

ML2R Coding Nuggets

Sorting as Linear Programming

Christian Bauckhage*
Machine Learning Rhine-Ruhr
Fraunhofer IAIS
St. Augustin, Germany

Pascal Welke†
Machine Learning Rhine-Ruhr
University of Bonn
Bonn, Germany

ABSTRACT

Linear programming is a surprisingly versatile tool. That is, many problems we would not usually think of in terms of a linear programming problem can actually be expressed as such. In this note, we show that sorting is such a problem and discuss how to solve linear programs for sorting using *SciPy*.

1 INTRODUCTION

We once again revisit the idea of [linear programming](#) [9, 10] and recall that it is all about solving optimization problems of the following form

$$\begin{aligned} z^* = \underset{z}{\operatorname{argmin}} \quad & c^\top z \\ \text{s.t.} \quad & Az \leq b \\ & Cz = d \end{aligned} \quad (1)$$

The practical application we consider in this note is the fundamental problem of sorting an unordered collection of real numbers in ascending order. Every first year computer science student knows that this is not difficult and that there are numerous algorithms (such as *quicksort*) which sort quickly and reliably.

A much lesser known fact is that the sorting problem can be expressed as a linear programming problem of the form in (1). Does this have practical advantages? No, at least not on commodity hardware! Good conventional sorting algorithms are fast enough ($O(n \log n)$) to deal with millions of numbers. Linear programs, on the other hand, are more demanding in that their solution requires polynomial effort ($O(n^d)$ where $d > 2$). In this sense, the method we discuss in this note is of theoretical rather than practical interest. Our goals are, first of all, to demonstrate that linear programming is very widely applicable and, second of all, to show that even well known problems often allow for different points of view. The resulting computational solutions may come in handy once next generation devices (such as quantum computers or neuromorphic computers) become more readily available.

Next, we show how to derive a linear program for the sorting problem (section 2). This then allows us to apply function `linprog` in *SciPy*'s `optimize` package to sort lists of numbers (section 3). Readers who would like to experiment with our exemplary code should be familiar with *NumPy* and *SciPy* [8] and only need to

```
import numpy as np
import numpy.random as rnd
import scipy.optimize as opt
```

* 0000-0001-6615-2128

† 0000-0002-2123-3781

2 THEORY

In what follows, we consider an arbitrary n -dimensional vector of real numbers

$$\mathbf{x} = [x_1, x_2, \dots, x_n]^\top \quad (2)$$

whose entries x_i are supposed to be sorted in ascending order.

Working with vectors allows us to unleash linear algebraic tools on the sorting problem. Indeed, sorting the entries of \mathbf{x} can also be understood as applying an $n \times n$ **permutation matrix** P such that the entries of the permuted vector

$$\mathbf{y} = P\mathbf{x} \quad (3)$$

obey $y_1 \leq y_2 \leq \dots \leq y_n$. This way, the sorting problem becomes the problem of finding an appropriate permutation matrix.

In order to devise an objective function whose minimization would yield the sought after permutation matrix, we next introduce a very specific, auxiliary n -dimensional vector, namely

$$\mathbf{n} = [1, 2, \dots, n]^\top \quad (4)$$

Given this vector, the **rearrangement inequality** due to Hardy, Littlewood, and Polya [5] implies that $-\mathbf{n}^\top P\mathbf{x} = -\mathbf{n}^\top \mathbf{y}$ is minimal, if $y_1 \leq y_2 \leq \dots \leq y_n$. In other words, the expression $-\mathbf{n}^\top P\mathbf{x}$ is minimal if the permutation matrix P causes the entries of \mathbf{x} to be sorted in the same order as the entries of \mathbf{n} .

Next, we recall that a **doubly stochastic matrix** is a square, non-negative matrix whose rows and columns all sum to one. That is, $M \in \mathbb{R}^{n \times n}$ is doubly stochastic, if $M \geq \mathbf{0}_{n \times n}$, $M\mathbf{1}_n = \mathbf{1}_n$, and $M^\top \mathbf{1}_n = \mathbf{1}_n$ where $\mathbf{0}_{n \times n}$ denotes the $n \times n$ matrix of all zeros and $\mathbf{1}_n$ is the n vector of all ones. We also recall that the **Birkhoff polytope** \mathcal{B}_n is the convex hull of all doubly stochastic matrices of size $n \times n$ and that the **Birkhoff-von Neumann theorem** establishes that the $n!$ vertices of \mathcal{B}_n are the $n \times n$ permutation matrices [12].

