Data-Driven Approximations to NP-Hard Problems

Anton Milan S. Hamid Rezatofighi Ravi Garg Anthony Dick Ian Reid

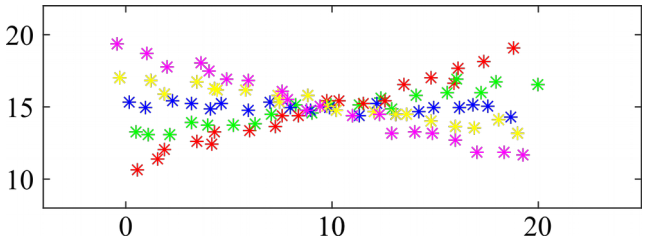

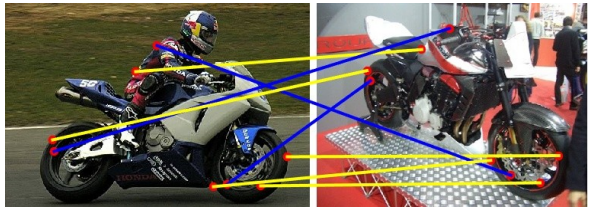

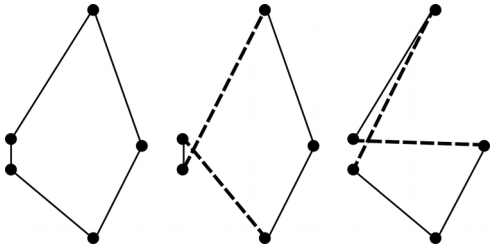

THE UNIVERSITY of ADELAIDE

Australian Centre for Visual Technologies
Innovation and education in visual information processing.

AAAI-17 San Francisco

Motivation

- Learn complex algorithms from data
- Efficient inference
- End-to-end learning

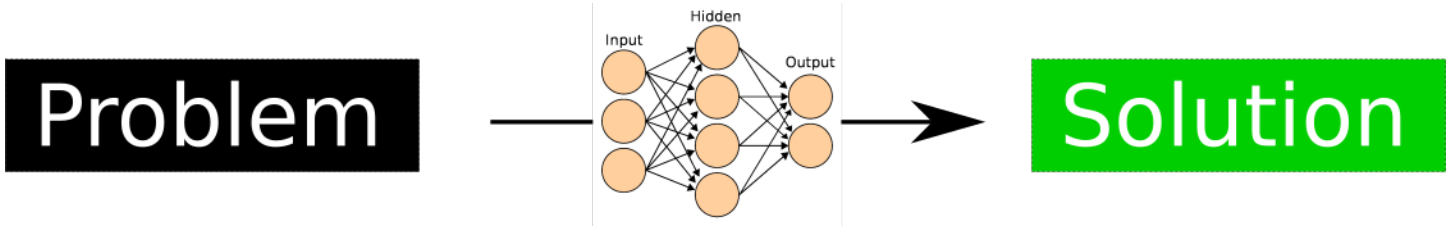

Contributions

1) Sequential (LSTM) bipartite matching
2) Training with “approximate” ground truth
3) Loss- vs. objective-based training

