Case study: RCNN

Object Detection

- ConvNets classify images well

- How can we use ConvNets to detect objects?

  - e.g. find object locations?

RCNN

- Object detection as region classification

- Object or not?

Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation
Girshick et al., CVPR 2014

RCNN - basic idea

- Crop and resize image regions
- Classify all regions
  - Object or not?

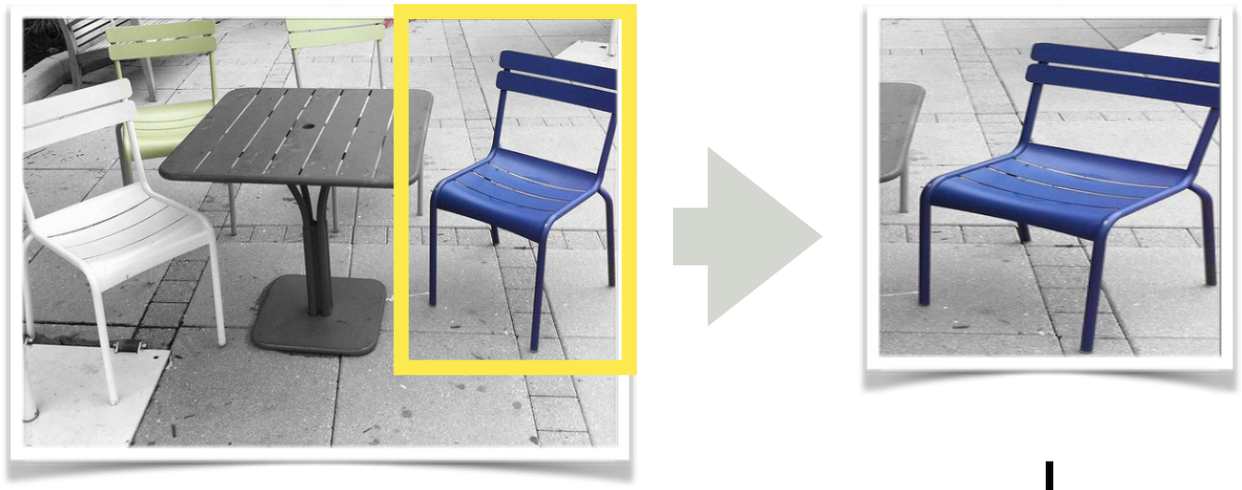

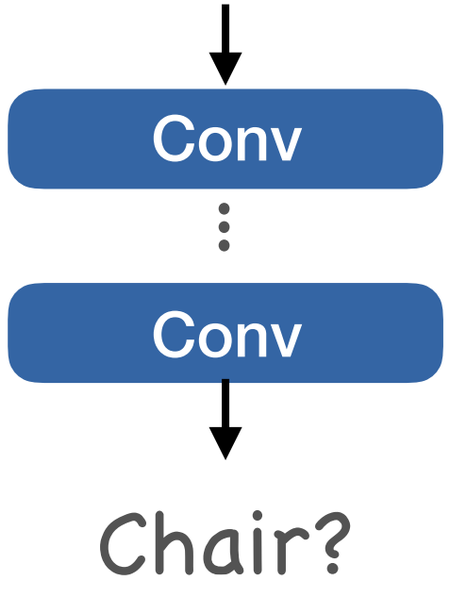

RCNN - Issues

- There are many potential regions

- Solution: region "proposals"