All of this is to say that the problem of sorting the elements of \mathbf{x} , i.e. the problem of finding an appropriate permutation matrix P , can be cast as a linear programming problem, namely

$$\begin{aligned} P = \underset{M \in \mathbb{R}^{n \times n}}{\operatorname{argmin}} \quad & -\mathbf{n}^\top M\mathbf{x} \\ & M \geq \mathbf{0}_{n \times n} \\ \text{s.t.} \quad & M\mathbf{1}_n = \mathbf{1}_n \\ & M^\top \mathbf{1}_n = \mathbf{1}_n \end{aligned} \quad (5)$$

Indeed, due to its constraints, the feasible set of this minimization problem is the Birkhoff polytope \mathcal{B}_n which is a convex set. Also, as shown in the appendix, the objective $-\mathbf{n}^\top M\mathbf{x}$ is linear in M . Since the minimum of a linear function over a convex set coincides with a vertex of said set and since the vertices of \mathcal{B}_n are permutation matrices, the solution to (5) will be a permutation matrix.

At this point, we have already reached our first major goal. That is, we have established that sorting can be accomplished via linear programming.

However, our linear program in (5) is a linear program over matrices and thus not of the commonly considered form in (1). Since most software tools for linear programming assume the problem to be formulated in terms of vectors (*SciPy*'s `linprog` is no exception), we still have work to do before we can turn our new insights into practical code. In other words, we have to rewrite the problem in (5) as a problem of the form in (1).

For this purpose, it will be more convenient to work with the transpose $-\mathbf{x}^\top \mathbf{M}^\top \mathbf{n}$ of the objective function in (5). In the appendix, we prove the following identity

$$\mathbf{M}^\top \mathbf{n} = \mathbf{N} \mathbf{m} \quad (6)$$

where the vector $\mathbf{m} \in \mathbb{R}^{n^2}$ and matrix $\mathbf{N} \in \mathbb{R}^{n \times n^2}$ are defined as

$$\mathbf{m} = \text{vec}(\mathbf{M}) \quad (7)$$

$$\mathbf{N} = \mathbf{I}_{n \times n} \otimes \mathbf{n}^\top \quad (8)$$

where $\mathbf{I}_{n \times n}$ is the $n \times n$ identity matrix and \otimes denotes the Kronecker product. In other words, we can express the product of the $n \times n$ matrix \mathbf{M}^\top and the n vector \mathbf{n} as the product of an $n \times n^2$ matrix \mathbf{N} and an n^2 vector \mathbf{m} . The vector $\mathbf{m} = \text{vec}(\mathbf{M})$ contains the stacked columns of \mathbf{M} and the matrix $\mathbf{N} = \mathbf{I}_{n \times n} \otimes \mathbf{n}^\top$ amounts to

$$\mathbf{N} = \begin{bmatrix} \mathbf{n}^\top & \mathbf{0}_n^\top & \mathbf{0}_n^\top & \cdots & \mathbf{0}_n^\top & \mathbf{0}_n^\top \\ \mathbf{0}_n^\top & \mathbf{n}^\top & \mathbf{0}_n^\top & \cdots & \mathbf{0}_n^\top & \mathbf{0}_n^\top \\ & & \cdots & & & \\ \mathbf{0}_n^\top & \mathbf{0}_n^\top & \mathbf{0}_n^\top & \cdots & \mathbf{0}_n^\top & \mathbf{n}^\top \end{bmatrix} \quad (9)$$

where $\mathbf{0}_n$ denotes the n vector of all zeros.