Applications

- **Data Association:**
  Marginalization

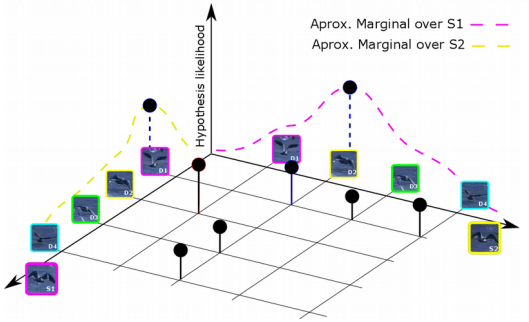

- **Keypoint Matching:**
  Quadratic Programming

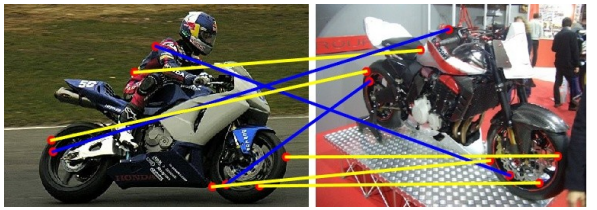

- **Travelling Salesman Problem:**
  Combinatorics

Applications

- **Data Association:**
  Marginalization

- **Keypoint Matching:**
  Quadratic Programming

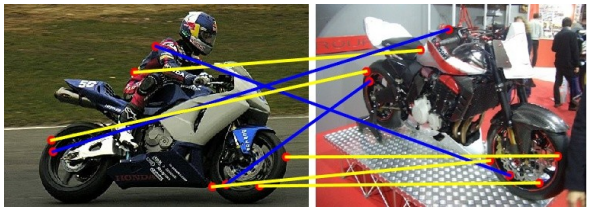

- **Travelling Salesman Problem:**
  Combinatorics

Bipartite Matching

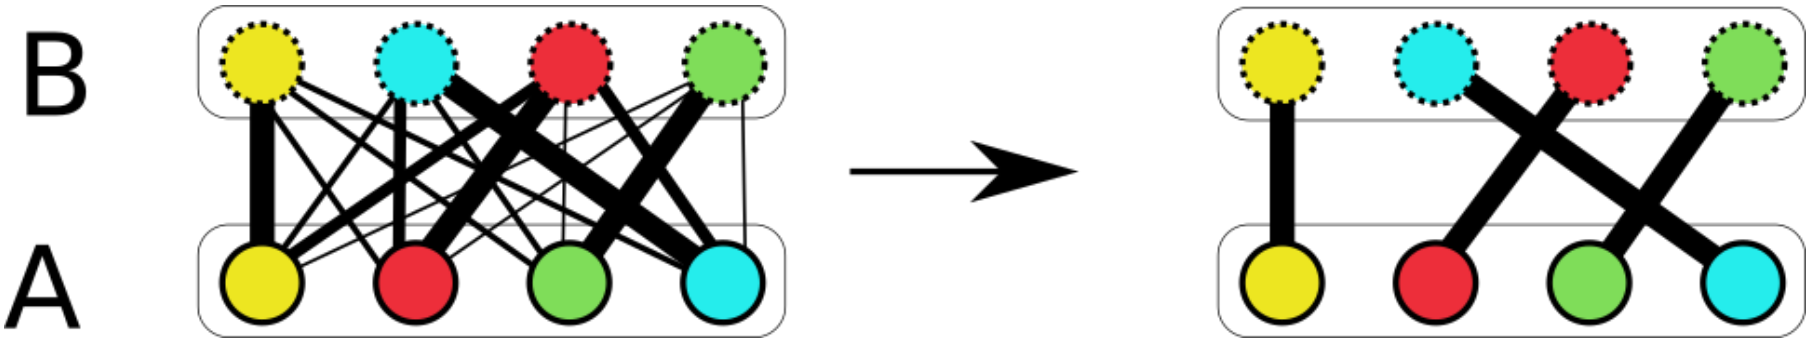

Linear Assignment: $X^* = \underset{X}{\operatorname{argmin}} \quad C^\top X$ s.t. *X* binary and one-to-one

→ Hungarian (Munkres) Algorithm

Quadratic Assignment Problem

$$X^* = \underset{X \in \mathcal{X}}{\operatorname{argmax}} \quad J(X) = X^\top K X$$

