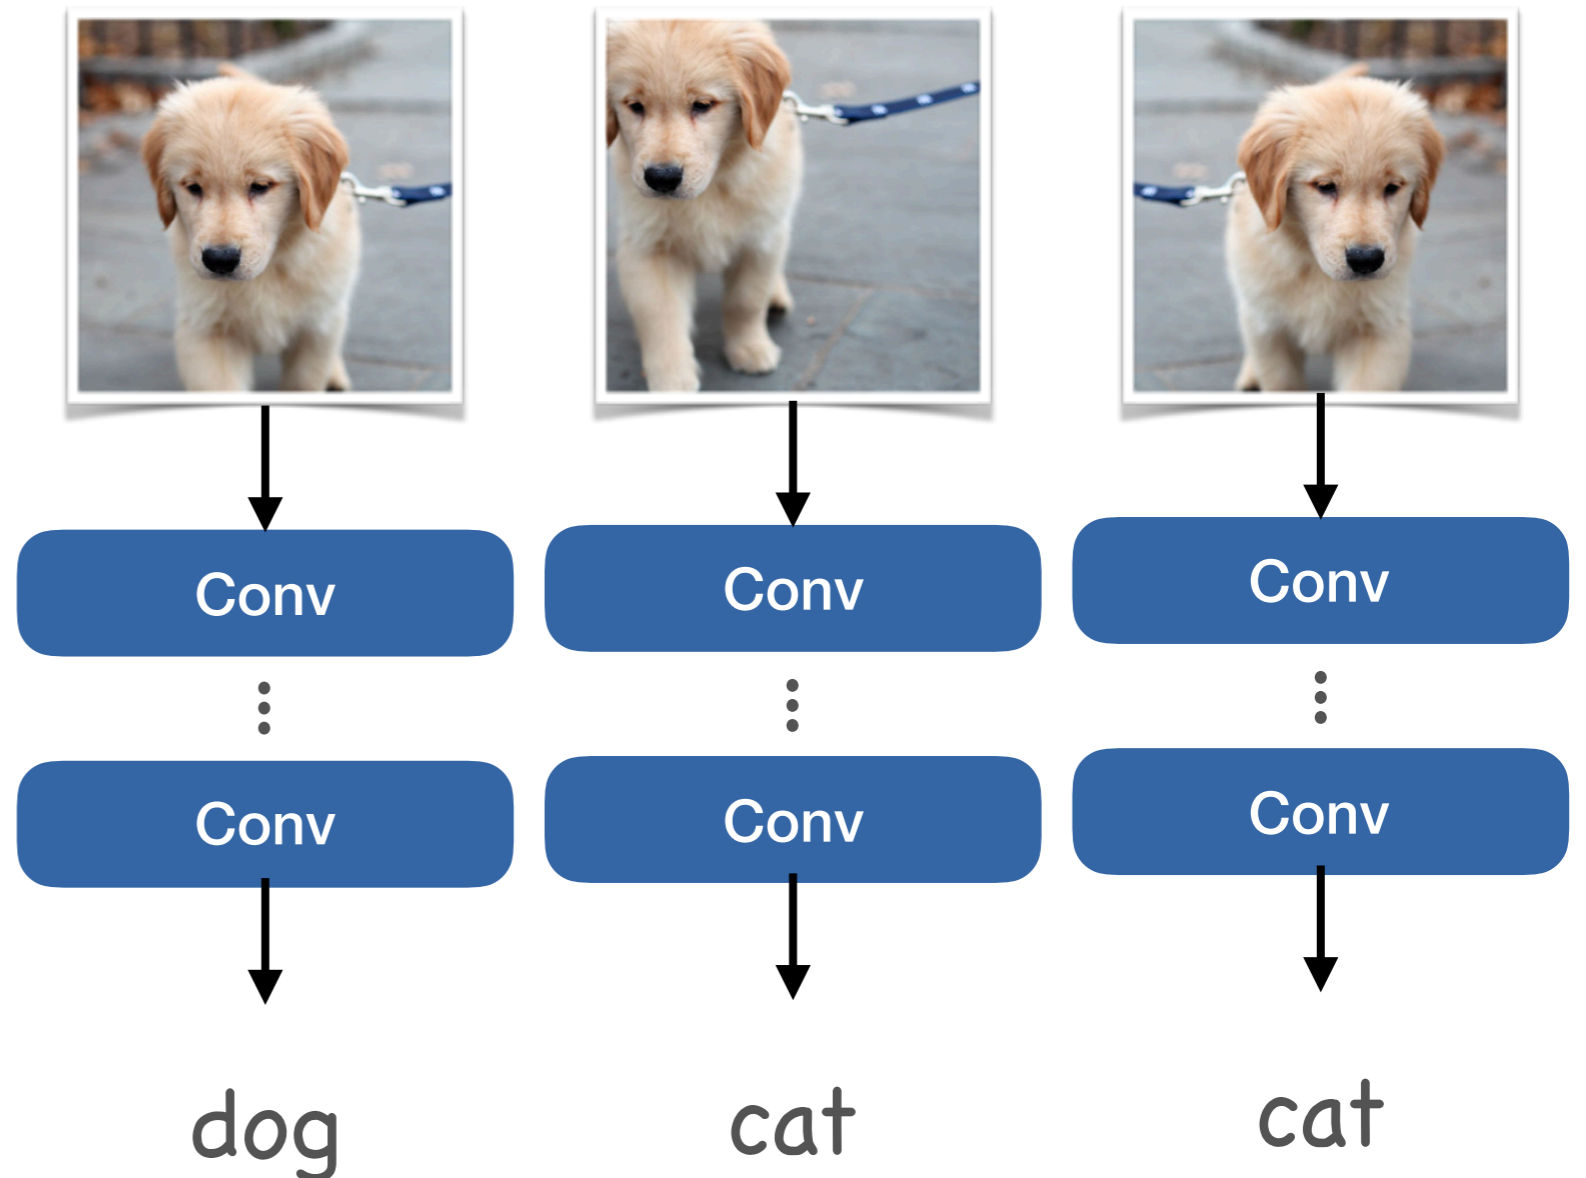


Data augmentation

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

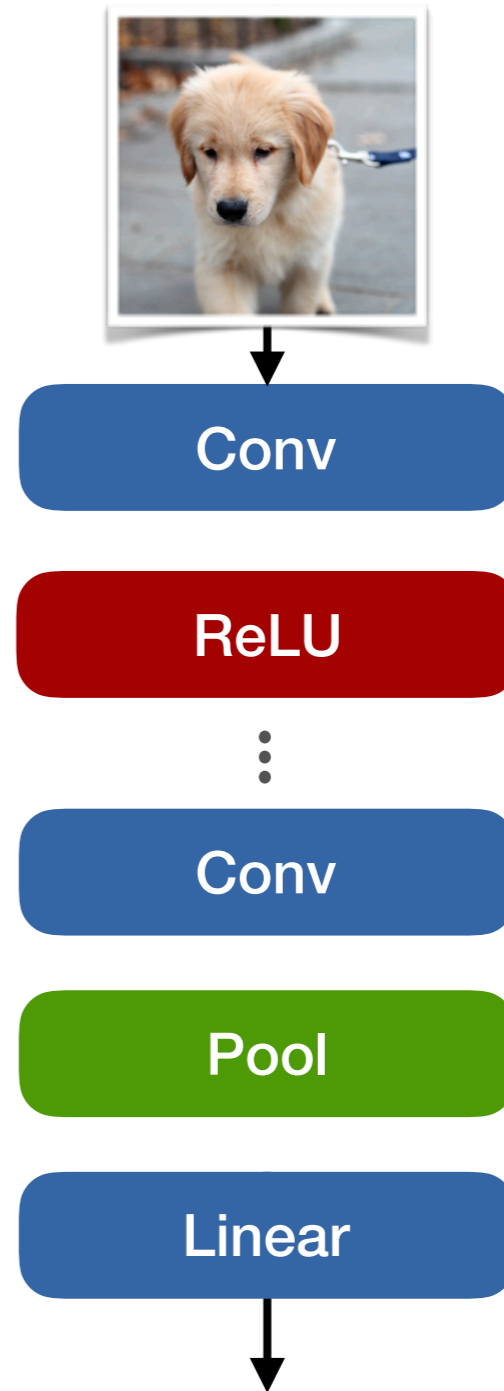
Signs of overfitting

- Does not capture invariances in data



How to capture invariances?

- Build them into the model
 - Convolutions
 - All-convolutional models
- Build them into the data
 - Data augmentation



Data augmentation

- Capture invariances in data
- (Randomly) transform data during training
- Reuse a label

