# THE AI DEVCON 2018



# APPLIED MACHINE LEARNING AT FACEBOOK: A DATACENTER INFRASTRUCTURE PERSPECTIVE



## Kim Hazelwood

Facebook AI Infrastructure

## **RE-EMERGENCE OF MACHINE LEARNING**

Gradient-Based Learning Applied to Document Recognition, LeCun et al., 1998



AIDC

## WHY NOW?





## MACHINE LEARNING EXECUTION FLOW











# **LET'S ANSWER SOME PRESSING QUESTIONS**

- How does Facebook leverage machine learning?
- Does Facebook design hardware? For machine learning?
- Does Facebook design machine learning platforms and frameworks?
- Are these hardware and software solutions available to the community?
- What assumptions break when scaling to 2B people?



## **HOW DOES FACEBOOK USE MACHINE LEARNING?**



## **MAJOR SERVICES AND USE CASES**





# WHAT ML MODELS DO WE LEVERAGE?



## HOW OFTEN DO WE TRAIN MODELS?





## **HOW LONG DOES TRAINING TAKE?**





## HOW MUCH COMPUTE DOES INFERENCE CONSUME?





# **LET'S ANSWER SOME PRESSING QUESTIONS**

- How does Facebook leverage machine learning?
- Does Facebook design hardware? For machine learning?
- Does Facebook design machine learning platforms and frameworks?
- Are these hardware and software solutions available to the community?
- What assumptions break when scaling to 2B people?

## FACEBOOK AI ECOSYSTEM

### Frameworks: Core ML Software Caffe2 / PyTorch / etc

### Platforms: Workflow Management, Deployment FB Learner

Large-Scale Infrastructure Servers, Storage, Network Strategy



## **THE INFRASTRUCTURE VIEW**





# **DOES FACEBOOK DESIGN HARDWARE?**

- Yes! All designs released through the Open Compute Project since 2010
- Facebook Server Design Philosophy
  - Identify a **small number** of **major services** with unique resource requirements
  - Design servers for those major services







#### For the web tier and other "stateless services"

Open Compute "Sleds" are 2U x 3 Across in an Open Compute Rack





#### For compute or memory-intensive workloads:







# **BRYCE CANYON AND LIGHTNING**

For storage-heavy workloads







## **BIG BASIN**

In 2017, we transitioned from Big Sur to Big Basin GPU Servers for ML training

**Big Sur** Integrated Compute 8 Nvidia M40 GPUs



**Big Basin** JBOG Design (CPU headnode) 8 Nvidia V100 GPUs per Big Basin





# **LET'S ANSWER SOME PRESSING QUESTIONS**

- How does Facebook leverage machine learning?
- Does Facebook design hardware? For machine learning?
- Does Facebook design machine learning platforms and frameworks?
- Are these hardware and software solutions available to the community?
- What assumptions break when scaling to 2B people?

# **FACEBOOK AI FRAMEWORKS**



- Used for Production
- Stability
- Scale & Speed
- Data Integration
- Relatively Fixed



- Used for Research
- Flexible
- Fast Iteration
- Debuggable
- Less Robust





## Vendor and numeric libraries





**INTEL AI DEVCON 201** 



# FACEBOOK AI ECOSYSTEM

## Frameworks: Core ML Software Caffe2 / PyTorch / etc

## **Platforms**: Workflow Management, Deployment FB Learner

### Large-Scale Infrastructure Servers, Storage, Network Strategy



# **FB LEARNER PLATFORM**

- AI Workflow
- Model Management and Deployment





## **TYING IT ALL TOGETHER**



## What changes when you scale to over

# **2 BILLION PEOPLE**



Santa Clara, California Ashburn, Virginia Prineville, Oregon Forest City, North Carolina Lulea, Sweden Altoona, Iowa Fort Worth, Texas Clonee, Ireland Los Lunas, New Mexico Odense, Denmark New Albany, Ohio Papillion, Nebraska 12



# **SCALING CHALLENGES / OPPORTUNITIES**





Lots of Data

Lots of Compute



## **SCALING CHALLENGES / OPPORTUNITIES: DATA**



Lots of Data



Data quality (and potentially quantity) correlates well with user experience

Network design matters Geographic locations matter Database configs matter



## **SCALING CHALLENGES / OPPORTUNITIES: COMPUTE**



Lots of Compute



Can leverage idle resources on nights and weekends for "free"



Must consider geographic resource distribution



## **SCALING OPPORTUNITY: FREE COMPUTE!**





# **SCALING CHALLENGES: DISASTER RECOVERY**

- Seamlessly handle the loss of an entire datacenter
- Geographic compute diversity becomes critical
- Delaying even the offline training portion of machine learning workloads has a measurable impact



## **KEY TAKEAWAYS**

**Facebook Al** 



Lots of Data







Full stack challenges



**Global scale** 





# THE AI DEVCON 2018