


ML2R Coding Nuggets

Hopfield Nets for Set Cover

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

ABSTRACT

We once again revisit the minimum set cover problem and show that the underlying integer linear program can be rewritten as a quadratic unconstrained binary optimization problem. This can then be cast as an energy minimization problem which we solve by running Hopfield nets. Using multiple restarts, our simple *NumPy* implementation consistently produces good results.

1 INTRODUCTION

Previously [1, 2], we looked at this basic¹ variant of the **set cover problem**: Given a set \mathcal{X} of m elements x_i and a set \mathcal{Y} of n subsets \mathcal{Y}_j of \mathcal{X} which are guaranteed to cover \mathcal{X} , determine a smallest subset \mathcal{Z} of \mathcal{Y} that covers \mathcal{X} .

To better grasp the nature of this problem, we considered a simple example where

$$\mathcal{X} = \{x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10}\}$$

$$\mathcal{Y} = \left\{ \begin{array}{l} \mathcal{Y}_1 = \{x_2, \} \\ \mathcal{Y}_2 = \{x_1, x_2, x_3, \} \\ \mathcal{Y}_3 = \{x_4, x_5, \} \\ \mathcal{Y}_4 = \{x_8, x_9, x_{10}, \} \\ \mathcal{Y}_5 = \{x_4, x_5, x_6, x_7, \} \\ \mathcal{Y}_6 = \{x_1, x_3, x_5, x_7, x_9, \} \end{array} \right\}$$

Looking at this example, it is clear that the $\mathcal{Y}_j \in \mathcal{Y}$ cover \mathcal{X} , because

$$\bigcup_{\mathcal{Y}_j \in \mathcal{Y}} \mathcal{Y}_j = \mathcal{X}$$

We further realize that the following smaller subset $\mathcal{Z} = \{\mathcal{Y}_2, \mathcal{Y}_4, \mathcal{Y}_5\}$ of \mathcal{Y} also covers \mathcal{X} , because

$$\bigcup_{\mathcal{Y}_j \in \mathcal{Z}} \mathcal{Y}_j = \mathcal{X}$$

Moreover, it is easy to verify that \mathcal{Z} is indeed the single smallest subset of \mathcal{Y} that covers \mathcal{X} and therefore solves our problem.

While this example makes set covering look easy, we discussed that it is NP-hard in general [2]. This was further confirmed by our observation that minimum set cover is actually an integer linear programming problem in disguise [1].

To see this once again, recall that, if we are given a certain set $\mathcal{X} = \{x_1, \dots, x_m\}$, we can represent any subset $\mathcal{Y}_j \subseteq \mathcal{X}$ in terms of a binary indicator vector $\mathbf{y}_j \in \{0, 1\}^m$ where

$$[\mathbf{y}_j]_i = \begin{cases} 1 & \text{if } x_i \in \mathcal{Y}_j \\ 0 & \text{otherwise} \end{cases}$$

We also recall that indicator vector representations extend to sets of subsets: If we are given a whole set $\mathcal{Y} = \{\mathcal{Y}_1, \dots, \mathcal{Y}_n\}$ of subsets $\mathcal{Y}_j \subseteq \mathcal{X}$, we can represent this collection as a matrix $Y \in \{0, 1\}^{m \times n}$ where $Y = [\mathbf{y}_1 \ \dots \ \mathbf{y}_n]$.

For instance, if we represent the subsets $\mathcal{Y}_1, \dots, \mathcal{Y}_6$ of our simple example using indicator vectors $\mathbf{y}_1, \dots, \mathbf{y}_6$, gather them as the columns of a matrix Y , and look at the transpose of this matrix, we find

$$Y^T = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}$$

which is easily recognized as an abstract representation of set \mathcal{Y} as presented above.

Finally, we recall that we may represent the given set \mathcal{X} in terms of an indicator vector $\mathbf{x} = \mathbf{1}_m$ of all ones and any subset $\mathcal{Z} \subseteq \mathcal{Y}$ as an indicator vector $\mathbf{z} \in \{0, 1\}^n$. But this means that we may think of set covering as a linear algebra problem: If we could find a vector \mathbf{z} such that the linear combination $\mathbf{x}(\mathbf{z}) = \mathbf{y}_1 \cdot z_1 + \dots + \mathbf{y}_n \cdot z_n = Y\mathbf{z}$ obeys $\mathbf{x}(\mathbf{z}) \geq \mathbf{1}_m$, we would have found an indicator vector which identifies a collection \mathcal{Z} of subsets \mathcal{Y}_j of \mathcal{Y} which cover \mathcal{X} .

Note, however, that the set cover problem does not simply ask for any cover of a given set but for a minimum cover of said set. Since we can express this objective by demanding that the sum of the entries of \mathbf{z} , i.e. the inner product $\mathbf{1}_n^T \mathbf{z}$, is to be minimized, we arrive at the following formulation of the minimum set cover problem in terms of an **integer linear program (ILP)**

$$\begin{aligned} \mathbf{z}^* &= \operatorname{argmin}_{\mathbf{z} \in \{0, 1\}^n} \mathbf{1}_n^T \mathbf{z} \\ &\text{s.t. } Y\mathbf{z} \geq \mathbf{1}_m \end{aligned} \tag{1}$$

Previously, we solved this ILP using correspondingly adapted earlier code for approximative greedy set covering. Our strategy in this note will be fundamentally different.

