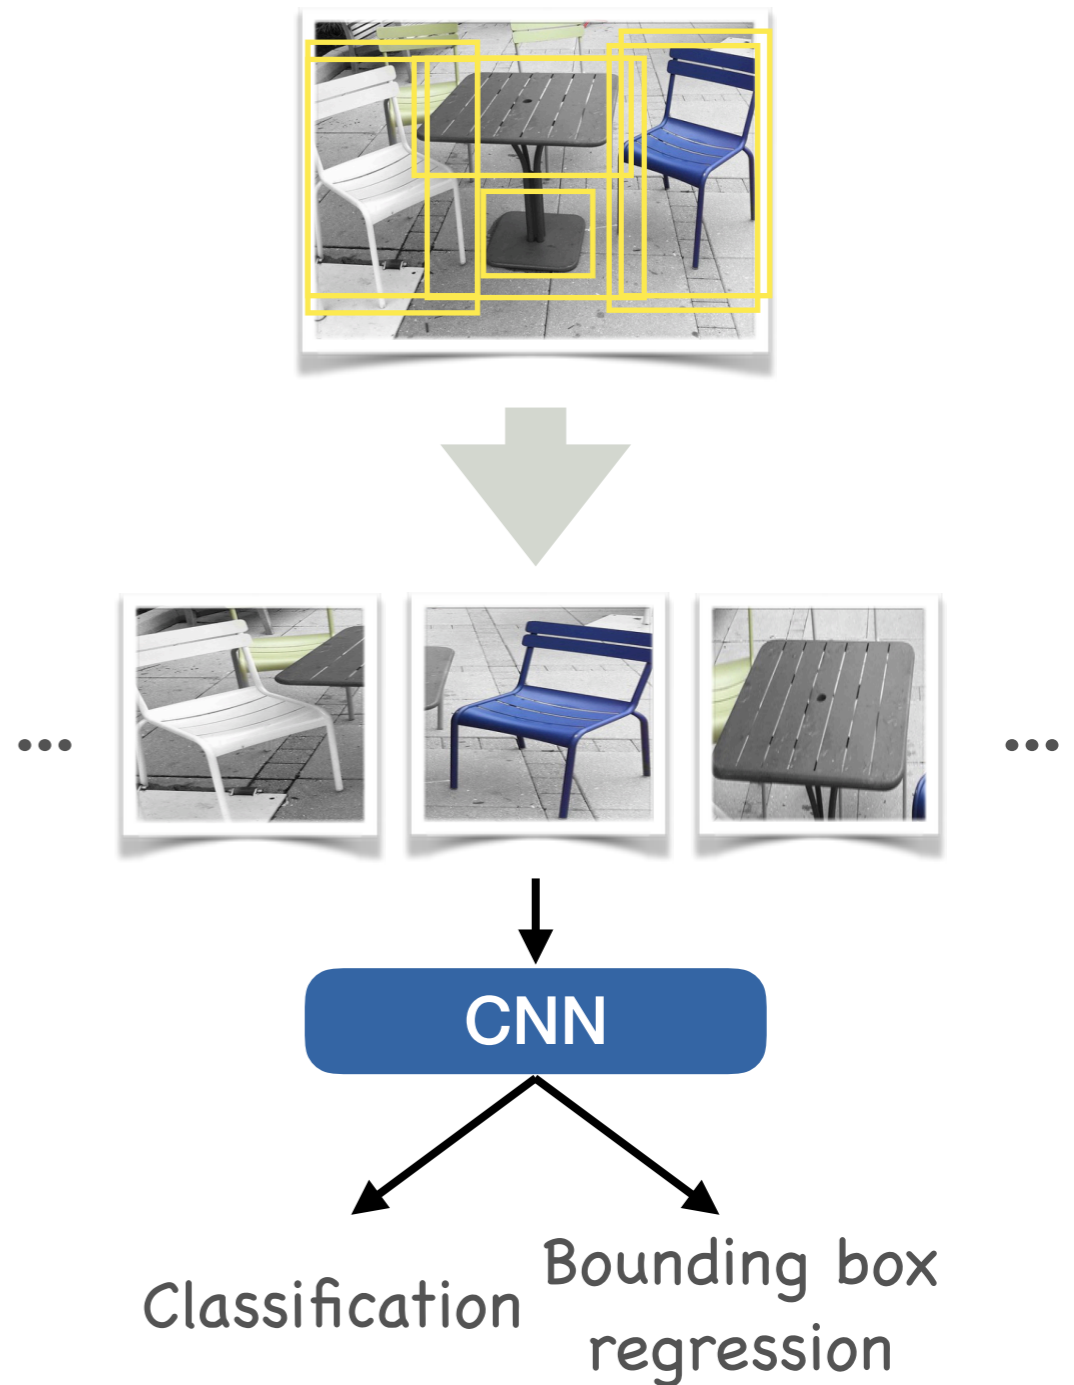


Case study: Faster RCNN

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

RCNN

- RCNN issues
 - Slow (1k forward passes)
 - Relies on good object proposals



Many forward passes

- FastRCNN
 - Forward the entire image through all conv layers
 - Crop activations
 - ROI Pooling, ROI Align



