



Discussion Paper

Redistribution of mass and probabilities

Léon Willenborg

20 January 2023

1 Introduction

This paper¹⁾²⁾ is about optimal transportation, of mass (real) or probability mass (virtual). Not so much about the solution of such problems, but about their formulation. Optimal transportation models have been considered since quite some time, starting with a contribution by Monge (see [12]), who considered the problem of transporting soil between several locations in an optimal way. Later, in the 20th century, Kantorovitch (see [8] and [9]) and Koopmans (see [10]) have made important pioneering contributions to this area.

The area of optimal transportation has turned out to be a connection point between economics and physics. In the latter area physical transport problems are studied in areas such as fluid mechanics and plasma physics. The Boltzmann equation describes the motion of particles in a fluid or a gas (a compressible fluid), and the Vlasov equation describes the motion of a plasma, which is an electrically charged, hot gas. Both economics and physics have an interest in transporting matter (goods in economics and mass in physics) in an optimal way, by minimizing costs or energy, distance, time, etc., under certain constraints which are particular for the problem at hand. See e.g. [19] and [20].

Applying similar ideas as in transporting real mass in an optimal way, one can also consider the problem how to transform a probability on a certain space into another problem on the same space. Probability mass is handled as if it is real matter. For theory and applications in this area see [14] and [15].

We are now in the position to explain what our interest in the problem is. It is not necessarily related to transportation economics (logistics) or physics. Our source of inspiration is (official) statistics. In certain areas of official statistics one runs into the problem of disaggregating two (or more) distributions. For example, two aggregate tables $S \geq 0$ and $D \geq 0$ are given and it is necessary to find a $2d$ -table U such that the marginals of U are S and D .

Assuming such a U exists, it has the property that

$$\begin{aligned} U(i, \cdot) &= S(i), \\ U(\cdot, j) &= D(j), \end{aligned} \tag{1}$$

where ‘ \cdot ’ denotes aggregation over appropriate categories.

To be consistent the sum of the elements of S and D need to be the same, which we express symbolically as

$$S(\cdot) = D(\cdot). \tag{2}$$

¹⁾ The views expressed in this paper are those of the author(s) and do not necessarily reflect the policies of Statistics Netherlands.

²⁾ The author is grateful to Sander Scholtus for reviewing a mature draft of the paper. He made several suggestions to improve the paper, which the author gladly accepted.

In general, such a U does not need to exist. And if it exists it does not need to be unique. Extra conditions or requirements are then needed to imply a unique U .

In principle, the sizes of S and T – viewed as vectors – need not be the same. In the application considered in the present paper they are the same, however. Not only the sizes of the vectors are the same, also the meaning of the categories. In the applications we have in mind they can be locations or sectors in the economy. In the first case, we are dealing with physical locations and in the second one with abstract entities that cannot be naturally associated with a metric space.

In case there is a solution to our disaggregation problem, and there are in fact several solutions, we wish to select a specific disaggregation method, based on the transportation model, which is a standard problem from operations research (see e.g. [5]). But this presupposes that we should have information about the (dis)similarity of the categories used in the model, or their distance. As the transportation model itself is fairly standard, the focus of the attention should be on these parameters in the model. In the present paper we consider several examples in which we try to specify reasonable values for them.

In the approaches we describe we can distinguish two kinds of mass: physical mass and probability mass. Changing the first requires transport of matter, so questions about doing this efficiently are raised automatically. Changing the latter does not require any physical activity. In this case we are dealing with information.

The remainder of the paper is organized as follows. In Section 2 transportation models are presented. These can be used to ship goods from a number of locations to other locations at minimal transportation costs. An important quantity in such models is the cost matrix, which defines the cost of transporting a unit of the good considered from location i to location j , for all possible location pairs (supply location, demand location). In Section 3 some examples of such cost matrices are shown. In Section 4 we use copulas as weight matrices. Copulas are devices that express the interdependency of random variables. In Section 5 we discuss the application of the well-known technique of Iterative Proportional Fitting (IPF) to a 2D table to find the interior of a table whose marginals are known, together with a weight matrix for the interior. In Section 6 there is a brief discussion of the approaches considered to solve the problem central in this paper. The paper is completed with a list of references and two appendices, one on the IPF method and the other on copula theory.

2 Transportation models

2.1 Base model

We present here a special case of the well-known optimization model from OR³⁾, the transportation model, in particular the Hitchcock transportation model⁴⁾, named after the author of [7]. This model is about optimally transporting a certain product from a number of factories to a number of shops, with minimal costs. Let V denote the set of factories and W the set of shops. For factory i and shop j the cost to ship the product from i to j is c_{ij} . Let $C = (c_{ij})$ be the $m \times n$ cost matrix, with c_{ij} denoting the cost of transporting one unit of the good from i to j .

The transportation model can be formulated as the following optimization problem:

$$\begin{aligned} \min \quad & \sum_{i \in V} \sum_{j \in W} c_{ij} x_{ij} \\ & \sum_{j \in W} x_{ij} = s_i \text{ for } i \in V, \\ & \sum_{i \in V} x_{ij} = d_j \text{ for } j \in W, \\ & x_{ij} \geq 0 \text{ for } (i, j) \in V \times W \end{aligned} \tag{3}$$

$(s_1, \dots, s_m) \geq 0$, with $m = |V|$ is the vector of supplies and $(d_1, \dots, d_n) \geq 0$ with $n = |W|$ is the vector of demands. We assume that total supply equals total demand, i.e.

$$\sum_{i \in V} s_i = \sum_{j \in W} d_j = t, \tag{4}$$

with $t > 0$ denoting the common total.

Remark The cost matrix C is not unique, in the sense that κC yields the same results, for any $\kappa > 0$. If we divide the supply and demand matrices by t as in (4), we obtain two probability densities:

$$\begin{aligned} \sigma_i &= s_i/t, \\ \delta_j &= d_j/t. \end{aligned} \tag{5}$$

We need this extension to probability distributions in the sequel as well. \square

³⁾ OR = operations research

⁴⁾ Named after Frank Hitchcock (1875–1957). More accurate is to call it the Hitchcock-Koopmans transportation problem, as Tjalling Koopmans (1910–1985) was another pioneer of such models.

Remark The implicit assumption in model (3) is that each supply location i is connected to each demand location j and vice versa. However, this does not need to be the case in any application. For instance, some sectors in an input output model of the economy may not be connected in the sense that there is no flow of goods or services (and hence money) from sector i to sector j . □

2.2 Generalizations

The model presented in Section 2.1 can be generalized in several ways, depending on the treatment of the masses, supplies, demands or transportation links. As to the masses it is important to know whether they can be split or not. In the latter case the entire supply at one site should be transported to another location (provided there is no demand at the supply site). As for the demands at the various locations it is important to know whether deliveries at that location should be equal to the demand at a location, or should not exceed the demand, or should be at least as big as the demand.