Next, in section 2, we will rewrite the ILP in (1) as a **quadratic unconstrained binary optimization problem (QUBO)**. This will require some work but also allows us to treat minimum set cover as a Hopfield energy minimization problem. We therefore also review Hopfield nets for problem solving. In section 3, we then put them to work and present corresponding *NumPy* implementations. Readers who would like to experiment with our code snippets should be familiar with *NumPy* / *SciPy* [16] and need to

```
import numpy as np
import numpy.random as rnd
```

¹There also exists the more general notion of *weighted* set cover problems.

2 THEORY

Our primary goal in this section is to establish that the set cover problem can be cast as a QUBO of the following form

$$\mathbf{s}^* = \operatorname{argmin}_{\mathbf{s} \in \{\pm 1\}^N} \mathbf{s}^\top \mathbf{Q} \mathbf{s} + \mathbf{q}^\top \mathbf{s} \quad (2)$$

where the matrix $\mathbf{Q} \in \mathbb{R}^{N \times N}$ and vector $\mathbf{q} \in \mathbb{R}^N$ contain problem specific parameters.

The fact that this can be done is not new. However, the frequently cited derivation of this result in [12] lacks details. We therefore get inspiration from the general recipes in [9], adapt them to the set cover problem, and present a step-by-step derivation starting from the ILP in (1).

Our secondary goal is to tie in the QUBO formulation of set cover with Hopfield nets for combinatorial problem solving [10, 11]. Here, our discussion will largely recap points we made in earlier notes (see, for instance, [3, 4, 6]).

2.1 Set Cover as a QUBO

Our first step in turning the ILP in (1) into a QUBO as in (2) will be to rewrite the inequality constraint

$$\mathbf{Y} \mathbf{z} \geq \mathbf{1}_m \quad (3)$$

as an equality constraint.

To this end, we introduce slack variables o_1, \dots, o_m , gather them in a vector $\mathbf{o} \in \mathbb{R}^m$, and consider

$$\mathbf{Y} \mathbf{z} - \mathbf{o} = \mathbf{1}_m \quad (4)$$

If we rearrange this equation, we obtain the following expression for the slack variables

$$\mathbf{o} = \mathbf{Y} \mathbf{z} - \mathbf{1}_m \quad (5)$$

Why do we call our slack variables o_i ? Because they they allow us to cope with “overshooting”. To see what we mean by this, we need to take a step back ...

From the point of view of linear algebra with indicator vectors, finding a set cover is tantamount to finding linear coefficients $z_j \in \{0, 1\}$ such that the linear combination $\mathbf{x}(\mathbf{z}) = \mathbf{y}_1 \cdot z_1 + \dots + \mathbf{y}_n \cdot z_n$ obeys $\mathbf{x}(\mathbf{z}) \geq \mathbf{1}_m$.

Since the \mathbf{y}_j form the columns of matrix \mathbf{Y} , we usually express this goal as having to find a coefficient vector \mathbf{z} such that $\mathbf{x}(\mathbf{z}) = \mathbf{Y} \mathbf{z} \geq \mathbf{1}_m$.

To better understand the rationale behind this idea, we shall briefly motivate it w.r.t. the exemplary set cover problem from the introduction.

There, we are given a set of subsets $\mathcal{Y} = \{\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_6\}$ which we represent in terms of an indicator matrix $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \dots \ \mathbf{y}_6]$ which reads

$$\mathbf{Y} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix}$$

We already know that the collection $\mathcal{Z}_{245} = \{\mathcal{Y}_2, \mathcal{Y}_4, \mathcal{Y}_5\}$ optimally solves our problem; it can be represented in terms of the indicator vector $\mathbf{z}_{245} = [0 \ 1 \ 0 \ 1 \ 1 \ 0]^\top$.

Another, sub-optimal cover is given by $\mathcal{Z}_{1456} = \{\mathcal{Y}_1, \mathcal{Y}_4, \mathcal{Y}_5, \mathcal{Y}_6\}$ which can be represented by the indicator vector $\mathbf{z}_{1456} = [1 \ 0 \ 0 \ 1 \ 1 \ 1]^\top$

If we compute the expression $\mathbf{x}(\mathbf{z}_{245}) = \mathbf{Y} \mathbf{z}_{245}$, we find

$$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 1 \\ 1 \\ 1 \\ 0 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

Since all the $m = 10$ entries of the vector $\mathbf{x}(\mathbf{z}_{245})$ are covered, i.e. are greater than 0, we have found a vector $\mathbf{z} = \mathbf{z}_{245}$ such that $\mathbf{Y} \mathbf{z} \geq \mathbf{1}_{10}$. In fact, in this particular case, we actually have equality $\mathbf{Y} \mathbf{z} = \mathbf{1}_{10}$.

