

# Deep Learning

Andrew Ng




## Opportunities for Deep Learning



Images



Speech




Behavior

Computer vision:  
Find coffee mug




## Computer vision: Find coffee mug






## Neurons in the brain




## Deep Learning: Neural network




# Computer vision



# What is a neural network?



# Supervised learning (learning from tagged data)



Data:

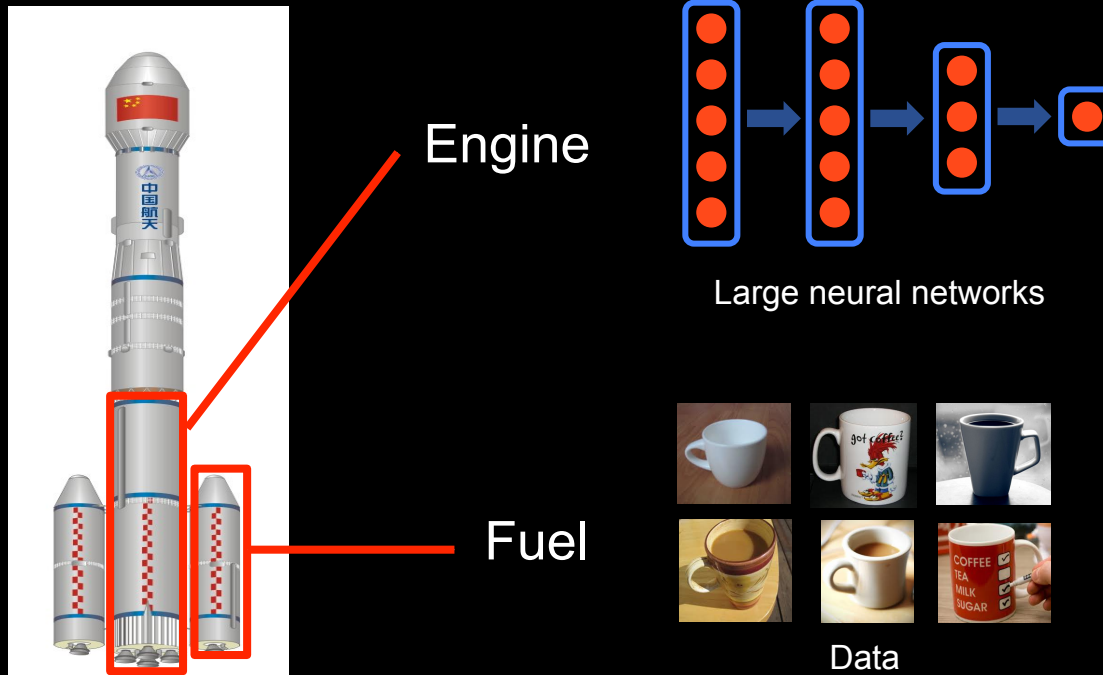


Yes



No

# Why is Deep Learning taking off?




# Why is Deep Learning taking off?

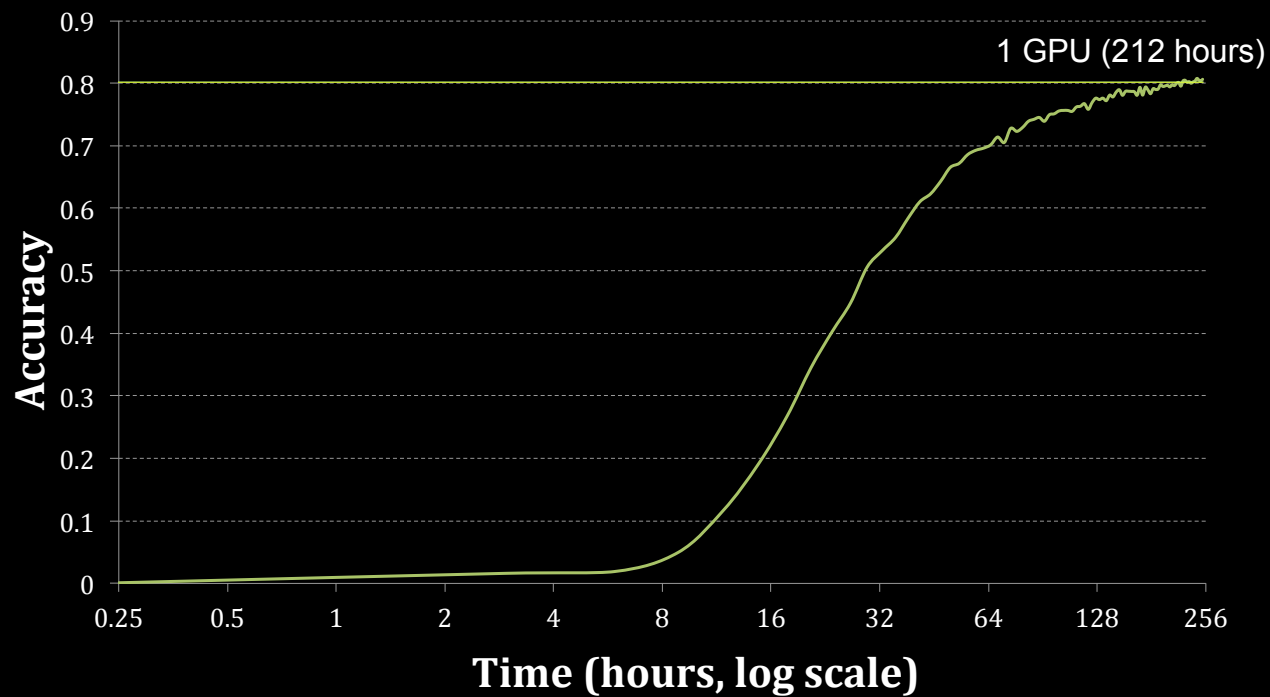





