## Graph Matching & Linear Programming

Part 1: Matchings in Graphs

Machine Learning for Structured Data Vlad Niculae · LTL, UvA · https://vene.ro/mlsd

## Graph Matching & Linear Programming

#### **1** Matchings in Graphs

**2** Finding a Maximum-Weight Matching

**3** Linear Programming

**4** Bipartite graphs and the assignment problem

## Today's agenda

- So far, we've seen two examples of structure prediction using dynamic programming: I wouldn't be too surprised if you thought DP is all you need.
- DP is widely applicable and, once you get the principle, you can come up with your own DP algorithms (e.g.: alignments with transitions, rewarding MM more)
- Today we see a problem that is:
  - relatively easy to state
  - widely applicable in practice
  - not tractable by DP: in fact, logsumexp is intractable!
- There are specialized algorithms for argmax for this problem. But you need to know them; sometimes they are hard to implement, etc.
- So, I will teach you a magic trick: formulate and solve (almost) any structured argmax problem using linear programming.

Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:



Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:

•  $M = \{(1,2), (3,4), (5,6)\}$ 



Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:

- $M = \{(1,2), (3,4), (5,6)\}$
- $M = \{(1,2), (4,6)\}$



Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:

- $M = \{(1,2), (3,4), (5,6)\}$
- $M = \{(1,2), (4,6)\}$
- $M = \{\}$



### Weighted Matchings in Graphs

Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:

- $M = \{(1,2), (3,4), (5,6)\}$
- $M = \{(1, 2), (4, 6)\}$

• 
$$M = \{\}$$

If the graph is **weighted** G = (V, E, w), where w(e) is the weight of edge  $e \in E$ , we define the weight of a matching:

 $w(M) = \sum_{e \in M} w(e)$ 



### Weighted Matchings in Graphs

Given an undirected graph G = (V, E), a **matching** (sometimes: coupling) is a subset of edges  $M \subseteq E$  such that no two edges in M touch.

Examples:

- $M = \{(1,2), (3,4), (5,6)\}$
- $M = \{(1, 2), (4, 6)\}$

• 
$$M = \{\}$$

If the graph is **weighted** G = (V, E, w), where w(e) is the weight of edge  $e \in E$ , we define the weight of a matching:

 $w(M) = \sum_{e \in M} w(e)$ 

A maximum weight matching is a matching M maximizing w(M).



#### **Applications in Biology: Protein Structure**



Image courtesy of dr. Laura Carroll, created with BioRender.com

#### **Applications in Biology: Protein Structure**

One of the many forms of *protein structure*: covalent bonds between parts of the chain.

Ex: disulfide bridges between Cysteine residues.

These grant stability, determine protein folding patterns, etc.

But which pairs will form bonds and which won't?



Disulfide bridges in human insulin. Source: https://www.chem.ucla.edu/~harding/IGOC/D/ disulfide\_bridge.html, public domain image.

#### **Applications in Biology: Protein Structure**

One of the many forms of *protein structure*: covalent bonds between parts of the chain.

Ex: disulfide bridges between Cysteine residues.

These grant stability, determine protein folding patterns, etc.

But which pairs will form bonds and which won't?





Research Article

Disulfide Bonds Play a Critical Role in the Structure and Function of the Receptor-binding Domain of the SARS-CoV-2 Spike Antigen

Andrey M. Grishin <sup>1</sup> 유 평, Nataliya V. Dolgova <sup>1, 2</sup>, Shelby Landreth <sup>4</sup>, Olivier Fisette <sup>5</sup>, Ingrid J. Pickering <sup>2, 3</sup>, Graham N. George <sup>2, 3</sup>, Darryl Falzarano <sup>4</sup> 유 평, Mirosław Cygler <sup>1</sup>

#### Show more 🗸

+ Add to Mendeley 🧠 Share 🍠 Cite

https://doi.org/10.1016/j.jmb.2021.167357

Get rights and content

Highlights

- The Receptor-Binding Domain of the Spike protein is dependent on its disulfide bonds.
- Reduction of the disulfide bonds in the RBD render it unstable.
- Reduction of the disulfide bonds of the RBD decreases its affinity to <u>ACE2</u>.
- The highly-reducing environment can stop SARS-CoV-2 replication in cell-based assays.

