

Deep Learning and Beyond

Scalability

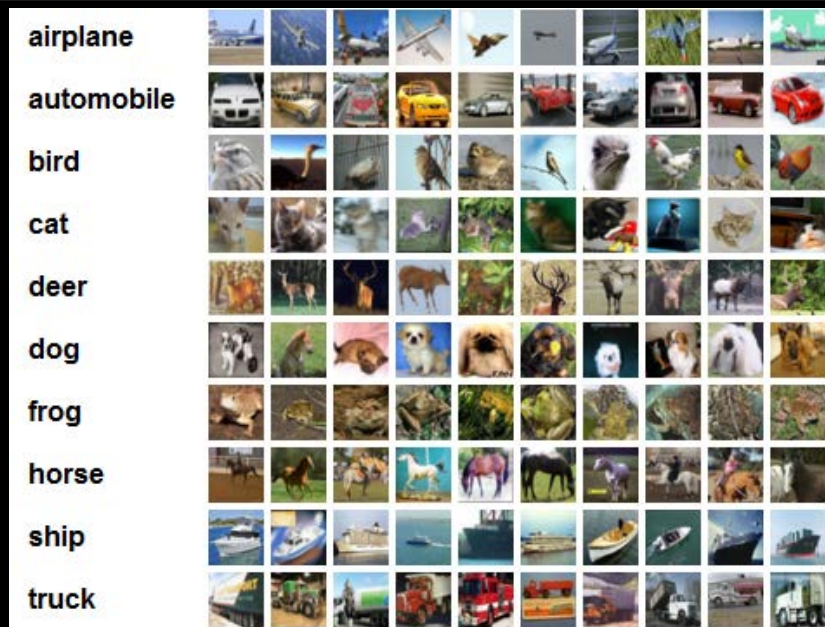
Adam Coates

adamcoates@baidu.com