# Rocket engines: Deep Learning driven by scale



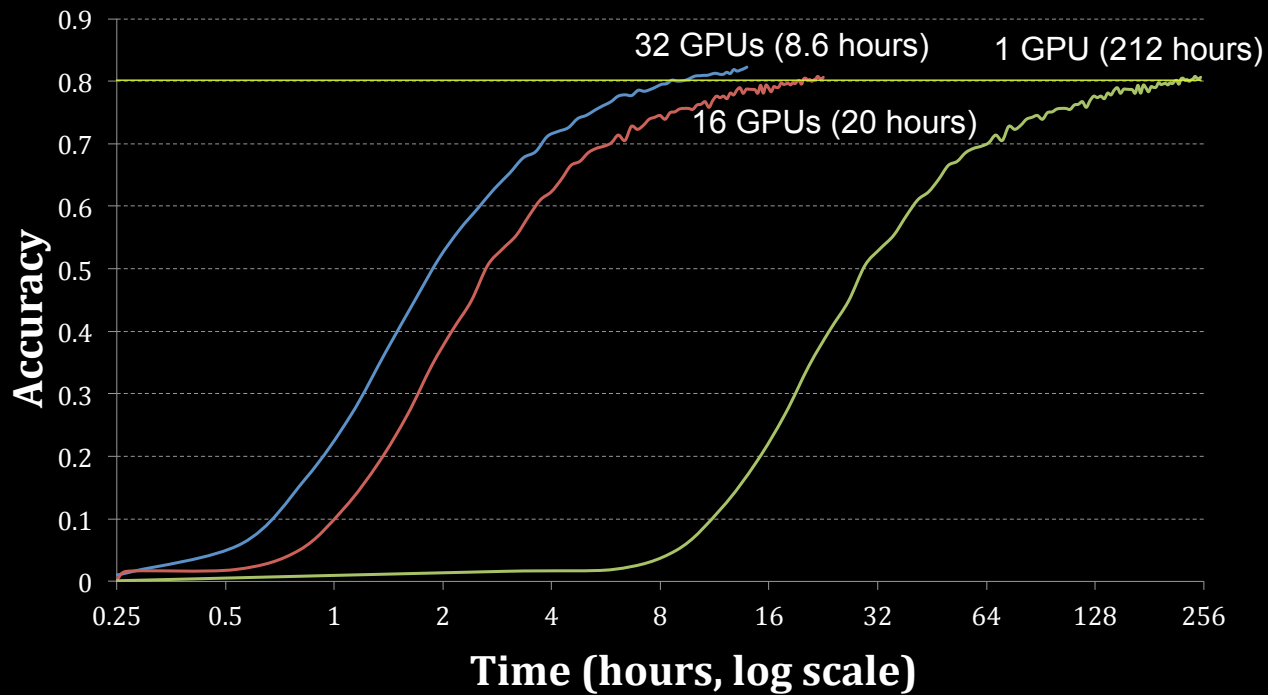
# The importance of strong scaling



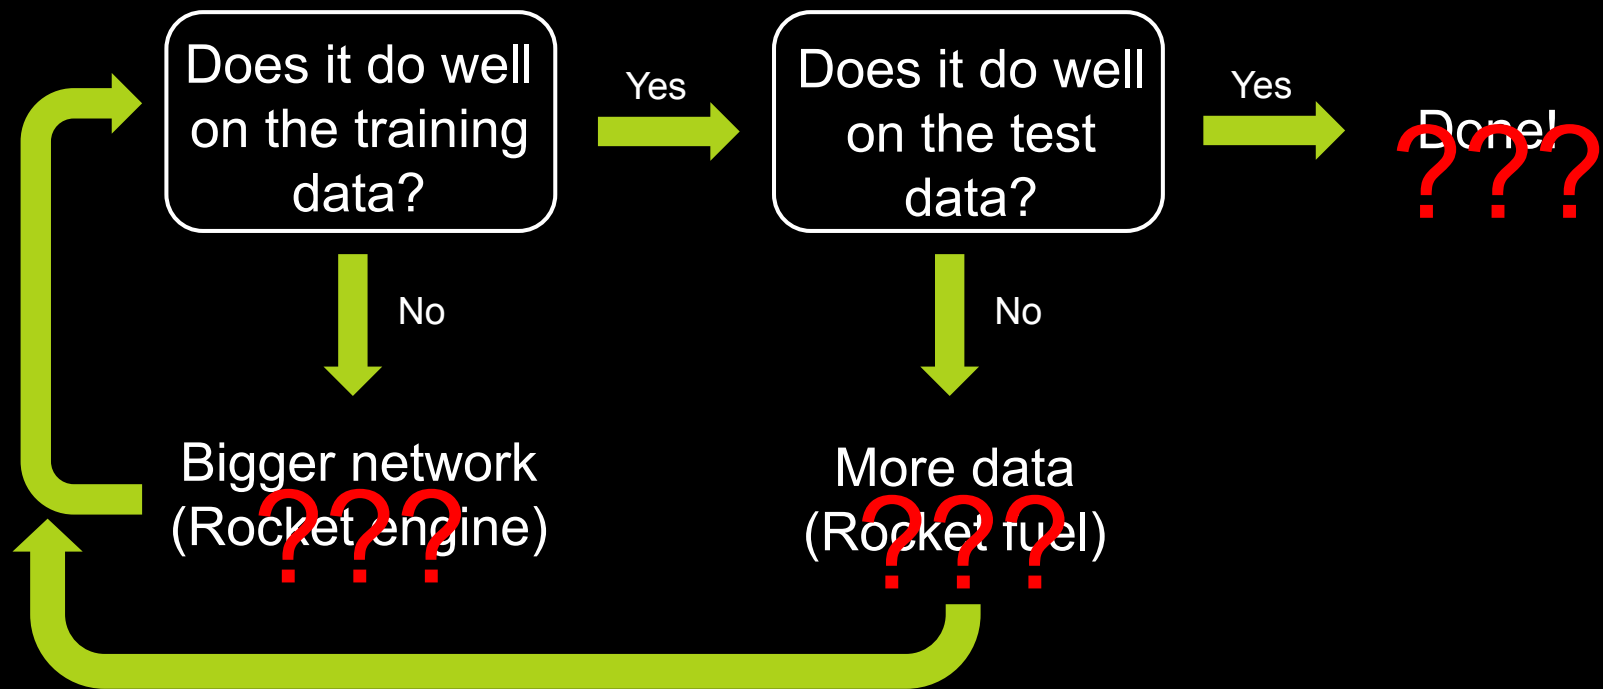
# The importance of strong scaling



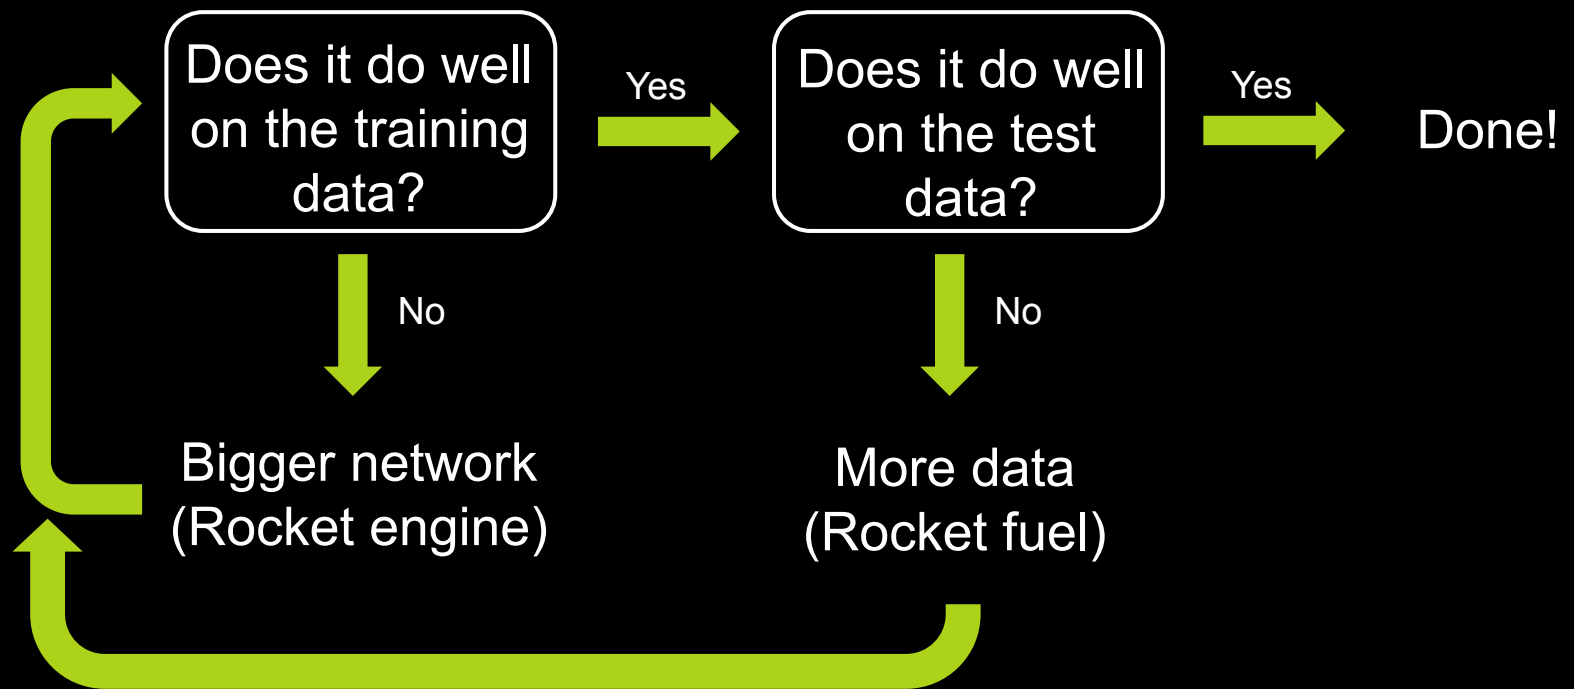
# The importance of strong scaling



# Basic recipe for machine learning



# Basic recipe for machine learning






# Learning to recognize faces

Are these the same person?



Same



Different

# Learning to recognize faces



Data:

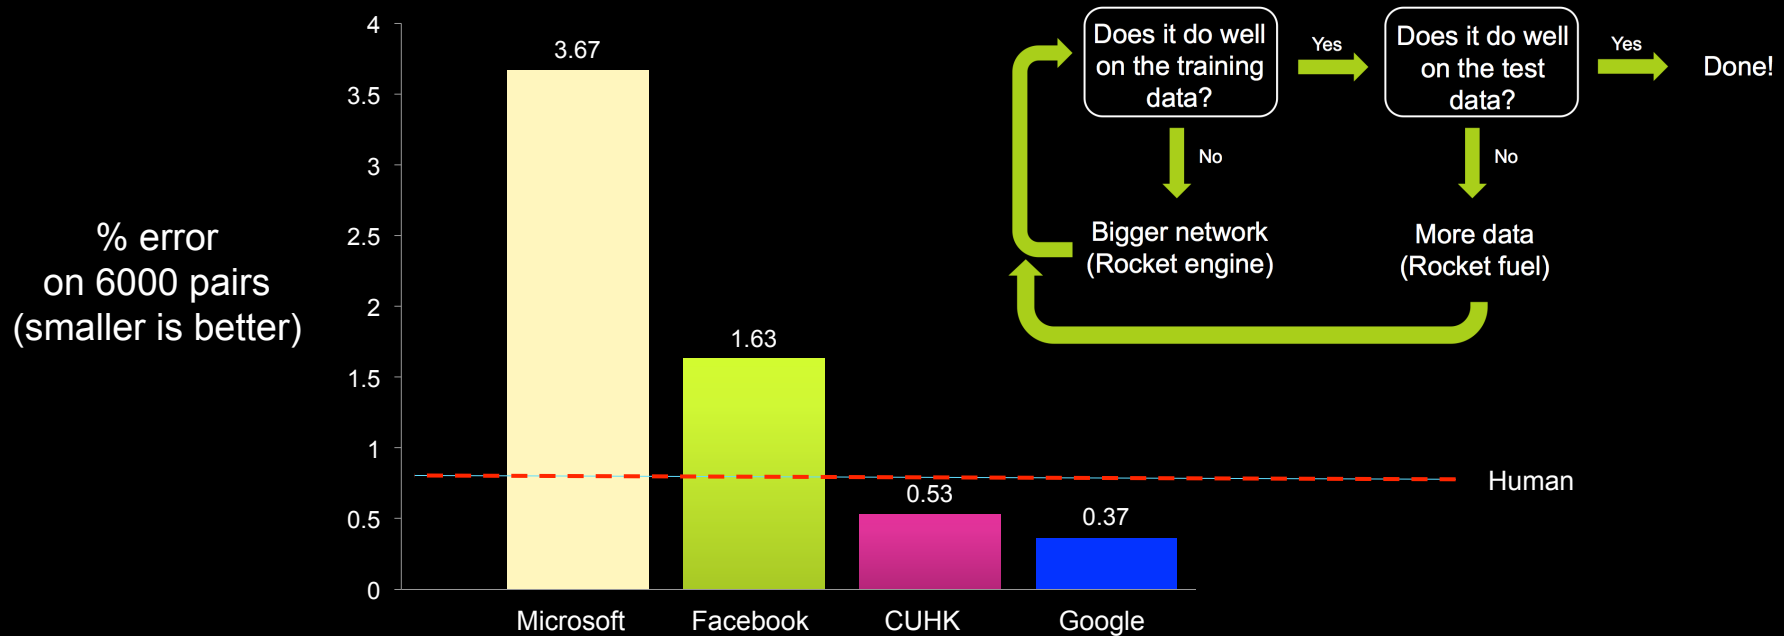


Same

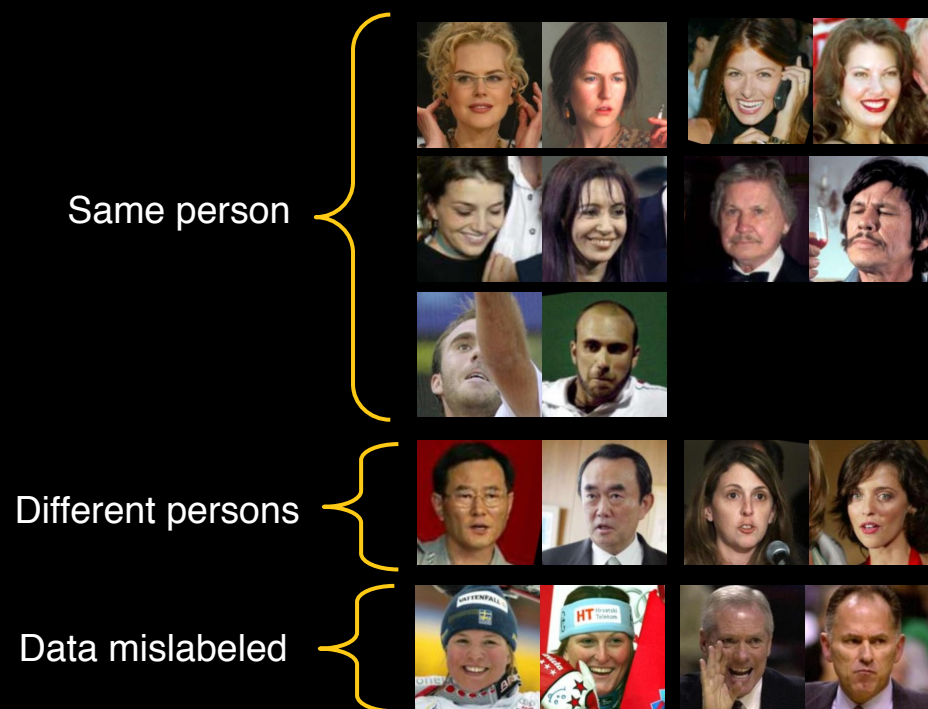


Different

# Face recognition errors (smaller is better)



# Out of 6000 test examples, 9 mistakes



## Can a computer understand these pictures?