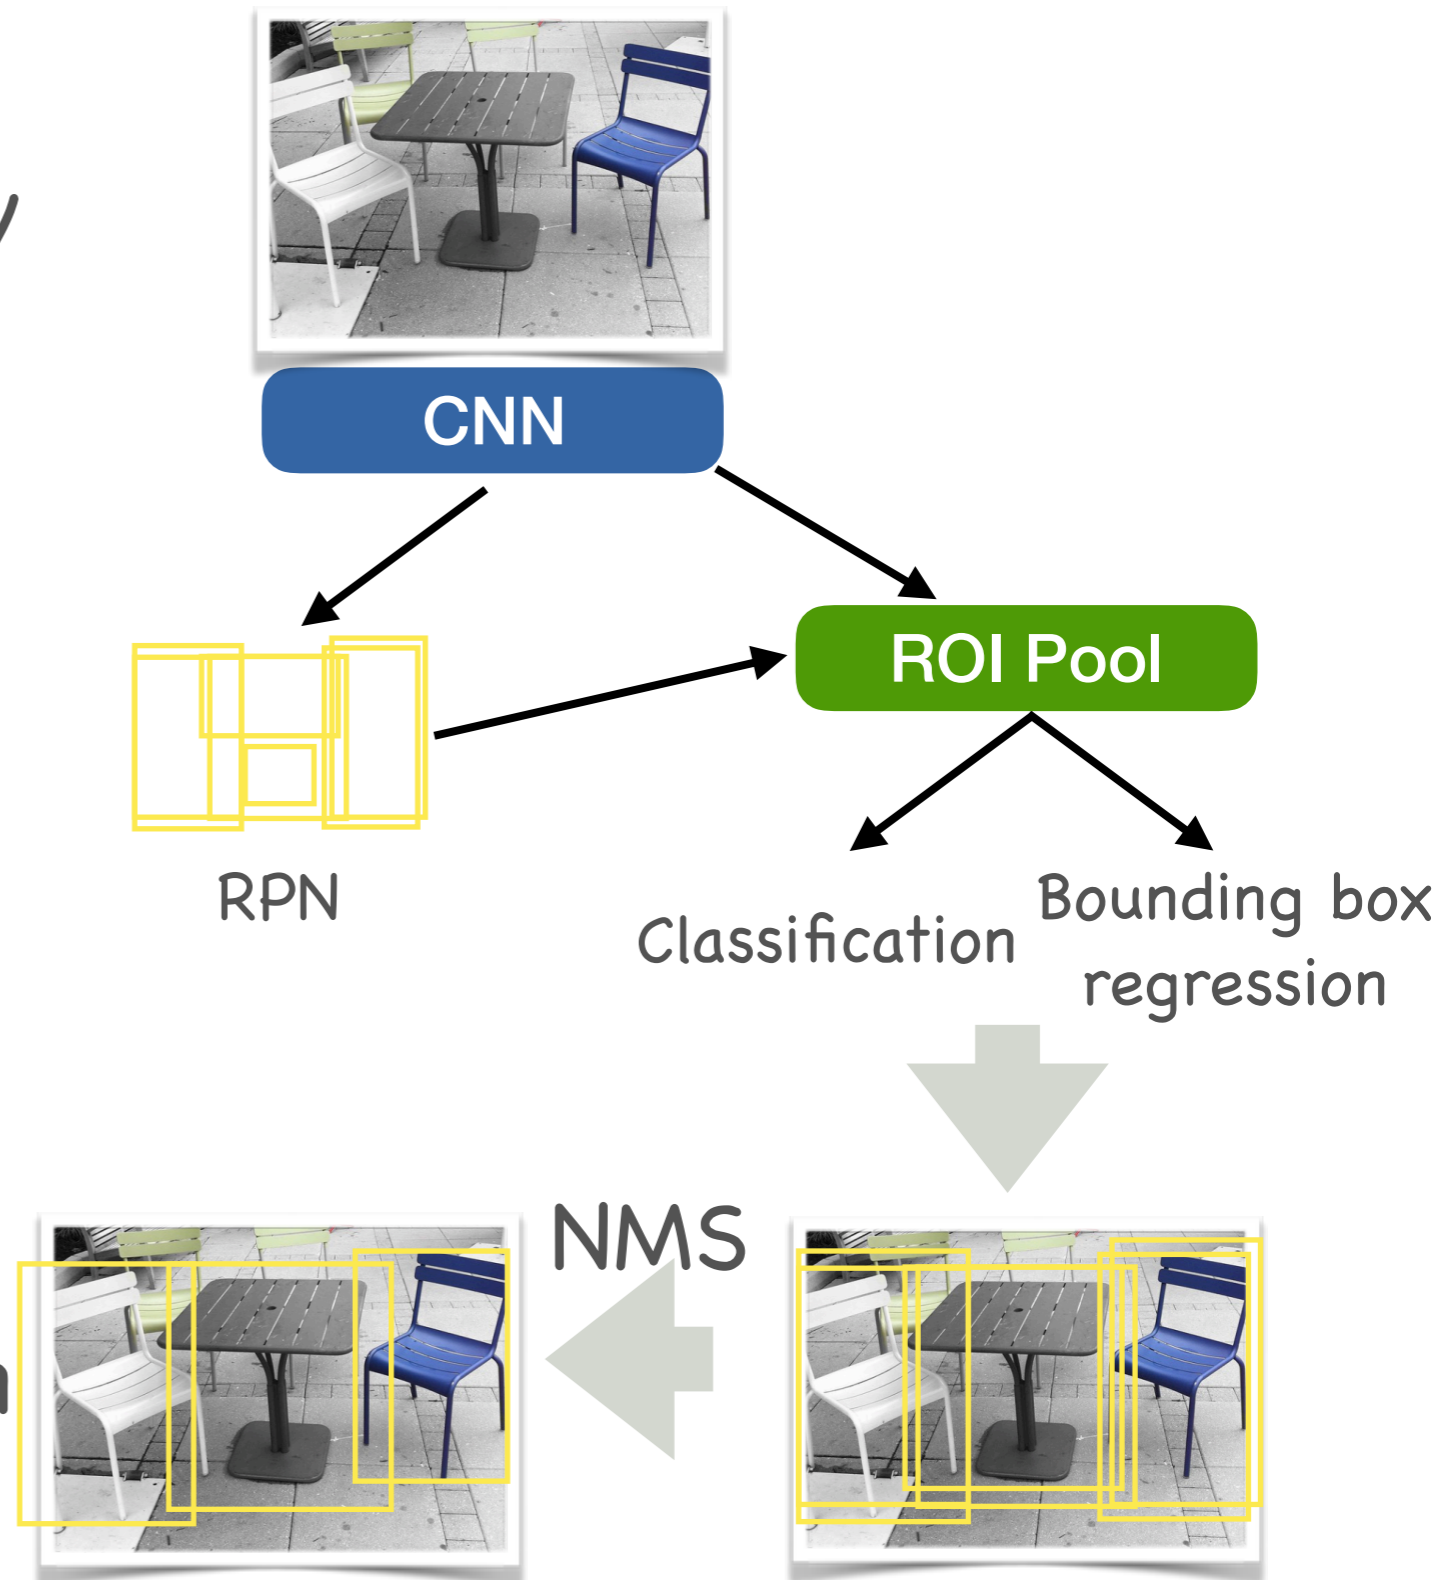
Region proposal network

- Use deep network for region proposals (RPN)
 - Region or not?
 - Try all spatial locations
 - Fully convolutional
 - Try multiple sizes
 - “Anchor boxes”



FasterRCNN

- Run image through fully convolutional network
- Classify Region
- Crop activation (ROI Pooling)
- Classify activation
- Non-maxima suppression



Summary

- Fast
 - 200ms / image
- Single network
 - Trained jointly
 - Assigning positive and negative boxes messy

