Doubly Stochastic Matrix Models for Estimation of Distribution Algorithms

Valentino Santucci
University for Foreigners of Perugia
Perugia, Italy
valentino.santucci@unistrapg.it

Josu Ceberio
University of the Basque Country
San Sebastian, Spain
josu.ceberio@ehu.eus

Abstract

Problems with solutions represented by permutations are very prominent in combinatorial optimization. Thus, in recent decades, a number of evolutionary algorithms have been proposed to solve them, and among them, those based on probability models have received much attention. In that sense, most efforts have focused on introducing algorithms that are suited for solving ordering/ranking nature problems. However, when it comes to proposing probability-based evolutionary algorithms for assignment problems, the works have not gone beyond proposing simple and in most cases univariate models. In this paper, we explore the use of Doubly Stochastic Matrices (DSM) for optimizing matching and assignment nature permutation problems. To that end, we explore some learning and sampling methods to efficiently incorporate DSMs within the picture of evolutionary algorithms. Specifically, we adopt the framework of estimation of distribution algorithms and compare DSMs to some existing proposals for permutation problems. Conducted preliminary experiments on instances of the quadratic assignment problem validate this line of research and show that DSMs may obtain very competitive results, while computational cost issues still need to be further investigated.

1 Introduction

Permutation problems have been a prominent research topic for the combinatorial optimization community. Although permutation problems belong to the family of combinatorial problems, the developments that have been made over the years have raised specific proposals for this type of problems due to the particularity that distinguishes them: the permutation codification. As stated by [1], permutations are probably among the richest combinatorial structures. Motivated principally by their versatility, permutations appear in a vast range of domains, such as graph theory, mathematical psychology or bioinformatics, but particularly, in logistic problems such as routing [2], scheduling [3] or assignment [4].

The literature presents a significant number of different permutation problems, however, it is possible to classify them based on the nature of what they represent. Santucci et al. [5], divide permutation problems into ordering and matching problems. The first class of problems aims to find an ordering/ranking of a given set of items, while the second pretends to match two given equally sized sets of items. It becomes obvious that, due to the different meaning that the permutations have in each case, those problems require different algorithmic approaches.

With the aim of solving permutation problems, the community of evolutionary computation has proposed many approaches, and has focused with special intensity in the development of probability-based strategies (on a variety of paradigms).

To enumerate some, Ceberio et al. investigated the usage of probability models that describe a probability distribution over permutations such as Plackett-Luce [6] and Mallows models under Kendall's-τ [7] and Ulam [8] distances, and evaluated their performance in the framework of Estimation of Distribution Algorithms (EDAs). Ayodele et al. [9] proposed the use of a transformation between real-coded vectors and permutations called "Random Key" and approached the permutation flowshop scheduling problem with an EDA. Santucci et al. [10] studied model-based gradient search under the Plackett-Luce model for optimizing the linear ordering problem.

A deep literature analysis reveals that when approaching permutation problems, the authors have focused on designing strategies that are compatible with the permutation codification [11, 12]. However, it becomes obvious to think that there is no model that works well for any permutation problem. In fact, we observe that the models enumerated above focus principally on modelling the ordering/rankings of the items in the permutations, and models that capture the assignment nature of the problems are less frequent. In this trend, we find the work by Irurozki et al. [13] that developed a Mallows model under the Cayley distance and reported to be suited for modelling permutations from the second class of problems. A similar consideration about the suitability of the so-called "exchange" or "interchange" moves for assignment/matching problems is provided also in [14] and [15].

In this paper, we aim to continue this trend and propose probability models that are suited to model permutations from matching problems, and then use such models to develop probability-based evolutionary algorithms. With this purpose, we study the use of Doubly Stochastic Matrices (DSMs) as probability models over the space of permutations. It is already known that a permutation can be represented as a 0/1 matrix where each row and column have exactly one 1-entry (and 0 for the rest of the entries). This makes them particular cases of DSMs that relax the constraints to "each row/column are formed by positive real numbers summing-up to 1". These properties (that will be made clear in Section 3) allows seeing the DSMs as synthesis of multiple permutations, thus making it possible to learn a DSM from permutations and to sample permutations from a DSM. DSMs specify the probability of an item appearing at a particular position and, intuitively, they look good models for matching/assignment problems like the well known Quadratic Assignment Problem (QAP) [4].

On the basis of previous literature on DSMs, mostly exogenous to evolutionary computation community, we investigate learning and sampling procedures and analyze their basic theoretical properties. Then, a number of EDA proposals using DSMs as a model are built. Finally, experiments on a selected set of QAP benchmark instances are carried out and the results of the different strategies proposed are compared among them and with respect to the most relevant EDA proposals for the QAP. Moreover, also empirical analysis of convergence are investigated. Results reveal good potential in the usage of DSMs to deal with permutation matching problems.

The remainder of the paper is organized as follows. Section 2 provides background on permutation problems, the encodings used and their nature. It also presents a distinction on the classes of problems. Next, in Section 3, theoretical aspects of Doubly Stochastic Matrices are introduced and, afterward, in Section 4 learning and sampling procedures are explored. Section 5 presents the EDA algorithmic scheme adopted and discuss its characteristics. Preliminary experiments are summarized in Section 6. The paper concludes in Section 7 with a summary of the contribution and by proposing a number of lines for future research.

2 Representations and Encodings in Permutation Problems

Formally, a permutation is a bijection function of the set of items $[n] = \{1, \ldots, n\}$ onto itself. Usually, the Greek letters σ, π or ρ are used to denote them. $\sigma(i)$ (also denoted as σ_i for readability) represents the item at position i, and $\sigma^{-1}(i)$ represents the position of item i. The set of all permutations of size n is denoted as $\mathbb{S}_n$, defines a group under the composition operation, and is known as the symmetric group in algebra.