    - Fast low level object classifiers with high recall

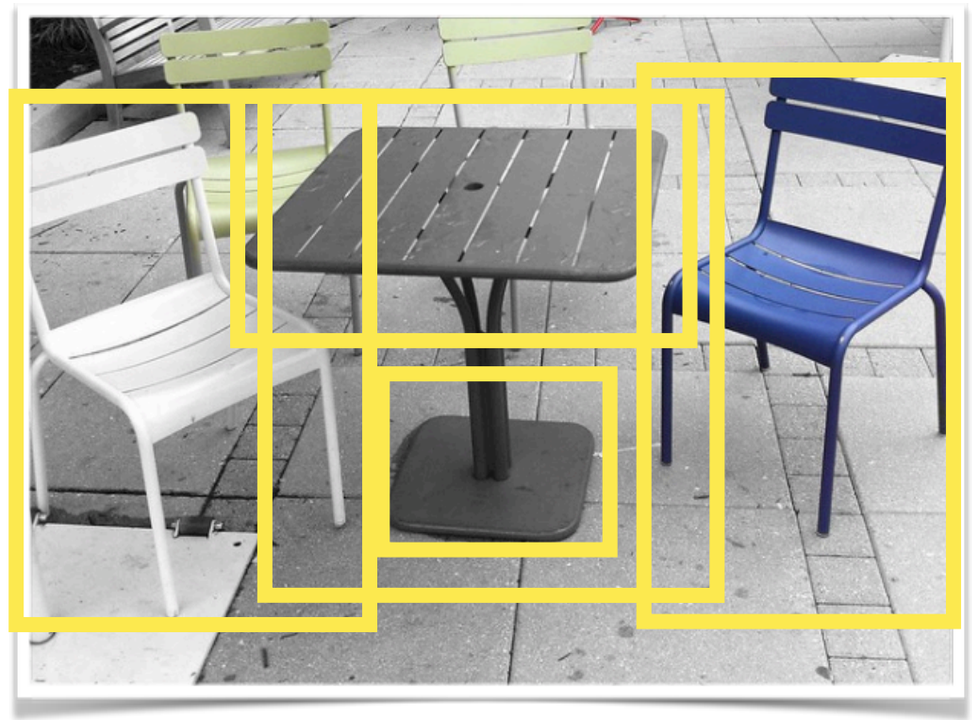

RCNN - Issues

- What do we do with overlapping regions?
  - Solution: Non-maxima suppression
    - Filter our all detections close to a strong detection

