 2016.
- [43] Luwei Yang, Heng Li, Jamal Ahmed Rahim, Zhaopeng Cui, and Ping Tan. End-to-end rotation averaging with multi-source propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11774–11783, June 2021.