The classical description of a permutation is the usual vector representation, where the items in the set $[n]$ appear in a particular order. Nevertheless, beyond an ordered set of items, a permutation can be represented uniquely (i.e., there exists a bijection between the vector representation and the following) by a collection of disjoint cycles, transpositions, pairwise precedence, 0/1 matrices, or even as graphs. Probably, due to the innumerable possibilities in which permutations can be encoded, they have served to represent solutions in combinatorial problems of different nature.

Santucci et al. [5] classified existing problems into two families (see also Fig. 1):

- *ordering problems*, where the goal is to find an optimal ordering of a given set of items (as e.g., in the permutation flowshop scheduling problem), and

- *matching problems*, where it is required to match, in the best possible way, two given equally sized sets of items (as e.g., in the quadratic assignment problem).

For both, ordering and matching problems, the fact that permutations are bijections of the first n integers has to be intended only as a genotypic encoding. In an ordering problem, a set A of n items to be optimally ordered (on the basis of a provided objective function) is given. Hence, an ordering of the items in A can be represented in two distinct ways: as a mapping from positions to items (*ordering representation*), or as a mapping from items to positions (*ranking representation*). Note that the semantic interpretation (ordering of items) is exactly the same for both the representations (ordering or ranking). Once chosen a representation, the next step is to find a genotypic encoding. Clearly, positions are
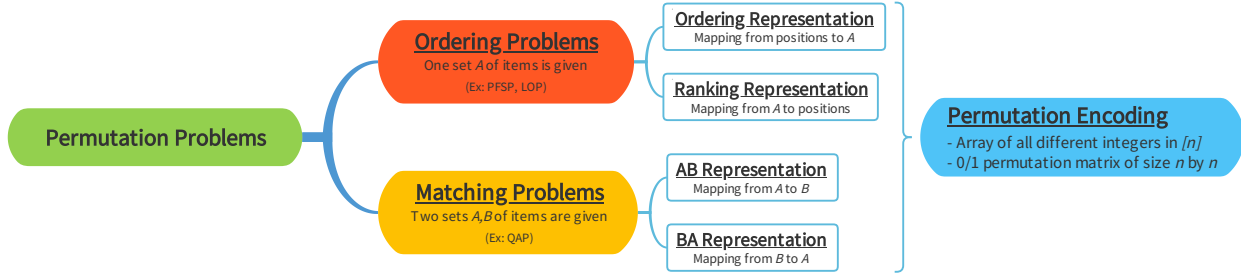
Figure 1: Relations among semantic interpretations, representations, and encodings in permutation problems.

integers in $[n]$, while the items in A can be arbitrarily assigned to (all-different) identification numbers in $[n]$. Therefore, both positions and items are encoded as elements of $[n]$, thus a mapping between them can be easily encoded by a permutation in $\mathcal{S}_n$. In other words, the ordering and ranking representations share the same genotypic encoding, but they still remain different representations of the same semantic interpretation. Note that in the case of ordering problems, under the ordering representation, the critical information about the solution is provided by the precedence between the items in the sequence (or the comparison of the magnitude of the rank under the ranking representation).

The genotype of matching problems describes a different idea. Specifically, the permutation describes a bijective function that matches the items in two discrete sets A and B of equal size n. Analogously to ordering problems, there are also two representations for matching problems: (1) *AB representation* and (2) *BA representation*. In the first case, each position in $[n]$ refers to the item from the set A, and the identification number placed in that position, denotes the item from B to which the matching has been made. The BA representation, as the reader can guess, swaps the sets A and B from position and identification numbers.

In either problem, ordering or matching, it is possible to easily convert one representation to the other under the inversion operation.

As stated in the introduction of this section, permutations can handle different descriptions beyond the usual vector representation. Two permutation encoding schemes which are relevant for this work are: (1) arrays of n all different integers from $[n]$ (vector representation), and (2) n-by-n permutation matrices, i.e., 0/1 matrices such that each row and each column has exactly one 1-entry (0/1 matrix representation). The two encoding schemes are clearly equivalent, and they can be converted to each other by simple conversion procedures. In the following, we denote by $\mathbb{S}_n$ and $\mathbb{P}_n$ the sets of linearly encoded n-length permutations and n-by-n permutation matrices, respectively. Hence, the conversion procedures realize an isomorphism between $\mathbb{S}_n$ and $\mathbb{P}_n$.

3 Doubly Stochastic Matrices

A Doubly Stochastic Matrix (DSM) is a matrix $D = [d_{ij}]_{n \times n}$ of non-negative real numbers such that each one of its rows and columns sums to 1, i.e., $d_{ij} \geq 0$ and $\sum_{i=1}^{n} d_{ij} = \sum_{j=1}^{n} d_{ij} = 1$, for all the items $i, j \in [n]$. We denote by $\mathbb{D}_n$ the set of all the DSMs of order n.