An example of the latter case is when a certain amount of bricks is needed for building a house. With less bricks supplied than required, the building work cannot be finished. But, although a surplus (say of bricks) is required for a job to be finished at a site, it should not be too big, as this counts as a loss. To require that demand and supply are exactly equal at a site, let alone at any site, may be a tall order. Consider the situation of the bricks again: they may be delivered pallet-wise, with each pallet containing a fixed number of bricks, say several hundreds of bricks. It is highly unlikely that the number of bricks supplied, pallet-wise, is exactly equal to the number required. Furthermore, the number supplied is preferably bigger than the number demanded, in order to have some spare bricks, in case some may be get broken, or cut wrongly by a bricklayer.

The case in which the demand should not be exceeded can be of importance when the materials distributed are quite expensive, and they are not needed all at once, but in due course, say during a building or construction process. If not enough of the material needed has been delivered at a particular time, the assumption is that additionally required materials can be ordered later. So in this case the tuning of supply and demand is not done in a single step, but in several steps.

Likewise, looking at the supply side, there may be various constraints and objectives. There may be a requirement that the entire supply at a location should be removed, as we mentioned before. Or the supply is divided into various batches (e.g. bags of sand, or pallets of bricks) and only entire numbers of discrete units are supplied to the demand sites. It may also be the case that not all the goods / materials can be sold. The amount that is left is considered a loss. The goal of each supplier is to minimize their loss.

The requirements of each supply location may be different. The same holds for each demand location. And this is for a single application, at the micro-level. The information as to which requirements apply to each location may be lacking. In this case one can only work with a simplified model, making assumptions about the requirements for the supply and demand sites.

We mention this variety of restrictions to indicate that the base model in Section 2.1 is an example of a particular choice of models. In an OR context alternative constraints and or target functions may be of interest. Such alternatives are, however, not so likely to be of importance for the application we are considering in the present paper. The reason is that, in the applications we consider, the problems are formulated at a macro-level, whereas the constraints we consider

(in the present section), are defined at a micro-level. And the requirements of each supplier may be different, even for each case at the micro-level. Similarly for each demand site. This prevents clear-cut constraints at the macro-level.

2.3 Wasserstein metric

The transportation model considered in 2.1 is discrete. However, the original model proposed by Monge (in [12]) is continuous. We can extend the approach in 2.1 to the continuous case. This is possible by using a Wasserstein metric, which yields a class of models parametrized by a parameter p . The generalization to general metric spaces is conceptually straightforward but is technically involved to ensure that certain quantities exist or converge. We do not wish to delve into these technicalities, so we present the formulas involved and assume that the necessary conditions have been imposed so that the quantities in which we are interested, exist.⁵⁾

Let σ, δ be two probability measures defined on a metric space (M, d) . Let $\Gamma(\sigma, \delta)$ be the set of all measures in $M \times M$ with marginals σ and δ , which is called the set of all couplings of σ and δ . The p^{th} Wasserstein distance $W_p(\sigma, \delta)$ between σ, δ is defined as

$$W_p(\sigma, \delta) = \left(\inf_{\mu \in \Gamma(\sigma, \delta)} \int_{M \times M} d(x, y)^p d\mu(x, y) \right)^{1/p}, \quad (6)$$

where μ is a measure on $M \times M$. The case $p = 1$ can be viewed as a direct generalization of the target function in (3):

$$W_1(\sigma, \delta) = \inf_{\mu \in \Gamma(\sigma, \delta)} \int_{M \times M} d(x, y) d\mu(x, y). \quad (7)$$

In the present paper we stick with discrete metric spaces. They are the ones that we need, which is fortunate as they are a lot easier to handle. But it is interesting to note that their generalization leads to (a specific case of) the Wasserstein metric.

2.4 Adopted model

For the application in the present paper we are interested in a special case of model (3), namely that in which $V = W$, so there is no distinction between factories and shops. Then, instead of model (3) we have:

⁵⁾ For a quick overview of the Wasserstein metric see https://en.wikipedia.org/wiki/Wasserstein_metric

$$\begin{aligned}
\min \sum_{i=1}^n \sum_{j=1}^n c_{ij} x_{ij} & \tag{8} \\
\sum_{j=1}^n x_{ij} &= s_i \text{ for } i = 1, \dots, n \\
\sum_{i=1}^n x_{ij} &= d_j \text{ for } j = 1, \dots, n \\
x_{ij} &\geq 0 \text{ for } i, j = 1, \dots, n.
\end{aligned}$$

with $s_i, d_i \geq 0$ for $i = 1, \dots, n$, $\sum_{i=1}^n s_i = \sum_{j=1}^n d_j = t$.

In OR applications the matrix C is supposed to be known and all the effort is devoted to solving the optimization problem. However, the choice of a suitable cost matrix is in a sense the key part of our problem. Solving the resulting optimization model is considered ‘routine’, given the knowledge available in the literature, the software available to solve such problems as well as the enormous experience in practice with such models.

We now consider the C -matrix in more detail. We assume that for any cost matrix that we consider in this paper the diagonal elements are all 0:

$$c_{ii} = 0 \text{ for } i \in V, \tag{9}$$

because shipping a good from a location to the same location does not require any shipping at all and so the shipping costs are then 0.

In practice (9) could be an over-idealization, as the categories could be small areas (instead of points) and the transportation distance within an area could be nonzero, although small, so that (9) has to be replaced by $c_{ii} = \zeta_i$ for $i \in V$, where the $\zeta_i > 0$ are small values, that is, small compared to the off-diagonal elements of C . Furthermore, $C \geq 0$, that is, $c_{ij} \geq 0$ for $i, j \in V$. C defines a metric if $c_{ij} = 0$ implies $i = j$ and vice versa (which is true because of (9)).

From the discussion above it is clear that cost functions and metrics are closely linked. The use of a cost matrix in the models we consider implies that the categories for which it is defined is a metric space if we adopt (9).

3 Examples of cost matrices

We now consider two examples. They differ in the choice of the cost matrix. C . In the first example it is a very crude distance matrix that only distinguishes between the same location or not. In the second example the distance matrix is more refined. Here, some locations are more distant than others to a given location. In fact, the distance matrix we consider corresponds to equidistant locations.

In both examples considered in the present section we use the situation as depicted in Figure 3.1. There are 6 locations and for each location (node) there is a supply (in red) of some stuff (say sand)⁶⁾ as well as a demand (in blue). Note that the total supply equals the total demand (100 units of sand). The goal is to redistribute the sand so that supply equals demand in each location, in such a way that the transportation costs are minimized. The difference between the examples is the cost matrices that are used. A cost matrix specifies the cost of transporting 1 unit of sand from location i to location j .




Figure 3.1 The locations with the supply (red) and demand (blue) of sand per location, measured in tons.

The problems in the Sections 3.1 and 3.2 can be solved with LP software, using the simplex algorithm. However, that is somewhat of an overkill as they can easily be solved by hand, which gives more insight in these kind of problems.

The aim of the first step in both examples is to simplify the problem by looking at supply and demand at each location and make sure that as much mass is kept ‘in situ’ to satisfy the local demand. This does not involve any costs. This gives a new picture of locations and adjusted supplies and demands. The next step is to eliminate those nodes (in the supply or demand set of locations) for which the adjusted supply or the demand is 0. Such locations cannot supply the good, nor do they demand this good, so they are of no interest anymore when the redistribution of the good is involved.

