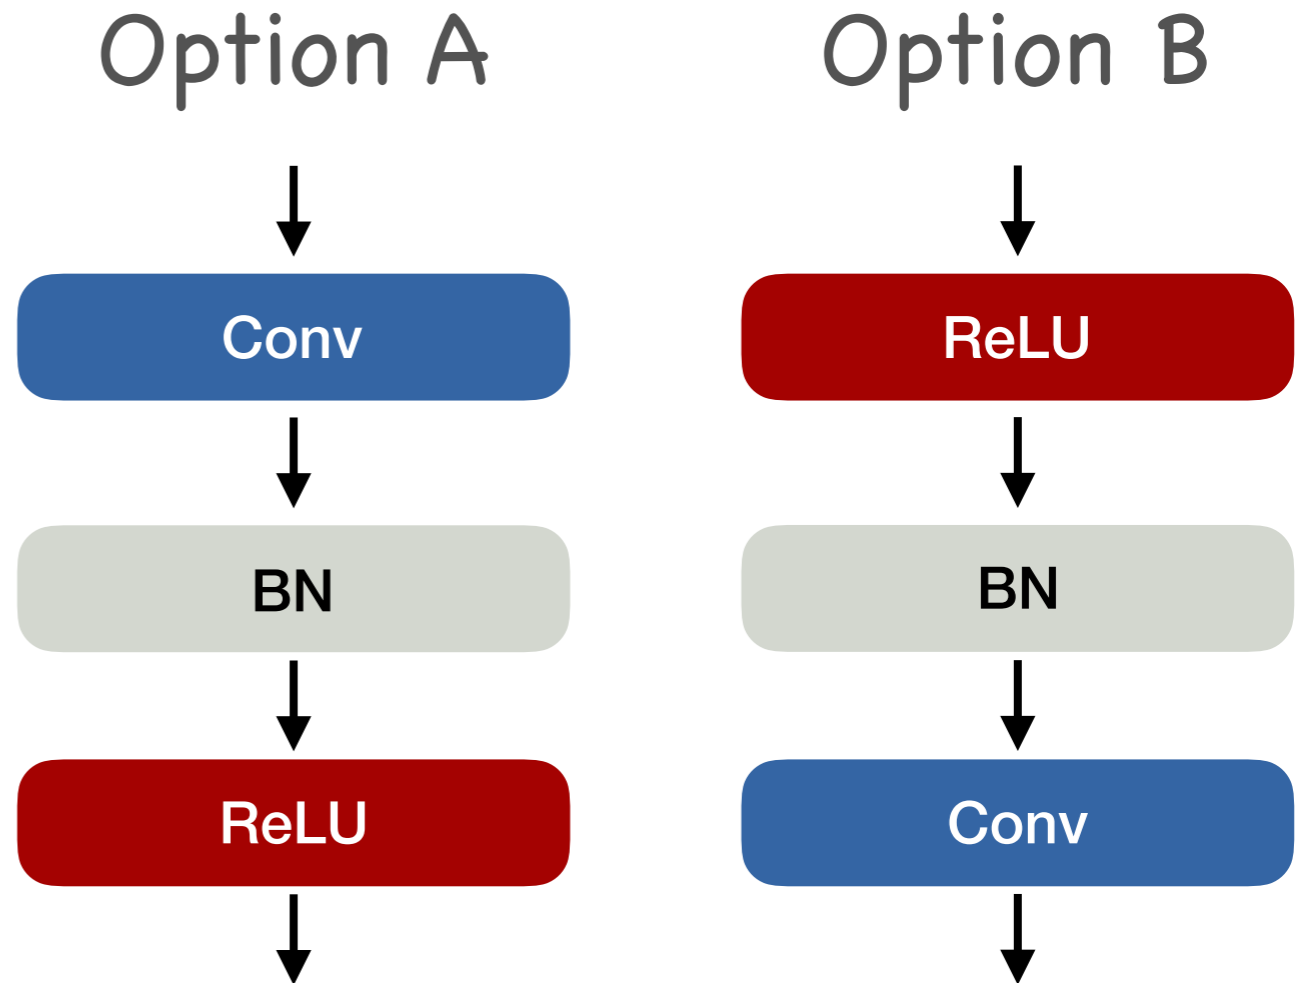


Where to add normalization?