It is easy to see that any permutation matrix is also a DSM but not vice versa, thus $\mathbb{P}_n \subset \mathbb{D}_n$. In particular, according to the Birkhoff-von Neumann (BvN) theorem [16], $\mathbb{D}_n$ defines a polytope which is the convex hull of $\mathbb{P}_n$. As a consequence, any DSM can be written as a convex combination of permutation matrices. Formally, given $D \in \mathbb{D}_n$, there exist k permutation matrices $P_1, \ldots, P_k \in \mathbb{P}_n$ and k weights $w_1, \ldots, w_k \in (0, 1]$ summing up to 1 (i.e., $\sum_{i=1}^{k} w_i = 1$) such that

$$D = w_1 P_1 + w_2 P_2 + \ldots + w_k P_k. \tag{1}$$

The right-hand side of Eq. (1) is said to be a decomposition of D and it can be computed by the so-called Birkhoff algorithm [17], which is known to return a decomposition whose length k is at most $n^2 - 2n + 2$, i.e., $k = O(n^2)$.

By definition, any row and any column of a DSM is a multinomial[1] distribution over the set $[n]$ of row/column indices. This property makes DSMs particularly appealing to build models for the AB (or BA) representation in the context of permutation matching problems.

---

[1]Also called "multinoulli distribution" in [18].

As a prominent example of matching problems, let us consider the QAP, where it is required to match a set of n facilities to a set of n locations (i.e., the sets *A* and *B* of the AB representation). Both the facilities and the locations can be encoded by the elements of $[n]$. However, note that QAP, being a matching problem, does not define any ordering, neither in the set of facilities nor in the set of locations, thus the elements in $[n]$ –though being cardinal numbers– have to be intended as identification numbers without relying on any ordering relation among them. In this context, we can design a DSM $D \in \mathbb{D}_n$ such that:

- the rows and columns of D represent, respectively, facilities and locations;
- the i-th row of D is a multinomial distribution representing the probabilities of assigning any location to the facility i;
- the j-th column of D is a multinomial distribution representing the probabilities of assigning any facility to the location j.

The points above clearly show that DSMs are suitable probability models for matching problems, being able to coherently encode both facilities and locations' distributions without relying on any ordering relation among them.

4 Learning and Sampling DSMs

The BvN theorem previously presented clearly states that a DSM can be seen as the aggregation of a number of permutations (encoded in the form of permutation matrices) such that only the matching or assignment nature of the permutations is considered. This property makes particularly appealing the use of DSMs as models for permutation matching problems. Therefore, here below we analyze different methodologies for learning a DSM from a set of permutations and, vice versa, sampling permutations from a DSM.

4.1 Learning DSMs

In order to learn a DSM from a set of permutations we devise two different strategies: *exact learning* and *smoothed learning*, of which the latter is a generalization of the former.

Exact Learning. Given m permutations $P_1, \ldots, P_m \in \mathbb{P}_n$, then a DSM $D \in \mathbb{D}_n$ can be learned as a convex combination of them. Formally,

$$D = w_1 P_1 + w_2 P_2 + \ldots + w_m P_m, \tag{2}$$

where $w_1, \ldots, w_m$ are m non-negative weights summing to 1, that can be uniformly set to $1/m$ or, alternatively, they can be made proportional to some importance measure of the corresponding permutations.

According to the BvN theorem (see Sect. 3), D can also be interpreted as the (weighted) centroid of $P_1, \ldots, P_n$ in the Birkhoff polytope.

As a simple consequence of the convex combination, if the generic entry (i, j) is null in all the permutation matrices $P_1, \ldots, P_m$, then also the entry (i, j) of D is null, and vice versa. As it will be clear from the next subsection, this property makes it impossible to sample from D, permutations which assign the facility i to the location j.

Smoothed Learning. In order to address the "null probability" drawback of the *exact learning* scheme, we slightly modify Eq. (2) by introducing –as a further term of the convex combination– the uniform DSM $U = [u_{ij}]_{n \times n}$ such that $u_{ij} = 1/n$ for every $i, j \in [n]$. Formally, a DSM D can now be learned as follows.

$$D = w_1 P_1 + w_2 P_2 + \ldots + w_m P_m + \alpha U, \tag{3}$$

where $\alpha \in (0, 1]$ is a smoothing factor which regulates the importance of U in the convex combination, while now $\sum_{i=1}^{m} w_i = 1 - \alpha$. Therefore, in the uniform setting, the weights are set to $(1 - \alpha)/m$.

The learning scheme of Eq. (3) produces a DSM without 0-entries, thus making possible to sample permutations containing items' assignments not present in the training set. Since this aspect is relevant for avoiding the premature convergence of an estimation-of-distribution algorithm based on the DSM model, in this work we will adopt the *smoothed learning* scheme.

4.2 Sampling DSMs

In order to sample permutations from a given DSM we devise three different strategies:

- *probabilistic sampling*, which allows to define a proper probability mass function;

- *algebraic sampling*, that exploits the multiplication of the DSM by a random vector;
- *geometric sampling*, which is based on the BvN theorem.

These strategies are described in the following.

Probabilistic Sampling (PS). For the sake of simplicity, in order to describe this sampling procedure, we use the linear permutation encoding. Therefore, given $D \in \mathbb{D}_n$, a permutation $\sigma \in \mathbb{S}_n$ is sampled from D as follows.

1. Select uniformly at random a row or a column of D and denote by $i \in [n]$ and $p \in [0,1]^n$ its index and the corresponding probability vector, respectively.
2. Since p is a multinomial probability distribution, sample an item $j \in [n]$ according to p.
3. If p is a row of D, then set $\sigma_i := j$, otherwise set $\sigma_j := i$.
4. Remove the row i (or j) and the column j (or i) of D when p is a row (or a column), then renormalize the remaining rows and columns.
5. If not all the rows and columns of D have been removed, go back to step (1).

It is easy to see that any iteration of this procedure sets exactly one entry of σ and removes exactly one row and one column from D. This guarantees that, after n iterations, σ is a proper permutation of $\mathbb{S}_n$. Moreover, since the computational cost of a single iteration is $\Theta(n)$, sampling a permutation with the PS costs $\Theta(n^2)$ time steps.

