Carl S. Marshall | April 25, 2022 | Eurographics

Practical Machine Learning for Rendering

From Research to Deployment

Course Goals

Give insights into recent neural models and help close the gap between taking a research neural model to deployment

Understand the challenges in data acquisition, development, training, deployment, and iteration of neural networks for rendering

Show practical use cases, neural models to start your path toward neural rendering in production software

Schedule

Introduction
Carl S. Marshall, Reality Labs Research at Meta
15 mins

ML for Graphics: A Brief Overview
Deepak Vembar, Intel Labs
40 mins

Synthetic Data For Computer Vision: Techniques, Challenges, and Tools
Sujoy Ganguly, Unity
40 mins

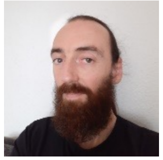

Machine Learning in Real-time
Florent Guinier, Unity Labs
40 mins

Conclusion
5 mins

Areas of Exploration

What are the latest techniques for Machine Learning in Rendering?

What types of neural network models have shown promising results?

Where can I get data to train my models?

How do I practically deploy my ML models into a rendering engine?

Challenges

Image Credit: Unity

Real-time vs. Offline

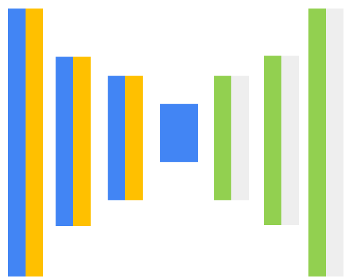

Performance vs. Quality
Model arch tradeoffs
Target HW tuning

Data Acquisition

Dataset availability
Models to curate datasets
Real world vs. synthetic

Deployment

ONNX RUNTIME

unity Barracuda

Single frame rendering
Integrated into engine
Hand tuning vs. API

Additional Rendering Attributes

Input Buffers | Resolution Changes | Temporal Stability | HDR/LDR Lighting

Reproducibility **Robustness** **Quality Comparison**

Simplified Practical ML for Rendering Workflow

Define Goals
- Desired output
- Performance/Quality
- Criteria for Success

Acquire Training Data
- Public datasets
- Real, synthetic, mixed
- 3D model availability

Model Development
- Define Input/Output
- Model architecture
  VAE, GAN, MLP, etc.
- Loss functions

Training & Deployment
- Train/Test iteration
- Deploy: ONNX Run-time, Unity Barracuda, DirectML
- Optimizations

Research Examples

Super Resolution

Image Credit: Meta Research

Frame Extrapolation

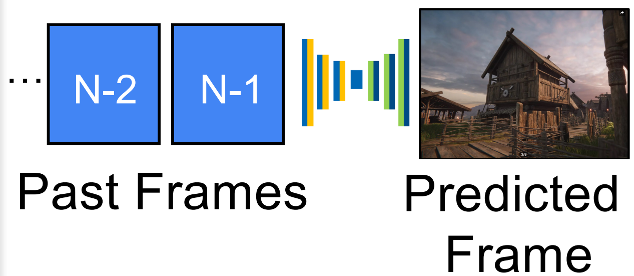

Image Credit: Unity Labs

Style Transfer

Image Credit: Intel Labs

Style Transfer: Goals

- Real-time for videos and 3D graphics scenes
- Temporally consistent
- High-Definition resolution
- Ability to segment objects for personalized style transfer
- Tradeoffs:
  - Training per Style versus Universal Style Transfer
  - Temporal stability through training or inferencing

Style Transfer:
Data Acquisition and Model Development

- Data Acquisition
  - *Style Transfer:* FlyingThings3D and Monkaa which provide Optical Flow and Motion boundaries for each consecutive frame
  - *Character Segmentation:* experimented with COCO, Supervisely Person Dataset and Carvana mask datasets
- Model Development:
  - Explored many different model architectures and started with ReCoNet as a base architecture with enhancement
  - Add ability for per-object style transfer