# Graph Matching & Linear Programming

#### Part 2: Finding a Maximum-Weight Matching

Machine Learning for Structured Data Vlad Niculae · LTL, UvA · https://vene.ro/mlsd

## Graph Matching & Linear Programming

**1** Matchings in Graphs

#### **2** Finding a Maximum-Weight Matching

**3** Linear Programming

**4** Bipartite graphs and the assignment problem

#### **Greedy Approaches Fail**



#### **Greedy Approaches Fail**



#### **Greedy Approaches Fail**



#### Finding a Maximum-Weight Matching

Spoiler: a good algorithm exists. But what if:

- you can't find it / don't know the right keyword to search for
- no good implementation is available?
- you want to modify the problem slightly?
- you just want to prototype something to make sure the rest of your model makes sense?

# Graph Matching & Linear Programming

Part 3: Linear Programming

Machine Learning for Structured Data Vlad Niculae · LTL, UvA · https://vene.ro/mlsd

## Graph Matching & Linear Programming

#### **1** Matchings in Graphs

**2** Finding a Maximum-Weight Matching

#### **3** Linear Programming

**4** Bipartite graphs and the assignment problem

## A crash course in linear optimization ("programming")

Here once again "programming" means "optimization".

Optimization problem (in general):

A problem of the form:

```
minimize F(\mathbf{y})
subject to G_i(\mathbf{y}) \le 0, \quad i = 1, \dots, p
\mathbf{y} \in \mathbb{R}^d
```

Notation/terminology:

- $y_j$  are called the variables,
- *F* is called the "objective",
- *G<sub>i</sub>* are constraints.

Here, y is just a variable.

Example: machine learning model training:

minimize  $L(\theta)$ 

where *L* is a total loss over a training set, and  $\theta$  are the model weights, is an optimization problem (but not a linear one usually).

#### A crash course in linear optimization

An optimization problem is **linear** if the objective *F* and all constraints *G<sub>i</sub>* are linear:

#### Linear program (LP):

A problem of the form:

```
minimize \boldsymbol{a} \cdot \boldsymbol{y}
subject to \boldsymbol{g}_i \cdot \boldsymbol{y} + b_j \leq 0, \quad i = 1, \dots, p
\boldsymbol{y} \in \mathbb{R}^d
```

General-purpose algorithms can solve this problem in polynomial time in expectation.

#### A crash course in linear optimization

An optimization problem is **linear** if the objective *F* and all constraints *G<sub>i</sub>* are linear:

#### Linear program (LP):

A problem of the form:

```
minimize \boldsymbol{a} \cdot \boldsymbol{y}
subject to \boldsymbol{g}_i \cdot \boldsymbol{y} + b_j \leq 0, \quad i = 1, \dots, p
\boldsymbol{y} \in \mathbb{R}^d
```

Integer linear program (ILP): A problem of the form: minimize  $\mathbf{a} \cdot \mathbf{y}$ subject to  $\mathbf{g}_i \cdot \mathbf{y} + b_j \le 0, i = 1, \dots, p$  $\mathbf{y} \in \mathbb{Z}^d$ 

General-purpose algorithms can solve this problem in polynomial time in expectation. This is discrete optimization!

Useful heuristics exist, but very hard in general. (NP-complete!)

### ILP for structure prediction

- ILPs with boolean variables *y<sub>i</sub>* ∈ {0, 1} can express many structured problems we care about.
- Think of this like a "domain-specific language":
- Just like pytorch can be used to express and train ML models, the language of LP can be used to express and solve argmax problems in structure prediction.
- Whenever possible, specialized algorithms are much better.
- But LP is great for prototyping, testing, exploring small changes, etc...

### ILP for max-weight matchings

Finding a max-weight matching in a weighted graph G = (V, E, w):

- A variable  $y_e \in \{0, 1\}$  for each  $e \in E$ .
- In a matching, for each node, at most one incident edge can be selected.
   For any *u* ∈ *V* define
   *E*(*u*) = {*e* ∈ *E* : *e* = (*u*, ·) or *e* = (·, *u*)}.
   Then we need:

$$\left(\sum_{e \in E(u)} y_e\right) \le 1, \text{ for every node } u \in V$$

• Maximize the sum of weights of the selected edges:

$$\sum_{e \in E} w(e) y_e$$

Putting it all together,

## Max-weight matching ILP maximize $\sum_{e \in E} w(e)y_e$ subject to $\left(\sum_{e \in E(u)} y_e\right) \le 1$ , for all $u \in V$ , $y_e \in \{0, 1\}$ for all $e \in E$ .