For a given n-by-n DSM $D = [d_{ij}]$, the probability of sampling a permutation $\sigma \in \mathbb{S}_n$, via PS, is given by

$$\Pr(\sigma|D) = \frac{\prod_{i=1}^n d_{i,\sigma_i}}{\sum_{\pi \in \mathbb{S}_n} \prod_{i=1}^n d_{i,\pi_i}} = \frac{\prod_{i=1}^n d_{i,\sigma_i}}{\mathrm{Perm}(D)}, \qquad (4)$$

where the denominator is the permanent of the matrix D [19].

The probability mass function of Eq. (4) allows to derive two interesting properties of PS as follows. By recalling that permutation matrices are particular cases of DSMs, when the current DSM is a permutation matrix P, then only P can be sampled. At the other extreme, when the current DSM is the uniform DSM U (introduced in Sect. 4.1), all the permutations are equiprobable.

Algebraic Sampling (AS). This sampling strategy is based on the "randomized rounding" methodology introduced in [20] (which, in turn, extends the methods described in [21] and [22]).

The idea is that, to sample a permutation matrix $P \in \mathbb{P}_n$ from a $D \in \mathbb{D}_n$, we first generate a vector $v \in [0,1]^n$ uniformly at random and then we obtain P by solving the equation

$$P \cdot \mathrm{rank}(v) = \mathrm{rank}(D \cdot v), \qquad (5)$$

where $\cdot$ is the usual matrix-vector multiplication, while the vector $\mathrm{rank}(v)$ is defined as $\mathrm{rank}(v)_i = j$, where v_i is the j-th smallest value in v. For example, if $v = (0.1, 0.5, 0.8, 0.2)$, then $\mathrm{rank}(v) = (1, 3, 4, 2)$. Practically, $\mathrm{rank}(v)$ returns the inverse permutation of $\mathrm{argsort}(v)$, as the well-known random key transformation for permutations [23].

Therefore, by using the linear permutation encoding, a permutation $\sigma \in \mathbb{S}_n$ is sampled from a DSM $D \in \mathbb{D}_n$ as follows.

1. Generate a vector $v \in [0,1]^n$ uniformly at random.
2. Calculate the permutations $\pi = \mathrm{rank}(D \cdot v)$ and $\rho = \mathrm{argsort}(v)$.
3. Returns the permutation $\sigma = \pi \circ \rho$ (where $\circ$ is the usual permutation composition).

In [20] it has been proved that the linear permutation σ, computed as above, corresponds to the permutation matrix P^* which solves the following optimization problem:

$$P^* = \mathrm{argmin}_{P \in \mathbb{P}_n} ||D \cdot v - P \cdot v||_F^2, \qquad (6)$$

where $|| \cdot ||_F$ is the usual Frobenius norm of a matrix.

Eq. (6) provides an intuitive justification to AS. Moreover, it also shows that, as for the PS case, when v entries are all different[2] and the current DSM is a permutation matrix, then only this permutation can be sampled. Conversely, when the DSM is the uniform one, all the permutations have the same probability to be sampled.

---

[2]Since v entries are sampled independently and uniformly at random, they are all different with probability 1.

Finally, the computational complexity of the procedure is $\Theta(n^2)$, which is given by the matrix-vector multiplication in step (2). However, it is worthwhile to note that the matrix-vector multiplication can benefit of very fast implementations, thus making AS very efficient in practice.

Geometric Sampling (GS). Given a DSM $D \in \mathbb{D}_n$, it is possible to sample a permutation $P \in \mathbb{P}_n$ –expressed in the form of a permutation matrix for convenience of description– by means of the following steps.

1. Execute the Birkhoff algorithm [17] and obtain a decomposition of D with length $k = O(n^2)$.
2. The decomposition has the form of Eq. (1), hence sample an index $i \in [k]$ according to the decomposition weights $w_1, \ldots, w_k$.
3. Select and return the permutation matrix P_i from the computed decomposition.

Substantially, the returned permutation is sampled from the vertices of the Birkhoff polytope by implicitly considering probabilities which are inversely proportional to the geometric distances of each vertex from the point corresponding to D.

As in the PS and AS cases, when the DSM is a permutation matrix, its Birkhoff decomposition is formed only by itself, thus it is the only permutation that can be sampled. However, when the DSM is the uniform DSM, since the Birkhoff algorithm is not randomized (or, at least, not yet), not all the permutations are equiprobable. In fact, an easy to see drawback of GS is that the sampling domain is limited by BvN theorem, i.e., no more than $n^2 - 2n + 2$ different permutations can be sampled.

Furthermore, GS is not efficient at all. Indeed, the computational complexity of GS is dominated by the Birkhoff algorithm which, as described in [17], is a greedy iterative method that, at any iteration, calculates a perfect matching in a bipartite graph formed by $2n$ vertices. The most used solver for bipartite matching is the Hopcroft-Karp algorithm [24], whose computational complexity is $O(n^{2.5})$. Note also that, due to the BvN theorem, the Birkhoff algorithm requires $O(n^2)$ iterations, therefore GS has the very high computational complexity of $O(n^{4.5})$.

5 The proposed EDA

In this section we introduce an Estimation of Distribution (EDA) for permutation matching problems whose probability model is designed as a DSM.

We adopt the *smoothed learning* scheme introduced in Sect. 4.1 and the three sampling strategies introduced in Sect. 4.2. Therefore, we have three different implementations of our EDA to which we refer with the terms DSM-PS, DSM-AS and DSM-GS, on the basis of the sampling scheme adopted.