We present a solution in each example in a number of steps. At each stage we show the current vector of supplies (first row) and demands (second row), for all locations. Below each step we mention the amount of mass moved. Some steps could be carried out simultaneously. But we are not interested in that aspect, at the moment; only the transportation costs count. But movement of masses from different locations are occasionally combined into a single step, for the sake of brevity of description. We compute the transportation costs at the end of each procedure, using the cost matrix (16) and the total transported mass between the various locations.

3.1 Example 1: discrete metric

In this case we consider the distance matrix in (10). It favours not to transport any goods, and if stuff has to be transported it does not matter to which location it is transported, as the transportation costs are equal for all destinations from any (other) supply location.

⁶⁾ We consider sand to be a continuous good, that can be divided in arbitrary portions.

$$C = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 \end{pmatrix} \quad (10)$$

Below we give a number of steps showing the state of affairs after some sand has been transported. We give several such steps in order to better show how the sand is actually moved between locations. In fact all these steps could have been combined in a single step. But then one does not see how the sand has actually been moved.

We start with:

$$\begin{pmatrix} 6 & 4 & 10 & 20 & 40 & 20 \\ 17 & 3 & 15 & 9 & 26 & 30 \end{pmatrix} \quad (11)$$

The mass that can remain in the same location is handled first, which means that supply and demand per location should be met as well as possible, as no transportation costs are involved. This means that if we denote demand and supply at a particular location as d and s , respectively, we have the following ‘zero cost’ adjustments:

$$d \leftarrow d - \min\{d, s\} \quad (12a)$$

$$s \leftarrow s - \min\{d, s\}, \quad (12b)$$

Applying (12) to (11) yields:

$$\begin{pmatrix} 0 & 1 & 0 & 11 & 14 & 0 \\ 11 & 0 & 5 & 0 & 0 & 10 \end{pmatrix} \quad (13)$$

The supply of sand in locations 2, 4 and 5 has to be redistributed in such a way that the demands in locations 1, 3 and 6 are satisfied. This can be done in many ways, which all yield the same transportation costs. For instance, we can ship 1 unit of sand from location 2 to location 6, and 11 units from location 4 to location 1. We then get the situation

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 14 & 0 \\ 0 & 0 & 5 & 0 & 0 & 9 \end{pmatrix} \quad (14)$$

In the next step we transport 5 units of sand from location 5 to location 3, and 9 units from location 5 to location 6, and we are finished. That it is possible to split a load of sand to be transported to different locations should be admissible.

We can summarize the totality of transportations in the following matrix:

$$U = \begin{pmatrix} 6 & 0 & 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 & 0 & 1 \\ 0 & 0 & 10 & 0 & 0 & 0 \\ 11 & 0 & 0 & 9 & 0 & 0 \\ 0 & 0 & 5 & 0 & 26 & 9 \\ 0 & 0 & 0 & 0 & 0 & 20 \end{pmatrix} \quad (15)$$

Remark In this example the amounts transported are taken to be integer, because it is convenient computationwise. But there is no need to do so. The transported amounts could just as well be values in \mathbb{Q}^+ . \square

Remark The first step, like the one in the present example resulting in (13), can often be made as an allowable step.⁷⁾ It simply gets rid of ‘auto-supplying’, where part of the goods are used at the same location. In such cases one may assume that the starting position for the transportation problem is such that for each location either supply or demand are non-zero, not both. \square

3.2 Example 2: equidistant locations

The present example is a bit more elaborate than the one in Section 3.1, in that it distinguishes in transportation costs between different locations. So if a good has to be transported from one location to another, one has to consider which pairs of (supply, demand)-locations are the most cost effective ones. In the example in Section 3.1 transporting to any other location is equally expensive, so one only has to consider the demand and supply constraints.

For simplicity, the 6 locations are supposed to be situated on a line, equally spaced, the neighboring locations 1 distance unit (say 1 km) apart. As a suitable cost matrix⁸⁾ we can take the distance matrix:

$$C = \begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 \\ 1 & 0 & 1 & 2 & 3 & 4 \\ 2 & 1 & 0 & 1 & 2 & 3 \\ 3 & 2 & 1 & 0 & 1 & 2 \\ 4 & 3 & 2 & 1 & 0 & 1 \\ 5 & 4 & 3 & 2 & 1 & 0 \end{pmatrix} \quad (16)$$

The goal is to redistribute the sand between locations in an economical way, that is with minimal transportation costs.

As in the previous example we start with:

$$\begin{pmatrix} 6 & 4 & 10 & 20 & 40 & 20 \\ 17 & 3 & 15 & 9 & 26 & 30 \end{pmatrix} \quad (17)$$

⁷⁾ But not in cases when loads can be broken up into smaller ones. In such cases one may be forced to obtain a load from elsewhere while not using the load at the location.

⁸⁾ A cost matrix is not unique: if we have such a matrix and we multiply all entries by the same factor $\lambda > 0$ we also have a cost matrix. And one that yields similar results as the initial one ($\lambda = 1$).

The mass that can remain in the same location is handled first, which means that supply and demand per location should be met as well as possible, as no transportation costs are involved. Applying the adjustments in (12) to (17) yields:

$$\begin{pmatrix} 0 & 1 & 0 & 11 & 14 & 0 \\ 11 & 0 & 5 & 0 & 0 & 10 \end{pmatrix} \quad (18)$$

Now we transport a mass 1 from location 2 to location 1, and mass 10 from location 5 to location 6. So in this step a total mass of 11 is transported with total costs $11 \cdot 1 = 11$. We then have the following new situation:

$$\begin{pmatrix} 0 & 0 & 0 & 11 & 4 & 0 \\ 10 & 0 & 5 & 0 & 0 & 0 \end{pmatrix} \quad (19)$$

In the next step we move a mass 10 from location 4 to location 1 at cost $10 \cdot 3 = 30$, a mass 4 from location 5 to location 3 at cost $4 \cdot 2 = 8$, so a total cost of 38 in this step. We then obtain:

$$\begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix} \quad (20)$$

In the final step we move a mass 1 at location 4 to location 3, at cost $1 \cdot 1 = 1$.

We can summarize the total mass being transported between the various locations in matrix form as:

$$U = \begin{pmatrix} 6 & 0 & 0 & 0 & 0 & 0 \\ 1 & 3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 10 & 0 & 0 & 0 \\ 10 & 0 & 1 & 9 & 0 & 0 \\ 0 & 0 & 4 & 0 & 26 & 10 \\ 0 & 0 & 0 & 0 & 0 & 20 \end{pmatrix} \quad (21)$$

Note that summing row-wise yields the supply vector, and column-wise the demand vector.

Multiplying C and U element-wise (\odot) yields the following total cost matrix T :

$$T = C \odot U = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 30 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 8 & 0 & 0 & 10 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (22)$$

Adding the entries in T yields the total cost: 50, as we found above when adding the costs found in each step.