A yellow bus driving down a road with green trees and green grass in the background.



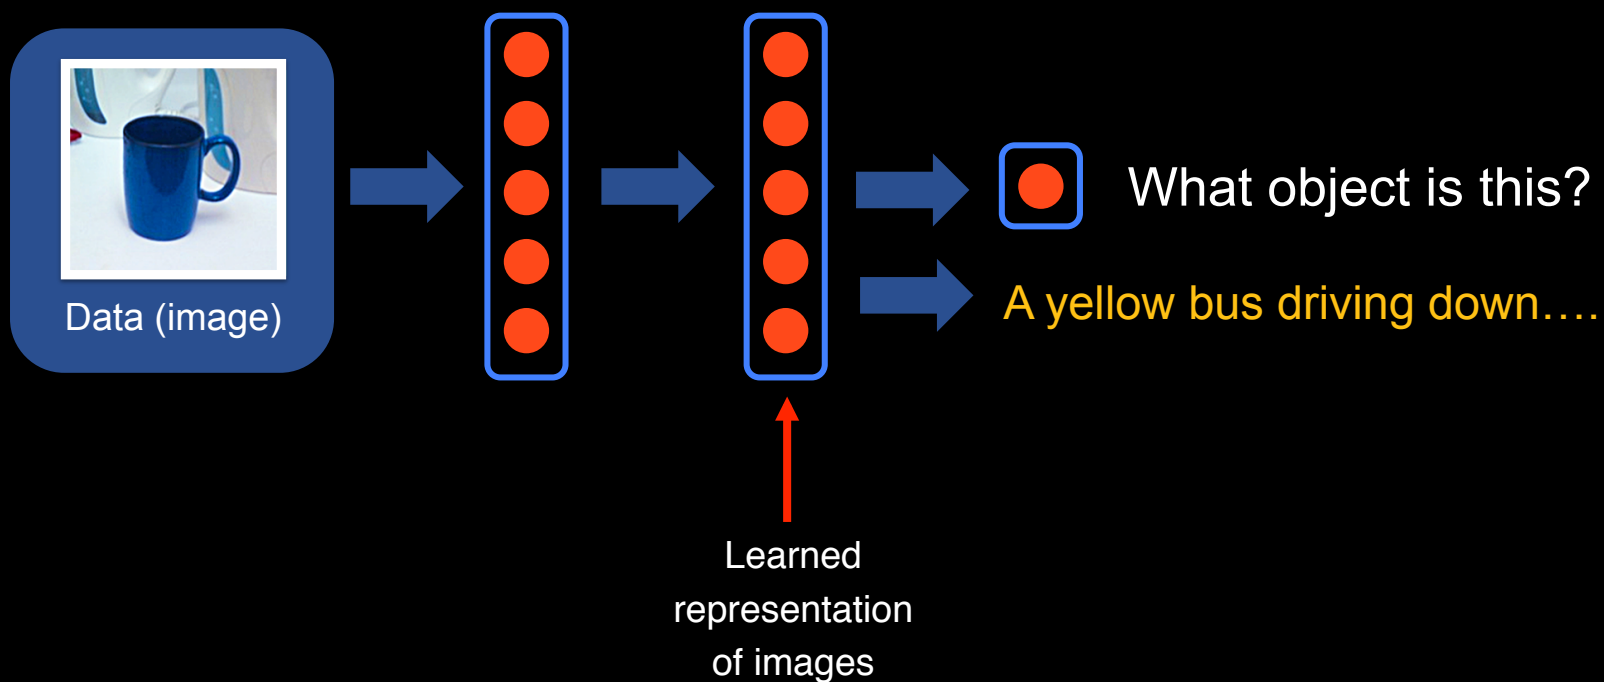
Living room with white couch and blue carpeting. The room in the apartment gets some afternoon sun.