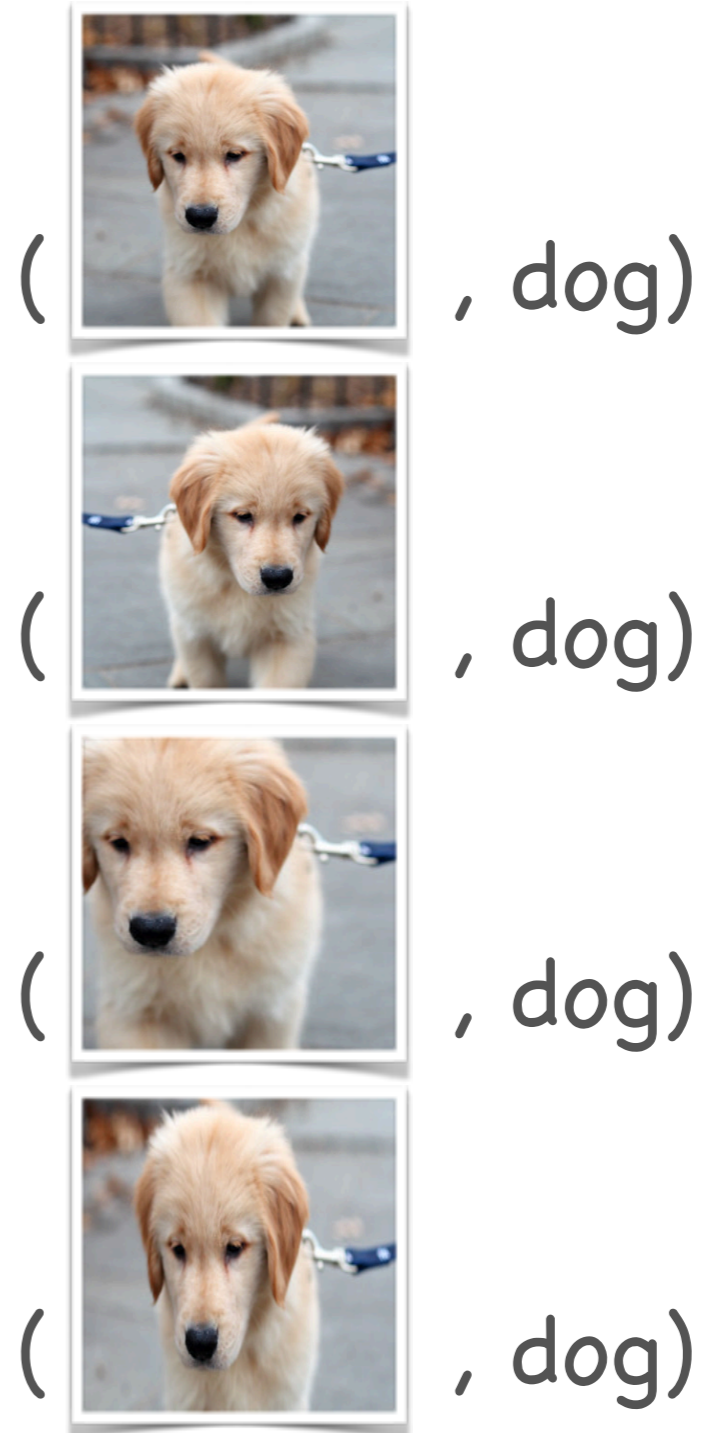


Image augmentations




flip




saturation



shift



brightness




scale



rotate

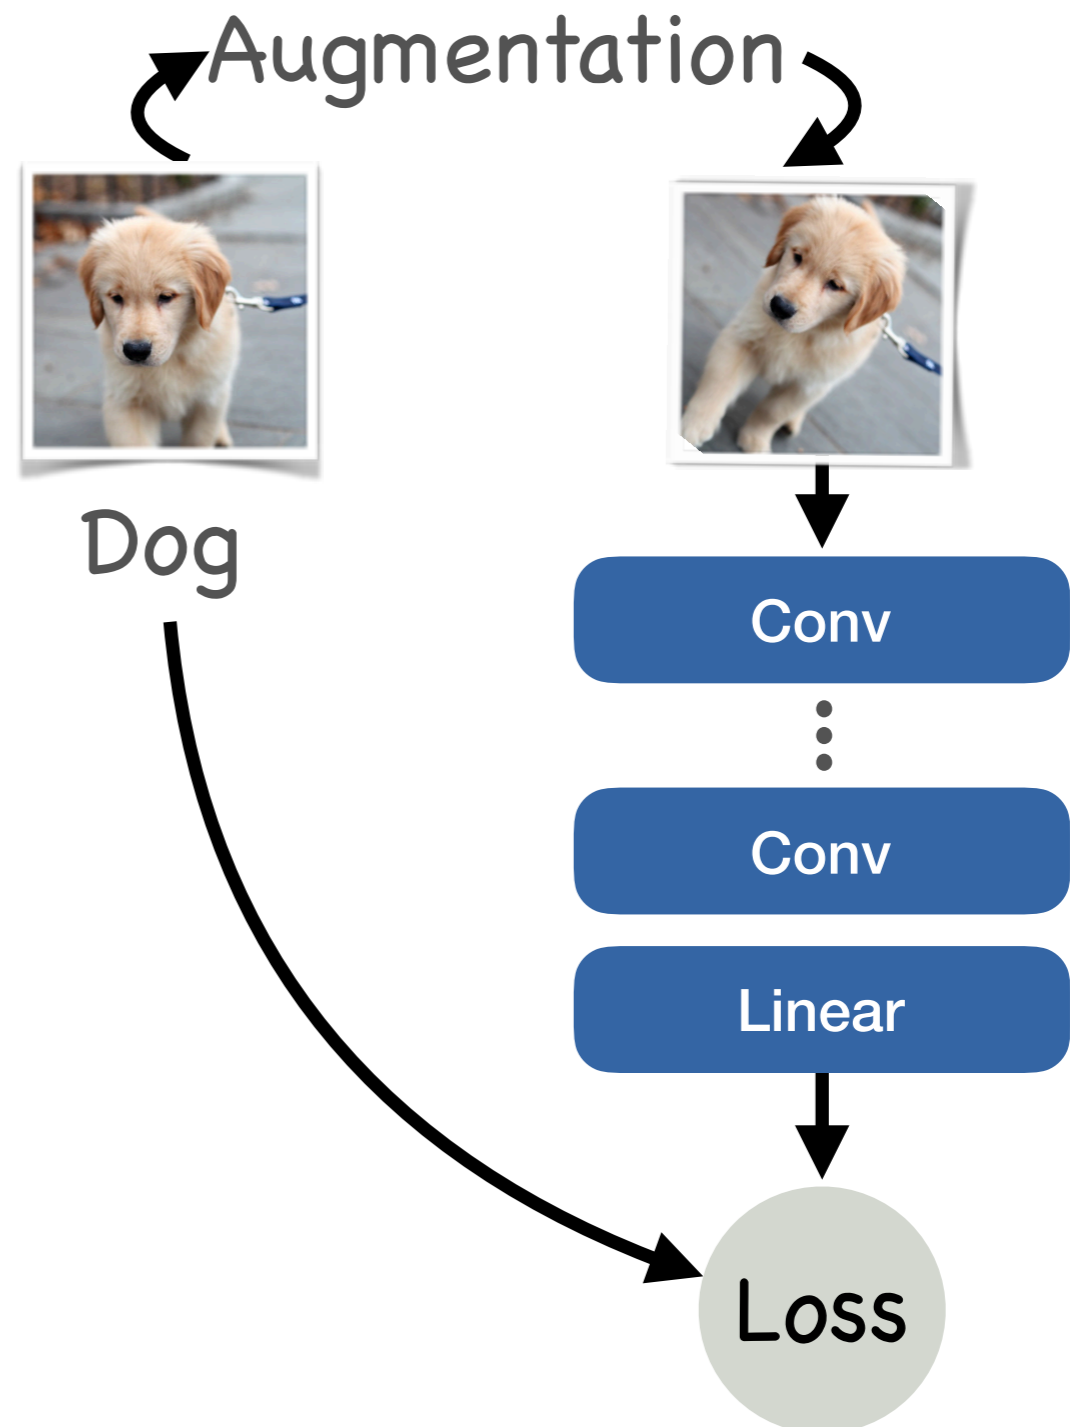