Remark In this example we again have chosen for transported values to be in \mathbb{N} , for computational ease. However, the values could have been chosen to be in \mathbb{Q}^+ or \mathbb{R}^+ just as well. \square

3.3 Hierarchy of categories

In some application the categories can be thought of as hierarchically ordered. This hierarchy can be used as a means to measure the distance between categories. In other words the hierarchy turns the set of categories into a metric space.




Figure 3.2 Categories ordered in a tree.

A simple metric associated with the (directed) tree in Figure 3.2 is that based on path length. Between each pair of nodes there is a unique path. Each arc is assumed to have length 1.⁹⁾ When traversing the tree we do not care about the direction of the arc.¹⁰⁾ In Table 3.1 a distance table for this (di)tree is shown, assuming that the distance between adjacent nodes is 1.

	a	b	c	d	e	f	g	h	i	j	k	l
a	0	1	1	2	2	2	2	3	3	3	3	3
b	1	0	2	1	1	3	3	2	2	2	4	4
c	1	2	0	3	3	1	1	4	4	4	2	2
d	2	1	3	0	2	4	4	3	3	3	5	5
e	2	1	3	2	0	4	4	1	1	1	5	5
f	2	3	1	4	4	0	2	5	5	5	3	3
g	2	3	1	4	4	2	0	5	5	5	1	1
h	3	2	4	3	1	5	5	0	2	2	6	6
i	3	2	4	3	1	5	5	2	0	2	6	6
j	3	2	4	3	1	5	5	2	2	0	6	6
k	3	4	2	5	5	3	1	6	6	6	0	2
l	3	4	2	5	5	3	1	6	6	6	2	0

Table 3.1 Distance table for the (di)tree in Figure 3.2.

⁹⁾ This is an assumption for simplicity's sake.

¹⁰⁾ The direction of the arrows is of no importance when determining the distance. The arrows indicate a semantic relationship.

If we assume that only the leaves of the (di)tree in Figure 3.2 are the ones that ‘count’, that is nodes d, f, h, i, j, k, l , we are dealing with a subtable of Table 3.1, which is represented in Table 3.2. These leaves correspond to the most detailed categories.

	d	f	h	i	j	k	l
d	0	4	3	3	3	5	5
f	4	0	5	5	5	3	3
h	3	5	0	2	2	6	6
i	3	5	2	0	2	6	6
j	3	5	2	2	0	6	6
k	5	3	6	6	6	0	2
l	5	3	6	6	6	2	0

Table 3.2 Distance table for the leaves of the (di)tree in Figure 3.2.

Table 3.1 can also be derived from the adjacency matrix of the underlying tree of the ditree in Figure 3.2 which is given in Table 3.3. If we denote the adjacency matrix in Table 3.3 by A and the matrix in Table 3.1 by \mathcal{T} , we have that $\mathcal{T}_{ij} = \min_q (A^q)_{ij} > 0$. The exponent 6 is the largest distance between any two points in the underlying tree of the ditree in Figure 3.2, as inspection of Table 3.1 shows. In other words, the diameter of the underlying tree is 6.

	a	b	c	d	e	f	g	h	i	j	k	l
a	0	1	1	0	0	0	0	0	0	0	0	0
b	1	0	0	1	1	0	0	0	0	0	0	0
c	1	0	0	0	0	1	1	0	0	0	0	0
d	0	1	0	0	0	0	0	0	0	0	0	0
e	0	1	0	0	0	0	0	1	1	1	0	0
f	0	0	1	0	0	0	0	0	0	0	0	0
g	0	0	1	0	0	0	0	0	0	0	1	1
h	0	0	0	0	1	0	0	0	0	0	0	0
i	0	0	0	0	1	0	0	0	0	0	0	0
j	0	0	0	0	1	0	0	0	0	0	0	0
k	0	0	0	0	0	0	1	0	0	0	0	0
l	0	0	0	0	0	0	1	0	0	0	0	0

Table 3.3 Adjacency matrix of the underlying tree of the ditree in Figure 3.2.

An assumption underlying the distance table in Table 3.2 is that the arcs in the ditree are all of the same length T . What does this mean? It is obviously not a distance function in the ordinary sense of the word. It can be interpreted as a dissimilarity measure of the corresponding categories. Such measures are also called proximities. See e.g. [4]. Chapter 6 of this book is devoted to obtaining proximities in practical applications.

One of these methods is to associate a number of variables with a classification. The categories in such a classification correspond with certain scores on these variables. These vectors of scores can be used to define a distance function of the set of categories in the classification. One could take a Hamming distance, for instance. Obviously, this distance depends on the variables chosen as well as the categorization for each variable. However it is a more refined method of obtaining a metric than the one used in Table 3.1 which assumes unit length for each arc. So we can define the Hamming distance Δ_H :

$$\Delta_H(v^a, v^b) = |\{i \in \{1, \dots, l\} | v_i^a \neq v_i^b\}|, \quad (23)$$

where v^a and v^b are two vectors with characteristics of length l , and where $|\cdot|$ denotes the count operator. The distance Δ_{ab} between categories a and b would then be measured by (23)

$$\Delta_{ab} = \Delta_H(v^a, v^b), \quad (24)$$

which, as the notation reflects, depends on the choice of the characteristic vectors v^a and v^b .

A distance table for the ditree in Figure 3.2 based on the Hamming distance is presented in Table 3.4.

	a	b	c	d	e	f	g	h	i	j	k	l
a	0	Δ_{ab}	Δ_{ac}	0	0	0	0	0	0	0	0	0
b	Δ_{ba}	0	0	Δ_{bd}	Δ_{be}	0	0	0	0	0	0	0
c	Δ_{ca}	0	0	0	0	Δ_{cf}	Δ_{cg}	0	0	0	0	0
d	0	Δ_{db}	0	0	0	0	0	0	0	0	0	0
e	0	Δ_{eb}	0	0	0	0	0	Δ_{eh}	Δ_{ei}	Δ_{ej}	0	0
f	0	0	Δ_{fc}	0	0	0	0	0	0	0	0	0
g	0	0	Δ_{gc}	0	0	0	0	0	0	0	Δ_{gk}	Δ_{gl}
h	0	0	0	0	Δ_{he}	0	0	0	0	0	0	0
i	0	0	0	0	Δ_{ie}	0	0	0	0	0	0	0
j	0	0	0	0	Δ_{je}	0	0	0	0	0	0	0
k	0	0	0	0	0	0	Δ_{kg}	0	0	0	0	0
l	0	0	0	0	0	0	Δ_{lg}	0	0	0	0	0

Table 3.4 Hamming distances for neighboring nodes of the ditree in Figure 3.2. The matrix is supposed to be symmetric.

From Table 3.4 we can compute (by employing matrix algebra) the distance matrix for any pair of categories. If \mathcal{T}_H denotes the matrix in Table 3.4, the distance table derived from this for any pair of nodes is $\mathcal{T}_{H,i,j}^d = \min_q \{q \mid \mathcal{T}_{H,i,j}^q > 0\}$. Note that any node in the ditree in Figure 3.2 can be reached from any other node in at most six steps.