## The limitations of supervised learning

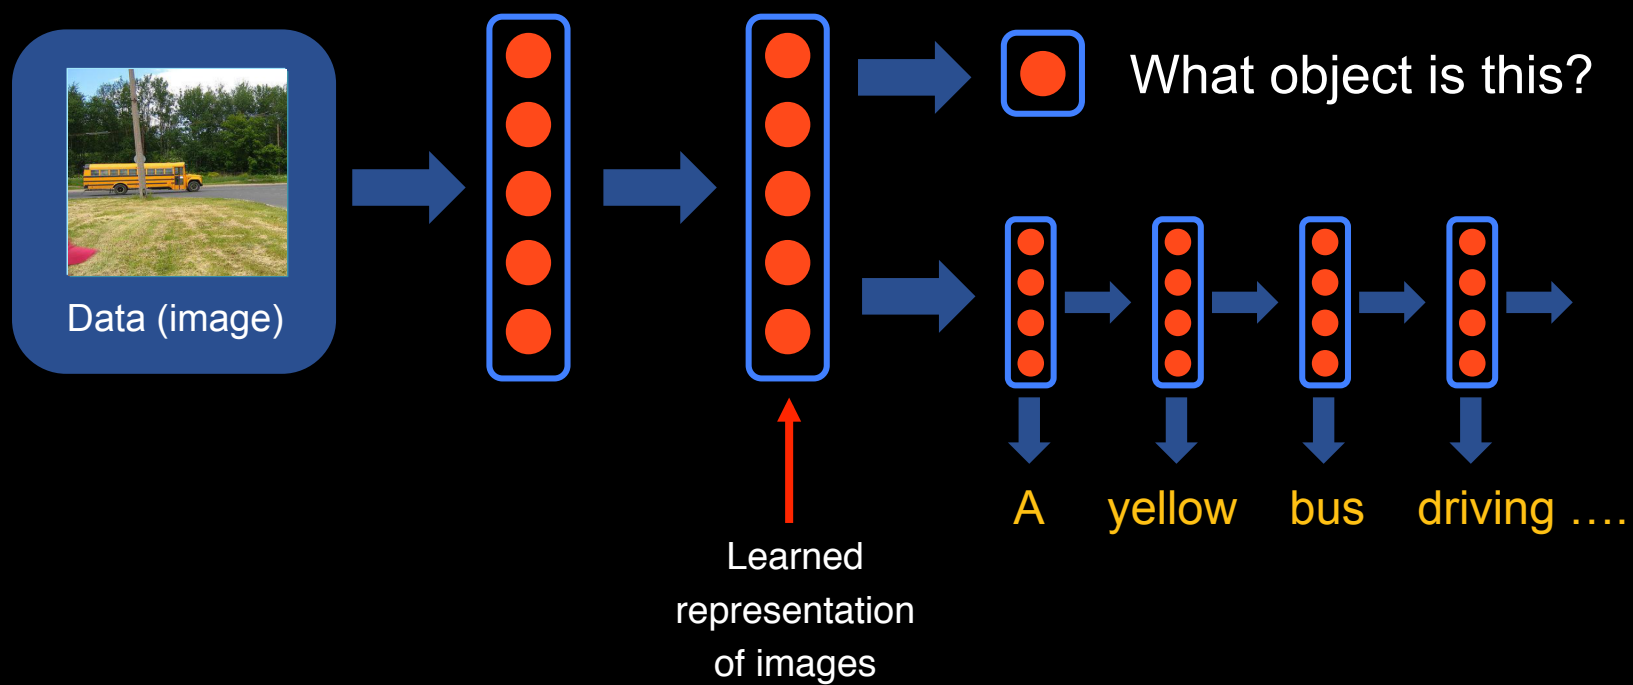




## Learning to caption (Multi-task learning)



## Learning to caption (Multi-task learning)



## Chinese captions



一辆双层巴士行驶在街道上。

(A double-decker bus driving on a street.)



一个棒球运动员  
准备击球。

(A baseball player  
getting ready to bat.)



一个人在海洋中冲浪。

(A person surfing on the ocean.)

## Image similarity search




一辆双层巴士行驶在街道上。

(A double-decker bus driving on a street.)



## Promising computer vision applications



Medical imaging



Retail (clothing search)