Apart from the sampling strategy, all the implementations share the same base algorithm whose pseudocode is provided in Alg. 1.

Algorithm 1 Algorithmic scheme of the proposed EDA

Input: $f : \mathbb{S}_n \to \mathbb{R}$, $\mathrm{SS} \in \{\mathrm{PS}, \mathrm{AS}, \mathrm{GS}\}$

1: $\lambda \leftarrow 10n$ ▷ Sample size
2: $\mu \leftarrow n$ ▷ Selection size
3: $\alpha \leftarrow 1/n^2$ ▷ Smoothing factor
4: $t \leftarrow 0$
5: $X_0 \leftarrow$ a set of λ permutations drawn uniformly at random
6: **while** the budget of evaluations is not exhausted **do**
7:      $t \leftarrow t + 1$
8:      $Y \leftarrow$ select the best μ permutations from X_{t-1}
9:      $D \leftarrow$ a DSM learnt from Y with smoothing factor α
10:      $Z \leftarrow$ a set of λ permutations sampled from D using SS
11:      evaluate $f(\sigma)$ for all $\sigma \in Z$
12:      $X_t \leftarrow Y \cup Z$
13: **end while**
14: **return** the best permutation in X_t

Alg. 1 receives in input the objective function to optimize and the desired sampling strategy to be chosen among PS, AS, or GS.

In lines 1–3, three algorithmic parameters are set on the basis of the problem size n and according to previous experiences in the field of EDAs [7, 8]. Namely, a high selection pressure is considered by setting the sample size λ

and the selection size μ to, respectively, $10n$ and n. The smoothing factor α is set to $1/n^2$ in order to make possible to escape stagnation states, though without making the search totally random.

Then, a set X of solutions is initialized with λ random permutations in line 5 and iteratively updated in the main loop of lines 6–13. At any iteration of the loop, the best μ solutions in X (line 8) are used to learn the DSM model D (line 9), from which λ new solutions are sampled (line 10). The sampled solutions are evaluated (line 11) and X is updated by merging the previously selected solutions with the new generated ones (line 12).

Note that, in the learning step, apart from α, all the weights for the solutions in the training set are set in a uniform way.

Moreover, the iterative update scheme ensures an elitist behaviour to the algorithm. Indeed, the set X always contains the best solution ever visited, which is finally returned in line 14. Furthermore, the elitism guarantees that, at any iteration t, the DSM model D is always learnt from the best μ solutions ever visited by the algorithm.

The elitist behaviour and the high selection pressure are counterbalanced by the smoothed learning which allows to generate unseen items' matchings with a non-null probability.

As discussed, learning and sampling are carried out by using, respectively, the smoothed learning scheme and the chosen sampling strategy. Therefore, the complexity of an iteration is dominated by the sampling and evaluation steps.

6 Experimental Study

In order to analyze the effectiveness and efficiency of the explored solutions for the sampling and learning methods of DSMs, in the following we present a thorough experimental study.

6.1 Experimental setting

As exposed in the introduction, the present work is oriented to matching/assignment type problems and, thus, for the purpose of the experimentation, we have chosen the Quadratic Assignment Problem (QAP) as the benchmark problem. In the QAP, we are given two square matrices $\mathbf{B} = [b_{ij}]_{n \times n}$ and $\mathbf{H} = [h_{ij}]_{n \times n}$ of parameters and the goal is to find the permutation $\sigma \in \mathbb{S}_n$ that minimizes the objective function

$$f(\sigma) = \sum_{i=1}^{n} \sum_{j=1}^{n} b_{i,j} h_{\sigma(i)\sigma(j)}.$$

Twelve instances of a variety of sizes from the QAPLIB [25] have been chosen to illustrate the experiments, and the selection of the instances was made prior to observing any performance results. In particular, the set of selected instances is formed by: *tai15{a/b}, tai20{a/b}, tai30{a/b}, tai50{a/b}, tai80{a/b}* and *tai100{a/b}*.

In order to evaluate the usability of DSMs with optimization purposes, the three sampling procedures described in Sect. 4.2 –namely, PS, AS, and GS– are considered. To have a real view of their performance, we will also consider three other competitor EDA algorithms from the literature:

- Mallows EDA under the Cayley distance [13]. This is the main competitor of the proposed algorithms in this work.
- Mallows EDA under the Kendall distance [7]. Due to the nature of the metric employed, this algorithm should not be suited to deal with matching problems.
- Plackett-Luce EDA [6] has been acknowledged for being an algorithm to deal with ordering problems, and thus, as the previous, should not be very effective in matching problems.

The labels used to denote them are EDA-MC, EDA-MK and EDA-PL, respectively. Similarly, we distinguish among the DSM algorithms, as DSM-PS and DSM-AS with respect to the sampling method.

Regarding the geometric sampling, preliminary runs point out that DSM-GS the method is not practicable in this framework. In fact, it was observed that the GS sampling strategy mostly produces the same permutations which have been used to learn the DSM, even with very large values for the smoothed factor α. This in turn makes DSM-GS produce very rare improvements in the best-so-far solution. Therefore, by also considering the very high computational complexity of GS (see Sect. 4.2), we decided to omit it in the presentation of the experimental results, deferring a more accurate analysis to future work.

