

# Style transfer

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

# Style transfer

- Content of source
- Style of target
- How do we measure style and content?




+



Image generative by [https://github.com/pytorch/examples/tree/master/fast\\_neural\\_style](https://github.com/pytorch/examples/tree/master/fast_neural_style)

# Content

- High level activations of a CNN
- e.g. VGG





# Style

- Low level activations?
- Contain both style and content
- Spatial arrangement



# Style

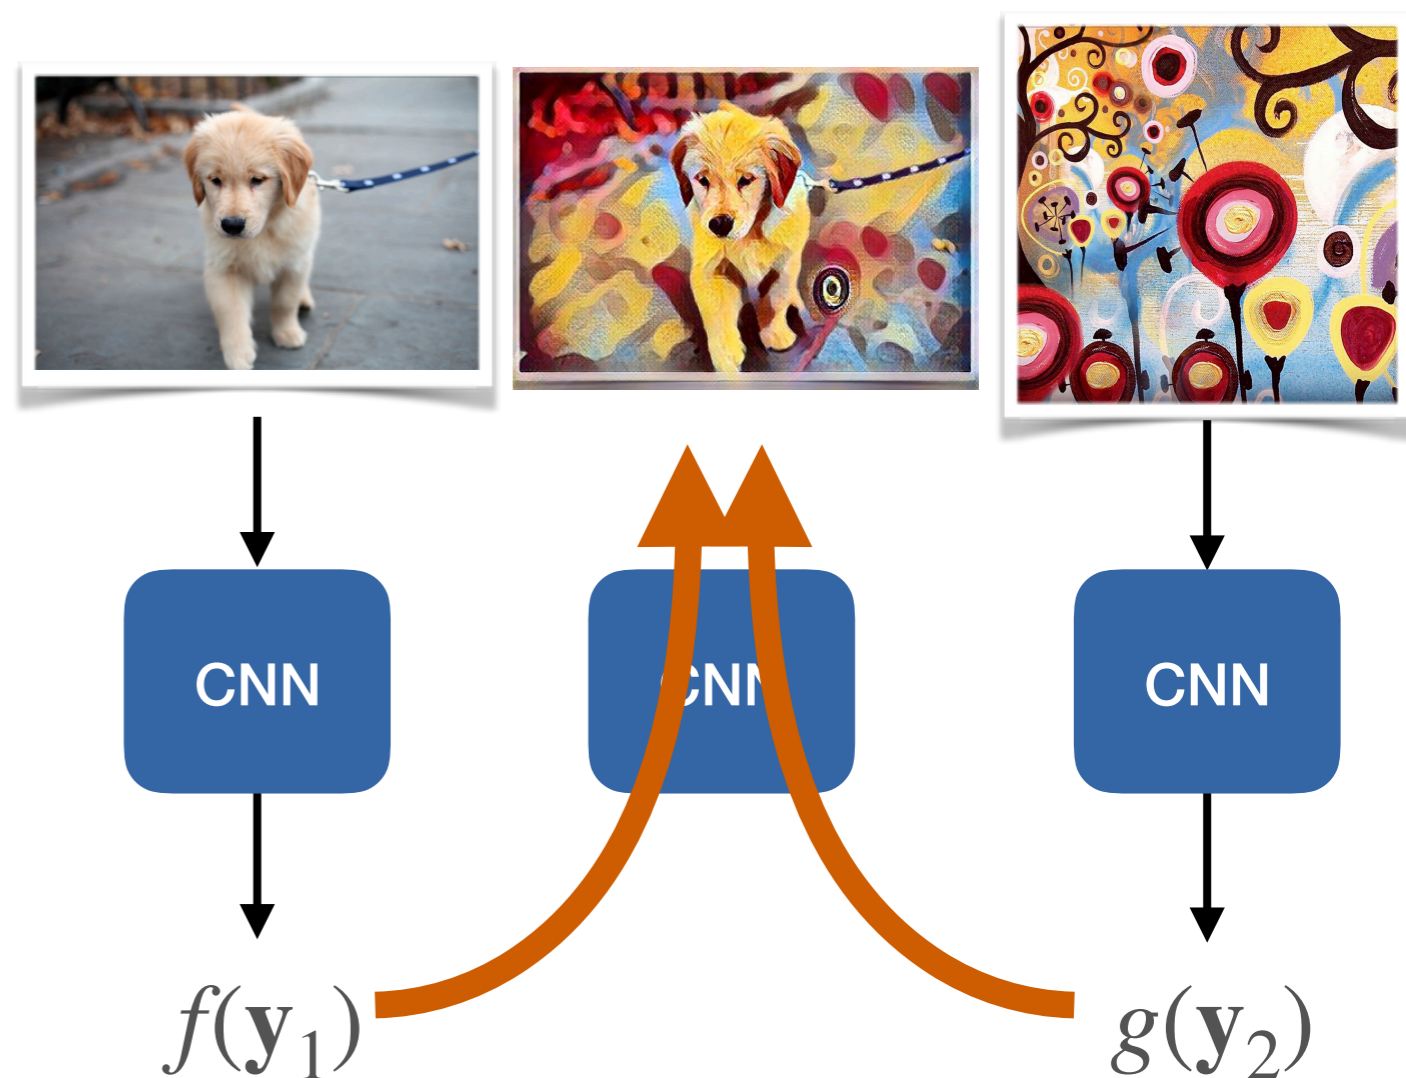
- Solution: Statistics of low level activations
- No (global) spatial information
- Local patterns only