Similar arguments apply to the expressions in the two equality constraints of problem (5). First of all, we have

$$\mathbf{M}^\top \mathbf{1}_n = \mathbf{C}_c \mathbf{m} \quad (10)$$

where the vector \mathbf{m} is given as above and matrix \mathbf{C}_c is defined as

$$\mathbf{C}_c = \mathbf{I}_{n \times n} \otimes \mathbf{1}_n^\top \quad (11)$$

Second of all, we observe that

$$\mathbf{M} \mathbf{1}_n = \mathbf{C}_r \mathbf{m} \quad (12)$$

where matrix \mathbf{C}_r is defined as

$$\mathbf{C}_r = \mathbf{1}_n^\top \otimes \mathbf{I}_{n \times n} \quad (13)$$

With respect to the inequality constraint of problem (5), we note that $\mathbf{M} \geq \mathbf{0}_{n \times n} \Leftrightarrow -\mathbf{M} \leq \mathbf{0}_{n \times n}$ as well as

$$-\mathbf{M} \leq \mathbf{0}_{n \times n} \Leftrightarrow -\mathbf{m} \leq \mathbf{0}_{n^2} \Leftrightarrow -\mathbf{I}_{n^2 \times n^2} \mathbf{m} \leq \mathbf{0}_{n^2} \quad (14)$$

Using all of this, we can therefore rewrite our matrix linear programming problem in terms of a vector linear programming problem, namely

$$\begin{aligned} \mathbf{p} = \operatorname{argmin}_{\mathbf{m} \in \mathbb{R}^{n^2}} & -\mathbf{x}^\top \mathbf{N} \mathbf{m} \\ & -\mathbf{I}_{n^2 \times n^2} \mathbf{m} \leq \mathbf{0}_{n^2} \\ \text{s.t.} & \quad \mathbf{C}_r \mathbf{m} = \mathbf{1}_n \\ & \quad \mathbf{C}_c \mathbf{m} = \mathbf{1}_n \end{aligned} \quad (15)$$

Assuming this problem has been solved, we can recover the originally sought after permutation matrix as $\mathbf{P} = \text{mat}(\mathbf{p})$ so as to then obtain the sorted version $\mathbf{y} = \mathbf{P} \mathbf{x}$ of \mathbf{x} .

However, while (15) is beginning to resemble the common form of a linear programming problem in (1), we are not quite there yet. The major difference between (1) and (15) is that the former involves only one equality constraint whereas the latter involves two.

But we can easily iron out these (mainly notational) differences. Indeed, the substitutions

$$\mathbf{c}^\top = -\mathbf{x}^\top \mathbf{N} \quad (16)$$

$$\mathbf{A} = -\mathbf{I}_{n^2 \times n^2} \quad (17)$$

$$\mathbf{b} = \mathbf{0}_{n^2} \quad (18)$$

and, crucially

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_r \\ \mathbf{C}_c \end{bmatrix} \quad (19)$$

$$\mathbf{d} = \begin{bmatrix} \mathbf{1}_n \\ \mathbf{1}_n \end{bmatrix} \quad (20)$$

allow us to also write the problem in (15) as follows

$$\begin{aligned} \mathbf{p} = \operatorname{argmin}_{\mathbf{m} \in \mathbb{R}^{n^2}} & \mathbf{c}^\top \mathbf{m} \\ \text{s.t.} & \quad \mathbf{A} \mathbf{m} \leq \mathbf{b} \\ & \quad \mathbf{C} \mathbf{m} = \mathbf{d} \end{aligned} \quad (21)$$

This is now a problem of the exact same form as in (1). In practice, we can therefore apply *SciPy*'s `linprog` function to solve it. In other words, we can practically use linear programming to sort a collection of numbers.

3 PRACTICAL COMPUTATION

Next, we look at how to use *SciPy* to solve the linear program in (21) and thus to sort the entries of a vector $\mathbf{x} \in \mathbb{R}^n$.

To work with a practical example, we first create a random vector \mathbf{x} of, say, $n = 5$ entries $0 \leq x_i \leq 100$ which we represent as a one-dimensional *NumPy* array. To keep things legible, we will force the x_i to be integers and proceed as follows

```
n = 5
vecX = rnd.randint(100, size=n)
```

In order to inspect the entries of this random vector, we simply use

```
print(vecX)
```

and may obtain something like this

```
[46 52 12 10 51]
```

Given \mathbf{x} , we next initialize a *NumPy* array which represent the auxiliary vector \mathbf{n} . To this end, we may use

```
vecN = np.arange(n) + 1
```

to obtain

```
[1 2 3 4 5]
```

Having completed all these preparatory steps, we next set up the ingredients of the linear program in (21). To this end, we may proceed as shown in Listing 1.