A general parameter settings has been decided for all the algorithms, following the guidelines of the EDA competitors in the literature. Hence, as also discussed in Sect. 5, the sample size and selection size are set to $\lambda = 10n$ and $\mu = n$, respectively.
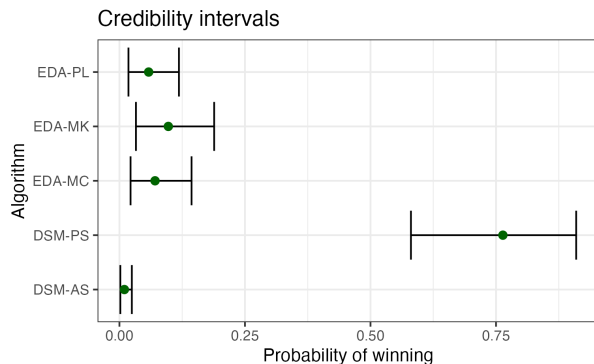
Figure 2: Credibility intervals of the evaluated algorithms on the set of benchmark instances. The intervals describe for each algorithm the probability of being the best option, based on the experimental data provided.

6.2 Effectiveness

In order to obtain a general view of the effectiveness of the DSM type algorithms when compared to the EDAs mentioned above, we executed each algorithm for 20 repetitions on the selected set of benchmark instances, while a budget of $100n^2$ evaluations is allowed for each execution. Results are summarized in Table 1 as Median Relative Deviations (MRD) with respect to the best known results reported in the QAPLIB website[3]. Surprisingly, results show that there is one algorithm that obtains much better results than the rest of the algorithms: DSM-PS. The rest of the algorithms seem to be worse than DSM-PS and have similar performance among them.

Table 1: Results of the EDA and DSM algorithms for the 12 QAP instances from the QAPLIB benchmark (from the Taillard set). The Median Relative Deviations (MRD) measures of the values found across the 20 repetitions of the best known results are reported. Results in bold highlight the algorithm that obtained the lowest MRD. A maximum number of $100n^2$ evaluations were performed by each of the algorithms.

Instance	Best Known	EDA-PL	EDA-MK	EDA-MC	DSM-PS	DSM-AS
tai15a	388214	0.09283	0.09039	0.06565	**0.04145**	0.09659
tai15b	51765268	0.01379	0.00960	0.00510	**0.00455**	0.01520
tai20a	703482	0.13623	0.13740	0.13293	**0.05471**	0.13805
tai20b	122455319	0.11341	0.05344	0.04272	**0.01807**	0.11903
tai30a	1818146	0.12511	0.12666	0.13055	**0.04783**	0.13016
tai30b	637117113	0.27583	0.15849	0.31548	**0.07563**	0.31814
tai50a	4938796	0.14096	0.14025	0.13820	**0.04965**	0.14048
tai50b	458821517	0.35262	0.09779	0.36341	**0.03915**	0.37153
tai80a	13499184	0.12401	0.12492	0.12404	**0.03945**	0.12495
tai80b	818415043	0.34160	0.31773	0.34531	**0.03729**	0.34593
tai100a	21052466	0.11712	0.11668	0.11664	**0.03480**	0.11713
tai100b	1185996137	0.32555	0.12570	0.33081	**0.01785**	0.33082

To statistically assess the differences among the different alternatives (combination of parameters), a Bayesian performance analysis was carried out [26]. The outcome of the analysis is summarized in Fig. 2 in the form of credibility intervals.

In the y-axis of the plot the different algorithms are listed, and for each case, a credibility interval is depicted. Such interval, formed with a green dot (the expectancy) and a range of values, describes the probability of that algorithm being the best alternative among the compared ones. The values needed to build the intervals are obtained by sampling the posterior distribution of the Bayesian model computed[4].

---

[3]https://coral.ise.lehigh.edu/data-sets/qaplib/qaplib-problem-instances-and-solutions/

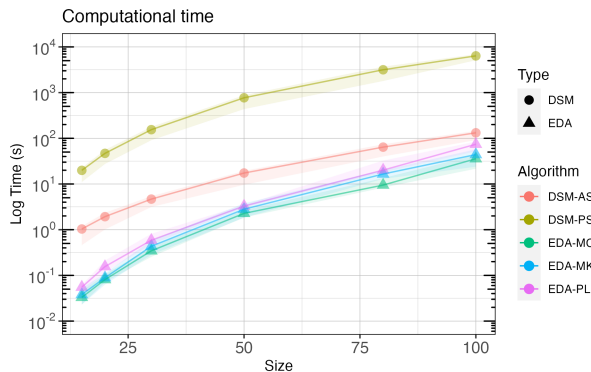[4]A detailed procedure for the analysis is provided in [26].

Figure 3: Average execution time (log seconds) of each run of the EDA-PL, EDA-MK, EDA-MC, and DSM-AS, DSM-PS algorithms for running $100n^2$ evaluations.

The analysis confirms that DSM-PS is one step ahead of the other algorithms, however, the width of the intervals, especially for EDA-PL and DSM-PS suggests that there is still much uncertainty related to previous statement. Taking into account that only 12 instances were used for the analysis, it becomes obvious that additional data is required to have more reliable conclusions. However, this results already point that DSMs may be very valuable for optimization purposes, at least in the context of matching problems.

6.3 Efficiency

We have seen that DSMs can offer superior performance to some competitors in the literature, but, what about their computational efficiency? In this section, we aim to evaluate the computational time required by the algorithms to run $100n^2$ evaluations and how it scales with the problem size. At this point, it is important to remark that EDAs were implemented in C++ while, and DSM versions were coded in Python, so this aspect needs to be taken into account when doing the analysis[5].