Figure 3.3 Hamming distances defined for the category tree in Figure 3.2.

In Figure 3.3 the distances in Table 3.4 have been associated to the arcs in Figure 3.2. The understanding is that these distances are equal for an arc as well as its counter-arc:

$\Delta_{(x,y)} = \Delta_{(y,x)}$ for arc (x, y) . So the distance between b and f equals $\Delta_{ba} + \Delta_{ac} + \Delta_{cf}$, and the distance between h and l is: $\Delta_{he} + \Delta_{eb} + \Delta_{ba} + \Delta_{ac} + \Delta_{cg} + \Delta_{gl}$. It should be borne in mind that in a tree there is a unique path between any pair of nodes, which is, trivially, also the shortest one for these nodes.

4 Using copulas as weight matrices

4.1 What are copulas?

The basic idea guiding copula theory is the separation of marginal distributions of a multivariate distribution and the dependence structure of derived variables. A theorem of Sklar (see [16]) is the basis for this separation. See also Appendix B.

Copulas¹¹⁾ are used to describe the dependence of random variables. A copula is a multivariate cumulative distribution function for which the marginal probability distribution of each variable is uniform.¹²⁾ In fact, the copula approach is directly usable if the marginal tables can be viewed as realizations of some random process. However, in the applications that we consider this is usually not the case. But still the basic idea of copula theory can be used if we are able to use a suitable dependency structure that can be used to redistribute the marginal values over the interior table.

In this approach, there are two steps to estimate the interior matrix. In the first step a so-called copula is estimated by applying IPF on the basis of available information about which nodes are connected. This copula is used as an approximation for the interior matrix. It is used in the second step, which consists of an IPF procedure with the two given univariate distributions as marginal distributions.

This approach first focuses on the dependence structures and then considers the redistribution of the probability mass. The dependence structure is assumed to be determined by an adjacency matrix only. This adjacency matrix is assumed to be provided by a specialist, in which case it can be assumed as subjective information. Take National Accounts as an example. In this case this 0 – 1 matrix A specifies which sector is related to which other sectors through the delivery of goods or services.¹³⁾ Suppose the vertices correspond to sectors in an economy. An arc pointing from sector i to j indicates that services or goods from i to j are offered.¹⁴⁾ If (i, j) is an arc, so node (say, sector) i is connected to node (say, sector j), and if there is no such arc we have $A_{ij} = 0$.

In Section 4.2 we consider methods to estimate copulas using the user specified adjacency matrix as well as the two marginal distributions. These marginal distributions are not used

¹¹⁾ ‘Copula’ is used here as a concept from probability theory and statistics. It is unrelated to a concept with the same name in linguistics. The Latin word ‘copula’, means ‘link’ or ‘tie’.

¹²⁾ For a quick introduction to copulas see [https://en.wikipedia.org/wiki/Copula_\(probability_theory\)](https://en.wikipedia.org/wiki/Copula_(probability_theory)). This entry also contains references in case one want to pursue the subject in more depth. In Appendix B some of this information is presented as well for easy reference.

¹³⁾ This implies a reverse relationship of money flowing in the opposite direction, as payments for the goods and services delivered.

¹⁴⁾ Or if one focuses on the payments: there is an arc in the opposite direction, from j to i .

directly. They are used to derive different marginals that are used in an IPF procedure to yield copulas. These copulas can then in turn be used as weight matrices in an IPF procedure to estimate a 2D-table with given marginals. See Section 5.1 for details.

4.2 Estimating copulas

We start with an $n \times n$ adjacency matrix A , estimated by a user, and two 1D marginals, vectors $\pi_c \geq 0$ and $\pi_r \geq 0$ of length n and with the same totals, i.e. $\pi_c \cdot \iota = \pi_r \cdot \iota$. First we order the elements of these marginals in increasing order and obtain $\bar{\pi}_c$ and $\bar{\pi}_r$. This implies permutations p_c and p_r of the columns and the rows, respectively. Apply p_c and p_r to the columns and rows of A , resulting in the doubly sorted matrix \bar{A} .

Now we generate two random vectors ρ_c and ρ_r of length n , in which each entry is independently drawn from the uniform distribution U on $[0, 1]$. Then sort the entries of these random vectors in increasing order, obtaining $\bar{\rho}_c$ and $\bar{\rho}_r$, respectively.

Next, apply IPF to \bar{A} , $\bar{\rho}_c$ and $\bar{\rho}_r$, resulting in the copula \bar{C} . Apply the inverses of p_c and p_r , that is p_c^{-1} and p_r^{-1} , to the columns and rows of \bar{C} respectively, to obtain the copula C_{ρ_c, ρ_r} , corresponding to A , where ρ_c and ρ_r correspond to the unsorted \bar{A} , $\bar{\rho}_c$ and $\bar{\rho}_r$.¹⁵⁾

If we wish, we can repeat this procedure many times, and average the resulting copulas, to obtain a new, averaged copula. We obtain a copula by repeatedly (say n times) drawing pairs of ρ_c and ρ_r vectors, and average the resulting $A_{\rho_c, \rho_r}^{(i)}$:

$$\bar{B}_{\rho_c, \rho_r}^{(n)} = \frac{1}{n} \sum_{i=1}^n A_{\rho_c, \rho_r}^{(i)}. \quad (25)$$

It should be remarked, that we can also use these results to obtain bootstrap estimates of the variance of the sampled copulas.

5 Using IPF in a 2D table

In Section 4 we used IPF to estimate copulas. These in turn are used in an approach described in Section 5.1 which is used to compute a 2D-matrix which can be seen as the result of two marginal 1D tables smeared out over the interior. In this case a copula is used as the weight matrix. This approach, which is a two-step procedure, is inspired by the copula approach described in Section 4. Typical for this approach is that first the dependence structure is estimated (see 4.2), which in turn is used as weight matrix for the second step, in which the marginal distributions are spread over the 2D-table. As input for the first step the adjacency matrix is used that is the result of an assessment of the user as to the interconnection of the sectors (in case of a national accounts example).

¹⁵⁾ Or rather, by applying p_c^{-1} and p_r^{-1} to $\bar{\rho}_c$ and $\bar{\rho}_r$, respectively.

The second approach we describe in the present section uses a one step IPF procedure to obtain a $2D$ table. It is described in Section 5.2. It uses the same ingredients as the two-step procedure (viz adjacency matrix A , marginal distributions π_c and π_r) but in a different, more direct way.

Using notation introduced in Section 4.2 it is possible to succinctly describe the results of the methods treated in Sections 5.1 and 5.2, and immediately see the differences, not numerically but symbolically. In Section (5.3) some properties of both approaches are brought forward.

5.1 Two-step procedure

The first step of this approach is the estimation of a suitable weight matrix W , for instance the one described in Section 4.2, which is a copula. In the second step W is used as a weight matrix to distribute the marginal distributions π_c and π_r over the interior of a $2D$ -matrix using IPF.

So we assume two marginal distributions of an unknown $2D$ -distribution to be given. In the first step we focus on the estimation of the dependence structure, using the adjacency matrix A as a weight matrix.

