
On the use of the Doubly Stochastic Matrix models for the Quadratic Assignment Problem

Valentino Santucci

University for Foreigners of Perugia,
Perugia, 06123, Italy

valentino.santucci@unistrapg.it

Josu Ceberio

University of the Basque Country UPV/EHU
Donostia-San Sebastian, Spain

josu.ceberio@ehu.eus

Abstract

Permutation problems have captured the attention of the combinatorial optimization community for decades due to the challenge they pose. Although their solutions are naturally encoded as permutations, in each problem, the information to be used to optimize them can vary substantially. In this article, we consider the Quadratic Assignment Problem (QAP) as a case study and propose using Doubly Stochastic Matrices (DSMs) under the framework of Estimation of Distribution Algorithms. To that end, we design efficient learning and sampling schemes that enable an effective iterative update of the probability model. Conducted experiments on commonly adopted benchmarks for the QAP prove doubly stochastic matrices to be preferred to four other models for permutations, both in terms of effectiveness and computational efficiency.

Reference: The following document is a brief summary of Santucci and Ceberio (2025). For further information, we refer the interested reader to the original paper.

1 Introduction

Permutation problems are a subset of combinatorial optimization problems characterized by representing their solutions naturally through permutations. Despite the fact that the encoding of the solutions is their point in common, in each case permutations can encode diverse information.

As a result, focusing on the optimization of permutation problems by means of general-purpose algorithms does not seem to be the most suitable strategy. In addition, to improve the performance of the algorithms, it is necessary for each problem to define those operators or strategies that are compatible with the encoding and, at the same time, are well aligned with the characteristics of the problem. That is, they are able to capture the features of the solutions that influence the objective function value.

In this article, we investigate the Doubly Stochastic Matrix (DSM) models for modeling permutations, and introduce it in the framework of Estimation of Distribution Algorithms (EDAs) to tackle the Quadratic Assignment Problem (QAP). To that end, we propose an efficient learning scheme, based on the well-known algebraic properties of DSMs, which captures the information of the item-to-item assignments appearing in a set of training permutations. With regard to the sampling scheme, we introduce and discuss several strategies. Two of them, namely the probabilistic and algebraic sampling strategies, are analyzed to understand how faithful they are to the learning algorithm. We also show that DSMs are able to capture and propagate the relevant information contained in the solutions of QAP instances more efficiently.

2 Learning and Sampling Doubly Stochastic Matrices

A DSM is a square matrix $\mathbf{D} = [d_{ij}]_{n \times n}$ of non-negative real numbers whose rows and columns sum to 1. Formally, we have $d_{ij} \geq 0$, $\sum_{k=1}^n d_{ik} = 1$, and $\sum_{k=1}^n d_{kj} = 1$, for all pairs $i, j \in [n]$. Moreover, we denote by $\mathbb{D}_n$ the set of all DSMs of order n .

It is easy to see that any permutation can be encoded as a special DSM with the additional constraint that the entries are either 0 or 1. In fact, there exists an isomorphism between $\mathbb{S}_n$ and the set $\mathbb{P}_n$ of so-called *permutation matrices*. Given $\sigma \in \mathbb{S}_n$, its associated permutation matrix $\mathbf{P} \in \mathbb{P}_n$ has the form $\mathbf{P} = [p_{ij}]_{n \times n}$, where, for all $i, j \in [n]$, p_{ij} is 1 if and only if $\sigma(i) = j$ and 0 otherwise. Hence, permutation matrices are a proper subset of DSMs, that is, $\mathbb{P}_n \subset \mathbb{D}_n$.

From this point of view, it is apparent that DSMs can effectively model probability distributions over permutations for assignment problems such as the QAP. In fact, in the QAP there are two sets A and B – of equal size and without any requirement for the internal order of their elements – that must be matched. Therefore, the row indices in $[n]$ can be used to encode the items of A , while the column indices in $[n]$ can be used to encode the items of B .

2.1 Learning

A well-known result in the field of DSMs is the Birkhoff-von Neumann (BvN) theorem, which states that $\mathbb{D}_n$ defines a polytope, embedded in the n^2 -dimensional Euclidean space, which is the convex hull of $\mathbb{P}_n$. In other words, the BvN theorem asserts that $\mathbb{D}_n$ is closed under a convex combination of its elements and that every DSM can be written as a convex combination of permutation matrices.

This result allows us to design a learning procedure based on the concept of convex combination. Given m training permutations $\mathbf{P}_1, \dots, \mathbf{P}_m \in \mathbb{P}_n$ (expressed as permutation matrices for presentation purposes), we define the learned DSM $\mathbf{D} \in \mathbb{D}_n$ as follows:

$$\mathbf{D} \leftarrow w_1 \mathbf{P}_1 + w_2 \mathbf{P}_2 + \dots + w_m \mathbf{P}_m + \alpha \mathbf{U}, \quad (1)$$

where $\mathbf{U} = [u_{ij}]_{n \times n}$ is the *uniform DSM* such that $u_{ij} = 1/n$ for all pairs $i, j \in [n]$, while α along with $w_1, w_2, \dots, w_m$ are non-negative weights summing to 1.