Plug this into a general-purpose solver and you have exact solutions for small problems. (In general, solving ILP is exponential in the problem size.)

#### Relaxing the integer constraints

ILP is hard, but LP (without I) is much easier.

(Continuous optimization generally easier than discrete optimization.)

This leads to the <u>fractional relaxation</u> that can be solved in polynomial time:

Max-weight matching relaxed LP

maximize 
$$\sum_{e \in E} w(e) y_e$$
  
subject to  $\left(\sum_{e \in E(u)} y_e\right) \le 1$ , for all  $u \in V$ ,  
 $0 \le y_e \le 1$  for all  $e \in E$ .

In fact, since every edge participates in at least one sum constraint, the constraints  $x_e \le 1$  are redundant and can be removed. But leaving them in is not a problem.

#### Relaxations generally introduce approximation

The ILP is exact, the relaxed LP is not:



The ILP constraints give:

 $\begin{cases} y_{12} + y_{13} \leq 1 \\ y_{12} + y_{23} \leq 1 \\ y_{13} + y_{23} \leq 1 \\ y_{ij} \in \{0, 1\} \end{cases}$ 

Write  $y = (y_{12}, y_{23}, y_{13})$ . Trying all  $2^3$  combinations shows that the only allowed configurations are:

y = (0, 0, 0)obj= 0y = (1, 0, 0)obj= 10y = (0, 1, 0)obj= 10y = (0, 0, 1)obj= 10

#### Relaxations generally introduce approximation

The ILP is exact, the relaxed LP is not:



The ILP constraints give:

 $\begin{cases} y_{12} + y_{13} \leq 1 \\ y_{12} + y_{23} \leq 1 \\ y_{13} + y_{23} \leq 1 \\ y_{ij} \in \{0, 1\} \end{cases}$ 

Relaxing the integer constraints still allows these configurations but also allows the fractional one

 $\boldsymbol{y} = \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$ 

But this "configuration" has higher objective:

Write  $y = (y_{12}, y_{23}, y_{13})$ . Trying all  $2^3$  combinations shows that the only allowed configurations are:

$$y = (0, 0, 0)$$
obj= 0 $y = (1, 0, 0)$ obj= 10 $y = (0, 1, 0)$ obj= 10 $y = (0, 0, 1)$ obj= 10

$$\frac{10}{2} + \frac{10}{2} + \frac{10}{2} = 15 > 10$$

In general, fractional relaxations can have spurious solutions.

# Graph Matching & Linear Programming

Part 4: Bipartite graphs and the assignment problem

Machine Learning for Structured Data Vlad Niculae · LTL, UvA · https://vene.ro/mlsd

## Graph Matching & Linear Programming

**1** Matchings in Graphs

**2** Finding a Maximum-Weight Matching

**3** Linear Programming



#### Bipartite weighted matching: the assignment problem

An undirected graph G = (V, E) is called bipartite if

- $V = V_A \cup V_B$  with  $V_A \cap V_B = \emptyset$
- every edge is from some node in  $V_A$  to some node in  $V_B$ , never within.

Write  $n = |V_A|$  and  $m = |V_B|$  and assume  $n \le m$  (otherwise, swap them.)

The **assignment problem**, e.g., optimally assigning tasks  $V_A$  to workers  $V_B$ : max-weight matchings in a bipartite graphs, with the additional constraint that every node in  $V_A$  gets assigned.

The weights and variables can be organized into  $n \times m$  matrices **A** and **Y**. The node-degree constraints are row and column sums:

$$\begin{cases} \sum_{j=1}^{m} y_{ij} = 1 & \text{for } i = 1, \dots, n \\ \sum_{i=1}^{n} y_{ij} \le 1 & \text{for } j = 1, \dots, m \end{cases}$$

In this bipartite case, it can be shown that the LP relaxation is exact!

#### Bipartite weighted matching LP





### Learning To Match

Matchings can be structured outputs in ML tasks: we train a model that predicts the edge weights. For example:

- Protein structure: which pairs of Cysteine residues are more compatible?
- Learning to match students with thesis projects (recommender system)
- Matching words in foreign language sentence translations (reordering)

Example architecture: parametrize edge weights  $a_{ij} = u_i \cdot v_j$  where  $u_i$  is an embedding of a student *i* and  $v_j$  embedding of the project *j* (each possibly passed through multiple hidden layers).