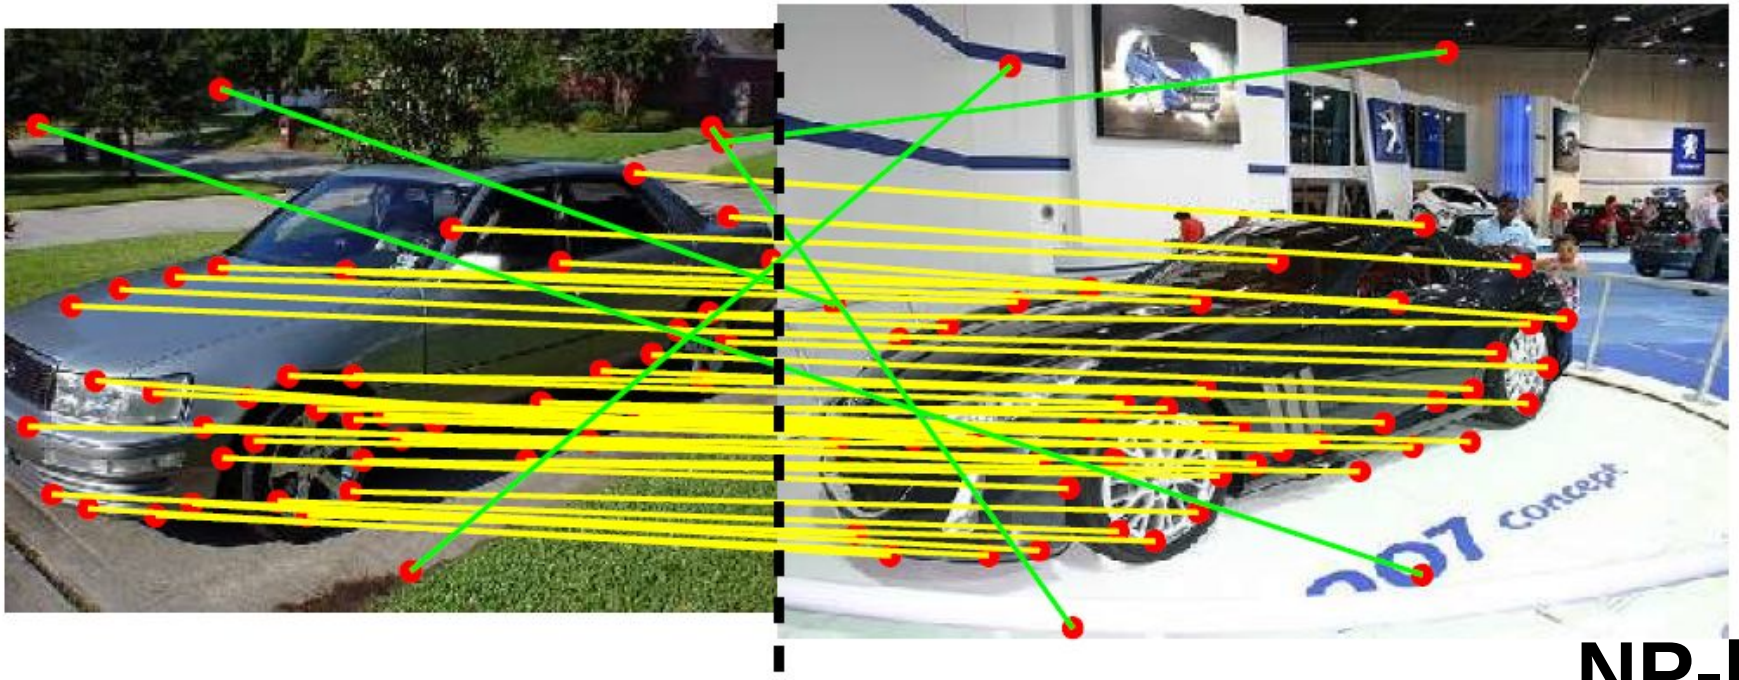

NP-hard

Our Model

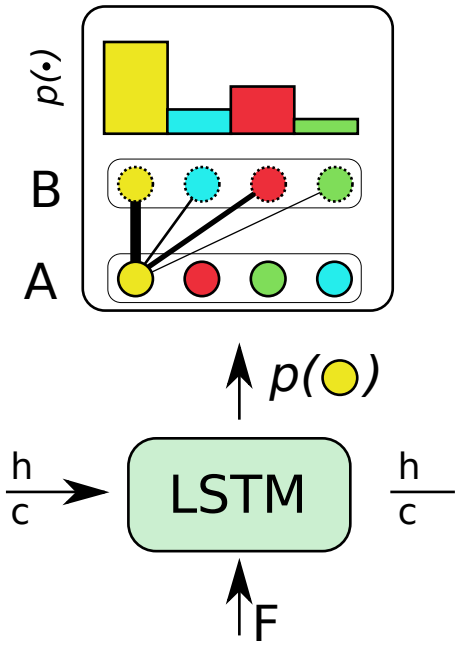

Our Model

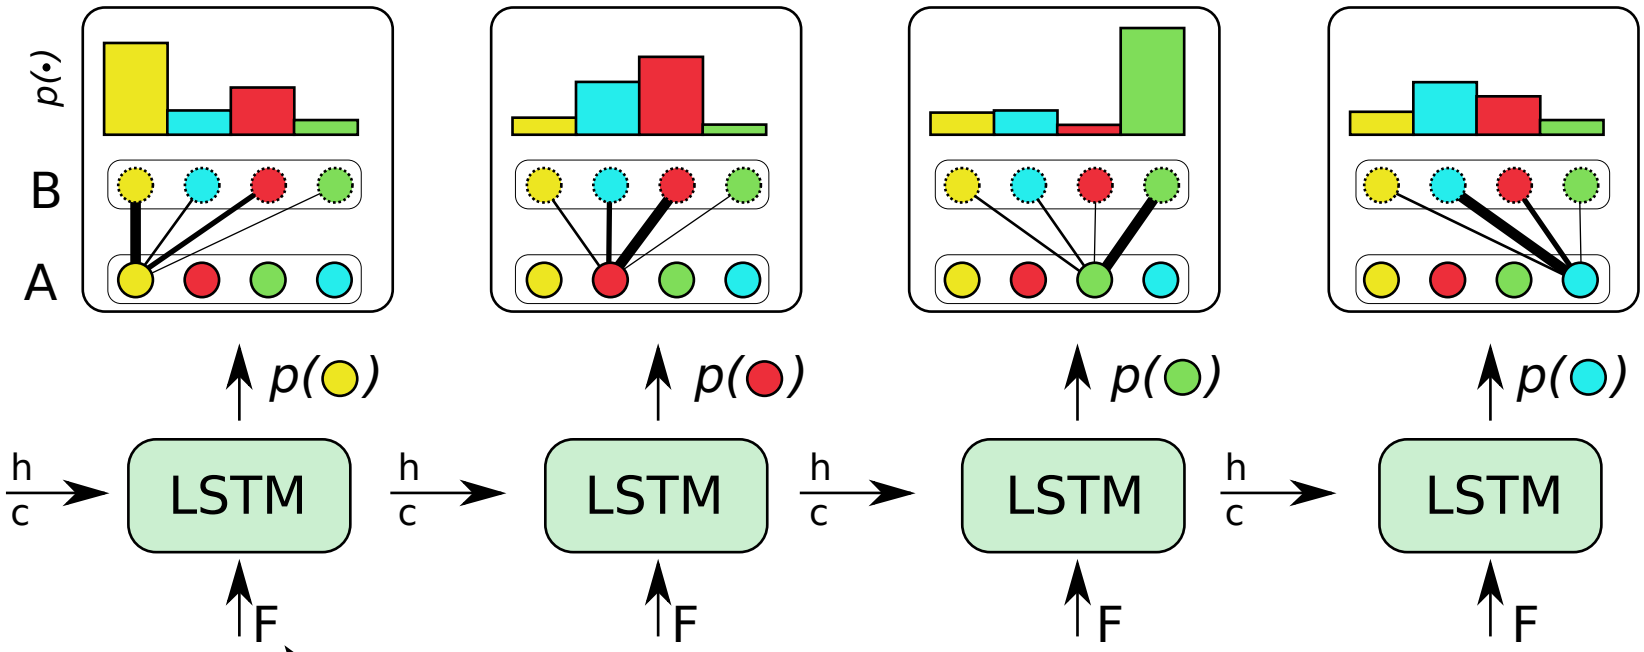

“Feature” matrix

$$X^\top K X$$

Bootstrapping the Training Set

- What if ground truth is ‘hard to obtain’?
- Start with what we have
- Improve over time

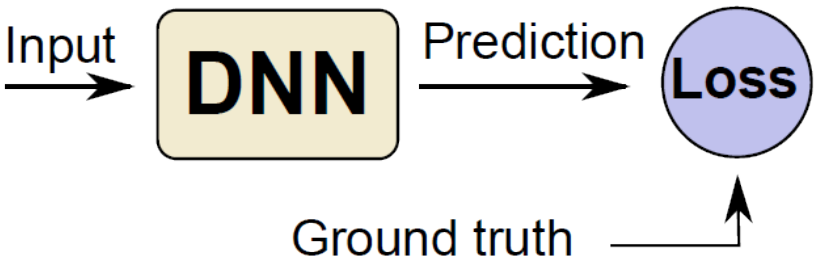

Bootstrapping the Training Set

- What if ground truth is ‘hard to obtain’?
- Start with what we have
- Improve over time