Style Transfer: Deployment

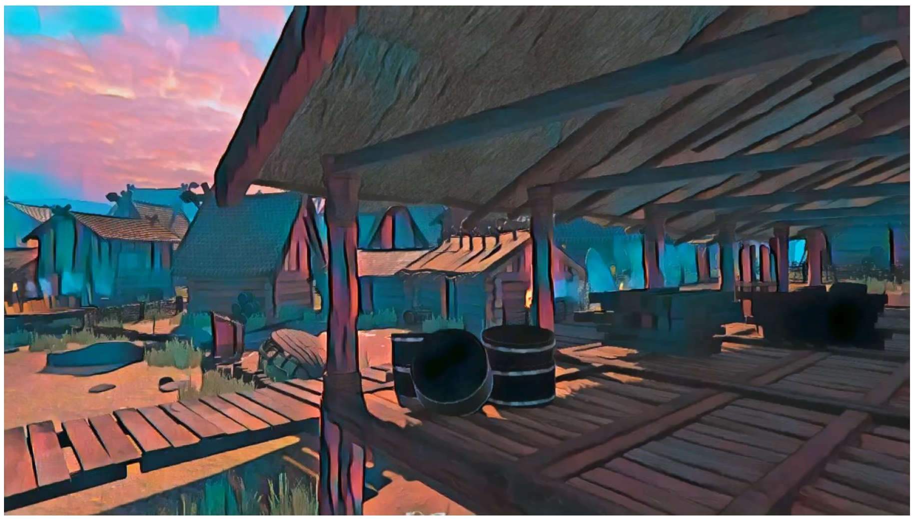

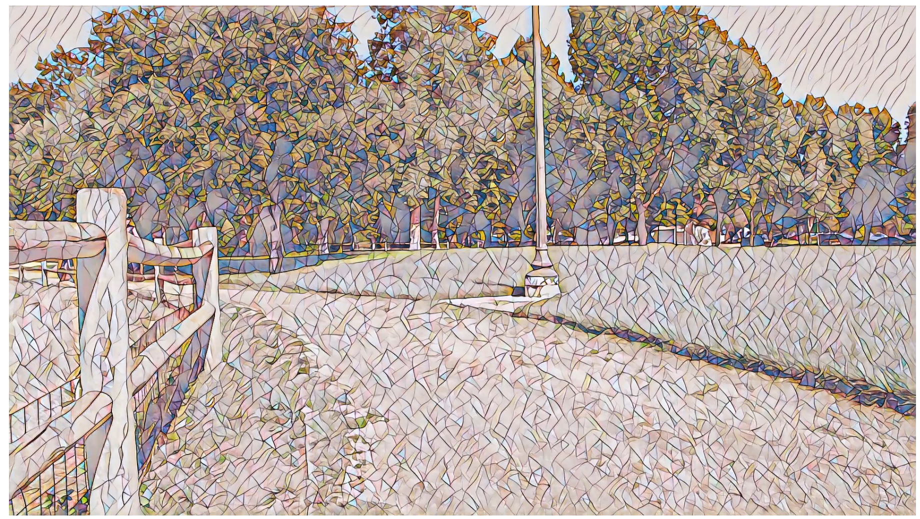

Videos Credit: Intel Labs

Style Transfer:
Segmentation Results

Image Credits: Intel Labs

Schedule

Introduction
Carl S. Marshall, Reality Labs Research at Meta
15 mins

ML for Graphics: A Brief Overview
Deepak Vembar, Intel Labs
40 mins

Synthetic Data For Computer Vision: Techniques, Challenges, and Tools
Sujoy Ganguly, Unity
40 mins

