


Faster Fundamental Graph Algorithms via Learned Predictions


Justin Chen, Sandeep Silwal, Ali Vakilian, Fred Zhang

Learning-Augmented Algorithms

New view of algorithm design: rigorous algorithmic analysis with ML help




+




For algorithmic tasks repeated day after day, **predictions** help overcome “worst-case” behaviour

Learning-Augmented Algorithms

New view of algorithm design: rigorous algorithmic analysis with ML help



+



For algorithmic tasks repeated day after day, **predictions** help overcome “worst-case” behaviour

Most results have focused on online, streaming, or sampling settings where predictions yield knowledge about the unseen part of the input

Faster Matchings via Learned Duals (Dinitz, Im, Lavastida, Moseley, Vassilvitskii '21)

- Learning-augmented framework applied to runtime analysis for batch graph problems
- Minimum cost matching via the Hungarian algorithm

Faster Matchings via Learned Duals (Dinitz, Im, Lavastida, Moseley, Vassilvitskii '21)

- Learning-augmented framework applied to runtime analysis for batch graph problems
- Minimum cost matching via the Hungarian algorithm
- Used learned predictions of the duals of the matching problem to “warm-start” the Hungarian algorithm

Faster Matchings via Learned Duals (Dinitz, Im, Lavastida, Moseley, Vassilvitskii '21)

- Learning-augmented framework applied to runtime analysis for batch graph problems
- Minimum cost matching via the Hungarian algorithm
- Used learned predictions of the duals of the matching problem to “warm-start” the Hungarian algorithm
- Theoretically analyze the complexity of learning a predicted dual, rounding it to feasibility, and the resulting runtime parameterized by the prediction error

Our Results

1. Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)

Our Results

1. Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)
2. Extend the learned duals approach to single-source shortest path with negative edge weights

Our Results

1. Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)
2. Extend the learned duals approach to single-source shortest path with negative edge weights
3. Establish a connection between reductions and learning-augmented algorithms (with case-study on foreign exchange data)

Our Results


1. Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)
2. Extend the learned duals approach to single-source shortest path with negative edge weights
3. Establish a connection between reductions and learning-augmented algorithms (with case-study on foreign exchange data)
4. Prove general learnability theorems, showing that the predicted duals used in our algorithms can be efficiently PAC-learned

Our Results

1. *Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)*
2. Extend the learned duals approach to single-source shortest path with negative edge weights
3. Establish a connection between reductions and learning-augmented algorithms (with case-study on foreign exchange data)
4. Prove general learnability theorems, showing that the predicted duals used in our algorithms can be efficiently PAC-learned

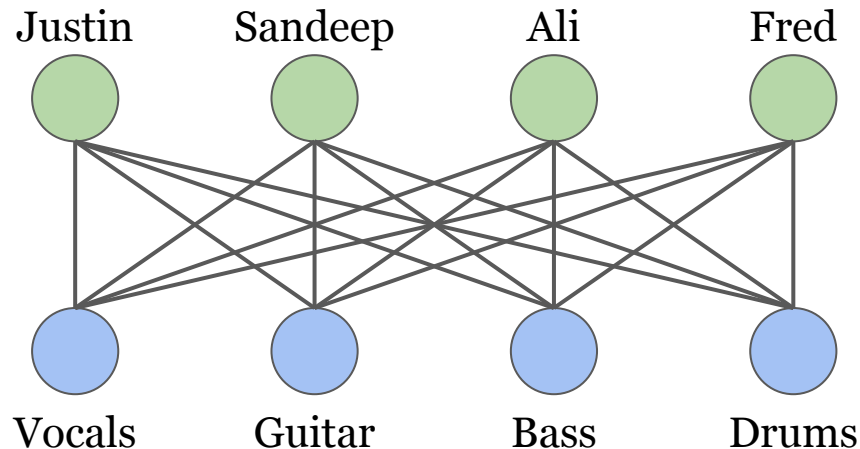
Minimum Cost Matching (Assignment Problem)