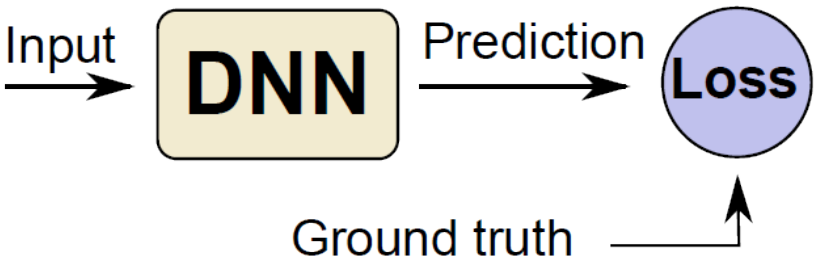

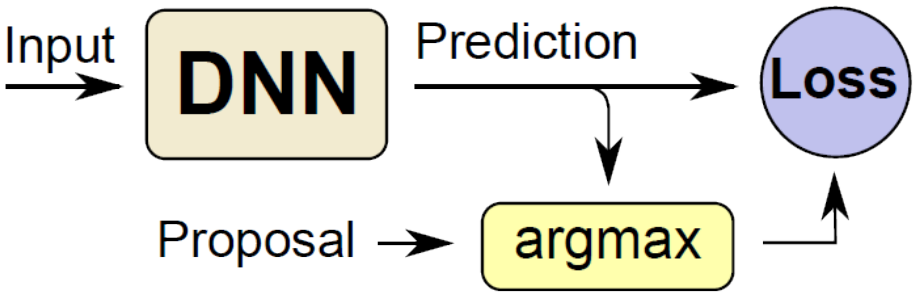

$$X^\top K X$$

Keypoint Matching

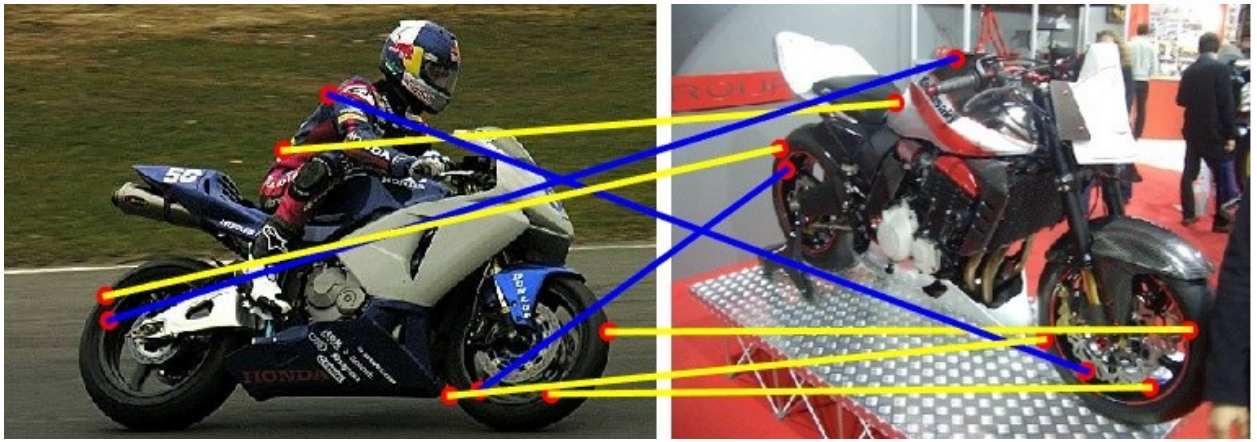

Name	Accuracy	Objective	Time [ms]
Branch-and-cut	0.90	10.99	7
IPFP-S [1] (best of 10) *	0.70	10.47	56
LSTM	**0.76**	**10.52**	**4**

* Used as ‘ground truth’

[1] Leordeanu et al., 2011

Travelling Salesman Problem

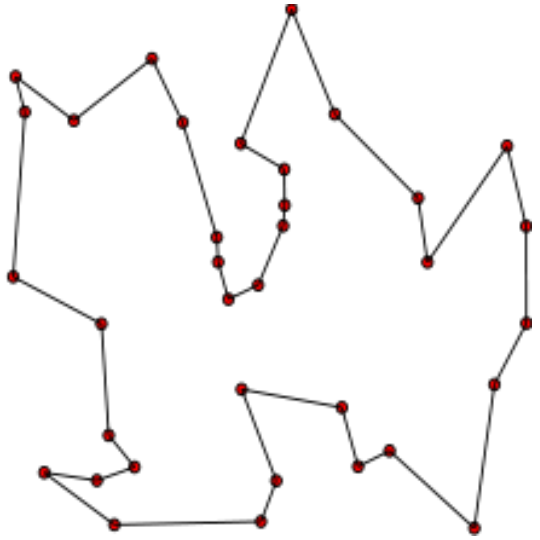

Given a set of nodes, visit each one exactly once and return to start.