Self-driving cars



Home security



Elder care



Wearables: Baidu Eye



# Baidu Eye





## Opportunities for Deep Learning



Images



Speech



Behavior


# Li Chongyang



# Li Chongyang

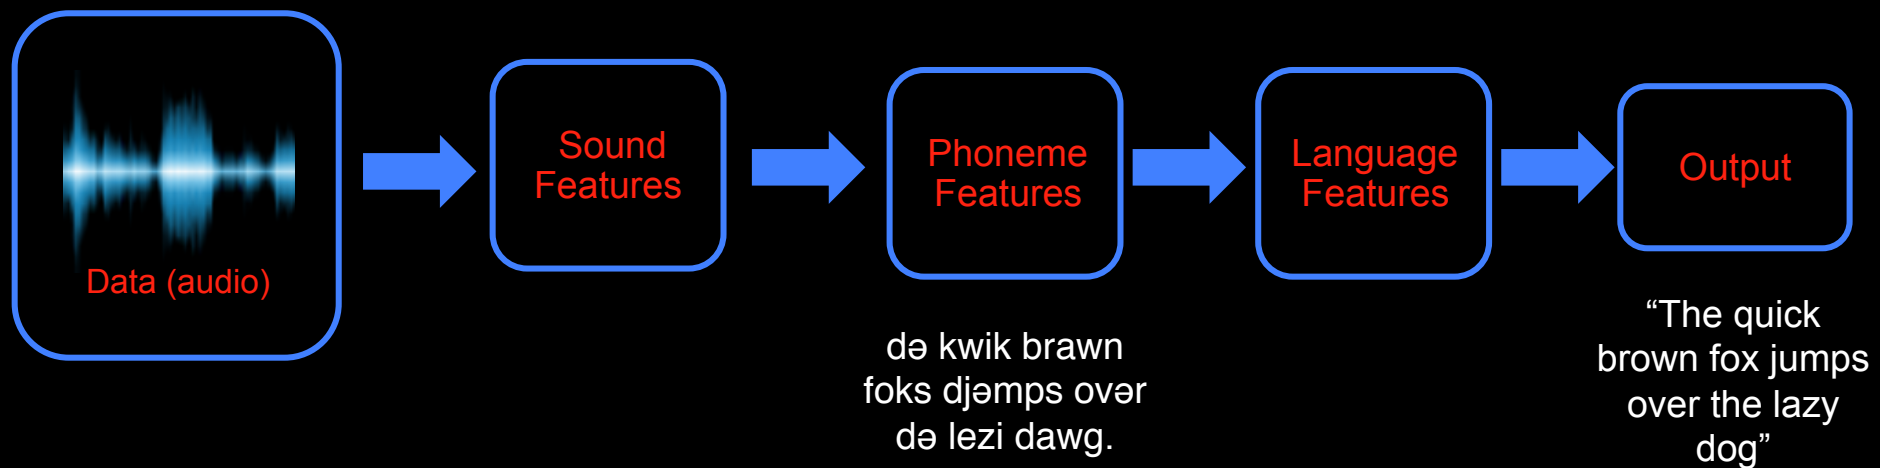


# Speech recognition

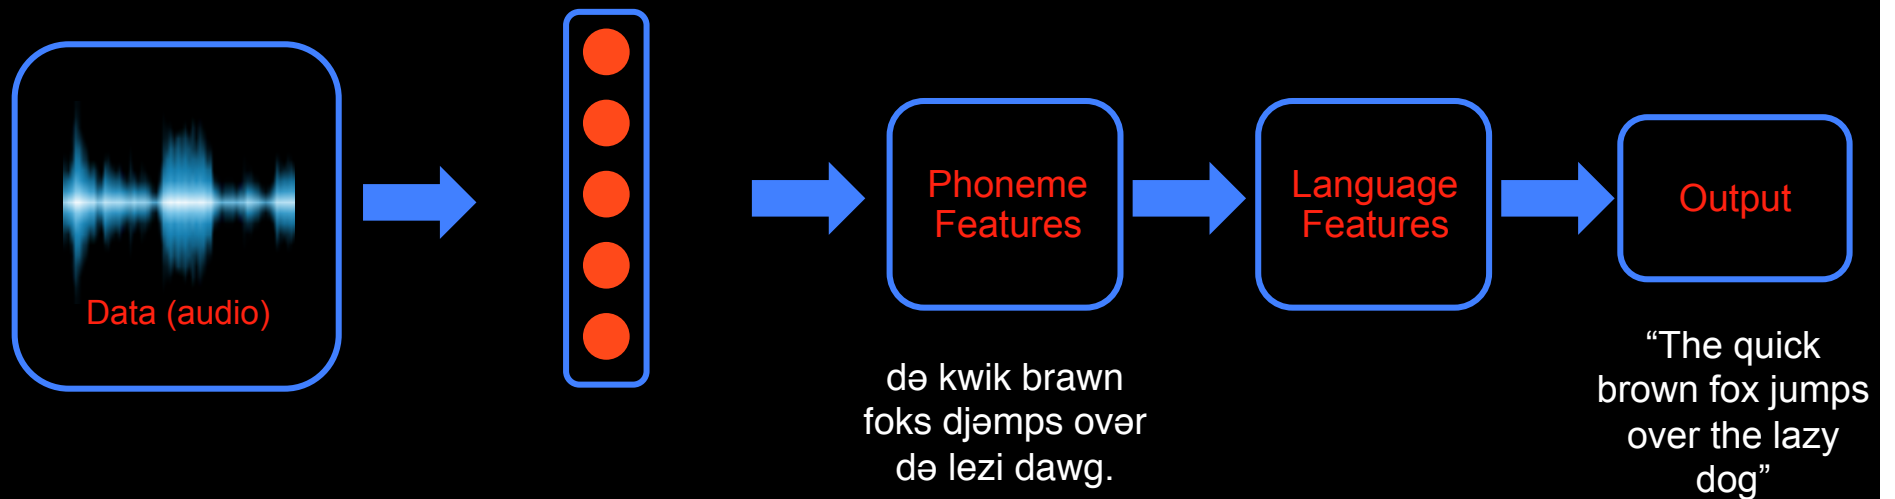


688,058


# Speech recognition



# Speech recognition

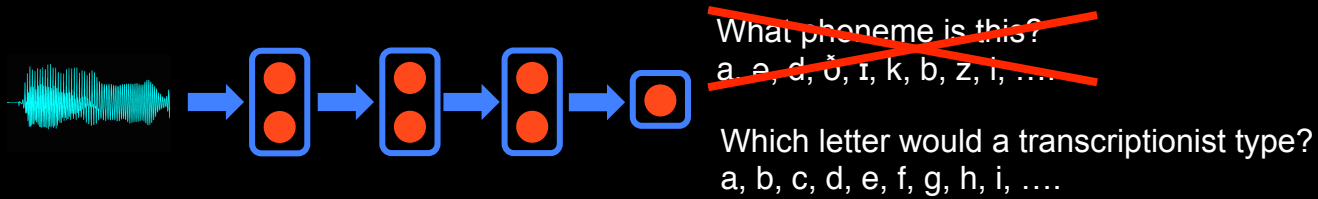


## Baidu Deep Speech: The rocket engine



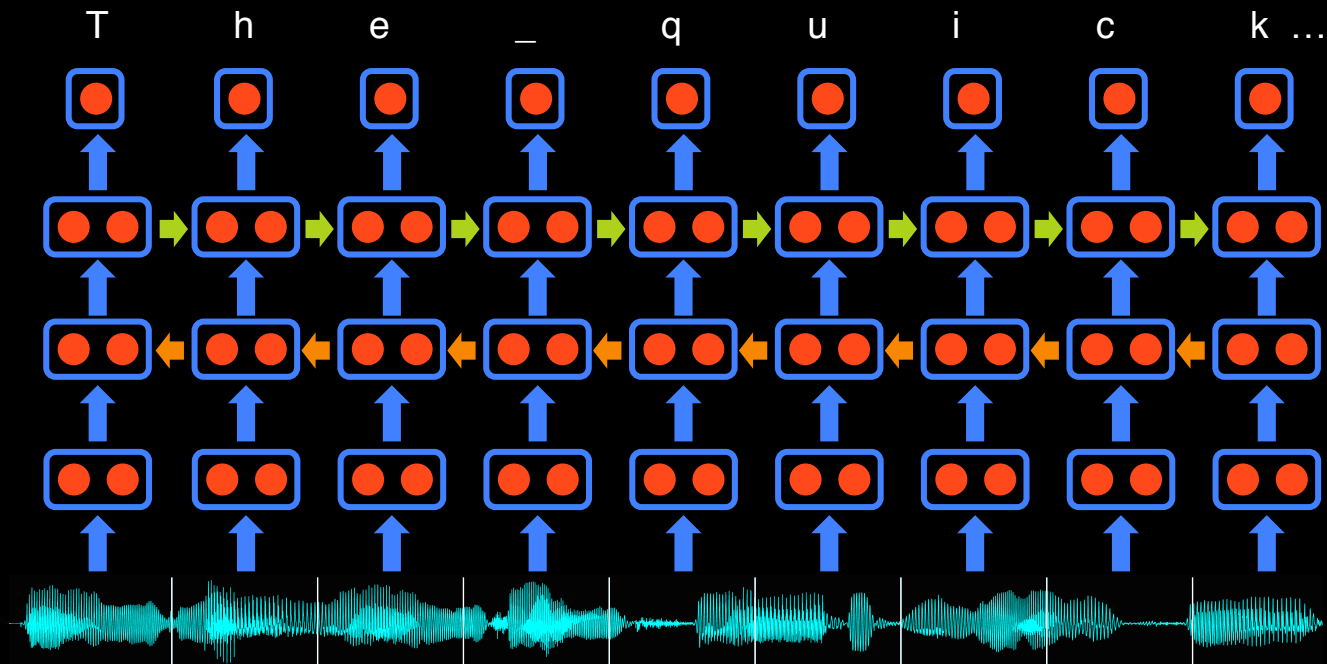