Silicon Valley AI Lab

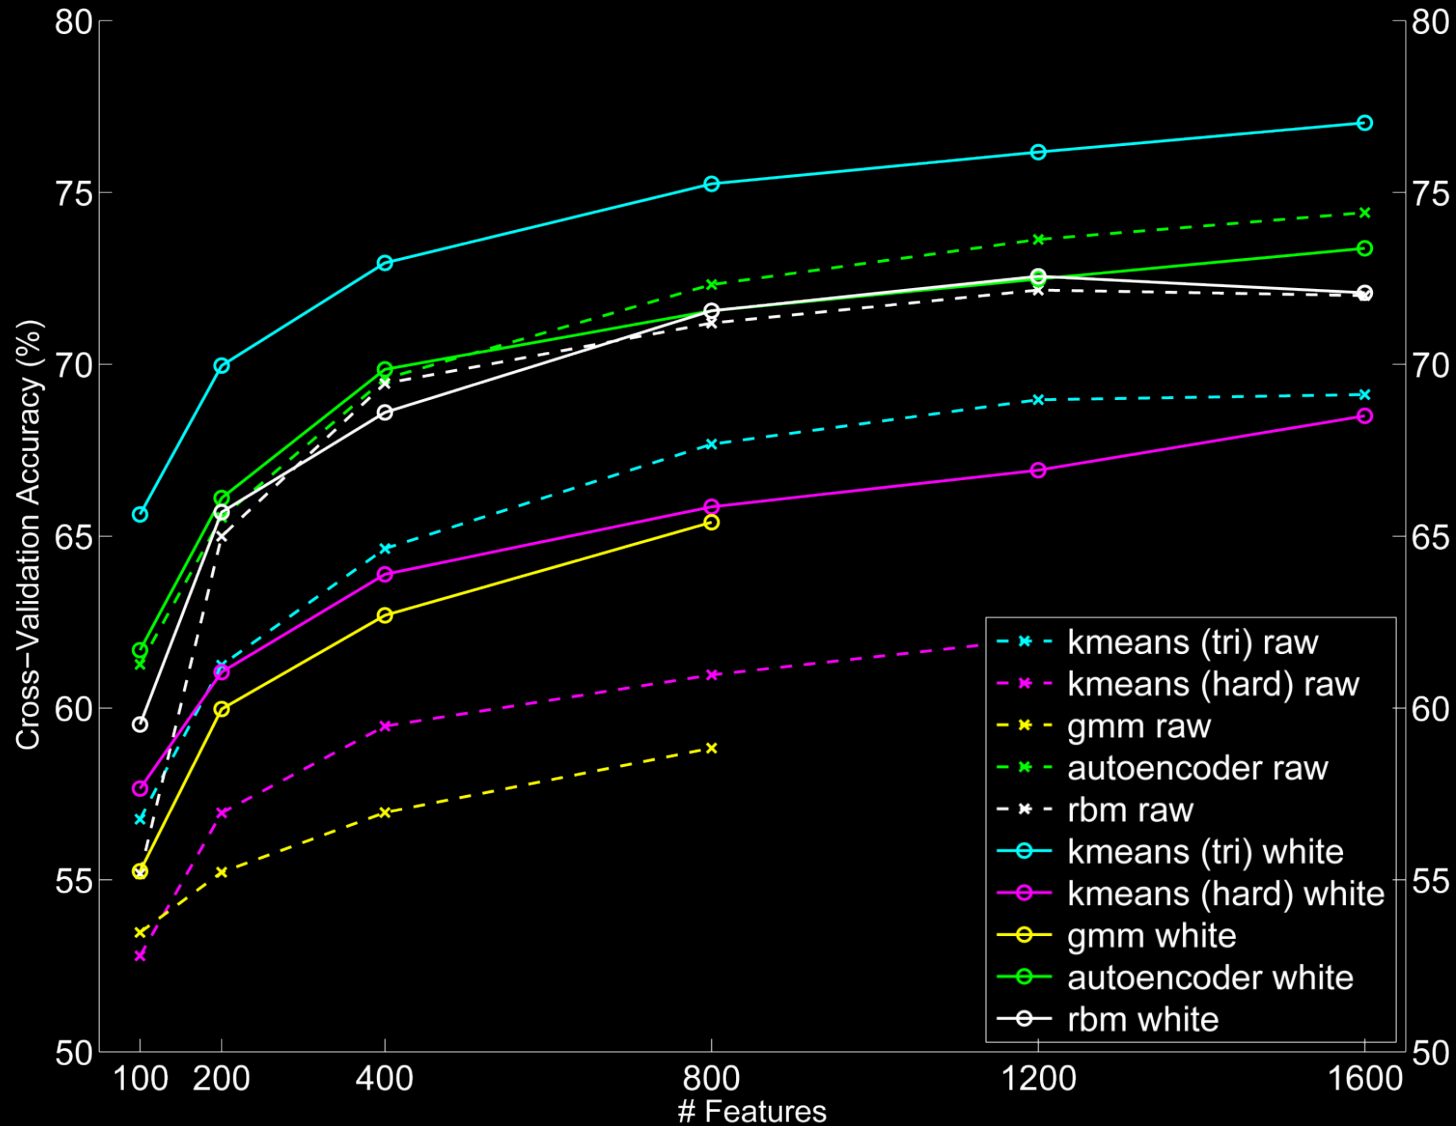
A Brief History of CIFAR-10 (2010-2012)