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

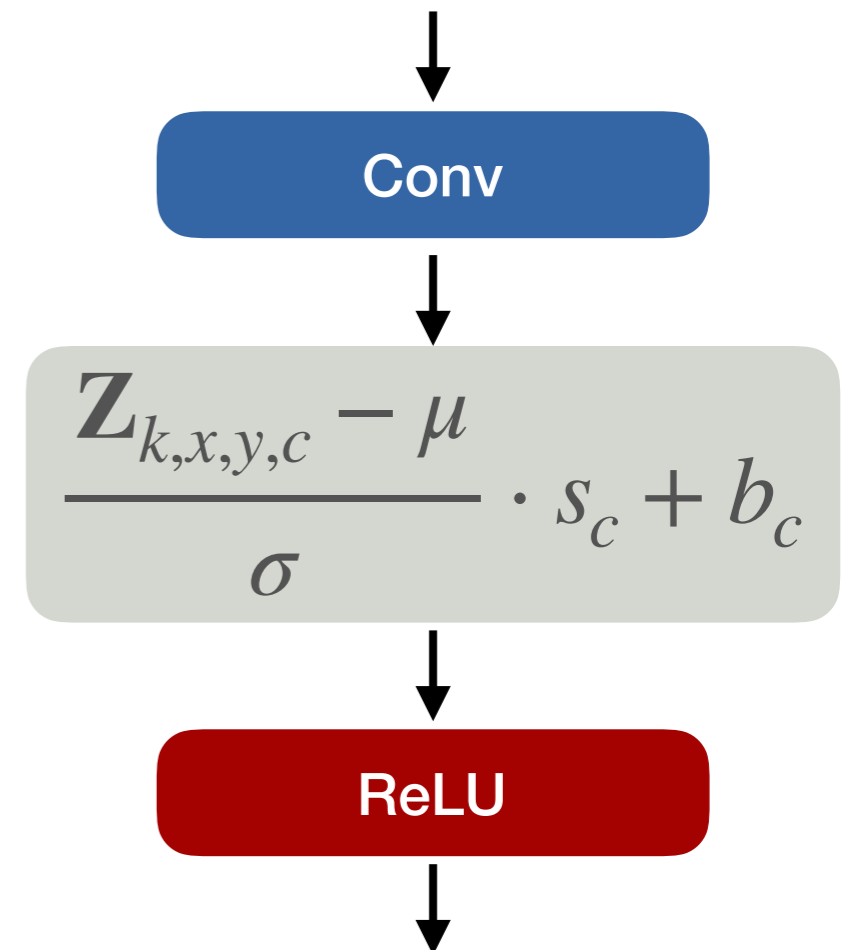
Where to add normalization?

- Option A
 - After convolution
- Option B
 - After ReLU (non-linearity)



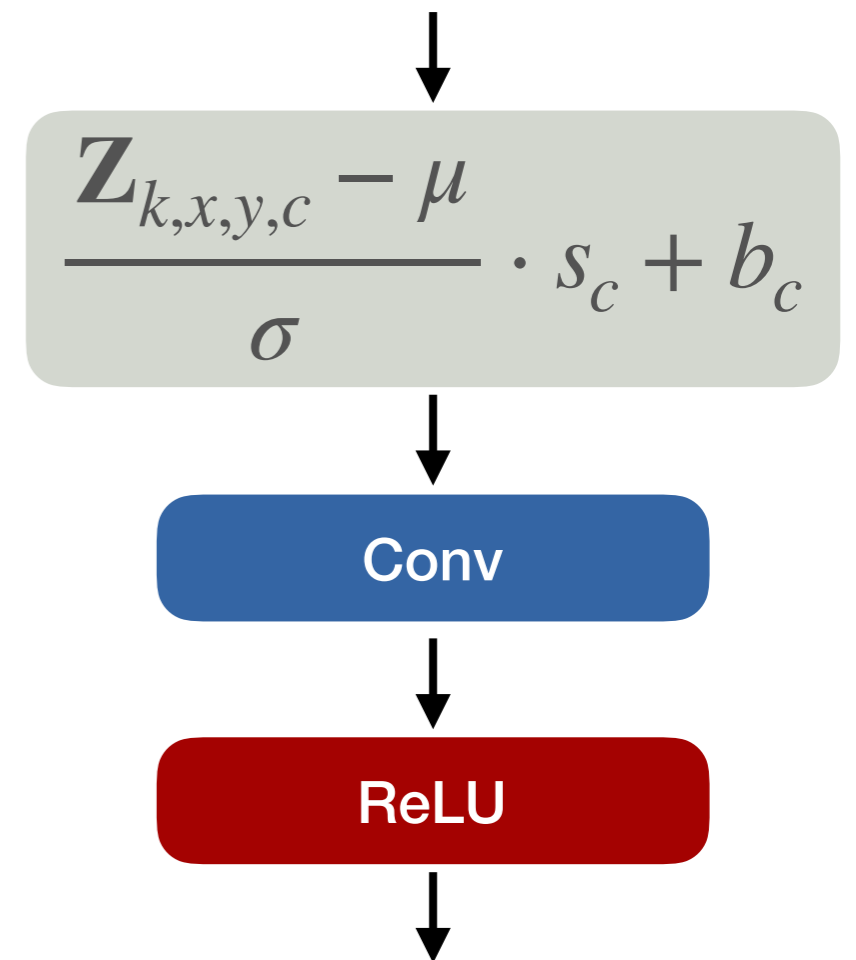
Option A

- No bias in conv
- Activations are zero mean
 - Half of activations zeroed out in ReLU
- Solution:
 - Learn a scale s_c and bias b_c after norm