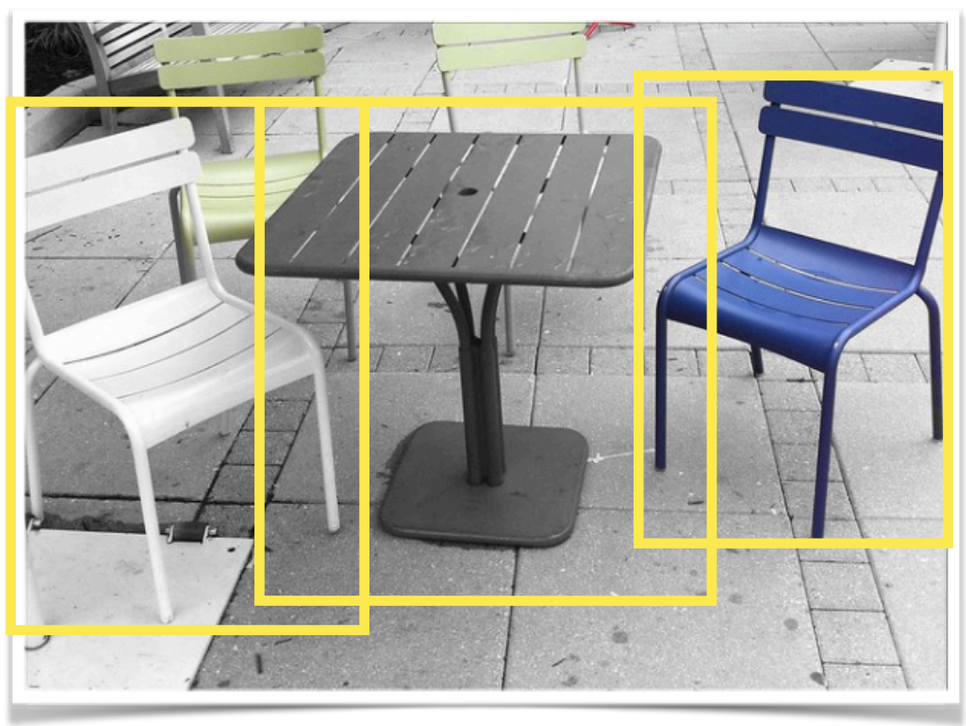

RCNN - Issues

- Proposals might not localize accurately enough

- Solution: Bounding box regression

  - Regress to precise object location

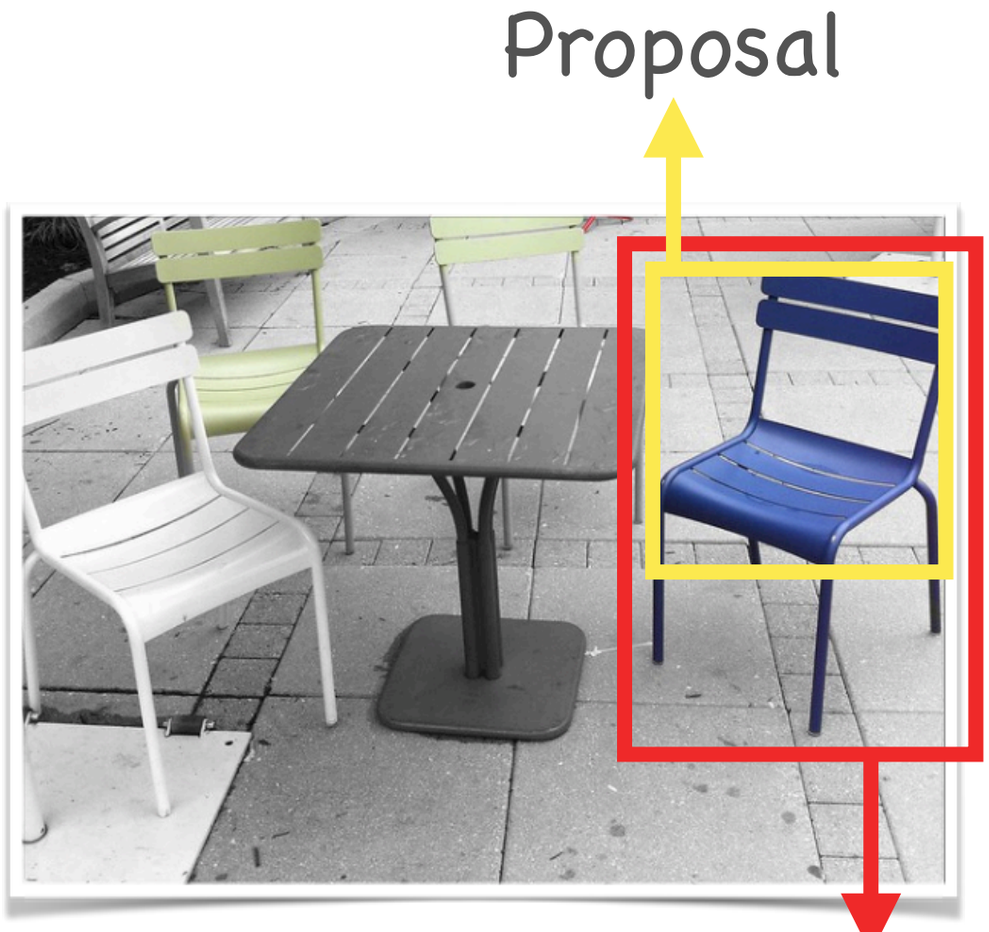

RCNN - Full algorithm

- Find region proposals
- Crop and resize regions
- Classify all regions
  - Object or not?
  - Regress exact location
- Extract detections using Non-maxima suppression