Pointer-Networks

[Vinyals et al., NIPS*2015]

- Loss: cross-entropy
- Better: objective-based training

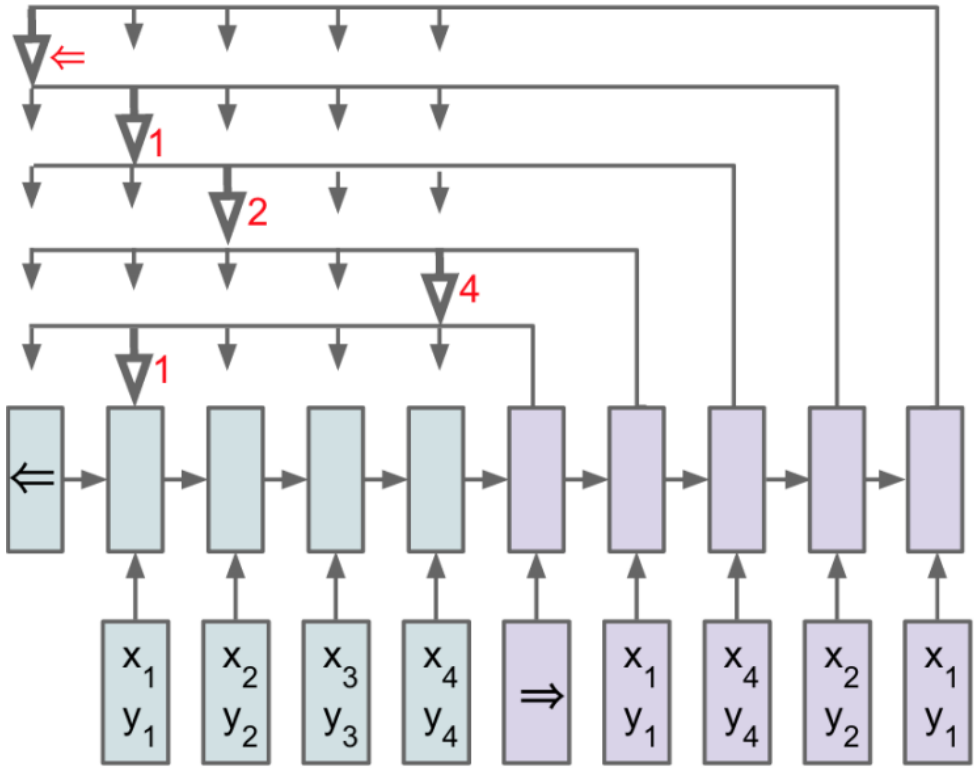

Objective-based training

Ground Truth

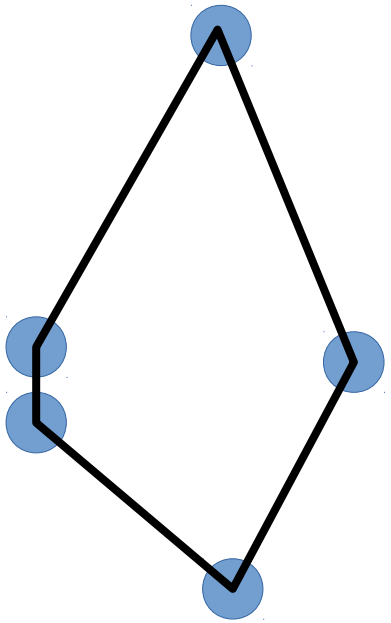

Errors: **0**
Length: **310**

Objective-based training

Ground Truth

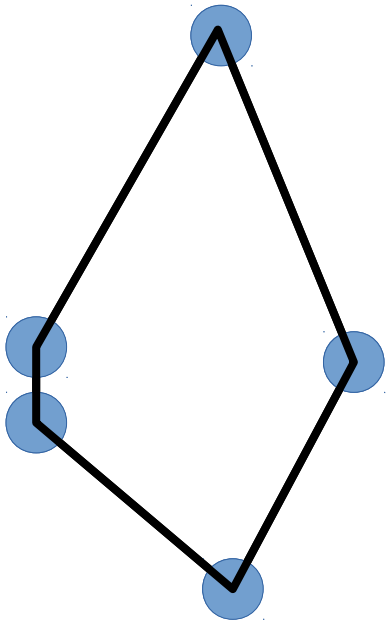

Errors: **0**
Length: **310**

Objective-based training

Ground Truth

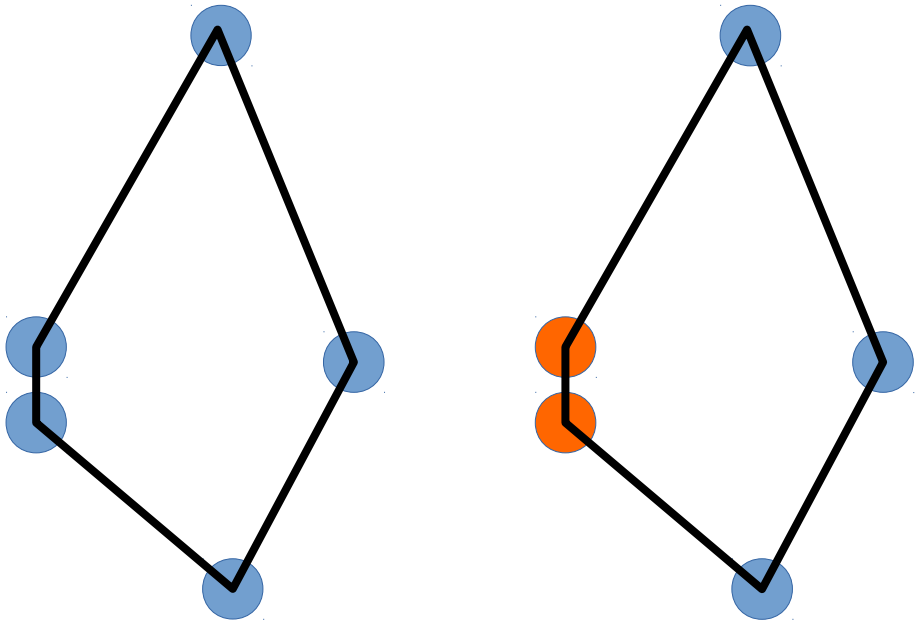

Errors: **0**
Length: **310**

Objective-based training

Ground Truth