If we compute the expression $\mathbf{x}(\mathbf{z}_{1456}) = \mathbf{Y} \mathbf{z}_{1456}$, we find

$$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 2 \\ 2 \\ 1 \\ 1 \\ 1 \\ 2 \\ 2 \end{bmatrix}$$

Again, all entries of $\mathbf{x}(\mathbf{z}_{1456})$ are covered so that we have found another vector $\mathbf{z} = \mathbf{z}_{1456}$ such that $\mathbf{Y} \mathbf{z} \geq \mathbf{1}_{10}$. However, the 5th and 9th entry of $\mathbf{x}(\mathbf{z}_{1456})$ are covered twice. Speaking in terms of the sets $\mathcal{Y}_1, \mathcal{Y}_4, \mathcal{Y}_5$, and \mathcal{Y}_6 , this is to say that they cover the given universe \mathcal{X} but are not disjoint.

In our silly example this means that the cover $\mathcal{Z}_{1456} = \{\mathcal{Y}_1, \mathcal{Y}_4, \mathcal{Y}_5, \mathcal{Y}_6\}$ is sub-optimal. However, in general, it may well happen that a minimum set cover involves overlapping subsets.

From the point of view of integer linear programming for set cover, this means that we have to incorporate inequalities into our considerations. Fundamentally, this is because the union (\cup) of sets behaves differently than the addition ($+$) of indicator vectors. For instance, the union of $\mathcal{Y}_2 = \{x_1, x_2, x_3\}$ and $\mathcal{Y}_6 = \{x_1, x_3, x_5, x_7, x_9\}$ is given by

$$\mathcal{Y}_2 \cup \mathcal{Y}_6 = \{x_1, x_2, x_3, x_5, x_7, x_9\}$$

whereas the sum of the corresponding indicator vectors $\mathbf{y}_2 = [1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]^\top$ and $\mathbf{y}_6 = [1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1]^\top$ amounts to

$$\mathbf{y}_2 + \mathbf{y}_6 = [1 \ 1 \ 2 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1]^\top$$

Given these observations it is clear why we say that an indicator vector \mathbf{z} that represents a subset $\mathcal{Z} \subseteq \mathcal{Y}$ also represents a cover of \mathcal{X} exactly if $\mathbf{x}(\mathbf{z}) = \mathbf{Y} \mathbf{z} \geq \mathbf{1}_m$. For those entries $[\mathbf{x}(\mathbf{z})]_i$ of $\mathbf{x}(\mathbf{z})$ where $[\mathbf{x}(\mathbf{z})]_i > 1$, we colloquially say that the linear algebraic solution “overshot” the target.

Our next step in turning the ILP in (1) into a QUBO as in (2) will be to treat the slack variables in \mathbf{o} as additional latent or hidden variables. In fact, equation (5) suggests that \mathbf{o} is not fixed but depends on the sought after indicator vector \mathbf{z} .

How small or large can the entries o_i of \mathbf{o} possibly be? Well, if \mathbf{z} solves the ILP in (1), then we have the following lower bound

$$o_{i,\min} = 0 \leq o_i = \sum_{j=1}^n Y_{ij} z_j - 1 \quad (6)$$

because at least one Y_{ij} must have a value of 1 and must have been selected into the set cover ($\exists j : z_j = 1$). By the same token, we have the following upper bound

$$o_i = \sum_{j=1}^n Y_{ij} z_j - 1 \leq \sum_{j=1}^n Y_{ij} - 1 = o_{i,\max} \quad (7)$$

because all Y_{ij} may have a value of 1 and may have been selected into the set cover ($\forall j : z_j = 1$)

Our next crucial observation is that the o_i must necessarily be integers. This is because $Y_{ij} \in \{0, 1\}$ and $z_j \in \{0, 1\}$ such that any sum of products of these binaries will be an element of \mathbb{Z}_+ . But

this is to say that each o_i can be written as a binary number or, more precisely, as a binary expansion. That is, letting

$$P_i = \max\left\{0, \lfloor \log_2(o_{i,\max}) \rfloor\right\} \quad (8)$$

we may write

$$o_i = \sum_{p=0}^{P_i} 2^p c_p \quad (9)$$

where the $c_p \in \{0, 1\}$ are binary expansion coefficients. Even more, if we introduce binary coefficient vectors $\mathbf{c}_i \in \{0, 1\}^{P_i+1}$ and ‘‘power of two’’ vectors

$$\mathbf{p}_i = [2^0 \ 2^1 \ \dots \ 2^{P_i}]^\top \quad (10)$$

we may write

$$o_i = \mathbf{p}_i^\top \mathbf{c}_i \quad (11)$$

Even better, if we define an overall ‘‘power of two’’ matrix and coefficient vector

$$\mathbf{P} = \begin{bmatrix} \mathbf{p}_1^\top & \mathbf{0}^\top & \dots & \mathbf{0}^\top \\ \mathbf{0}^\top & \mathbf{p}_2^\top & \dots & \mathbf{0}^\top \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^\top & \mathbf{0}^\top & \dots & \mathbf{p}_n^\top \end{bmatrix} \quad (12)$$