Machine Learning in Real-time
Florent Guinier, Unity Labs
40 mins

Conclusion
5 mins

Carl S. Marshall | April 25, 2022 | Eurographics

Conclusion:
Practical Machine Learning for Rendering

From Research to Deployment

Brief Recap Continued:
ML in Rendering Overview: Workflow and Challenges

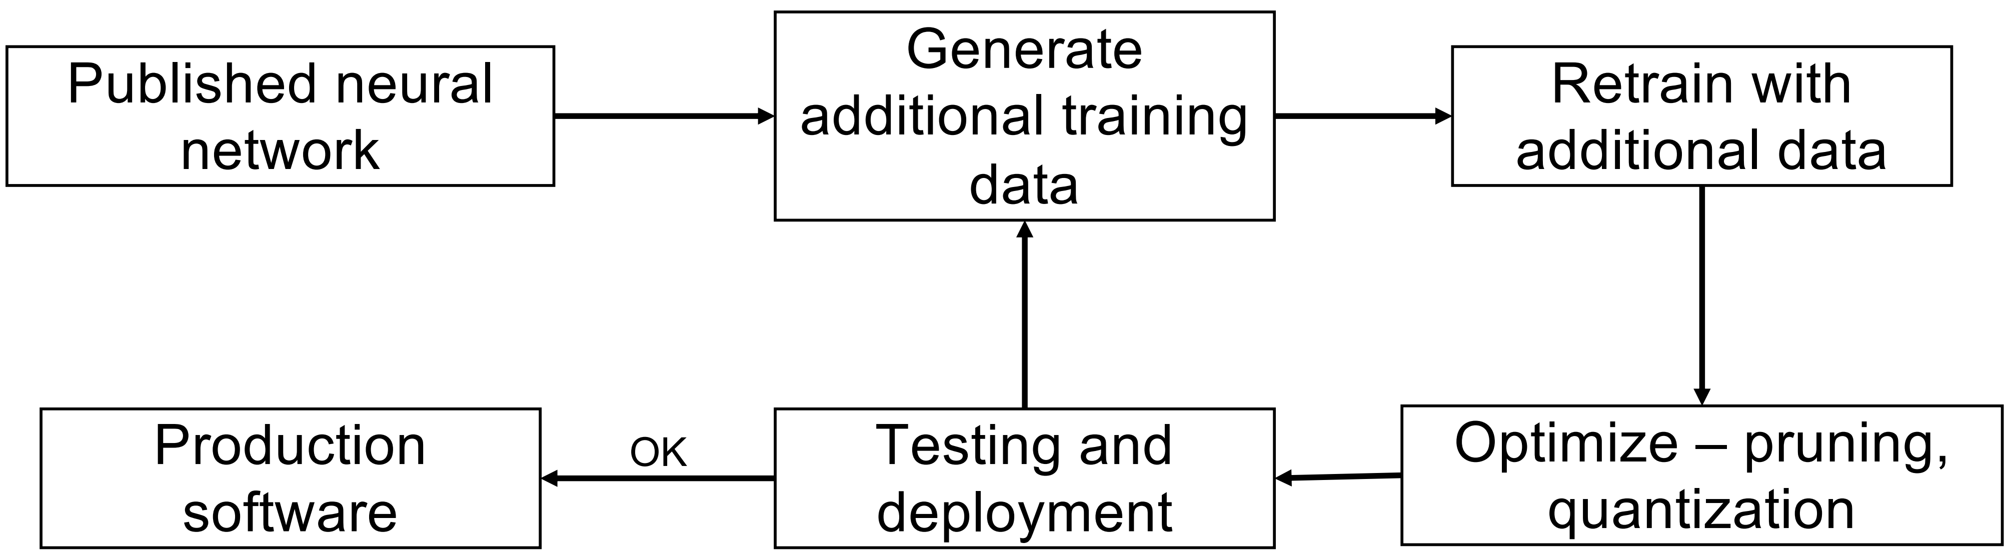

Brief Recap Continued: Synthetic Data Generation

- Synthetic data generation
  - Described methods to bridge the Sim-to-Real gap
  - Burdens of Domain Randomization
  - Sensor and Perception SDKs
  - Benchmark environments: SynthDet, SynthCOCO-18, PeopleSansPeople