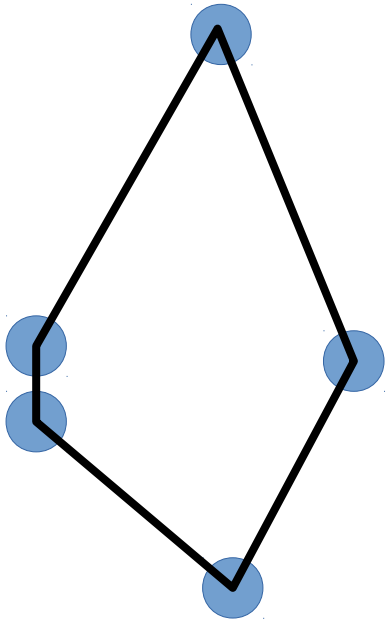

Errors: **0**
Length: **310**

Solution 1

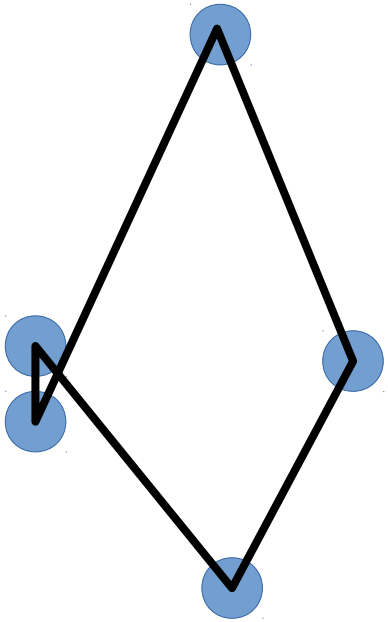

Errors: **2**
Length: **340**

Objective-based training

Ground Truth

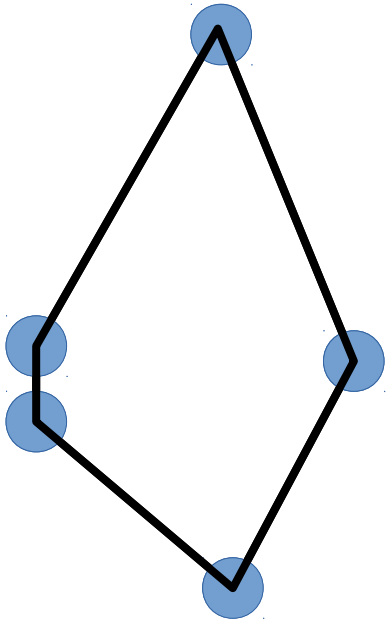

Errors: **0**
Length: **310**

Solution 1

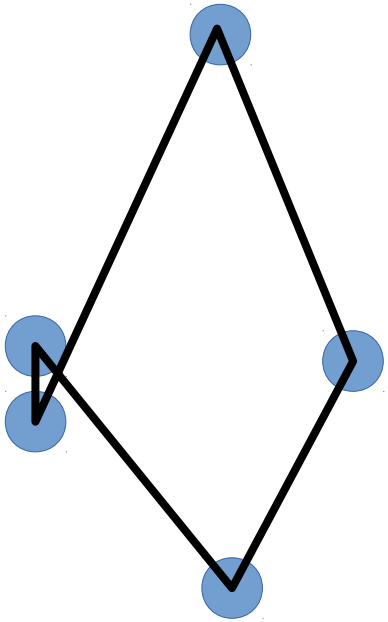

Errors: **2**
Length: **340**

Objective-based training

Ground Truth

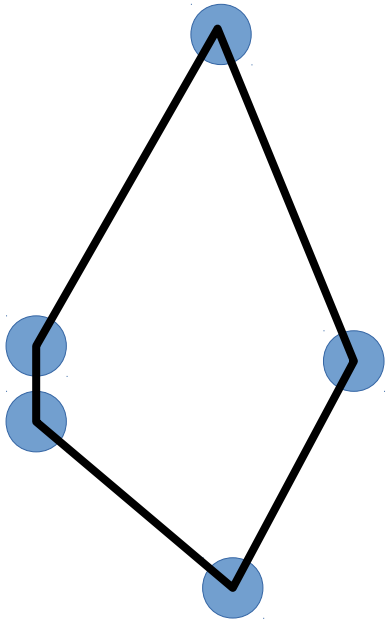

Errors: **0**
Length: **310**

Solution 1

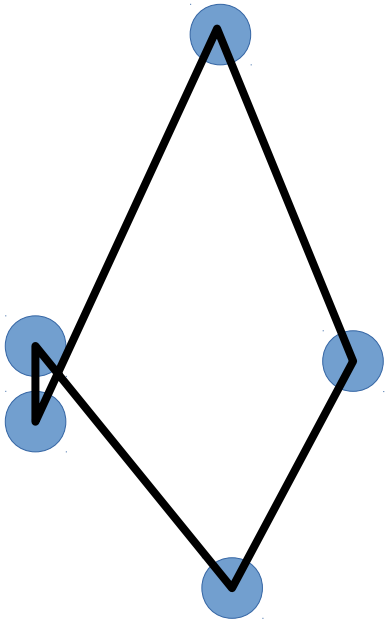

Errors: **2**