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \vdots \\ \mathbf{c}_n \end{bmatrix} \quad (13)$$

we can express the whole vector \mathbf{o} of slack variables o_i as

$$\mathbf{o} = \mathbf{P}\mathbf{c} \quad (14)$$

Note: the number of columns of \mathbf{P} and the corresponding number of rows of \mathbf{c} is a variable that depends on the specific set cover problem we are dealing with. It is given by

$$n' = \sum_{i=1}^m (P_i + 1) \quad (15)$$

Plugging our new expression for \mathbf{o} into equation (4), we obtain

$$\mathbf{Y}\mathbf{z} - \mathbf{P}\mathbf{c} = \mathbf{1}_m \quad (16)$$

Next, we will further rewrite this equation to make it more compact. To this end, we horizontally stack matrices \mathbf{Y} and $-\mathbf{P}$ and vertically stack vectors \mathbf{z} and \mathbf{c} , i.e. we introduce

$$\hat{\mathbf{Y}} = [\mathbf{Y} \quad -\mathbf{P}] \quad (17)$$

$$\hat{\mathbf{z}} = \begin{bmatrix} \mathbf{z} \\ \mathbf{c} \end{bmatrix} \quad (18)$$

because this allows us to write (16) as

$$\hat{\mathbf{Y}}\hat{\mathbf{z}} = \mathbf{1}_m \quad (19)$$

As an important intermediate result our our efforts so far, we therefore have that inequality constrained ILP in (1) can also be written as an equality constrained ILP, namely

$$\hat{\mathbf{z}}^* = \underset{\hat{\mathbf{z}} \in \{0,1\}^N}{\operatorname{argmin}} \mathbf{1}_N^\top \hat{\mathbf{z}} \quad (20)$$

$$\text{s.t. } \hat{\mathbf{Y}}\hat{\mathbf{z}} = \mathbf{1}_m$$

where $N = n + n'$. Since $N \geq n$, the price we have to pay for this ‘‘simplification’’ is that we now must solve for N binary variables \hat{z}_j^* rather than for n binary variables z_j^* as before. However, this will pay off because it is now easy to turn the ILP in (20) into a QUBO as in (2)

Continuing with our derivation, we next note the equivalence

$$\hat{\mathbf{Y}}\hat{\mathbf{z}} = \mathbf{1}_m \Leftrightarrow \|\hat{\mathbf{Y}}\hat{\mathbf{z}} - \mathbf{1}_m\|^2 = 0 \quad (21)$$

and the fact that we may write this squared Euclidean distance as

$$\|\hat{\mathbf{Y}}\hat{\mathbf{z}} - \mathbf{1}_m\|^2 = \hat{\mathbf{z}}^\top \hat{\mathbf{Y}}^\top \hat{\mathbf{Y}}\hat{\mathbf{z}} - 2 \mathbf{1}_m^\top \hat{\mathbf{Y}}\hat{\mathbf{z}} + \mathbf{1}_m^\top \mathbf{1}_m \quad (22)$$

Since the inner product $\mathbf{1}_m^\top \mathbf{1}_m$ is a constant independent of $\hat{\mathbf{z}}$, we therefore have the following Lagrangian for the problem in (20)

$$\begin{aligned} L(\hat{\mathbf{z}}, \lambda) &= \mathbf{1}_N^\top \hat{\mathbf{z}} + \lambda [\hat{\mathbf{z}}^\top \hat{\mathbf{Y}}^\top \hat{\mathbf{Y}}\hat{\mathbf{z}} - 2 \mathbf{1}_m^\top \hat{\mathbf{Y}}\hat{\mathbf{z}}] \\ &= \hat{\mathbf{z}}^\top [\lambda \hat{\mathbf{Y}}^\top \hat{\mathbf{Y}}] \hat{\mathbf{z}} + [\mathbf{1}_N^\top - 2 \lambda \mathbf{1}_m^\top \hat{\mathbf{Y}}] \hat{\mathbf{z}} \\ &\equiv \hat{\mathbf{z}}^\top \mathbf{R} \hat{\mathbf{z}} + \mathbf{r}^\top \hat{\mathbf{z}} \end{aligned} \quad (23)$$

Here, λ is a Lagrange multiplier which we henceforth treat as a parameter that needs to be set manually.

