


Selecting training examples

© 2019 Philipp Krähenbühl and Chao-Yuan Wu


Sampling

- How do we select positives and negatives?




All pairs / triples?

- Bad idea
- very slow
- Pairs $O(N^2)$
- Triples $O(N^3)$




Random pairs / triples?

- Random positives
 - Fast
 - Good gradient
- Random negatives
 - Far apart
 - Small loss
 - Small gradient




Hard negatives

- Pick one negative
- Closed to each positive




Hard negatives

- Too noisy
- No meaningful gradient direction
- Too hard
- Stronger gradient than positives



Semi-hard negatives

- Fine one negatives
- at same distance as a positive



Semi-hard negatives

- Works well enough in practice
- A bit hacky
- Alternative: Weighted random sampling