Length: **340**

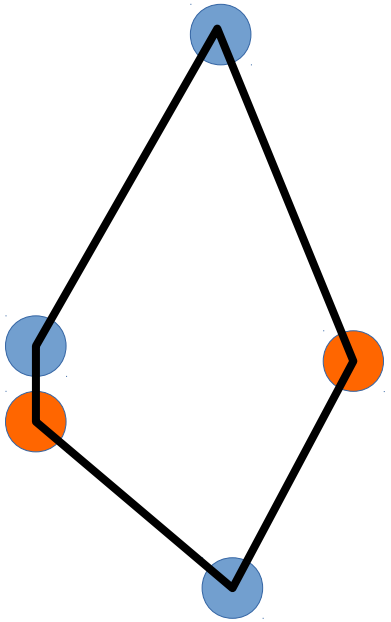

Objective-based training

Ground Truth

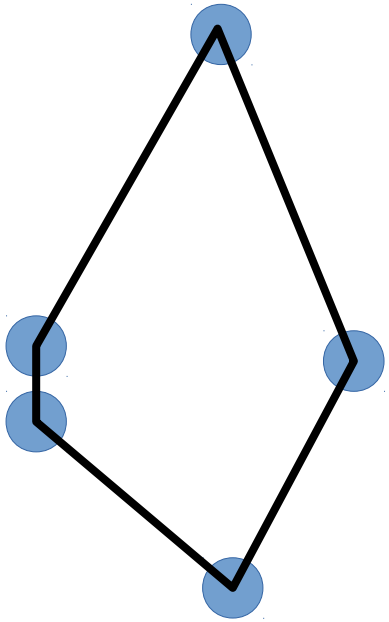

Errors: 0
Length: 310

Solution 1

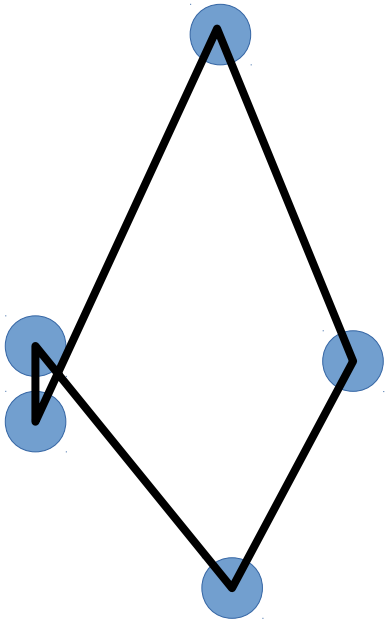

Errors: 2
Length: 340

Solution 2

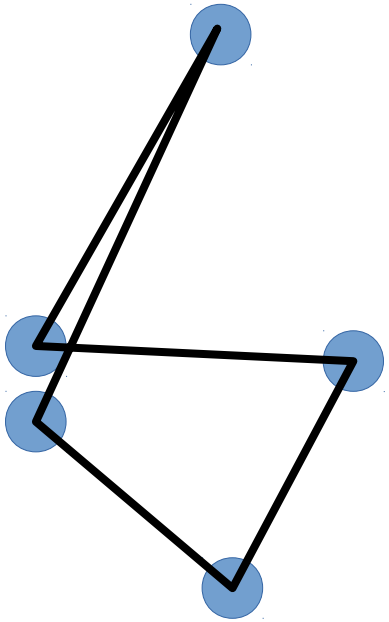

Errors: 2
Length: 400

Non-differentiable Loss

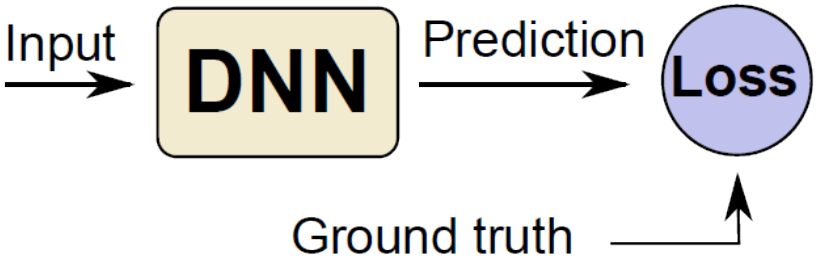

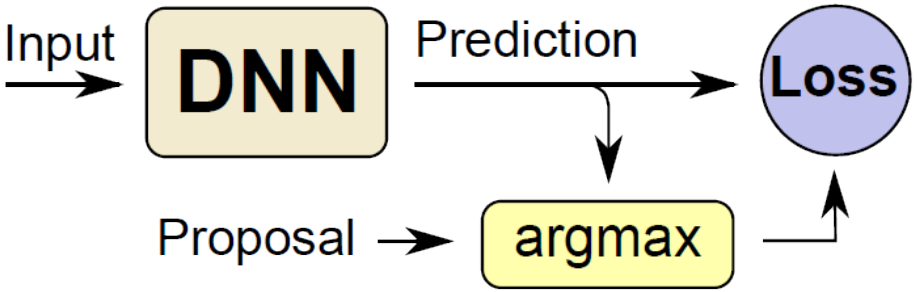

Tour length

Results

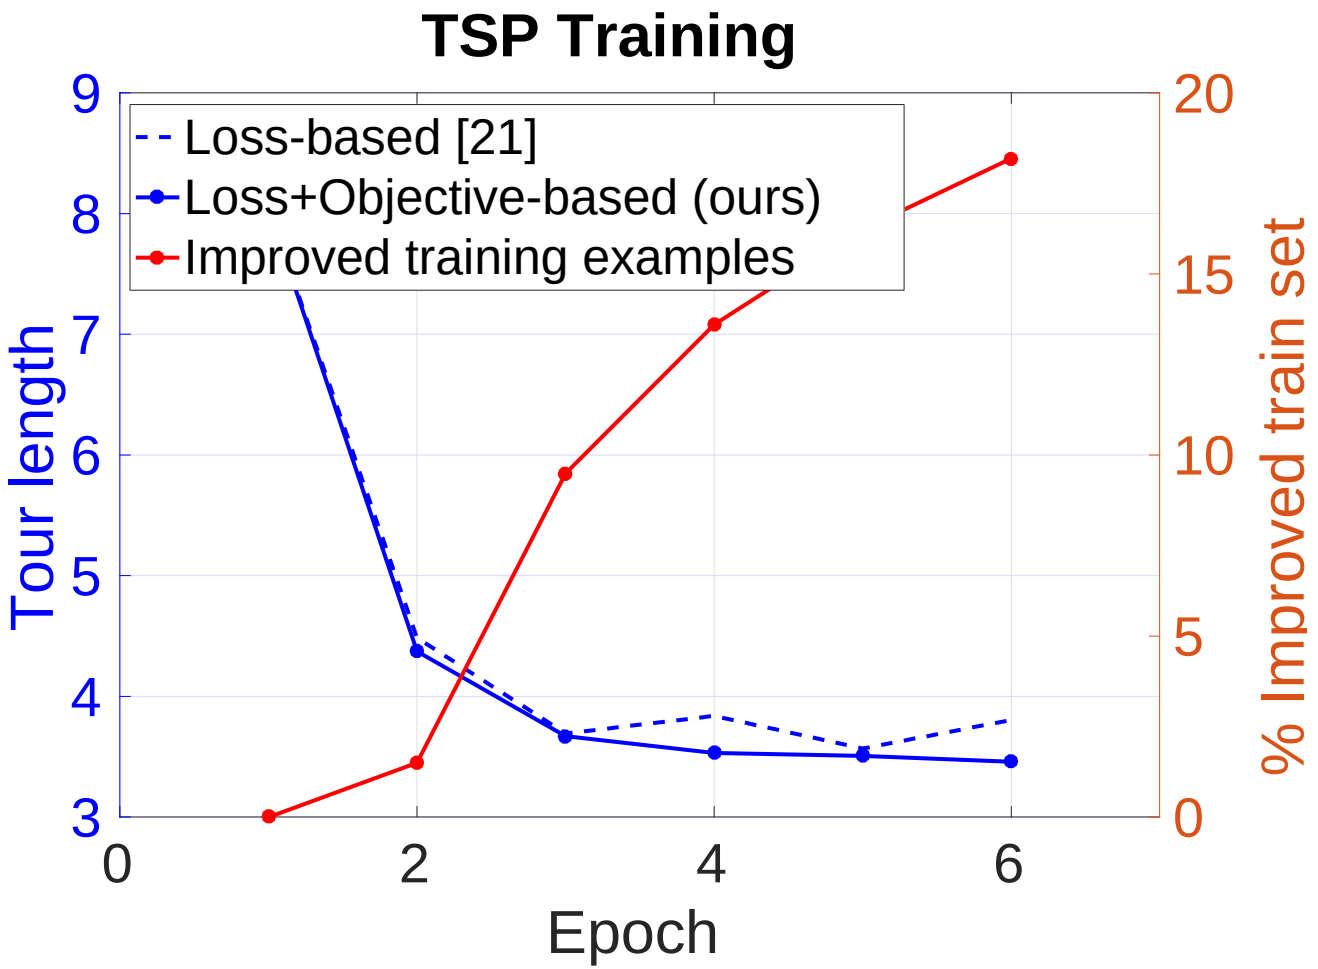
TSP Training
Loss-based [21]