In the two-step approach we deal with the copula. The input for this is the adjacency matrix. The marginals are the random vectors ρ_c and ρ_r , where each of their entries are drawn from a uniform distribution. With the drawing procedure we should consider the order of the entries of π_c and π_r . So use independent drawings from the uniform distribution and sort according to the distributions of ρ_c and ρ_r . Using IPF¹⁶⁾ we can find a copula, which we denote by A_{ρ_c, ρ_r} . This is used as the weight matrix for the second step, where it serves to guide the distribution of the marginals π_c and π_r over the interior. We then have the result B_{π_c, π_r} , where $B = A_{\rho_c, \rho_r}$ is the weight matrix obtained in the first step.

Like in Section 4.2 it is possible to extend the method to obtain a copula in the first step, is by repeatedly drawing pairs of ρ_c and ρ_r vectors and average the resulting $A_{\rho_c, \rho_r}^{(i)}$, to obtain the expression (25).

5.2 One-step procedure

The approach we consider in the present subsection is direct: it uses the adjacency matrix as a crude approximation of the interior matrix. IPF is used to spread the marginal distributions over the interior, respecting the information in the adjacency matrix. This means that zero cells remain zero cell and only non-zero cells can acquire mass.

In this approach we directly use the univariate distributions as the marginals to be fitted, with the adjacency matrix as an initial proxy for the bivariate distribution we want to estimate. Again IPF is applied to compute the result.

Example If A is the $n \times n$ ‘all 1 matrix’

¹⁶⁾ See Appendix A.

$$A = \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & & \vdots \\ 1 & \cdots & 1 \end{pmatrix}. \quad (26)$$

Then the resulting weight matrix is as follows:

$$W = \begin{pmatrix} \pi_1^c \pi_1^r & \cdots & \pi_n^c \pi_1^r \\ \vdots & & \vdots \\ \pi_1^c \pi_n^r & \cdots & \pi_n^c \pi_n^r \end{pmatrix} = \pi^c \cdot (\pi^r)'. \quad (27)$$

where π^c is the column marginal represented as a column vector and $(\pi^r)'$ is the row marginal represented as a row vector (the prime ' denotes transposition).

The solution (27) is also obtained if the maximum entropy principle is used (see e.g. [3], Section 3.4). It represents a solution that corresponds to independence of the marginals (when they are normalized so that they can be viewed as random variables). \square

5.3 Properties of the two approaches

Before we start with considering the properties of the two approaches in more detail, we introduce some useful notation.

Symbolically we can write the computation of a $2D$ -table starting with a weight matrix $W \geq 0$ and marginals $\pi_r > 0$ and $\pi_c > 0$ ¹⁷⁾, where r refers to 'rows' and c to 'columns', as a mapping

$$(W, \pi_r, \pi_c) \mapsto W_{\pi_r, \pi_c}, \quad (28)$$

assuming that the value of this mapping exists, that is, is an element of $\mathbb{M}_{n,n}$, the set of $n \times n$ real-valued matrices. To be consistent, the marginal vectors need to add to the same total $\pi_c \cdot \iota = \pi_r \cdot \iota$.

It is interesting to consider some properties of this mapping. An easy property to verify is that for any scalar $\gamma > 0$ it holds that $W_{\gamma \pi_r, \gamma \pi_c} = W_{\pi_r, \pi_c}$. In view of this property we may also assume that $\pi_c \cdot \iota = \pi_r \cdot \iota = 1$, although we do not always choose to do so.

The next set of properties concern the iteration of the IPF operation. Let τ_r and τ_c be other row and column vectors. A purely mathematically motivated question is about the commutativity of the mapping: are $(W_{\pi_r, \pi_c})_{\tau_r, \tau_c}$ and $(W_{\tau_r, \tau_c})_{\pi_r, \pi_c}$ the same or are they different? A general answer cannot be given, depending on the choice of the marginal vectors and the weight matrix.

¹⁷⁾ The requirement is actually weaker: $\pi_r \geq 0$ and $\pi_c \geq 0$. But the rows and columns corresponding to 0-values can be skipped. This leads to a reduced matrix and correspondingly reduced marginal vectors with strictly positive components. So without loss of generality we may assume that the original row and columns vectors are strictly positive. But then the number of rows and columns may be different.

If $\pi_r = \gamma \tau_r$ and $\pi_c = \gamma \tau_c$ for some constant $\gamma > 0$ the results are the same. Also we have $(W_{\pi_r, \pi_c})_{\pi_r, \pi_c} = W_{\pi_r, \pi_c}$, for all marginal vectors $\pi_r, \pi_c > 0$ of size n .

We can describe our one-step IPF procedure in term of this mapping as

$$(A, \pi_r, \pi_c) \mapsto A_{\pi_r, \pi_c}. \quad (29)$$

and our two-step procedure as

- step 1: $(A, \iota_r, \iota_c) \mapsto A_{\iota_r, \iota_c}$, where $\iota_r = (1, \dots, 1)$ the all one vector of size n and $\iota_c = \iota_r'$.
- step 2: $(A_{\iota_r, \iota_c}, \pi_r, \pi_c) \mapsto (A_{\iota_r, \iota_c})_{\pi_r, \pi_c}$

The resulting doubly stochastic matrix A_{ι_r, ι_c} is a copula. It is used as a weight matrix in a second IPF application, but this time with the supply and demand vectors π_r and π_c , respectively (possibly normalized, i.e. such that $\iota_r \cdot \pi_r = \iota_c \cdot \pi_c = 1$, where \cdot denotes the standard inner product in \mathbb{R}^n). In general, $A_{\pi_r, \pi_c} \neq (A_{\iota_r, \iota_c})_{\pi_r, \pi_c}$.

6 Discussion

In this paper several methods were considered to derive a $2D$ measure given two $1D$ measures. Different approaches yield different solutions. There is no way to tell which of these methods is the best one. Which model can be applied depends on the problem at hand, and the information available.

The best approach to a problem of the type discussed is to try several methods and see how much (or little) their solutions differ. This gives a good indication of model selection variability (or sensitivity). This is actually an approach that is commendable in many applications.

The $1D$ measures are assumed to be known, exactly or approximately. In case they can be viewed as realizations of a random process, they can be replicated. For each combination of such uni-variate distributions one can compute a $2D$ solution. This gives an idea of the sensitivity of the result (the $2D$ measure) on these quantities, for a chosen model.

To produce good estimates of the $2D$ distribution, one should use as much prior information about such a measure as is available, in particular about the 0 cells, the structural zeroes.

References

- [1] M. Bacharach (1970). *Biproportional Matrices & Input-Output Change*. Cambridge University Press.