### Learning To Match

Matchings can be structured outputs in ML tasks: we train a model that predicts the edge weights. For example:

- Protein structure: which pairs of Cysteine residues are more compatible?
- Learning to match students with thesis projects (recommender system)
- Matching words in foreign language sentence translations (reordering)

Example architecture: parametrize edge weights  $a_{ij} = u_i \cdot v_j$  where  $u_i$  is an embedding of a student *i* and  $v_j$  embedding of the project *j* (each possibly passed through multiple hidden layers).

For predicting assignment, can solve the LP.

However, for learning, logsumexp over all matchings is intractable!

#### Non-Probabilistic Structure Learning

What can we do when we cannot compute the probabilistic loss:

$$-\log P(\mathbf{y}) = L_{\mathsf{NLL}}(\mathbf{y}) = -\operatorname{score}(\mathbf{y}) + \log \sum_{\mathbf{y}' \in \mathcal{Y}} \exp \operatorname{score}(\mathbf{y}')$$

#### Non-Probabilistic Structure Learning

What can we do when we cannot compute the probabilistic loss:

$$-\log P(\mathbf{y}) = L_{\mathsf{NLL}}(\mathbf{y}) = -\operatorname{score}(\mathbf{y}) + \log \sum_{\mathbf{y}' \in \mathcal{Y}} \exp \operatorname{score}(\mathbf{y}')$$

Recall: An alternative training objective is the perceptron loss:

$$L_{\mathsf{Perc}}(\boldsymbol{y}) = -\operatorname{score}(\boldsymbol{y}) + \max_{\boldsymbol{y}' \in \mathcal{Y}} \operatorname{score}(\boldsymbol{y}')$$

The Perceptron loss can sometimes be computed using (I)LP!

#### **Structured Perceptron Learning**

$$L_{\mathsf{Perc}}(\boldsymbol{y}) = -\operatorname{score}(\boldsymbol{y}) + \max_{\boldsymbol{y}' \in \mathcal{Y}} \operatorname{score}(\boldsymbol{y}')$$

When we can

- represent the structures y as vectors and  $\mathcal{Y}$  as linear constraints,
- write score( $\boldsymbol{y}$ ) =  $\boldsymbol{a} \cdot \boldsymbol{y}$ ,

(e.g., our graph matching applications). Then:

- solve the ILP  $y^* = \arg \max_{y \in \mathcal{Y}} \operatorname{score}(y)$  or the LP  $y^* = \arg \max_{y \in \tilde{\mathcal{Y}}} \operatorname{score}(y)$
- $L_{\text{Perc}}(\boldsymbol{y}) = -\boldsymbol{a} \cdot \boldsymbol{y} + \boldsymbol{a} \cdot \boldsymbol{y}^*$

In the relaxed perceptron, fractional solutions are allowed: there is theory that shows this is ok and can even work better.

In both cases, as long as you have any code (non-pytorch) to find  $y^*$ , plug it in and pytorch gives the correct gradient  $\nabla_a L_{Perc}(y) = y^* - y$ .

#### Practical Example: Knapsack problems

Say we have a knapsack that can hold *B* liters, and we are packing for a trip.

We have *n* objects, each with a volume  $v_i$  liters, and with a score (value)  $a_i$ . We want to pack the most valuable items.

#### Weighted Knapsack (I)LP

maximize 
$$\sum_{i=1}^{n} a_i y_i$$
  
subject to 
$$\sum_{i=1}^{n} v_i y_i \le B$$
  
(if ILP:)  $y_i \in \{0, 1\}$  for all  $i$   
(if LP:)  $y_i \in [0, 1]$  for all  $i$ 

import numpy as np import cvxpy as cp

```
v = np.array([ .5, .1, .9, .2, 3.0])
a = np.array([-.2, -.5, .3, .2, 1.0])
B = 4.0
```

```
n = v.shape[0]
y = cp.Variable(n, integer=False) # integer=True for ILP
```

problem = cp.Problem(cp.Maximize(objective), constraints)
problem.solve()

```
print(y.value.round(2))
# if LP: [0. 0. 0.95 1. 0.98]
# if ILP: [0. 0. 1.0 0. 1.0]
```

#### **Ease of Prototyping Extensions**

Say you are now asked to solve the Knapsack problem with an additional requirement: the first two items are mutually exclusive.

Can right away add a constraint  $y_1 + y_2 \le 1$ .

A dedicated algorithm can be much faster, exact, but not as easy to modify.

And if you come up with a better specialized algorithm, you should still  $\underline{\text{test it}}$  against the ILP.