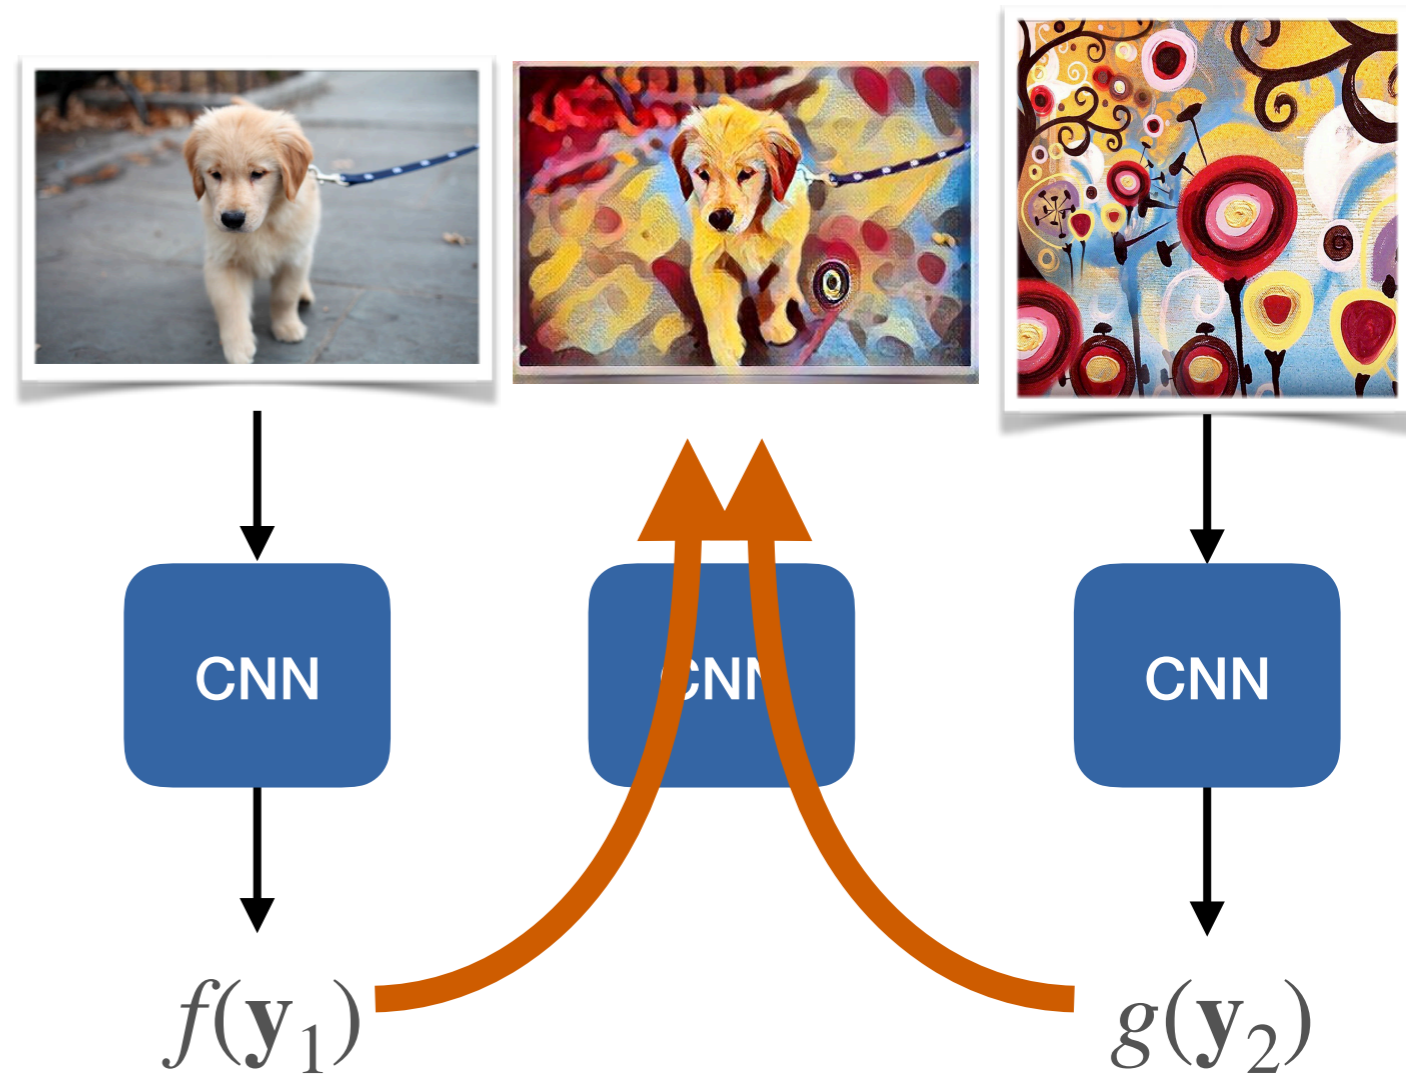
# Neural style transfer

- Optimize for image  $\mathbf{x}$
- With high level activations  $f(\mathbf{y}_1)$
- and Gram statistic of low level activation  $g(\mathbf{y}_2)$



# Neural style transfer - Issues

- Optimization is slow
- 100-1000 backpropagation steps per image



# Fast neural style transfer

- Train a network to do style transfer
  - Objective
    - Match activations
    - Match Gram statistics
- Fast inference
- Slower training