Now, all of this is to say that the original ILP in (1) or its modified version in (20) can just as well be cast a QUBO over binary decision variables $\hat{\mathbf{z}} \in \{0, 1\}^N$, namely

$$\hat{\mathbf{z}}^* = \underset{\hat{\mathbf{z}} \in \{0,1\}^N}{\operatorname{argmin}} \hat{\mathbf{z}}^\top \mathbf{R} \hat{\mathbf{z}} + \mathbf{r}^\top \hat{\mathbf{z}} \quad (24)$$

Finally, in order to turn this into a QUBO over bipolar decision variables $\mathbf{s} \in \{\pm 1\}^N$ such as in (2), we recall that binary and bipolar vectors are isomorphic in the sense that

$$\mathbf{s} = 2\hat{\mathbf{z}} - \mathbf{1}_N \Leftrightarrow \hat{\mathbf{z}} = \frac{1}{2} [\mathbf{s} + \mathbf{1}_N] \quad (25)$$

Hence, if we plug $\hat{\mathbf{z}} = \frac{1}{2} [\mathbf{s} + \mathbf{1}_N]$ into the objective function of our QUBO in (24), we obtain

$$\begin{aligned} \hat{\mathbf{z}}^\top \mathbf{R} \hat{\mathbf{z}} + \mathbf{r}^\top \hat{\mathbf{z}} &= \frac{1}{4} [\mathbf{s} + \mathbf{1}_N]^\top \mathbf{R} [\mathbf{s} + \mathbf{1}_N] + \frac{1}{2} \mathbf{r}^\top [\mathbf{s} + \mathbf{1}_N] \\ &= \frac{1}{4} \mathbf{s}^\top \mathbf{R} \mathbf{s} + \frac{1}{2} [\mathbf{R} \mathbf{1}_N + \mathbf{r}]^\top \mathbf{s} + \text{const} \\ &\equiv \mathbf{s}^\top \mathbf{Q} \mathbf{s} + \mathbf{q}^\top \mathbf{s} + \text{const} \end{aligned} \quad (26)$$

In other words, we find that set covering can also be understood as the problem of solving

$$\mathbf{s}^* = \underset{\mathbf{s} \in \{\pm 1\}^N}{\operatorname{argmin}} \mathbf{s}^\top \mathbf{Q} \mathbf{s} + \mathbf{q}^\top \mathbf{s} \quad (27)$$

where

$$\mathbf{Q} = \frac{1}{4} \mathbf{R} \quad (28)$$

$$\mathbf{q} = \frac{1}{2} [\mathbf{R} \mathbf{1}_N + \mathbf{r}] \quad (29)$$

and

$$\mathbf{R} = \lambda \hat{\mathbf{Y}}^\top \hat{\mathbf{Y}} \quad (30)$$

$$\mathbf{r} = \mathbf{1}_N - 2 \lambda \hat{\mathbf{Y}}^\top \mathbf{1}_m \quad (31)$$

That is it! We now have a QUBO formulation of the set cover problem and may run Hopfield nets to try to solve it. If we do this, we will obtain a bipolar vector $\mathbf{s}^* \in \{-1, +1\}^N$ which we turn into a binary vector $\hat{\mathbf{z}}^* = \frac{1}{2} [\mathbf{s} + \mathbf{1}_N]$. The actually sought after indicator vector \mathbf{z}^* is then contained in the first n components of $\hat{\mathbf{z}}^*$.

2.2 Hopfield Nets for Set Cover

Recall that a Hopfield net is a recurrent neural network of N neurons s_1, \dots, s_N which are bipolar threshold units. Hence, if $\mathbf{w}_j \in \mathbb{R}^N$ and $\theta_j \in \mathbb{R}$ denote the input weights and threshold value of neuron s_j and if $\mathbf{s} = [s_1, \dots, s_N]^\top$ denotes the current global state of a Hopfield net, then neuron s_j computes its activation or local state as

$$s_j = \begin{cases} +1 & \text{if } \mathbf{w}_j^\top \mathbf{s} - \theta_j \geq 0 \\ -1 & \text{otherwise} \end{cases} \quad (32)$$

$$= \text{sign}(\mathbf{w}_j^\top \mathbf{s} - \theta_j)$$

Consequently, the global state of a Hopfield net always coincides with a bipolar vector $\mathbf{s} \in \{-1, +1\}^N$.

Also recall that we may gather all the weight- and threshold parameters of a Hopfield net in a matrix and a vector

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1^\top \\ \vdots \\ \mathbf{w}_N^\top \end{bmatrix} \quad \text{and} \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \vdots \\ \theta_N \end{bmatrix} \quad (33)$$

Given this weight matrix and threshold vector, the energy of a Hopfield net in state \mathbf{s} can be written as

$$E(\mathbf{s}) = -\frac{1}{2} \mathbf{s}^\top \mathbf{W} \mathbf{s} + \boldsymbol{\theta}^\top \mathbf{s} \quad (34)$$

Finally, recall the following crucial fact: *If the weight matrix of a Hopfield net is symmetric ($\mathbf{W} = \mathbf{W}^\top$) and hollow ($\text{diag}[\mathbf{W}] = \mathbf{0}$) and if its neurons update asynchronously (compute their activation one at a time), then the energy of a Hopfield net can never increase. As there are “only” 2^N distinct global states a Hopfield net can be in, it will converge to a state of (locally) minimum energy after finitely many updates [10].*

All of this suggests that a Hopfield net can (approximately) solve a QUBO. Apparently, we only need to let $\mathbf{W} = -2\mathbf{Q}$ and $\boldsymbol{\theta} = \mathbf{q}$ to obtain a Hopfield net whose energy function corresponds to the QUBO’s minimization objective.