- [2] J. Barthélemy & T. Suesse (no date). mipfp: Multidimensional Iterative Proportional Fitting. CRAN.
- [3] J. Berger (1985). *Statistical Decision Theory and Bayesian Analysis* (2nd ed.). Springer.
- [4] I. Borg & P. Groenen (2005). *Modern Multidimensional Scaling*. Springer.
- [5] L. Ford & D. Fulkerson (1962). *Flows in Networks*. Princeton University Press.
- [6] I. Csiszár (1975). I-Divergence of Probability Distributions and Minimization Problems. *Annals of Probability*. Vol. 3, pp. 146–158.
- [7] F. Hitchcock (1941). The Distribution of a Product from Several Sources to Numerous Localities, *J. Math. Phys.*, 20, pp. 224–230.
- [8] L. Kantorovich (1942). On the Transfer of Masses. *Dokl. Akad. Nauk USSR*, Vol. 37, pp. 7–8.
- [9] L. Kantorovich (1948). On a Problem of Monge. *Uspekhi Mat. Nauk*, Vol. 3, pp. 225–226. (In Russian)
- [10] T. Koopmans (1942). Exchange Ratios between Cargoes on Various Routes (Non-Refrigerated Dry Cargoes). Memorandum for the Combined Shipping Adjustment Board, Washington, D.C., 1942. Publ. in *Scientific Papers of Tjalling C. Koopmans*, Springer Verlag, 1970. pp. 77–86.
- [11] A. Marshall & I. Olkin (1967). Scaling of Matrices to Achieve Specified Row and Column Sums. *Numerische Mathematik*. Vol. 12 (1), pp. 83–90.
- [12] G. Monge (1781). Mémoire sur la Théorie de Déblais et de Remblais. *Mémoires de l'Académie Royale des Sciences (France)*, pp. 666–704.
- [13] F. Pukelsheim & B. Simeone (2009). On the Iterative Proportional Fitting Procedure: Structure of Accumulation Points and L_1 -Error Analysis. Preprint, Department of Mathematics, University of Augsburg.
- [14] S. Rachev & L. Rüschendorf (1998a). *Mass Transportation Problems*, Vol. I: Theory, Springer.
- [15] S. Rachev & L. Rüschendorf (1998b). *Mass Transportation Problems*, Vol. II: Applications, Springer.
- [16] A. Sklar (1959). Fonctions de Répartition à n Dimensions et Leurs Marges, *Publ. Inst. Statist. Univ. Paris*, Vol. 8, pp. 229–231.
- [17] R. Sinkhorn (1964). A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist.*, Vol. 35, pp. 876–879.
- [18] R. Sinkhorn & P. Knopp (1967). Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math.*, Vol. 21, pp. 343–348.
- [19] C. Villani, (2003). *Topics in Optimal Transportation*, Graduate Studies in Mathematics, Vol. 58, American Mathematical Society.
- [20] C. Villani (2008). *Optimal Transport, Old and New*, Springer.

Appendix

A IPF

According to Wikipedia: Iterative Proportional Fitting (IPF) is an iterative algorithm for estimating cell values of a contingency table such that the marginal totals remain fixed and the estimated table decomposes into an outer product.¹⁸⁾ It is an algorithm that is known for a long time. Convergence in practice of IPF is usually not a problem. However, to provide necessary and sufficient conditions for convergence in case the contingency table has 0-cells is not so easy. Such conditions are known and are of relative recent origin, compared to the age of the IPF-algorithm (see [6] and [13]). It was used by the Cambridge economist Richard Stone as a tool in input-output analysis. See also [1].

It should be remarked that an R package is available for solving multivariate IPF problems: `mipf` (see [2] on the CRAN website). But to write an IPF script in R is not such a difficult task, particularly not for a $2D$ -table.

IPF appears in different guises, which we discuss below.¹⁹⁾ In some cases we assume that the variables used in the algorithm have a particular probability distribution (like a Poisson or multinomial distribution). But generally such an assumption is not made, nor necessary.

In fact we are especially interested in those case where some of the cells of the $2D$ -matrix are 0, as this is often the case in practice. These are beyond the traditional scope of the problem when proving convergence. However [13] has results on this case, which appears quite often in practice. In Section A.4.2 this issue is discussed in the context of Sinkhorn's theorem.

A.1 Classical IPF

Given a two-way $(I \times J)$ -table of counts x_{ij} , where the cell values are assumed to be Poisson or multinomially distributed, we wish to estimate a decomposition $\hat{m}_{ij} = a_i b_j$ for all i and j such that \hat{m}_{ij} is the maximum likelihood estimate (MLE) of the expected values m_{ij} leaving the marginals $x_{i+} = \sum_j x_{ij}$ and $x_{+j} = \sum_i x_{ij}$ fixed. The assumption that the table factorizes in such a manner is known as the 'model of independence' (I-model). Written in terms of a log-linear model, we can write this assumption as $\log m_{ij} = u + v_i + w_j + z_{ij}$, where $m_{ij} = \mathbb{E}(x_{ij})$, $\sum_i v_i = \sum_j w_j = 0$ and the interaction term vanishes, that is $z_{ij} = 0$ for all i and j .

Choose initial values $\hat{m}_{ij}^{(0)} = 1$ (different choices of initial values may lead to changes in convergence behavior), and for $\eta \geq 1$ set

¹⁸⁾ See: https://en.wikipedia.org/wiki/Iterative_proportional_fitting

¹⁹⁾ The discussion is almost verbatim taken from the Wikipedia page https://en.wikipedia.org/wiki/Iterative_proportional_fitting on this item.

$$\hat{m}_{ij}^{(2\eta-1)} = \frac{\hat{m}_{ij}^{(2\eta-2)} x_{i+}}{\sum_{k=1}^J \hat{m}_{ik}^{(2\eta-2)}}, \quad (\text{A.1a})$$

$$\hat{m}_{ij}^{(2\eta)} = \frac{\hat{m}_{ij}^{(2\eta-1)} x_{+j}}{\sum_{k=1}^I \hat{m}_{kj}^{(2\eta-1)}}. \quad (\text{A.1b})$$

Notes:

- Convergence does not depend on the actual distribution. Distributional assumptions are necessary for inferring that the limit $\hat{m}_{ij} = \lim_{\eta \rightarrow \infty} (\hat{m}_{ij}^{(\eta)})$ is an MLE indeed.
- IPF can be manipulated to generate any positive marginals by replacing x_{i+} by the desired row marginal u_i (analogously for the column marginals).
- IPF can be extended to fit the 'model of quasi-independence' (Q-model), where $m_{ij} = 0$ is known a priori for $(i, j) \in S$. Only the initial values have to be changed: Set $\hat{m}_{ij}^{(0)} = 0$ if $(i, j) \in S$ and 1 otherwise.

A.2 Factor estimation

Assume the same setting as in Section A.1. Alternatively, we can estimate the row and column factors separately: Choose initial values $\hat{b}_j^{(0)} = 1$, and for $\eta \geq 1$ set

$$\hat{a}_i^{(\eta)} = \frac{x_{i+}}{\sum_j \hat{b}_j^{(\eta-1)}}, \quad (\text{A.2a})$$

$$\hat{b}_j^{(\eta)} = \frac{x_{+j}}{\sum_i \hat{a}_i^{(\eta)}}. \quad (\text{A.2b})$$

Setting $\hat{m}_{ij}^{(2\eta)} = \hat{a}_i^{(\eta)} \hat{b}_j^{(\eta)}$, the two variants of the algorithm are mathematically equivalent (can be seen by formal induction).