## Most speech systems today

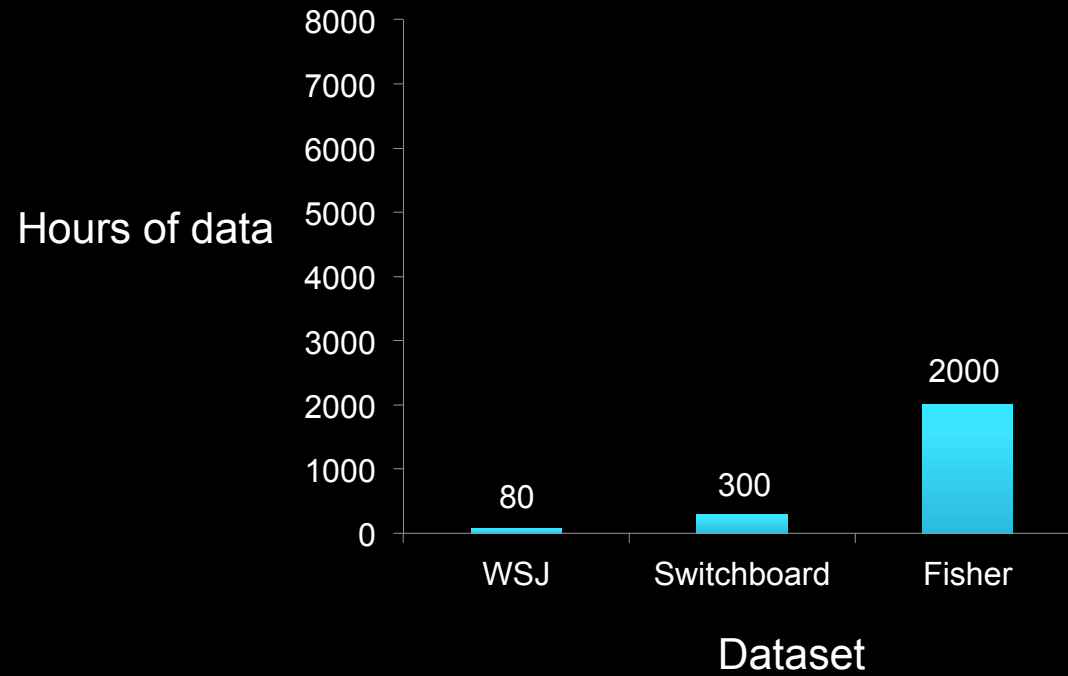


# Baidu Deep Speech


Bi-directional Recurrent  
Neural Network (BDRNN)



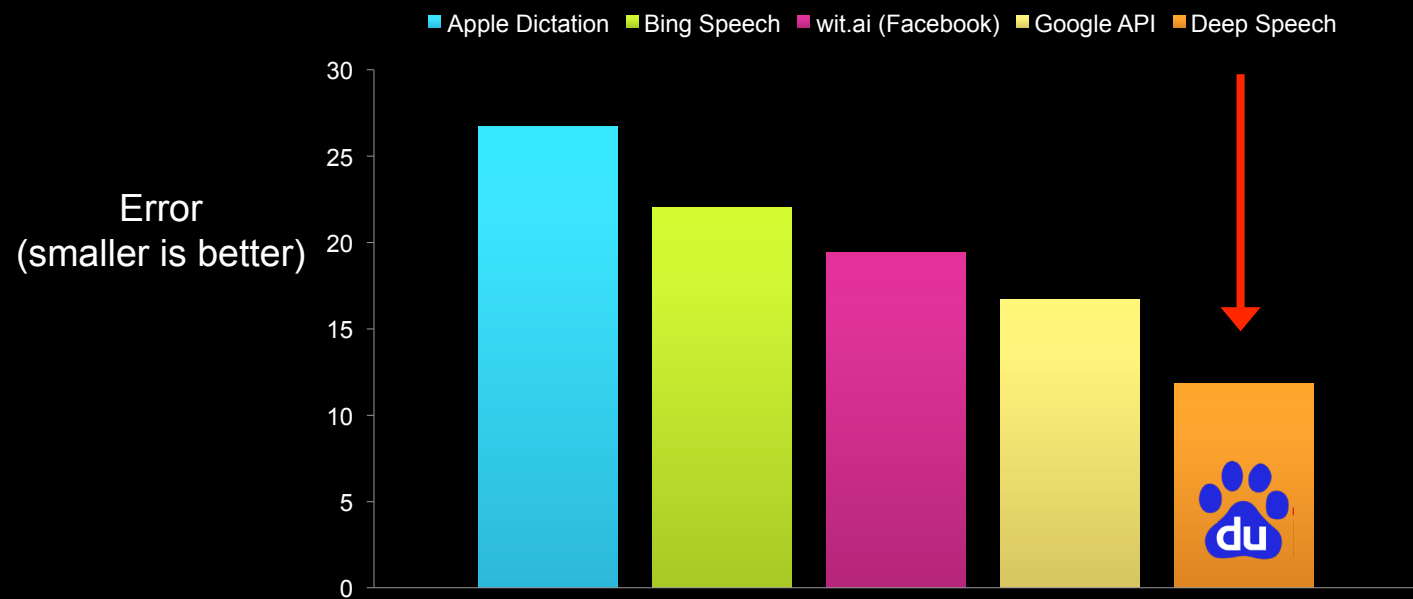
## Baidu Deep Speech: The rocket fuel (data)



## Baidu Deep Speech: The rocket fuel (data)



## Speech recognition errors



# Live Demo

# The Deep Speech team



Most people don't understand the difference  
between 95% accuracy and 99% accuracy.