However, there is a crucial subtlety we must pay attention to: Matrix \mathbf{Q} in (27) may not be hollow, i.e. may have non-zero diagonal elements. Yet, hollowness of the weight matrix \mathbf{W} of a Hopfield net is required to be able to guarantee that its global state eventually converges to a (local) minimum of its energy function.

We therefore note that we can write matrix \mathbf{Q} in (27) as a sum of a diagonal matrix \mathbf{D} and an off-diagonal matrix \mathbf{O} . The former represents the diagonal entries of \mathbf{Q} and the latter represents the off-diagonal entries of \mathbf{Q} . This way, the quadratic term in (27) becomes

$$\mathbf{s}^\top \mathbf{Q} \mathbf{s} = \mathbf{s}^\top [\mathbf{D} + \mathbf{O}] \mathbf{s} = \mathbf{s}^\top \mathbf{D} \mathbf{s} + \mathbf{s}^\top \mathbf{O} \mathbf{s} \quad (35)$$

Next, we point out that the term $\mathbf{s}^\top \mathbf{D} \mathbf{s}$ necessarily amounts to

$$\mathbf{s}^\top \mathbf{D} \mathbf{s} = \sum_j \sum_i s_j D_{ij} s_i = \sum_j D_{jj} s_j^2 \quad (36)$$

because \mathbf{D} is a diagonal matrix. Moreover, we will always have $s_j^2 = 1$, because $s_j \in \{-1, +1\}$. But this is to say that the term

$$\mathbf{s}^\top \mathbf{D} \mathbf{s} = \sum_j D_{jj} = \text{const} \quad (37)$$

has no impact on the solution of (27). Consequently, we may let

$$\mathbf{W} = -2[\mathbf{Q} - \mathbf{D}] \quad (38)$$

$$\boldsymbol{\theta} = \mathbf{q} \quad (39)$$

to obtain a Hopfield net with a hollow weight matrix that can (approximately) solve the problem in (27).

With respect to the update mechanism that governs the temporal evolution of a Hopfield net, we note that neuron s_u which updates its state in round t is typically chosen randomly. While this random selection mechanism is guaranteed to converge to a global state of (locally) minimum energy, it typically entails very slow convergence. We can do better than this.

Let \mathbf{s}_t and \mathbf{s}_{t+1} be two consecutive states during the temporal evolution of a Hopfield net and assume that neuron s_u updated at time t . It is well known that the energy difference in this situation amounts to $\Delta E = E(\mathbf{s}_{t+1}) - E(\mathbf{s}_t) = 2 \cdot [\mathbf{s}_t]_u \cdot (\mathbf{w}_u^\top \mathbf{s}_t - \theta_u) \leq 0$.

This can be used to determine which neuron should update its activation in iteration t so as to maximally decrease the network’s current energy $E(\mathbf{s}_t)$. Letting \odot denote the Hadamard product, we may compute a vector

$$\Delta \mathbf{e} = 2 \cdot \mathbf{s}_t \odot [\mathbf{W} \mathbf{s}_t - \boldsymbol{\theta}] \quad (40)$$

of expected energy decrements, determine the index u of its most negative entry

$$u = \underset{j}{\text{argmin}} [\Delta \mathbf{e}]_j \quad (41)$$

and then update neuron s_u to obtain a new global state where

$$[\mathbf{s}_{t+1}]_j = \begin{cases} [\mathbf{s}_t]_u \cdot \text{sign}[\Delta \mathbf{e}]_u & \text{if } j = u \\ [\mathbf{s}_t]_j & \text{otherwise} \end{cases} \quad (42)$$

This greedy selection of the updating neurons will obviously minimize the network’s energy but does so much faster than random selections. Alas, this mechanism may get trapped more easily in local minima of the energy landscape. However, set cover problems as considered in this note come with very rugged energy landscapes anyway so that random updates, too, are very likely to miss global minima. A valid strategy is thus to restart a Hopfield net many times from random global states and to keep track of the best solution found so far. Within this meta-algorithm, greedy updates of the states of the network will definitely be superior as they considerably reduce the duration of each round.

3 NUMPY IMPLEMENTATIONS

In this section, we will implement the above theory in *NumPy*. As a simple practical example, we once again consider the exemplary set cover problem in [2]. There we implemented the set \mathcal{X} as

$$\mathcal{X} = \{\emptyset, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$$

and realized the set of subsets \mathcal{Y} among which to find a cover \mathcal{Z} as

$$\mathcal{Y} = \text{dict}(\text{enumerate}([\{1\}, \{0, 1, 2\}, \{3, 4, 9\}, \{7, 8, 9\}, \{3, 4, 5, 6\}, \{0, 2, 4, 6, 8\}]))$$

Listing 1: converting sets to indicator vectors matrices

```
def sets2mat(setX, dctY):
    m, n = len(setX), len(dctY),
    matY = np.zeros((n,m))
    for key, vals in dctY.items():
        matY[key, list(vals)] = 1

    return matY.T
```