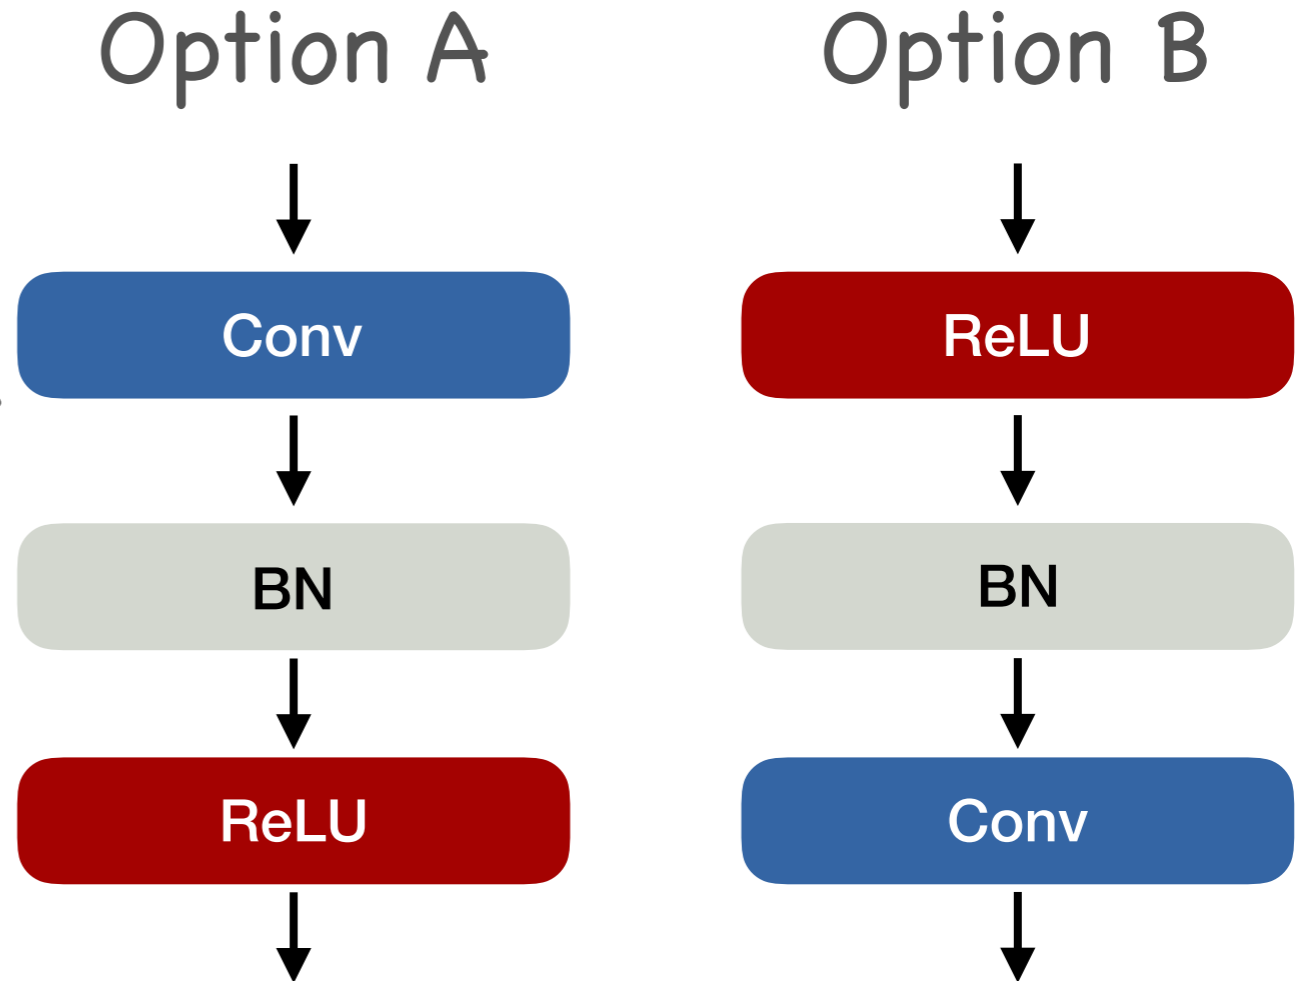
Option B

- Scale s_c and bias b_c optional



Where to add normalization?

- Both work
- Option A is more popular
- Option B is easier
 - Scale and bias optional
 - Conv unchanged



Where not to add batch norm?

- After fully connected layers
- Mean and standard deviation estimates too unstable

Why does normalization work?

- Regularizes the network
- Handles badly scaled weights
- Single parameter to learn scale s_c