Listing 1: setting up the linear programming problem in (21)

```

1 matI = np.eye(n)
2 vec1 = np.ones(n)
3
4 matN = np.kron(matI, vecN)
5 matCc = np.kron(matI, vec1)
6 matCr = np.kron(vec1, matI)
7
8 matC = np.vstack((matCr, matCc))
9 vecD = np.hstack((vec1, vec1))
10
11 vecC = -vecX @ matN

```

Arrays `matI` and `vec1` represent the $n \times n$ identity matrix and the n -dimensional vector of all ones, respectively.

Arrays `matN`, `matCc`, and `matCr` represent the matrices N , C_c , and C_r which we defined in equations (8), (11), and (13). In order to compute these arrays, we apply the *NumPy* function `kron` which implements the Kronecker product. In other words, lines 4–6 in Listing 1 directly correspond to the mathematical expressions in (8), (11), and (13).

In line 8, we initialize an array `matC` which represents the $2n \times n^2$ matrix C defined in (19) and array `vecD` in line 9 represents vector d defined in (20).

Finally, in line 11, we compute an array `vecC` which represents vector c in (16). Long-time readers of the series will have noticed our use of

```
-vecX @ matN
```

instead of a more cumbersome expression such as

```
-np.dot(vecX, matN)
```

Both expressions are indeed equivalent since the infix operator `@` realizes matrix-vector- or matrix-matrix multiplication. While we rarely used `@` in previous notes, [PEP 465](#) recommends its use and we have lately come to appreciate the convenience it provides.

Attentive readers will have also noticed that we did not implement matrix A and vector b in (17) and (18), respectively. In our current setting, this is indeed unnecessary. To understand why, we recall our discussion in [9] where we pointed out that the questionable default behavior of *SciPy*'s `linprog` function is to assume that the solution of a linear program is non-negative. In our current setting, this is actually convenient as we do not have to care about the non-negativity enforcing inequality constraint in (21).

Hence, we are good to go and can finally invoke function `linprog` in *SciPy*'s `optimize` module. To solve our problem, we simply call

```
result = opt.linprog(vecC, A_eq=matC, b_eq=vecD)
```

Note that, in contrast to our examples in [9, 10], we now set the parameters `A_eq` and `b_eq` which define equality constraints. Another difference to our earlier examples is that `linprog` now issues the following warning

```
OptimizeWarning: A_eq does not appear to be of full
row rank. To improve performance, check the problem
formulation for redundant equality constraints.
```

This is certainly correct; matrix C passed in the argument `A_eq` is rank deficient but there is nothing we can do about it. Indeed, the warning appears to be inconsequential, because printing `result` for inspection yields

```

con: array([-3.96327415e-12, ...])
fun: -636.0000000020979
message: 'Optimization terminated successfully.'
nit: 7
slack: array([], dtype=float64)
status: 0
success: True
x: array([4.88698700e-13, ...])

```

which indicates that `linprog` succeeded in finding a solution.

The following statements therefore determine the solution p of our linear program and the corresponding permutation matrix P

```
vecP = result.x
matP = vecP.reshape(n, n).T
```

and subsequent (rounded) printing of the array `matP` results in

```

[[0.  0.  0.  1.  0.]
 [0.  0.  1.  0.  0.]
 [1.  0.  0.  0.  0.]
 [0.  0.  0.  0.  1.]
 [0.  1.  0.  0.  0.]]

```

Finally, in order to verify that the above permutation matrix solves our problem, we use

```
print ('vecX = ', vecX.astype(float))
print ('vecY = ', matP @ vecX)
```

and obtain

```
vecX = [46.  52.  12.  10.  51.]
vecY = [10.  12.  46.  51.  52.]
```

Success! We have solved a linear programming problem to get a sorted version $y = Px$ of an unordered vector x .