Listing 2: setting the parameters of a set cover Hopfield net

```
1 def hnet_init_params(matY, l=1.):
2     m, n = matY.shape
3
4     omax = np.sum(matY, axis=1) - np.ones(m)
5     pwrs = [np.floor(np.log2(o)) if o>0 else 0 for o in omax]
6
7     nprm = np.sum(pwrs).astype(int) + m
8     matP = np.zeros((m, nprm))
9
10    j = 0
11    for i, p in enumerate(pwrs):
12        vecP = 2**np.arange(p+1)
13        dimP = len(vecP)
14        matP[i, j:j+dimP] = vecP
15        j += dimP
16
17    matY = np.hstack((matY, -matP))
18    m, N = matY.shape
19
20    vec1m = np.ones(m)
21    vec1N = np.ones(N)
22
23    matR = l * matY.T @ matY
24    vecR = vec1N - 2 * l * matY.T @ vec1m
25
26    matQ = 0.25 * matR
27    vecQ = 0.50 * (matR @ vec1N + vecR)
28
29    matW = -2 * matQ; np.fill_diagonal(matW, 0)
30    vecT = vecQ
31
32    return matW, vecT
```

For our purposes in this note, we need to translate the dictionary representing \mathcal{Y} into a binary matrix Y . To this end, we may use function `sets2mat` in Listing 1 which resembles a procedure we discussed in [1]. Indeed,

```
matY = sets2mat(X, Y)
```

will produce a corresponding *NumPy* array `matY` of shape

```
m, n = matY.shape
```

Given this array, we can compute arrays `matW` and `vecT` which represent the weight matrix W and the threshold vector θ of a set cover Hopfield net. To this end we call

```
matW, vecT = hnet_init_params(matY)
```

where function `hnet_init_params` is defined in Listing 2. Its first parameter `matY` is self explanatory; its second parameter `l` represent the Lagrange multiplier λ which, when in doubt, defaults to 1.

Line 4 computes the vector \mathbf{o}_{\max} according to what has been worked out in (7). Line 5 computes the P_i according to (8) and line 7 computes n' according to (15).

Lines 8–15 compute matrix P defined in (12). First, we initialize an $m \times n'$ array `matP` of all zeros. We then iterate over the elements of list `pwrs` to compute arrays `vecP` representing the vectors \mathbf{p}_i in (10). These are inserted at appropriate locations into `matP`.

Listing 3: greedy state updates for Hopfield nets

```
def signum(x):
    return np.where(x >= 0, +1, -1)

def hnet_run_greedy(vecS, matW, vecT, tmax=100):
    for t in range(tmax):
        grad = matW @ vecS - vecT
        updt = np.argmin(vecS * grad)
        vecS[updt] = signum(grad[updt])

    return vecS
```

Listing 4: (re-)running Hopfield nets for set cover

```
1 ebst = np.inf
2 for trial in range(100):
3     vecS = rnd.binomial(n=1, p=0.5, size=len(vecT)) * 2 - 1
4     vecS = hnet_run_greedy(vecS, matW, vecT, tmax=2*len(vecT))
5     eact = -0.5 * vecS @ matW @ vecS + vecT @ vecS
6     if eact < ebst:
7         sbst = vecS
8         ebst = eact
```

Line 17 creates an array representing matrix \hat{Y} in (17) and lines 20–21 implement the vectors $\mathbf{1}_m$ and $\mathbf{1}_N$. Lines 23–24 realize R and \mathbf{r} in (30) and (31). These are then used to implement Q and \mathbf{q} in (28) and (29), respectively. Finally, lines 29–30 implement *NumPy* arrays `matW` and `vecT` which represent our Hopfield net parameters as defined in (38) and (39). Regarding the former, we point out our use of `np.fill_diagonal` which is simple (and correct!) way of making sure that network's weight matrix is hollow.

Code for greedy activation updates for evolving a Hopfield net is provided in function `hnet_run_greedy` in Listing 3. This is an almost verbatim *NumPy* implementation of the above mathematics; however, we note our use of the custom made sign function `signum` which we have applied several times before [3, 4, 6].

Having all required ingredients and mechanisms in place, we can now run the meta-algorithm for solving set cover problems which we sketched above. Corresponding code is shown in Listing 4.

In line 1, we initialize a variable `ebst` to keep track of energy values at the end of each run of our Hopfield net. In this example, we perform 100 runs. In each run, we first randomly initialize an array `vecS` representing the initial state s_0 of the Hopfield net. We then run `hnet_run_greedy` to determine a state s_t (where we here choose $t = 100$). Variable `eact` gets assigned the energy value $E(s_t)$; if it is less than the current best energy value in `ebst`, we assign s_t to `sbst` and update our bookkeeping variable `ebst`.

For our running example, we may inspect the outcome of the whole procedure using

```
vecZ = (sbst+1) / 2
vecZ = vecZ[:n]
print (vecZ.astype(int))
print ((matY @ vecZ).astype(int))
```

This produces

```
>>> [0 1 0 1 1 0]
>>> [1 1 1 1 1 1 1 1 1]
```

and therefore indicates that our approach was successfully able to solve an (admittedly very simple) set cover problem.

