

# Detecting Aliasing Artifacts in Image Sequences Using Deep Neural Networks

Anjul Patney and Aaron Lefohn

NVIDIA

# Motivation

- We build a lot of renderers, and we evaluate them by their output image quality
- Manual evaluation is tedious and subjective
- Automatic evaluation is fast, but involves comparison against a known reference

# Aliasing





# Noise





# Compression Artifacts



# Image Quality Evaluation

- You **do not** need a reference for most artifacts
- You **do** need higher level visual processing



# Key Idea: Use Machine Learning

- Recent deep neural networks have demonstrated remarkable visual processing skills
  - Object classification
  - Face recognition
  - 3D reconstruction
- Similar networks can likely learn to identify image quality artifacts
- This short paper studies aliasing artifacts

# Network Architecture




# Start with a well-studied deep neural network classifier

- We picked VGG-19
  - High accuracy on classification and localization
  - Demonstrated success as a perceptual loss function

# Start with a well-studied deep neural network classifier (VGG-19)


224 x 224



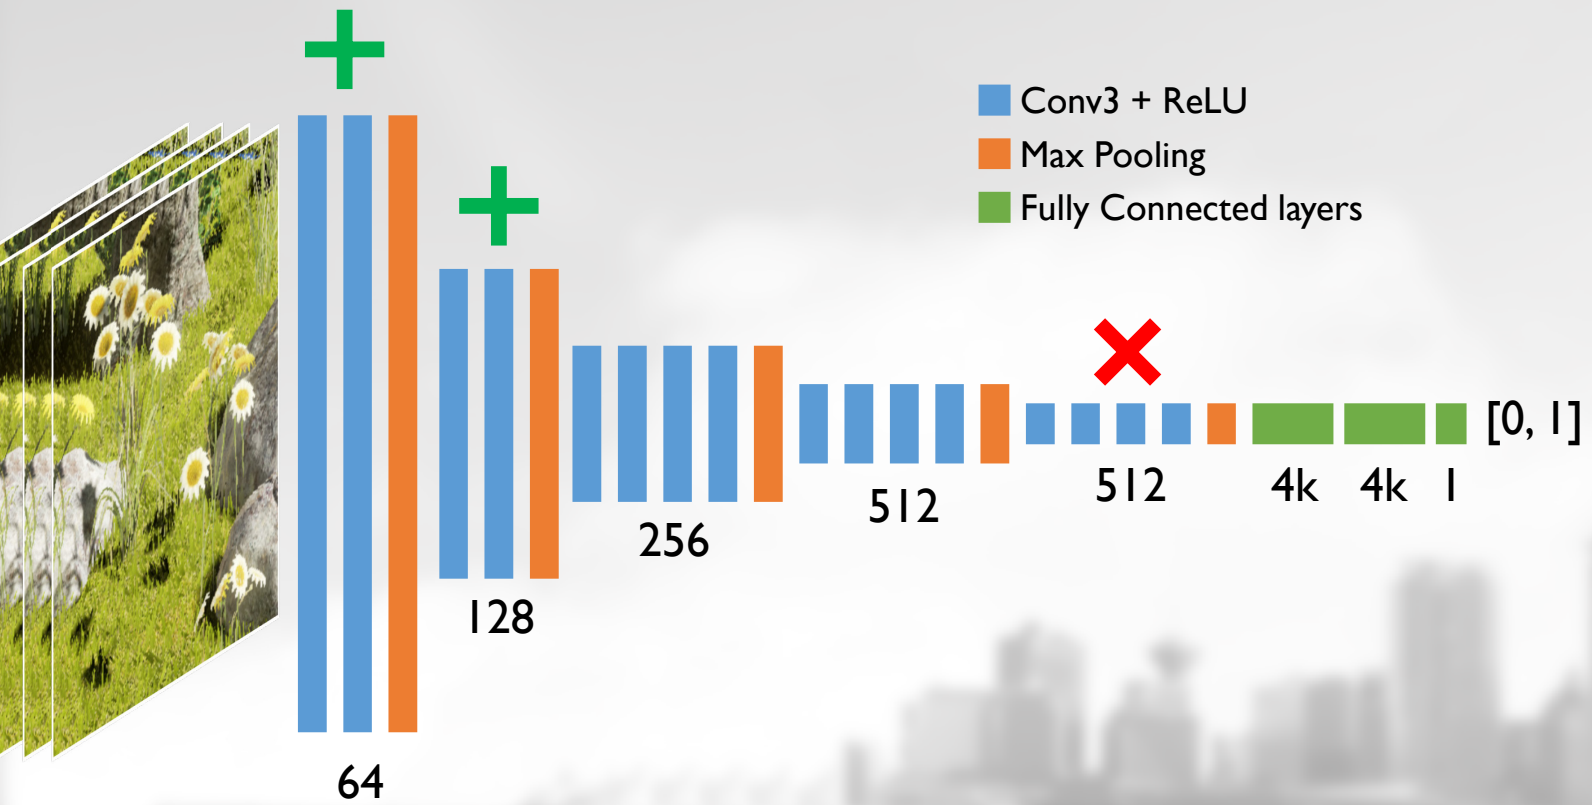


# Add support for multi-frame sequences

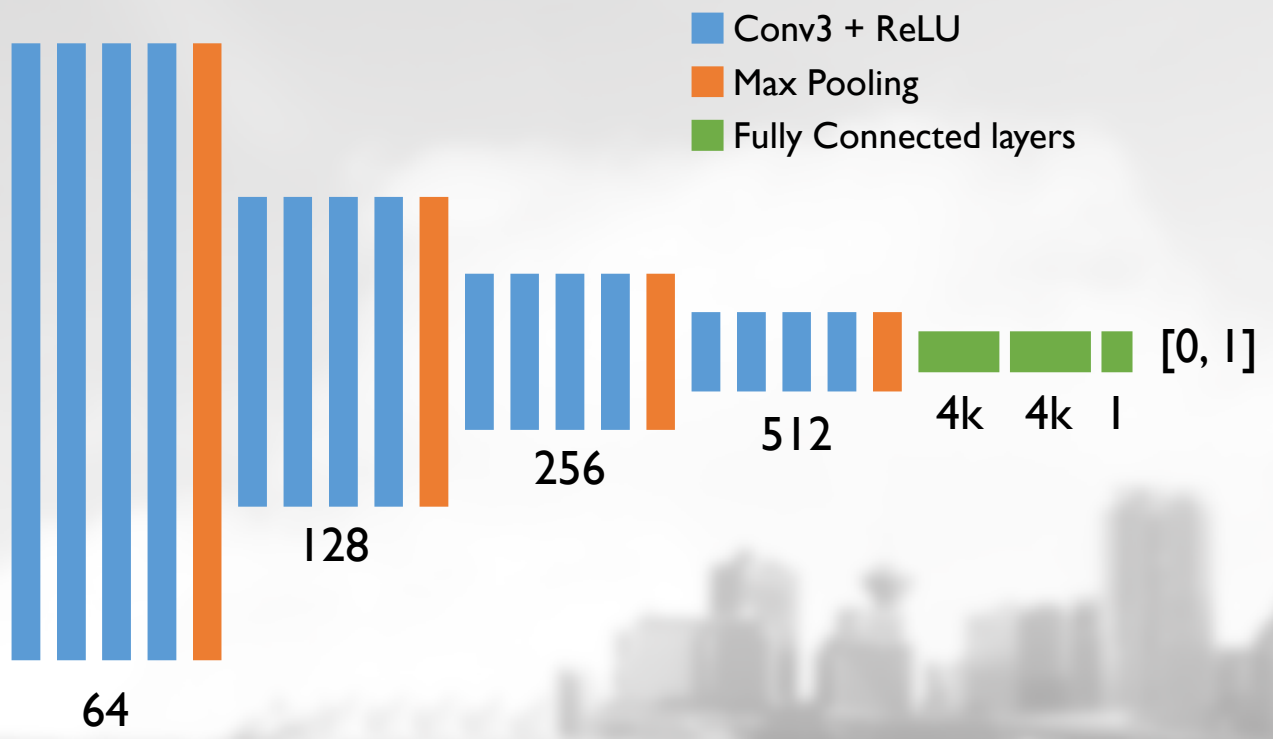
64 x 64 x 4