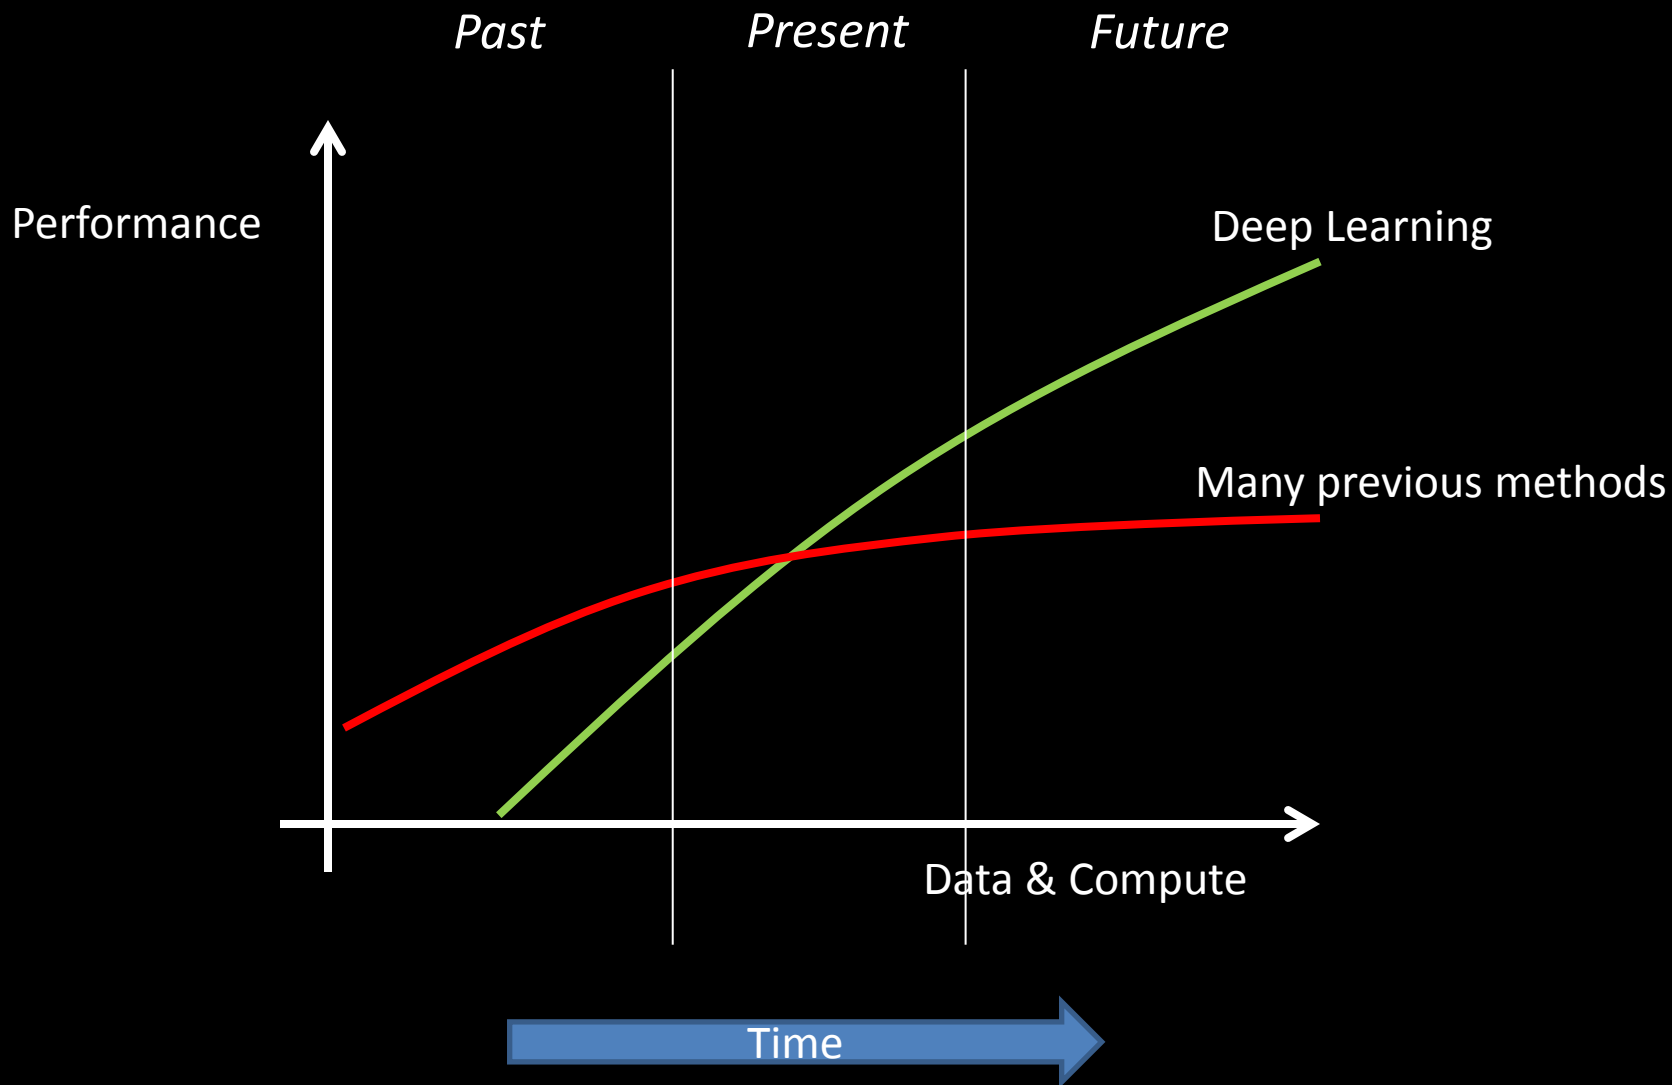
10-class image classification problem



Scale matters: big models often best

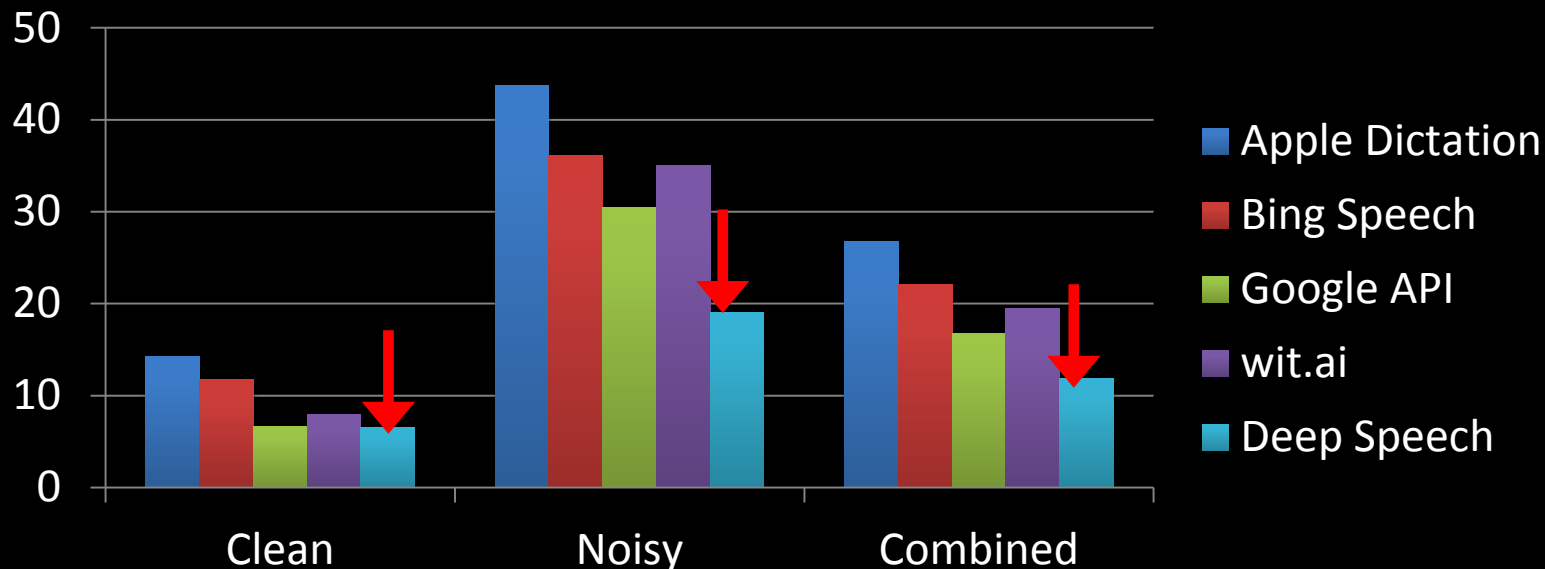


Deep learning improves with scale



Deep learning improves with scale

- New results coming from scale.
 - “Deep Speech” [Hannun et al., 2014, *arXiv preprint*]:



- Increasingly enabled by systems research & HPC:
 - Clusters of dense compute + high-speed networks.
 - [E.g., “Deep Learning with COTS HPC”, 2013]