Loss+Objective-based (ours)
Improved training examples
Tour length
% Improved train set
Epoch
9
8
7
6
5
4
3
20
15
10
5
0
0
2
4
6

Limitations and Discussion

- Time: Training vs. algorithm design
- Input/Output size is fixed
- Problem’s objective is not always clear (but if it is, use it!)

Conclusions

- LSTM model for matching

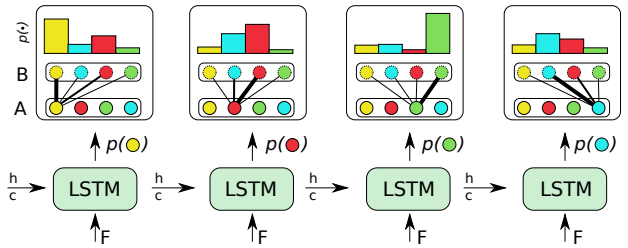

- Improving ‘approximate’ training set

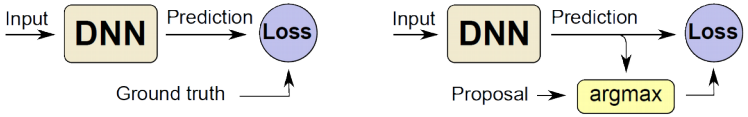

- Objective vs. loss-based training

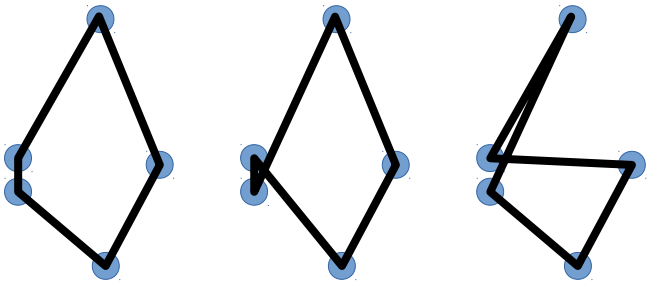

Thank you