tint/hue



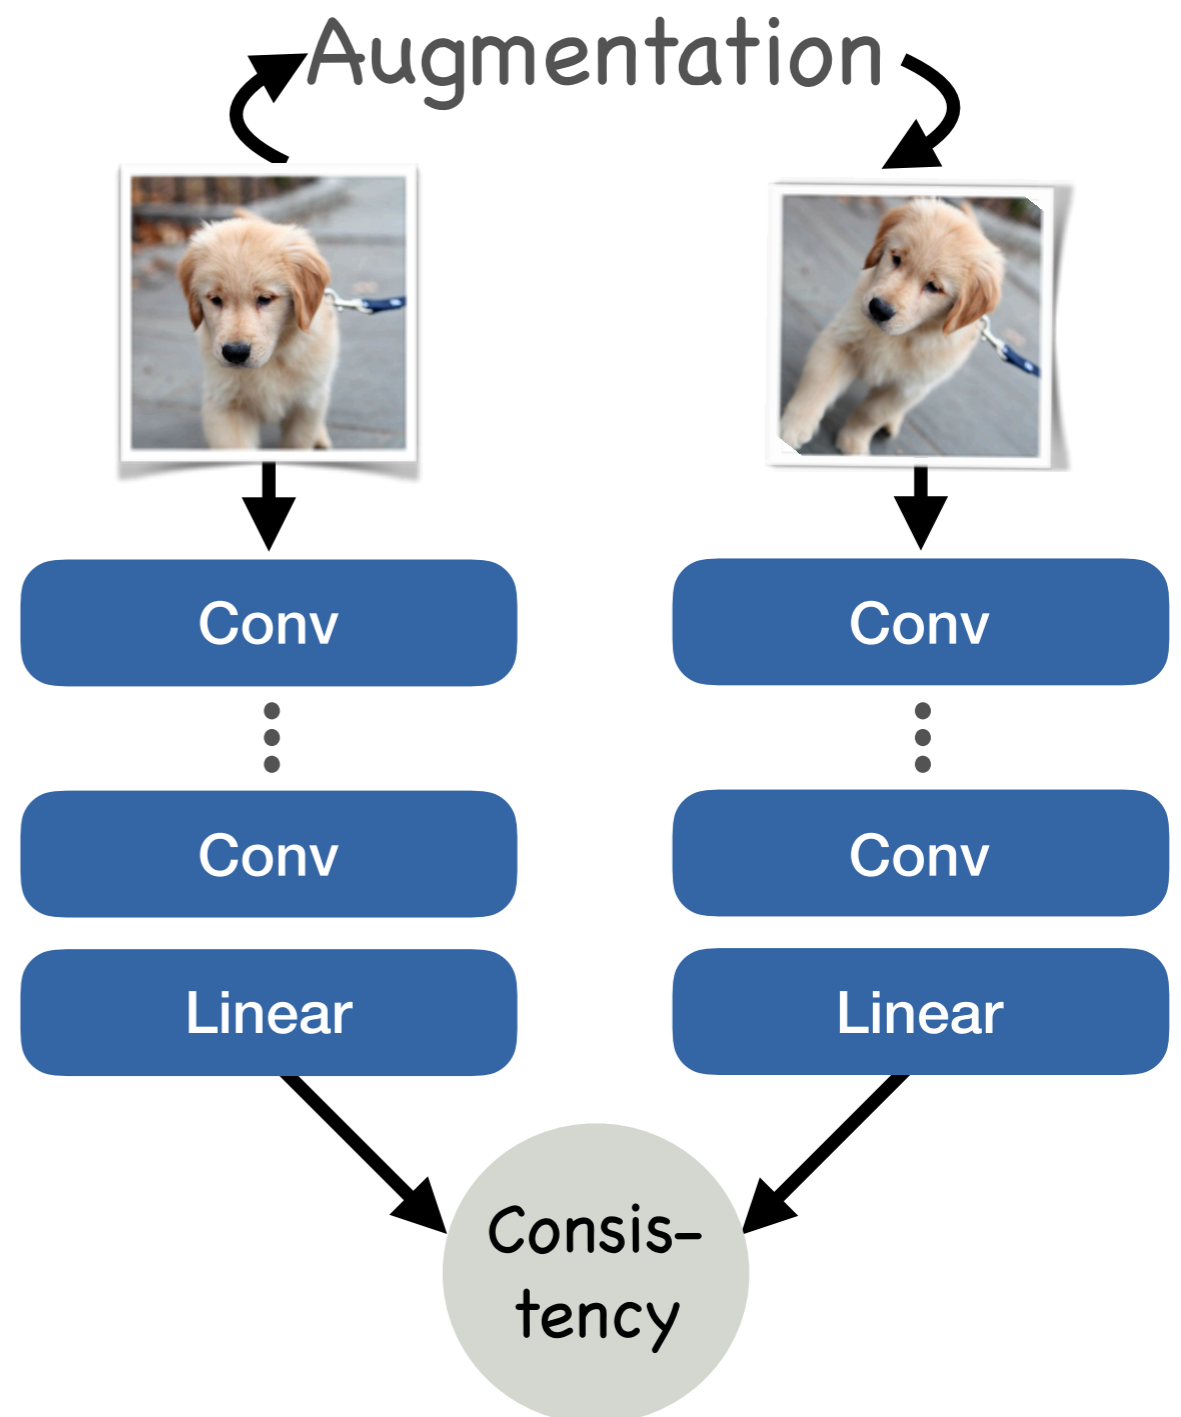
Training with data augmentation

- (Randomly) augment every single iteration
- Network never sees exact same data twice



Unsupervised data augmentation

- Captures invariances on unseen and unlabeled data



Xie, Dai, Hovy, Luong, Le,
"Unsupervised Data Augmentation",
arXiv 2019

Data augmentation

- Always use data augmentation if possible
- Some augmentations require augmentation of labels
 - e.g. for dense prediction tasks