Notes:

- In matrix notation, we can write $(\hat{m}_{ij}) = \hat{a} \hat{b}^T$, where $\hat{a} = (\hat{a}_1, \dots, \hat{a}_I)^T = \lim_{\eta \rightarrow \infty} \hat{a}^{(\eta)}$ and $\hat{b} = (\hat{b}_1, \dots, \hat{b}_J)^T = \lim_{\eta \rightarrow \infty} \hat{b}^{(\eta)}$.
- The factorization is not unique, since $m_{ij} = a_i b_j = (\gamma a_i) (\frac{1}{\gamma} b_j)$ for any $\gamma > 0$.
- The factor totals remain constant, i.e. $\sum_i \hat{a}_i^{(\eta)} = \sum_i \hat{a}_i^{(1)}$ for any $\eta \geq 1$ and $\sum_j \hat{b}_j^{(\eta)} = \sum_j \hat{b}_j^{(0)}$ for any $\eta \geq 0$.
- To fit the Q-model, where $m_{ij} = 0$ a priori for $(i, j) \in S$, set $\delta_{ij} = 0$ if $(i, j) \in S$ and $\delta_{ij} = 1$ otherwise. Then

$$\hat{a}_i^{(\eta)} = \frac{x_{i+}}{\sum_j \delta_{ij} \hat{b}_j^{(\eta-1)}}, \quad (\text{A.3a})$$

$$\hat{b}_j^{(\eta)} = \frac{x_{+j}}{\sum_i \delta_{ij} \hat{a}_i^{(\eta)}}, \quad (\text{A.3b})$$

$$\hat{m}_{ij}^{(2\eta)} = \delta_{ij} \hat{a}_i^{(\eta)} \hat{b}_j^{(\eta)}. \quad (\text{A.3c})$$

A.3 RAS

The Problem: Let $M = (m_{ij}^{(0)}) \in \mathbb{R}^{I \times J}$ be the initial matrix with nonnegative entries, $u \in \mathbb{R}^I$ a vector of specified row marginals (or, row sums) and $v \in \mathbb{R}^J$, a vector of column marginals. We wish to compute a matrix $\hat{M} = (\hat{m}_{ij}) \in \mathbb{R}^{I \times J}$ similar to M and consistent with the predefined marginals, meaning

$$\hat{m}_{i+} = \sum_{j=1}^n \hat{m}_{ij} = u_i, \quad (\text{A.4a})$$

$$\hat{m}_{+j} = \sum_{i=1}^m \hat{m}_{ij} = v_j. \quad (\text{A.4b})$$

Define the diagonalization operator $\text{diag} : \mathbb{R}^k \rightarrow \mathbb{R}^{k \times k}$, which produces a (diagonal) matrix with its input vector on the main diagonal and zeroes elsewhere. Then, for $\eta \geq 0$, set

$$M^{(2\eta+1)} = \text{diag}(r^{(\eta+1)})M^{(2\eta)}, \quad (\text{A.5a})$$

$$M^{(2\eta+2)} = M^{(2\eta+1)}\text{diag}(s^{(\eta+1)}). \quad (\text{A.5b})$$

where

$$r_i^{\eta+1} = \frac{u_i}{\sum_j m_{ij}^{(2\eta)}}, \quad (\text{A.6a})$$

$$s_j^{\eta+1} = \frac{v_j}{\sum_i m_{ij}^{(2\eta+1)}}. \quad (\text{A.6b})$$

A.4 Sinkhorn's theorem and beyond

A.4.1 The theorem

Sinkhorn's theorem can be viewed as an application of IPF. This theorem is as follows. Let A be a square matrix (say $n \times n$ with $A > 0$, that is, with only positive entries. Then there exist diagonal matrices D_1 and D_2 with strictly positive diagonal elements such that $D_1 A D_2$ is doubly stochastic. The matrices D_1 and D_2 are unique modulo scale factors.²⁰⁾ See [17] and [11]. The result also follows from the application of IPF to the matrix A and marginal distributions $\iota = (1, \dots, 1)$ (of length n). This obvious solution method was suggested in [18], which also analyzed the convergence of the algorithm.

²⁰⁾ Multiplying D_1 by $\lambda > 0$ and D_2 by $1/\lambda$ yields the same result.

A.4.2 Beyond the theorem

The restriction $A > 0$ is rather limiting, as in practice one often encounters situations where some of the entries of A are 0. An obvious approach is to replace each of these zeroes by small values $\epsilon > 0$, in which case a solution exists according to Sinkhorn's theorem, and then take the limit $\epsilon \downarrow 0$. Such an approach is suggested in [13].

B Copulas

Let X be a real-valued random variable, and F its cumulative distribution function (CDF). Then it is well-known that the random variable $F(X)$ is uniformly distributed on $[0, 1]$ ²¹⁾. This idea can be used in a multivariate setting as well. Let (X_1, X_2, \dots, X_n) be a random vector with values in \mathbb{R}^n . Let the marginal CDF be $F_i(x) = P[X_i \leq x]$. Then

$$(U_1, \dots, U_n) = (F(X_1), \dots, F(X_n)) \quad (\text{B.1})$$

is a distribution where each of the marginals has a uniform distribution on $[0, 1]$. The copula of $X = (X_1, X_2, \dots, X_n)$ is defined as the joint CDF of $U = (U_1, \dots, U_n)$:

$$C(u_1, \dots, u_n) = P[U_1 \leq u_1, \dots, U_n \leq u_n]. \quad (\text{B.2})$$

So (B.2) contains the dependence structure of a kind of normalized variables, which are independent of the marginal distributions of the various components of the random vector. The marginal distributions are separated off from the original problem.

The nice thing is that the steps can be reversed. For a given copula function C one can generate pseudo-random samples (U_1, \dots, U_n) and one can generate (X_1, X_2, \dots, X_n) as follows

$$(X_1, \dots, X_n) = (F_1^{-1}(U_1), \dots, F_n^{-1}(U_n)) \quad (\text{B.3})$$

It should be remarked that if F is a continuous function, the inverse F^{-1} exists. In general this need not be the case. However, a CDF F is always right continuous and a generalized inverse F^- can always be defined, with $F \circ F^-(x) = x$ and $F^- \circ F(y) = y$.

²¹⁾ Or $[0, 1)$ or $(0, 1)$ or $(0, 1]$. For the sake of the argument the exact set is not so important, as we assume that single point sets are assumed to have zero probability.

Colophon

Publisher

Statistics Netherlands
Henri Faasdreef 312, 2492 JP The Hague
www.cbs.nl

Prepress

Statistics Netherlands, Grafimedia

Design

Edenspiekermann

Information

Telephone +31 88 570 70 70, fax +31 70 337 59 94
Via contact form: www.cbs.nl/information

© Statistics Netherlands, The Hague/Heerlen/Bonaire 2018.
Reproduction is permitted, provided Statistics Netherlands is quoted as the source