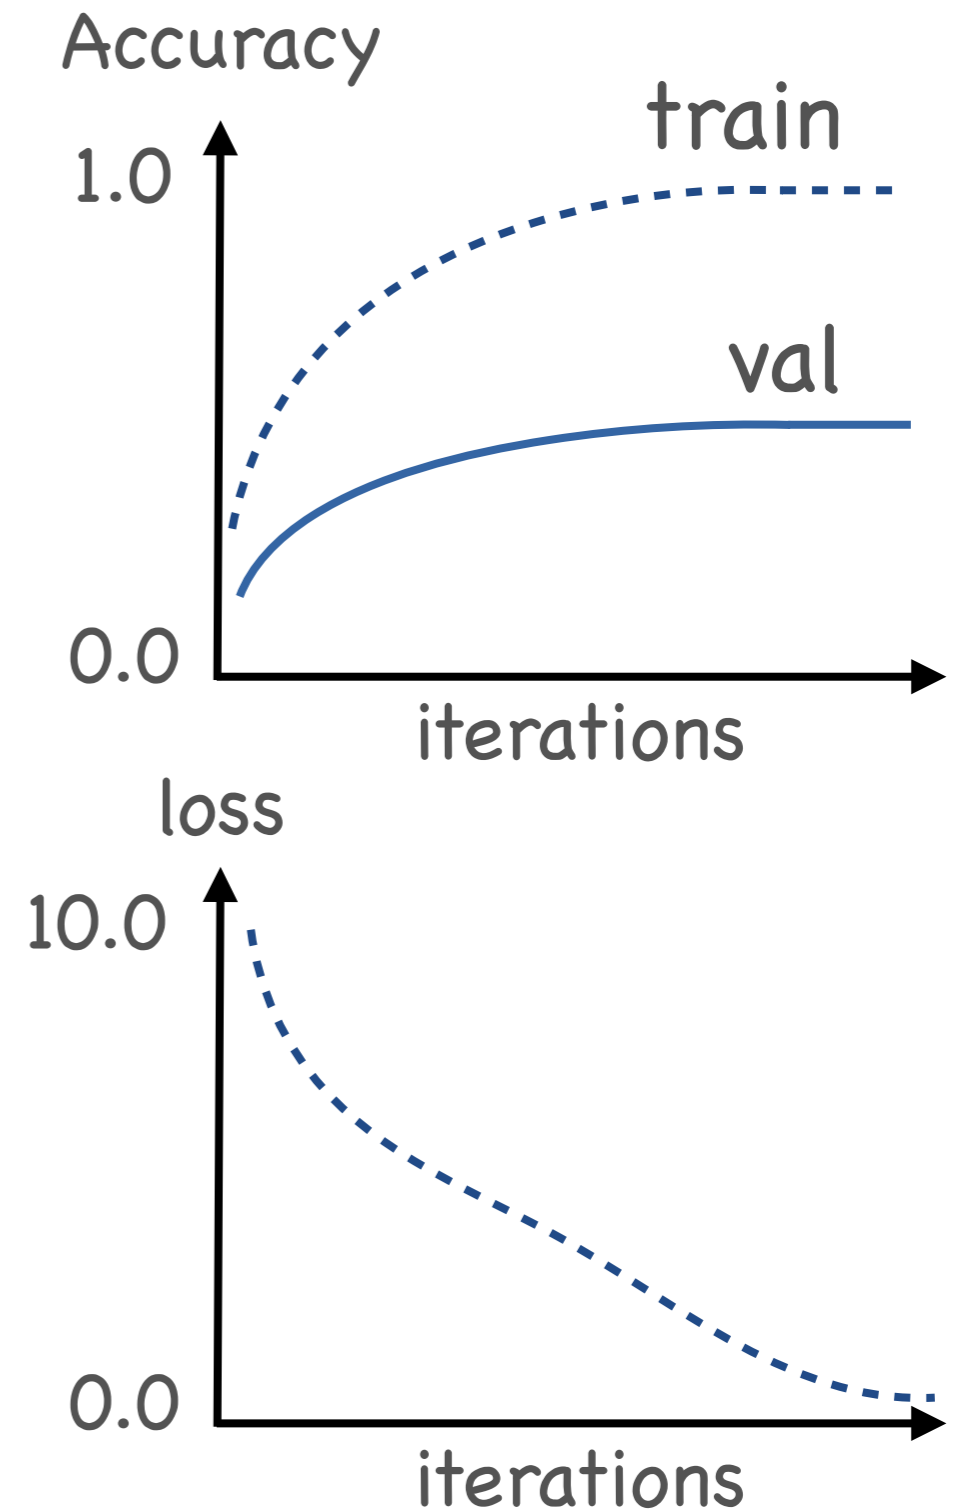


Overfitting and how to detect it

© 2019 Philipp Krähenbühl and Chao-Yuan Wu


Overfitting

- Fit model to training set
- Model does not work on unseen data (e.g. validation set)




How to detect overfitting?

- Plot training and validation accuracy




Is overfitting always bad?



Why do we overfit?

- Sampling bias
- Optimization fits patterns that only exist in training set



Do we overfit with infinite
training data?

How do we prevent overfitting?

- Collect more data
- Make the model simpler
 - Regularize
- Transfer learning