# Prioritize learning at local scale (nearby pixels)




# Prioritize learning at local scale (nearby pixels)







# Other changes, determined empirically



# Other changes, determined empirically



# Training




Antialiased  
0.0

1 spp: aliased

16-64spp: antialiased




# Training



1 spp: aliased

16-64spp: antialiased

# Training

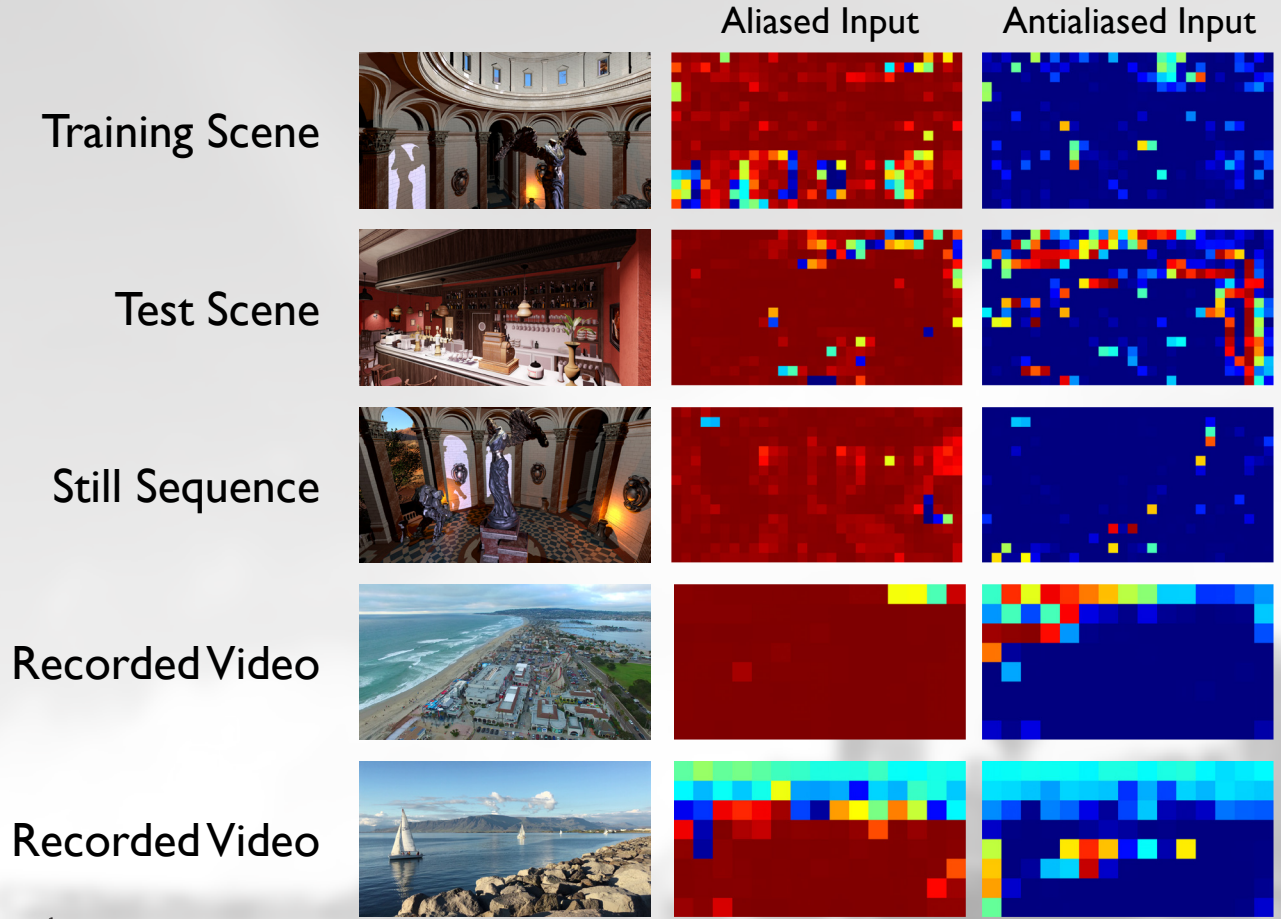


1 spp: aliased

16-64spp: antialiased

# Results

# Network Output





# Sun Temple (included in training)


Bistro (excluded from training)

# Still Sequence

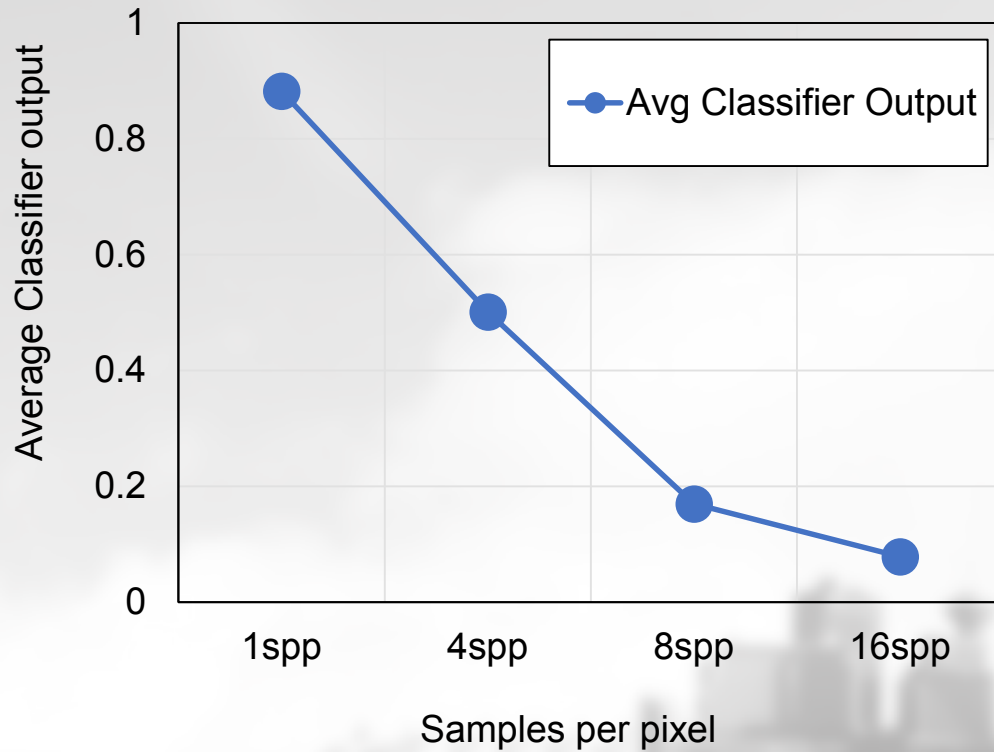
# Real-life Video I




# Real-life Video 2



# Output Trend




# Output Trend (follows 1-SSIM)




# Trained network is selective to Aliasing Artifacts

1 spp




SSIM 0.9709

4 spp




SSIM 0.9915

8 spp




SSIM 0.9966

16 spp



SSIM 1.000


TAA



SSIM 0.9269



# Trained network is selective to Aliasing Artifacts



SSIM indicates low quality, but our network correctly estimates low aliasing





# Conclusion

- Deep neural networks enable reference-free image quality assessment
- Selective classification helps isolate artifacts
- Future work: detect higher-level visual artifacts
  - Does this image/sequence contain light leaks?
  - Are the shadows in this image/sequence consistent?
  - Is this image/sequence realistic?

Thank you!