99% is game changing.



With 99% accuracy, we could  
redesign your cellphone using a  
speech interface.



# Speech will transform the Internet of Things



Car interfaces



Home appliances



Wearables

(e.g., TV, microwave,  
music player, ....)

## Opportunities for Deep Learning



Images

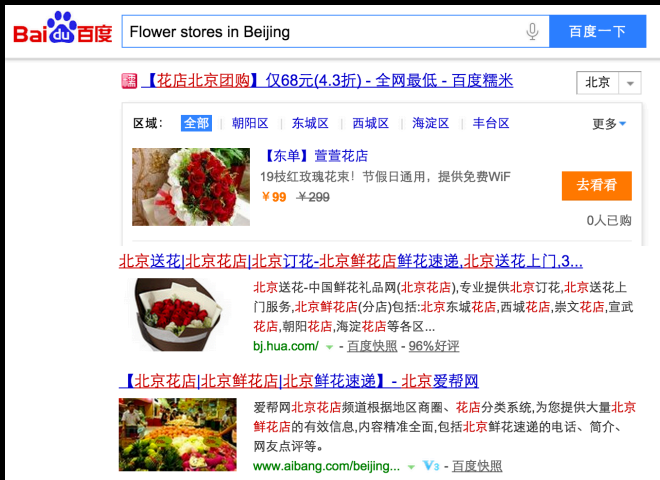


Speech



Behavior

# Deep Learning and big data



Web search/Advertising



Datacenter management



Computer security

## Opportunities for Deep Learning



Images

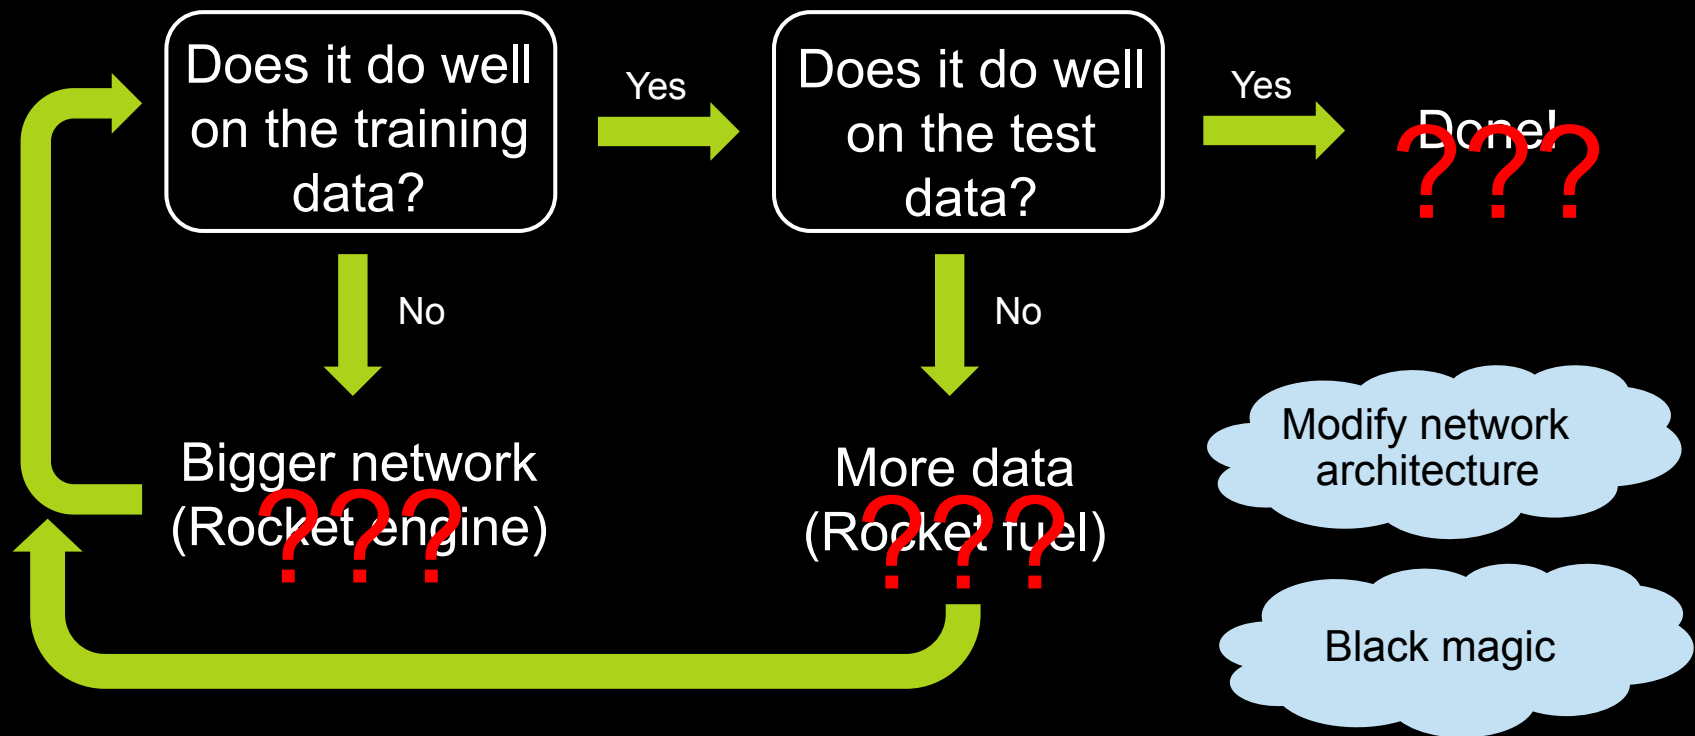


Speech




Behavior


## Developing new applications




# Network architectures



Basic neural network




Captioning Network




Deep Speech

# The process of machine learning research



## The future of Deep Learning



The AI rocket



Images




Speech



Behavior



What about the hype?



The AI rocket



“The Transformers” movie

We have superpowers



Thank you.