Brief Recap Continued:
Machine Learning in Real-time

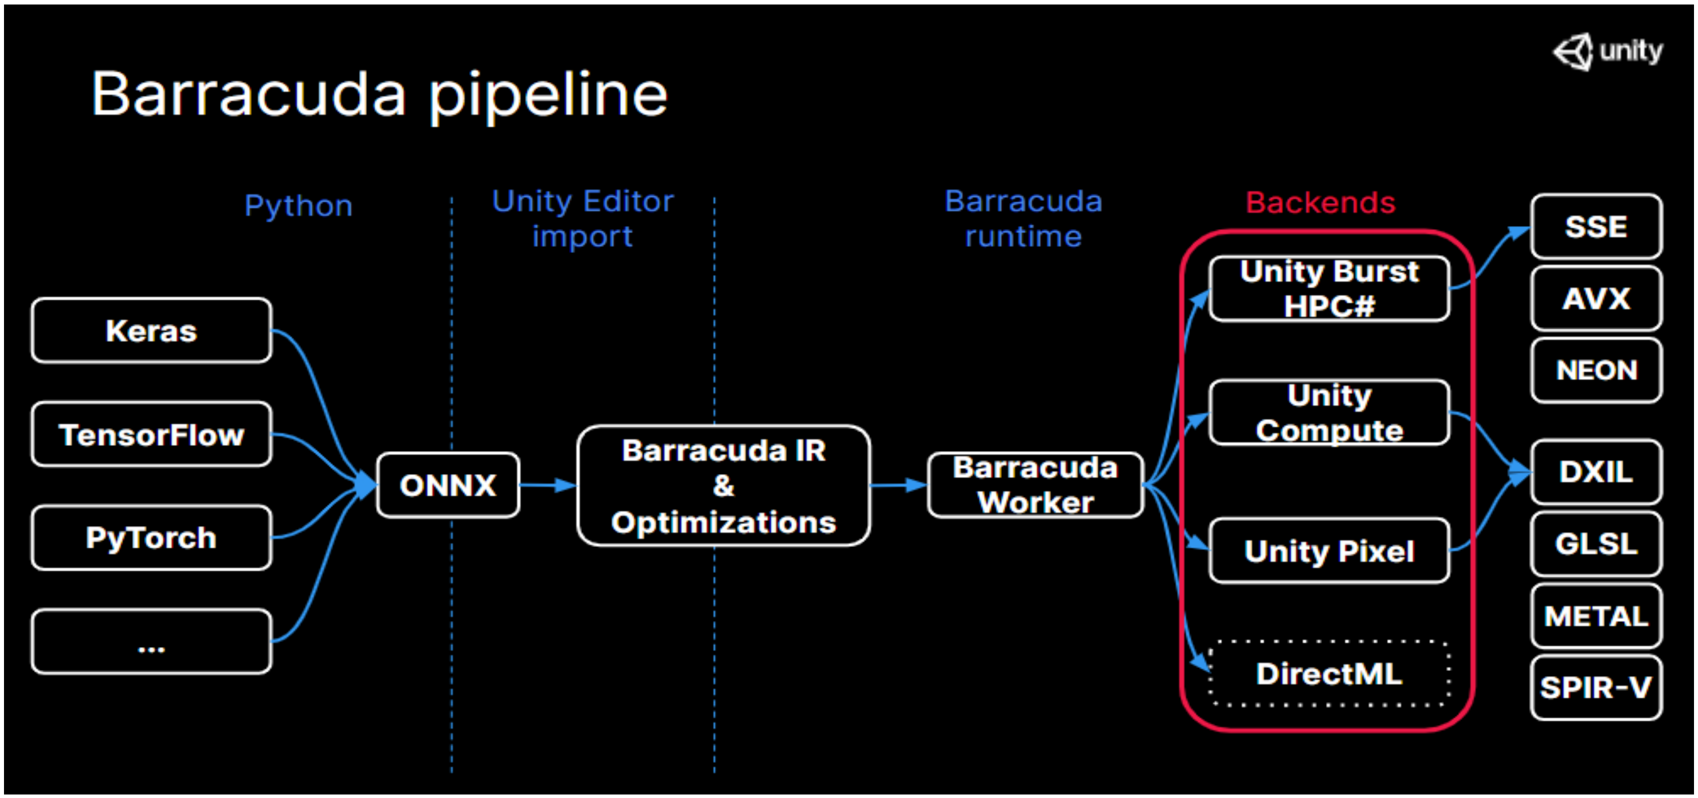

Call to Action

- Download and try out
  - Unity Barracuda
- Links for tools, renderers, etc. are listed in course notes
  - ML frameworks
  - Rendering engines
  - Tools
  - Deployment frameworks
  - Dataset links
  - Lab links

Contact Info

Carl S. Marshall, Reality Labs Research at Meta
csmarshall@fb.com

Deepak Vembar, Intel Labs
deepak.s.vembar@intel.com

Sujoy Ganguly, Unity
sujoy.ganguly@unity3d.com

Florent Guinier, Unity Labs
florent@unity3d.com

April 25, 2022

Thank you.