4 CONCLUSION

In this note, we made good on an earlier promise and showed that the combinatorial problem of finding minimum set covers can be tackled using Hopfield nets.

The main work we had to do was to transform the set cover ILP we derived in [1] into a QUBO. For those who have never seen the use of slack variables in QUBOs before, this may have been a rough ride. We therefore note that a more gentle introduction to this truly important technique can be found in an excellent tutorial paper by Glover and colleagues [9].

Once we derived a QUBO formulation of the set cover problem, it was straightforward to turn it into an energy minimization problem that can be solved by Hopfield nets. This is of practical interest, because set covering problems frequently arise in industry (examples include job scheduling, shift planning, location planning, route selection, and other logistics tasks).

The crux is that the NP-hardness of set cover requires (very) special purpose software to be able to solve such problems at scale. Hopfield nets, on the other hand, are simple omnipurpose tools that can be implemented on modern high performance architectures and therefore allow for highly efficient number crunching [7, 8, 14, 15].

Indeed, as we never fail to mention when discussing Hopfield nets, if a problem can be solved by running a Hopfield nets, it can also be solved on (adiabatic) quantum computers [5]. While these have not yet reached the level of technological maturity required for large scale industrial applications, there is considerable progress with respect to practical applications of quantum computing [13, 17, 18]. Should these developments continue at their current pace, it seems likely that quantum set cover solvers will be deployed in practice soon.

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. 01IS18038C). The authors gratefully acknowledge this support.

REFERENCES

- [1] C. Bauckhage. 2022. *ML2R Coding Nuggets: Greedy Set Cover with Binary Numpy Arrays*. Technical Report. MLAI, University of Bonn.
- [2] C. Bauckhage. 2022. *ML2R Coding Nuggets: Greedy Set Cover with Native Python Data Types*. Technical Report. MLAI, University of Bonn.
- [3] C. Bauckhage, F. Beaumont, and S. Müller. 2021. *ML2R Coding Nuggets: Hopfield Nets for Bipartition Clustering*. Technical Report. MLAI, University of Bonn.
- [4] C. Bauckhage, F. Beaumont, and S. Müller. 2021. *ML2R Coding Nuggets: Hopfield Nets for Max-Sum Diversification*. Technical Report. MLAI, University of Bonn.
- [5] C. Bauckhage, R. Sanchez, and R. Sifa. 2020. Problem Solving with Hopfield Networks and Adiabatic Quantum Computing. In *Proc. IJCNN*. IEEE.
- [6] C. Bauckhage and P. Welke. 2021. *ML2R Coding Nuggets: Hopfield Nets for Sorting*. Technical Report. MLAI, University of Bonn.
- [7] D. Biesner, R. Sifa, and C. Bauckhage. 2022. Solving Subset Sum Problems using Binary Optimization with Applications in Auditing and Financial Data Analysis. *TechRxiv* (2022).
- [8] S. Buschjäger, L. Pfahler, J. Buss, K. Morik, and W. Rhode. 2021. On-Site Gamma-Hadron Separation with Deep Learning on FPGAs. In *Proc. ECML PKDD*.
- [9] F. Glover, G. Kochenberger, R. Hennig, and Y. Du. 2022. Quantum Bridge Analytics I: A Tutorial on Formulating and Using QUBO Models. *Annals of Operations Research* 314 (2022).
- [10] J.J. Hopfield. 1982. Neural Networks and Physical Systems with Collective Computational Abilities. *PNAS* 79, 8 (1982).
- [11] J.J. Hopfield and D.W. Tank. 1985. "Neural" Computation of Decisions in Optimization Problems. *Biological Cybernetics* 52 (1985), 141–152.
- [12] A. Lucas. 2014. Ising Formulations of Many NP Problems. *Frontiers in Physics* 2, 5 (2014).
- [13] S. Mücke, R. Heese, S. Müller, M. Wolter, and N. Piatkowski. 2022. Quantum feature Selection. *arXiv:2203.13261 [quant-ph]* (2022).
- [14] S. Mücke, N. Piatkowski, and K. Morik. 2019. Hardware Acceleration of Machine Learning Beyond Linear Algebra. In *Proc. ECML/PKDD*.
- [15] S. Mücke, N. Piatkowski, and K. Morik. 2019. Learning Bit by Bit: Extracting the Essence of Machine Learning. In *Proc. KDML-LWDA*.
- [16] T.E. Oliphant. 2007. Python for Scientific Computing. *Computing in Science & Engineering* 9, 3 (2007).
- [17] N. Piatkowski, T. Gerlach, R. Hugues, R. Sifa, C. Bauckhage, and F. Barbaresco. 2022. Towards Bundle Adjustment for Satellite Imaging via Quantum Machine Learning. In *Proc. FUSION*.
- [18] N. Piatkowski and C. Zoufal. 2022. On Quantum Circuits for Discrete Graphical Models. *arXiv:2206.00398 [quant-ph]* (2022).