Average results of the computational time consumed by each of the compared algorithms are depicted in Figure 3. As the number of evaluations to perform is a function of the problem size, we have plotted a line for each algorithm for the different problem sizes considered in the instance set. Results reveal that EDAs (all of them implemented in C++), have much lower computational cost compared to the proposed DSM-based algorithms (note that y-axis is in log scale). In fact, the best performing algorithm in the previous section, DSM-PS, is by far the most time-consuming method. The DSM algorithms are prototype versions that have been implemented in Python and the room for improvement is high, and therefore the chances of approaching the performance of EDAs is real.

This last statement can be further substantiated by the fact that Figure 3 shows that the scalability of the algorithms is similar, as the slopes of the line plots are comparable.

6.4 Sampling convergence

As a way to analyze how effective are the different algorithms when modelling and capturing the relevant information of the candidate solutions for the problem, in this section, we analyze the quality of the solutions sampled at each iteration of the algorithm. We do not focus on the population itself, since its quality is usually monotonically increasing, but we aim to see the solutions sampled. In an appropriate model, we expect to see that the quality of the sampled solutions improves across the iterations.

To that end, we have conducted an alternative experiment on the instance *tai50a* where every solution sampled during the execution by each algorithm has been recorded. Then, the objective value of the samples at each iteration is averaged, and to avoid randomness, 5 repetitions for each algorithm were performed. Results are presented in the form of scatter and smooth line plots in Figure 4.

On the one hand, we can observe that the EDAs show almost no convergence, and they do not move from (normalized) objective values within the range $[0.19, 0.2]$. With respect to the DSMs, the results are very variable. DSM-PS shows a

---

[5]A comparison of the energy efficiency of 27 programming languages is carried out in [27]. The analysis reveals that while C and C++ are at the top of the ranking, Python usually is the one of the worst ranked, if not the worst.
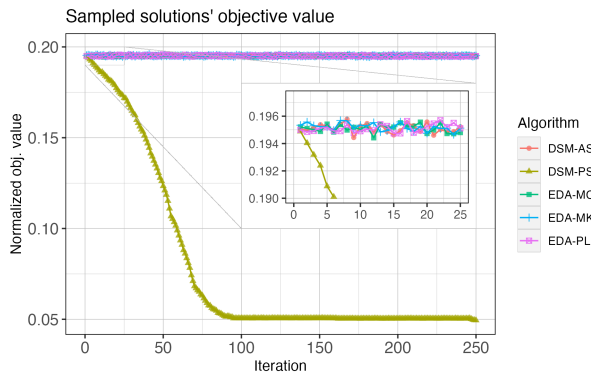
Figure 4: Average objective function value of the solutions sampled at each iteration of EDA-PL, EDA-MK, EDA-MC, and DSM-AS, DSM-PS algorithms for running $100n^2$ evaluations. Each algorithm was run 5 repetitions. The instance used in this plot, is *tai50a*. x-axis has been limited to 250 iterations.

very particular shape where we see that almost since the first iterations, the algorithm starts progressively improving the results. This cannot be generalized to DSMs, as the behaviour of DSM-AS seems equivalent to other EDAs.

7 Conclusion and Future Work

In this paper, we investigated the use of Doubly Stochastic Matrices (DSMs) in the framework of evolutionary algorithms and focused on Estimation of Distribution Algorithms (EDAs) for permutation-based combinatorial optimization problems.

We consider that existing EDAs are not suited for solving matching or assignment-type permutation problems, as most of them assume an ordering nature of the permutations. In this context, we think that DSMs can have good performance. To that end, we analyzed the potential use of DSMs within EDAs, by exploring different learning and sampling methods. In particular, we designed a simple learning process by exploiting the Birkhoff-von Neumann theorem, which is a well known result in the algebraic field. Regarding the sampling of candidate solutions (permutations) from a DSM, we found multiple ways to do so. In this paper, we studied three sampling strategies –probabilistic, algebraic, and geometric– with characteristics and variable time complexities.

The conducted experiments on a set of instances of the quadratic assignment problem reveal an interesting scenario. When compared with already published EDA algorithms for permutations problems, we observe that our DSM-based EDAs, particularly, under the probabilistic sampling, obtain very good, and even better, results.

This work intends to explore the potential of doubly stochastic matrices within the framework of estimation of distribution algorithms, and thus, there is a lot to improve in their design, integration, and fine-tuning. On the one hand, the computational time required by these algorithms needs still further research to make them more competitive. The fact that the DSM-based algorithms were coded in Python, while the competitor EDAs were implemented in C++, suggests that the efficiency gap observed can be narrowed down. On the other hand, DSM-based algorithms can be further refined by conducting a deeper analysis of the parameters' setting (now hard-coded in the algorithm) and by also studying a novel weighting scheme (now uniform) for the learning phase.Moreover, it is also interesting to explore other sampling methodologies, such as the one based on the Sinkhorn-Knopp algorithm which has been recently adopted in the field of machine learning [28].

Finally, in this paper, we exclusively focused on implementing EDAs, as it was straightforward, but doubly stochastic matrices can have other uses such as, in the context of genetic algorithms, for designing probabilistic crossover operators tailored to permutation matching problems or, in the context of model-based gradient search algorithms, for designing differentiable Mallows-like models where the mode permutation is relaxed to a doubly stochastic matrix.

Acknowledgments

References

[1] Richard P. Stanley. *Enumerative Combinatorics*. Wadsworth Publ. Co., Belmont, CA, USA, 1986.

[2] Paolo Toth and Daniele Vigo. *The vehicle routing problem*. Society for Industrial and Applied Mathematics, 2001.

[3] Jatinder N. D. Gupta and Edward F. Stafford. Flow shop scheduling research after five decades. *European Journal of Operational Research*, 169(3):699–711, 2006.