NOTE: While our example corroborates that sorting can be done via linear programming, we must point out that this approach is rather slow and typically suffers from numerical instabilities. Above, we see (from looking at `result.con`) that `linprog` found a solution with non-zero residuals for the equality constraint of our problem. As a consequence, the entries of `result.x` are very close to 0 or 1 but not exactly 0 or 1. When working with a small ($n = 5$) input vector x , this (floating point) imprecision does not have any noticeable impact. However, tiny errors may add up. Readers are encouraged to verify this for themselves and see what happens when considering a vector x of, say, $n = 100$ elements.

4 SUMMARY AND OUTLOOK

This note showed that linear programming is a surprisingly versatile tool. That is, many familiar problems we would not usually consider as a linear programming problem can actually be expressed as such and therefore be solved using linear programming solvers.

The particular problem we explored was the problem of sorting a lists of real numbers. We derived a corresponding linear program and solved it using function `linprog` in *SciPy*'s `optimize` module. In contrast to the practical applications in our previous notes [9, 10], the sorting problem required us to work with equality constraints which was easily done. Yet, for the particular equality constraint we had to consider, `linprog` issued a rank deficiency warning. While this was of little consequence for our toy example of sorting $n = 5$ numbers, such warnings are there for a reason.

Indeed, we cannot recommend the above approach as a practical solution. For larger n , it is much slower than conventional sorting algorithms and will suffer from floating point imprecision.

On the other hand, the fact that we may conceptualize sorting as a linear algebraic optimization problem is of interest to those working on next generation computing paradigms. Indeed, (21) can be further rewritten and cast as a binary unconstrained quadratic optimization problem for quantum computing [1, 2, 4]. This, too, is not really remarkable because sorting is not the kind of problem that calls for heavy machinery. However, the way of thinking we sketched in this note promises new solutions for much, much more demanding permutation problems [3, 6, 7, 11] and we will get back to it later.

A APPENDIX

In this appendix, we prove the not so obvious claims we made in section 2.

LEMMA A.1. *If $\mathbf{n}, \mathbf{x} \in \mathbb{R}^n$ are two given vectors and $\mathbf{M} \in \mathbb{R}^{n \times n}$ denotes a variable matrix, then the function $f : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ with*

$$f(\mathbf{M}) = \mathbf{n}^\top \mathbf{M} \mathbf{x}$$

is linear in \mathbf{M} .

PROOF. Consider $a, b \in \mathbb{R}$ and $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$, then

$$\begin{aligned} f(a\mathbf{A} + b\mathbf{B}) &= \mathbf{n}^\top (a\mathbf{A} + b\mathbf{B}) \mathbf{x} \\ &= a\mathbf{n}^\top \mathbf{A} \mathbf{x} + b\mathbf{n}^\top \mathbf{B} \mathbf{x} = a f(\mathbf{A}) + b f(\mathbf{B}) \end{aligned}$$

□

LEMMA A.2. *Let $\mathbf{n} \in \mathbb{R}^n$ and $\mathbf{M} \in \mathbb{R}^{n \times n}$. Then there exist a vector $\mathbf{m} \in \mathbb{R}^{n^2}$ and a matrix $\mathbf{N} \in \mathbb{R}^{n \times n^2}$ such that*

$$\mathbf{M}^\top \mathbf{n} = \mathbf{N} \mathbf{m}$$

In particular, the claim holds true for

$$\begin{aligned} \mathbf{m} &= \text{vec}(\mathbf{M}) \\ \mathbf{N} &= \mathbf{I} \otimes \mathbf{n}^\top \end{aligned}$$

where \mathbf{I} is the $n \times n$ identity matrix and \otimes denotes the Kronecker product.

PROOF. Consider the vector $\mathbf{v} = \mathbf{M}^\top \mathbf{n}$. In terms of column \mathbf{m}_i of matrix $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_n]$, entry v_i of \mathbf{v} is given by $v_i = \mathbf{m}_i^\top \mathbf{n}$ or, equivalently, $v_i = \mathbf{n}^\top \mathbf{m}_i$.