- Bipartite graph $G = (V, E)$ with non-negative integral edge-costs
- Find a perfect matching that minimizes the summed cost



Minimum Cost Matching (Assignment Problem)

- Bipartite graph $G = (V, E)$ with non-negative integral edge costs
- Find a perfect matching that minimizes the summed cost
- Hungarian algorithm (primal-dual): $O(mn)$ runtime



Warm Starting with Learned Duals

- *Dinitz et al*: Warm-start the Hungarian algorithm with y' , bounding the runtime by $O(m\sqrt{n}\|y - y^*\|_1)$ vs. $O(mn)$ worst-case time

Warm Starting with Learned Duals

- *Dinitz et al*: Warm-start the Hungarian algorithm with y' , bounding the runtime by $O(m\sqrt{n}\|y - y^*\|_1)$ vs. $O(mn)$ worst-case time
- *Our algorithm*: $O(m\sqrt{n} + (m + n \log n)\|y - y^*\|_0)$
 - Note: $\|y' - y^*\|_0 < \|y' - y^*\|_1$ as duals are integral

Warm Starting with Learned Duals

- *Dinitz et al*: Warm-start the Hungarian algorithm with y' , bounding the runtime by $O(m\sqrt{n}\|y - y^*\|_1)$ vs. $O(mn)$ worst-case time
- *Our algorithm*: $O(m\sqrt{n} + (m + n \log n)\|y - y^*\|_0)$
 - Note: $\|y' - y^*\|_0 < \|y' - y^*\|_1$ as duals are integral
- Improvement over worst-case bound when dual errors are $o(n)$ vs. $o(\sqrt{n})$

Warm Starting with Learned Duals

- *Dinitz et al:* Warm-start the Hungarian algorithm with y' , bounding the runtime by $O(m\sqrt{n}\|y - y^*\|_1)$ vs. $O(mn)$ worst-case time
- *Our algorithm:* $O(m\sqrt{n} + (m + n \log n)\|y - y^*\|_0)$
 - Note: $\|y' - y^*\|_0 < \|y' - y^*\|_1$ as duals are integral
- Improvement over worst-case bound when dual errors are $o(n)$ vs. $o(\sqrt{n})$
- In the regime where dual errors are $o(\sqrt{n})$, faster runtime by a factor of at least \sqrt{n}

Warm Starting with Learned Duals

- *Dinitz et al:* Warm-start the Hungarian algorithm with y' , bounding the runtime by $O(m\sqrt{n}\|y - y^*\|_1)$ vs. $O(mn)$ worst-case time
- *Our algorithm:* $O(m\sqrt{n} + (m + n \log n)\|y - y^*\|_0)$
 - Note: $\|y' - y^*\|_0 < \|y' - y^*\|_1$ as duals are integral
- Improvement over worst-case bound when dual errors are $o(n)$ vs. $o(\sqrt{n})$
- In the regime where dual errors are $o(\sqrt{n})$, faster runtime by a factor of at least \sqrt{n}
- Minor change to the algorithm combined with improved analysis

Our Results

- *Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)*
- Extend the learned duals approach to single-source shortest path with negative edge weights
- Establish a connection between reductions and learning-augmented algorithms (with case-study on foreign exchange data)
- Prove general learnability theorems, showing that the predicted duals used in our algorithms can be efficiently PAC-learned

Our Results

- *Develop an improved algorithm and analysis for minimum cost matching with learned duals (by at least a \sqrt{n} factor if the predictions are useful)*
- Extend the learned duals approach to single-source shortest path with negative edge weights
- Establish \tilde{O} bounds for \tilde{O} and learning-augmented algorithms (e.g., \tilde{O} on foreign exchange data)
- Prove general learnability theorems, showing that the predicted duals used in our algorithms can be efficiently PAC-learned

Check out our paper and poster!