Eq. 1 describes precisely the DSM $\mathbf{D}$ as a convex combination of $m + 1$ terms: the m training permutations provided in input, and the uniform DSM $\mathbf{U}$. This formulation serves two key purposes: (i) it allows summarizing in $\mathbf{D}$ all item-to-item assignments on the basis of their observed frequency as encoded in the training permutations, and (ii) the inclusion of the uniform DSM allows the smoothing of the multinomial distributions in $\mathbf{D}$, thus permitting positive probabilities also for those item-to-item assignments which were not observed in the training permutations.

Regarding the coefficients in the convex combination of Eq. 1, the values w_i can be set according to the utility of each training permutation – or, simply $w_i = (1 - \alpha)/m$ for $i = 1, \dots, m$, when no permutation is preferred over the others –, and the α parameter serves as a *smoothing factor* ($\alpha \in [0, 1]$) that plays a crucial role in regulating the exploration behavior of subsequent samplings performed with the learned model. In the two extreme cases: when $\alpha = 0$, only item-to-item assignments observed in the training permutations can be sampled, while when $\alpha = 1$, the learned DSM is formed only by uniform multinomial distributions, thus spreading evenly over the entire permutation space. In practical scenarios, α can be set to small values according to the space dimensionality n , such as $1/n^2$ or $1/n$, or it can be adjusted during iterations to balance the exploration-exploitation behavior of the search.

So we present a learning algorithm that takes as input m permutations (equally weighted) and the smoothing factor α . Once the DSM is initialized, the algorithm accounts for the item-to-item assignments in all input permutations (for an extensive explanation, check the original paper).

2.2 Sampling

Sampling permutations from a DSM in an exact manner is a #P-hard complexity process. In light of this, we resort to heuristic sampling approaches that align as closely as possible to exactly sampling the probability distribution induced by $\mathbb{D}$. In particular, we present two sampling procedures, namely *Probabilistic Sampling* (PS) and *Algebraic Sampling* (AS), which are empirically analyzed.

Probabilistic Sampling (PS). Given a DSM $\mathbf{D} \in \mathbb{D}_n$, the PS algorithm samples a permutation $\sigma \in \mathbb{S}_n$ by iteratively selecting item-to-item assignments based on the probabilities of rows and columns of $\mathbf{D}$. At each iteration, after the item-to-item assignment $\sigma(i) = j$ is set, the entries of $\mathbf{D}$ are updated by zeroing out both the row i and the column j . The process is repeated until all the n item-to-item assignments in σ are set. It is interesting to observe that, when the input DSM model is a permutation matrix, PS can only sample its corresponding permutation, while when the input DSM is $\mathbf{U}$, all permutations have the same probability of being sampled.

Algebraic Sampling (AS). The AS procedure is based on the “randomized rounding” methodology outlined by Wolstenholme and Walden (2016). Given a DSM $\mathbf{D} \in \mathbb{D}_n$, the AS procedure generates a random vector $v \in [0, 1]^n$ and selects the permutation matrix $\mathbf{P} \in \mathbb{P}_n$ that solves the equation

$$\mathbf{P} \cdot \text{rank}(v) = \text{rank}(\mathbf{D} \cdot v), \quad (2)$$

where $\cdot$ is the usual matrix-vector multiplication, while the vector $\text{rank}(v)$ is defined in such a way that $\text{rank}(v)_i$ is the rank of v_i among all elements in v . For example, if $v = (0.2, 0.6, 0.8, 0.4)$, then $\text{rank}(v) = (1, 3, 4, 2)$. Actually, $\text{rank}(v)$ is the inverse permutation of $\text{argsort}(v)$. Additionally, due to the finite precision of computer arithmetic, there is a tiny probability of observing ties in the values of v that are randomly broken in our implementation.

3 Experiments

To analyze the effectiveness of the proposed DSM model for optimizing QAP instances within EDAs, we designed an experiment that compares the ability of the proposed model to capture good features of assignment problem solutions with that of the Plackett-Luce (PL) model (known to be good for ordering problems). With this aim, for each instance/model pair, we plotted the recorded objective values of each iteration (using the DSM and PL models, respectively) as a density plot. The curves representing the different iterations, from 0 to 5, are colored with an increasing color gradient.




Figure 1: Density plots of the QAP objective values recorded in each iteration for both the DSM model (left plot) and the PL model (right plot) on a selected QAP instance.

From Fig. 1, it is apparent that the DSM model effectively allowed us to improve the objective values of the sampled solutions as the iterations progressed, whereas this improvement is much less evident for the PL model. Furthermore, considering that the initial models in the experiment were learned from a selection of good samples, Fig. 1 also shows that the DSM model allows us to quickly reach and improve the quality of these initial samples.

In conclusion, this experiment empirically demonstrates that a model suitable for one type of permutation problem, such as the PL model, which has been shown to be effective for ordering problems, may not necessarily be suitable when applied to another class of permutation problems. Moreover, the comparison also reveals that the proposed DSM model is particularly well-suited for the assignment class of problems, thereby supporting its incorporation into EDAs.

References

- Santucci, V. and Ceberio, J. (2025). On the use of doubly stochastic matrix models for the quadratic assignment problem. *Evolutionary Computation Journal*.
- Wolstenholme, R. and Walden, A. T. (2016). A sampling strategy for projecting to permutations in the graph matching problem. *arXiv preprint arXiv:1604.04235*.