If we introduce two (much) larger vectors $\mathbf{m} \in \mathbb{R}^{n^2}$ and $\mathbf{n}_i \in \mathbb{R}^{n^2}$ where $\mathbf{m} = \text{vec}(\mathbf{M})$ contains the stacked columns of \mathbf{M} and

$$\mathbf{n}_i^\top = \left[\underbrace{0^\top \dots 0^\top}_{i-1 \text{ times}} \mathbf{n}^\top \underbrace{0^\top \dots 0^\top}_{n-i \text{ times}} \right]$$

with $\mathbf{0} \in \mathbb{R}^n$, we can also write $v_i = \mathbf{n}_i^\top \mathbf{m}$.

Hence, if we gather the n different vectors \mathbf{n}_i^\top as the rows of a matrix $\mathbf{N} \in \mathbb{R}^{n \times n^2}$, i.e. if we consider

$$\mathbf{N} = \begin{bmatrix} \mathbf{n}^\top & \mathbf{0}^\top & \mathbf{0}^\top & \dots & \mathbf{0}^\top & \mathbf{0}^\top \\ \mathbf{0}^\top & \mathbf{n}^\top & \mathbf{0}^\top & \dots & \mathbf{0}^\top & \mathbf{0}^\top \\ \dots & \dots & \dots & \dots & \dots & \dots \\ \mathbf{0}^\top & \mathbf{0}^\top & \mathbf{0}^\top & \dots & \mathbf{0}^\top & \mathbf{n}^\top \end{bmatrix} = \mathbf{I} \otimes \mathbf{n}^\top$$

we find that \mathbf{v} can also be written as $\mathbf{v} = \mathbf{N} \mathbf{m}$. □

ACKNOWLEDGMENTS

This material was produced within the Competence Center for Machine Learning Rhine-Ruhr (ML2R) which is funded by the Federal Ministry of Education and Research of Germany (grant no. 01S18038C). The authors gratefully acknowledge this support.

REFERENCES

- [1] C. Bauckhage, E. Brito, K. Cvejovski, C. Ojeda, R. Sifa, and S. Wrobel. 2017. Ising Models for Binary Clustering via Adiabatic Quantum Computing. In *Proc. EMM-CVPR*. Springer.
- [2] C. Bauckhage, R. Sanchez, and R. Sifa. 2020. Problem Solving with Hopfield Networks and Adiabatic Quantum Computing. In *Proc. IJCNN*. IEEE.
- [3] C. Bauckhage, R. Sifa, and S. Wrobel. 2020. Adiabatic Quantum Computing for Max-Sum Diversification. In *Proc. SDM*. SIAM.
- [4] F. Glover, G. Kochenberger, and Y. Du. 2018. A Tutorial on Formulating and Using QUBO Models. *arXiv:1811.11538 [cs.DS]* (2018).
- [5] G.H. Hardy, J.E. Littlewood, and G. Polya. 1952. *Inequalities*. Cambridge University Press.
- [6] J. Kunegis, D. Fay, and C. Bauckhage. 2010. Network Growth and the Spectral Evolution Model. In *Proc. CIKM*. ACM.
- [7] A. Nowak, S. Villar, A.S. Bandeira, and J. Bruna. 2017. Revised Note on Learning Algorithms for Quadratic Assignment with Graph Neural Networks. *arXiv:1706.07450 [stat.ML]* (2017).
- [8] T.E. Oliphant. 2007. Python for Scientific Computing. *Computing in Science & Engineering* 9, 3 (2007).
- [9] P. Welke and C. Bauckhage. 2020. *ML2R Coding Nuggets: Linear Programming for Robust Regression*. Technical Report. MLAI, University of Bonn.
- [10] P. Welke and C. Bauckhage. 2020. *ML2R Coding Nuggets: Solving Linear Programming Problems*. Technical Report. MLAI, University of Bonn.
- [11] M.M. Zavlanos and G. J. Pappas. 2008. A Dynamical Systems Approach to Weighted Graph Matching. *Automatica* 44, 11 (2008).
- [12] G.M. Ziegler. 1995. *Lectures on Polytopes*. Springer.