[4] Tjalling C. Koopmans and Martin J. Beckmann. Assignment Problems and the Location of Economic Activities. Cowles Foundation Discussion Papers 4, Cowles Foundation for Research in Economics, Yale University, 1955.

[5] Valentino Santucci and Marco Baioletti. A fast randomized local search for low budget optimization in black-box permutation problems. In *2022 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8, 2022.

[6] Josu Ceberio, Alexander Mendiburu, and Jose A. Lozano. The plackett-luce ranking model on permutation-based optimization problems. In *2013 IEEE Congress on Evolutionary Computation (CEC), Cancun, Mexico*, pages 494–501. IEEE, 2013.

[7] Josu Ceberio, Ekhine Irurozki, Alexander Mendiburu, and Jose A. Lozano. A Distance-based Ranking Model Estimation of Distribution Algorithm for the Flowshop Scheduling Problem. *IEEE Transactions on Evolutionary Computation*, 18(2):286 – 300, April 2014.

[8] Ekhine Irurozki, Josu Ceberio, Borja Calvo, and Jose A. Lozano. Mallows model under the ulam distance: a feasible combinatorial approach. In *Workshop Book 2014 NIPS Conference*, Montreal, Quebec, Canada, December 2014.

[9] Mayowa Ayodele, John McCall, and Olivier Regnier-Coudert. Rk-eda: A novel random key based estimation of distribution algorithm. In Julia Handl, Emma Hart, Peter R. Lewis, Manuel López-Ibáñez, Gabriela Ochoa, and Ben Paechter, editors, *Parallel Problem Solving from Nature – PPSN XIV*, pages 849–858, Cham, 2016. Springer International Publishing.

[10] Valentino. Santucci, Josu Ceberio, and Marco Baioletti. Gradient search in the space of permutations: an application for the linear ordering problem. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, pages 1704–1711, 2020.

[11] Marco Baioletti, Alfredo Milani, and Valentino Santucci. Variable neighborhood algebraic differential evolution: An application to the linear ordering problem with cumulative costs. *Information Sciences*, 507:37–52, 2020.

[12] Arthur Guijt, Ngoc Hoang Luong, Peter AN Bosman, and Mathijs de Weerdt. On the impact of linkage learning, gene-pool optimal mixing, and non-redundant encoding on permutation optimization. *Swarm and Evolutionary Computation*, 70:101044, 2022.

[13] Ekhine Irurozki, Borja Calvo, and Jose A Lozano. Sampling and learning mallows and generalized mallows models under the cayley distance. *Methodology and Computing in Applied Probability*, 20(1):1–35, 2018.

[14] Marco Baioletti, Alfredo Milani, Valentino Santucci, and Marco Tomassini. Search moves in the local optima networks of permutation spaces: the qap case. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1535–1542, 2019.

[15] Marco Baioletti, Gabriele Di Bari, Alfredo Milani, and Valentino Santucci. An experimental comparison of algebraic crossover operators for permutation problems. *Fundamenta Informaticae*, 174(3-4):201–228, 2020.

[16] Fanny Dufossé and Bora Uçar. Notes on birkhoff–von neumann decomposition of doubly stochastic matrices. *Linear Algebra and its Applications*, 497:108–115, 2016.

[17] Diane M Johnson, AL Dulmage, and NS Mendelsohn. On an algorithm of g. birkhoff concerning doubly stochastic matrices. *Canadian Mathematical Bulletin*, 3(3):237–242, 1960.

[18] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

[19] David G Glynn. The permanent of a square matrix. *European Journal of Combinatorics*, 31(7):1887–1891, 2010.

[20] RJ Wolstenholme and Andrew T Walden. A sampling strategy for projecting to permutations in the graph matching problem. *arXiv preprint arXiv:1604.04235*, 2016.

[21] Fajwel Fogel, Rodolphe Jenatton, Francis Bach, and Alexandre d'Aspremont. Convex relaxations for permutation problems. *Advances in neural information processing systems*, 26, 2013.

[22] Alexander Barvinok. Approximating orthogonal matrices by permutation matrices. *Pure and Applied Mathematics Quarterly*, 2(4):943–961, 2006.

[23] Valentino Santucci, Marco Baioletti, and Alfredo Milani. An algebraic framework for swarm and evolutionary algorithms in combinatorial optimization. *Swarm and Evolutionary Computation*, 55:100673, 2020.

[24] John E Hopcroft and Richard M Karp. An n^5/2 algorithm for maximum matchings in bipartite graphs. *SIAM Journal on computing*, 2(4):225–231, 1973.

[25] Rainer E. Burkard, Stefan E. Karisch, and Franz Rendl. QAPLIB – a quadratic assignment problem library. *Journal of Global Optimization*, 10(4):391–403, 1997.

[26] Jairo Rojas-Delgado, Josu Ceberio, Borja Calvo, and Jose A. Lozano. Bayesian performance analysis for algorithm ranking comparison. *IEEE Transactions on Evolutionary Computation*, pages 1–1, 2022.

[27] Rui Pereira, Marco Couto, Francisco Ribeiro, Rui Rua, Jácome Cunha, João Paulo Fernandes, and João Saraiva. Ranking programming languages by energy efficiency. *Science of Computer Programming*, 205:102609, 2021.

[28] Gonzalo Mena, David Belanger, Gonzalo Munoz, and Jasper Snoek. Sinkhorn networks: Using optimal transport techniques to learn permutations. In *NIPS Workshop in Optimal Transport and Machine Learning*, volume 3, 2017.