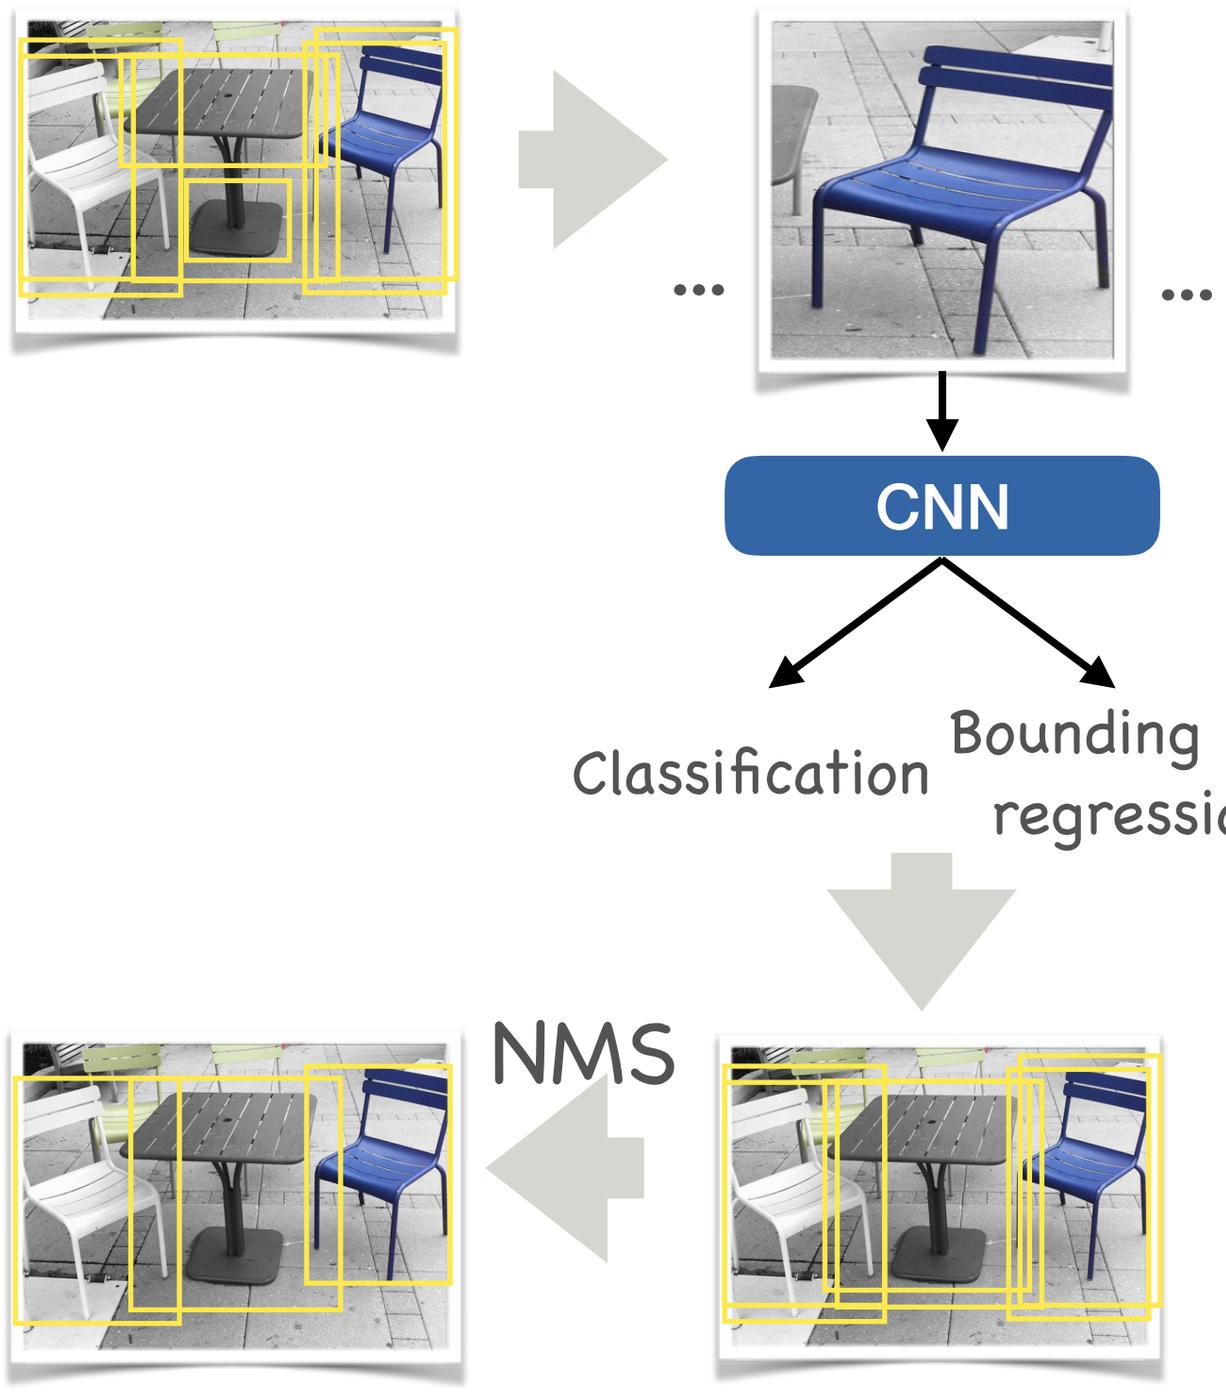

Summary

- State of the art performance 2014
- Slow
  - 1 min / image
  - Too many forward passes (1k+)
- Relies on region proposals