

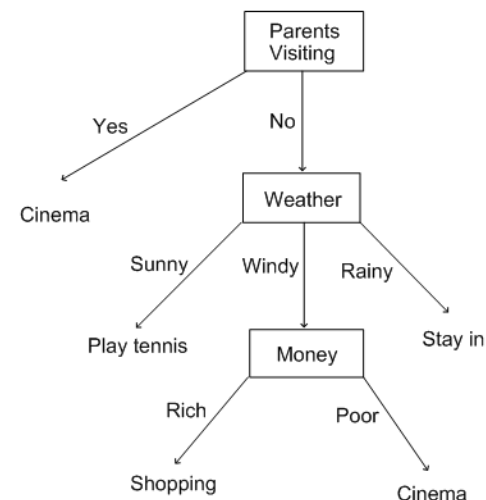
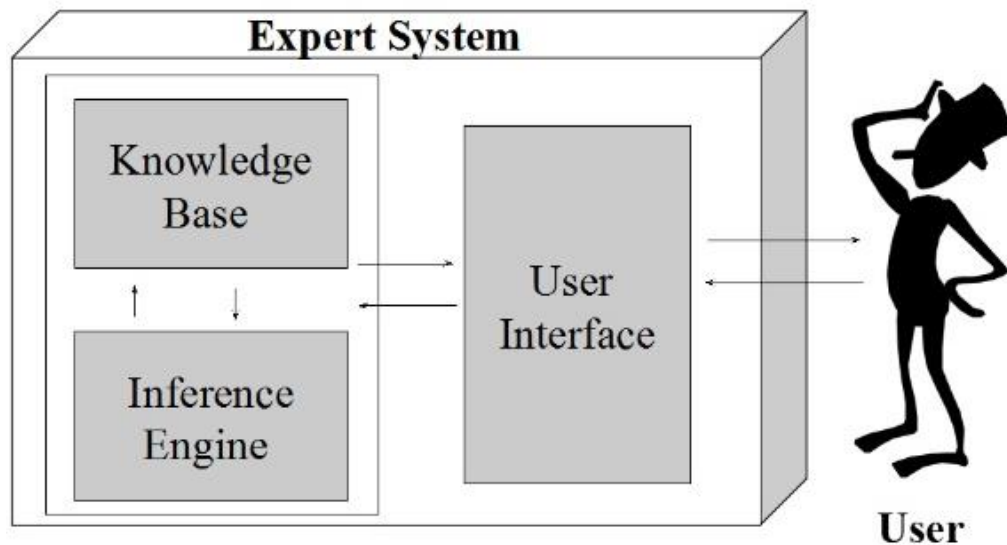
DEEP LEARNING INTRODUCTION

Bryan Catanzaro, 1 March 2017




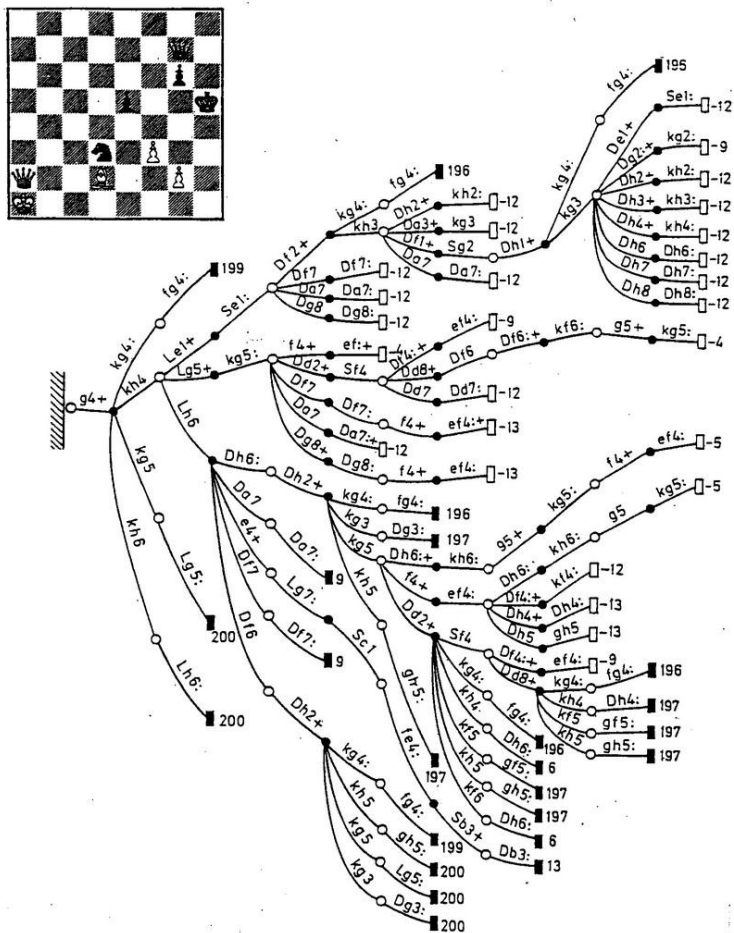
WHAT IS AI TO YOU?

Rules, scripts




```
> (doctor)
  tell me your troubles. please terminate input with an enter.
<-my computer dreams of you
  do computers worry you.
<-they are scary
  did you think they might not be scary.
<-i don't
  don't you really.
<-no, i don't
  why don't you.
<-because they might talk to me
  is that the real reason.
<-well, they might be smarter than i am
  is it because you are that you came to me.
<-yes
  you're being somewhat short with me.
<-i came to you because i am afraid that my computer is smarter than i am
  why do you mention computers.
<-because i am afraid
  don't any other reasons come to mind.
<-there are too many to count
```

Solvers



WHAT IS AI TO YOU?

Statistical methods, Machine Learning, Deep Learning



WHAT IS AI TO YOU?


All of these are AI

So why are we focused on Deep Learning?





DEEP LEARNING

Huge progress in many fields

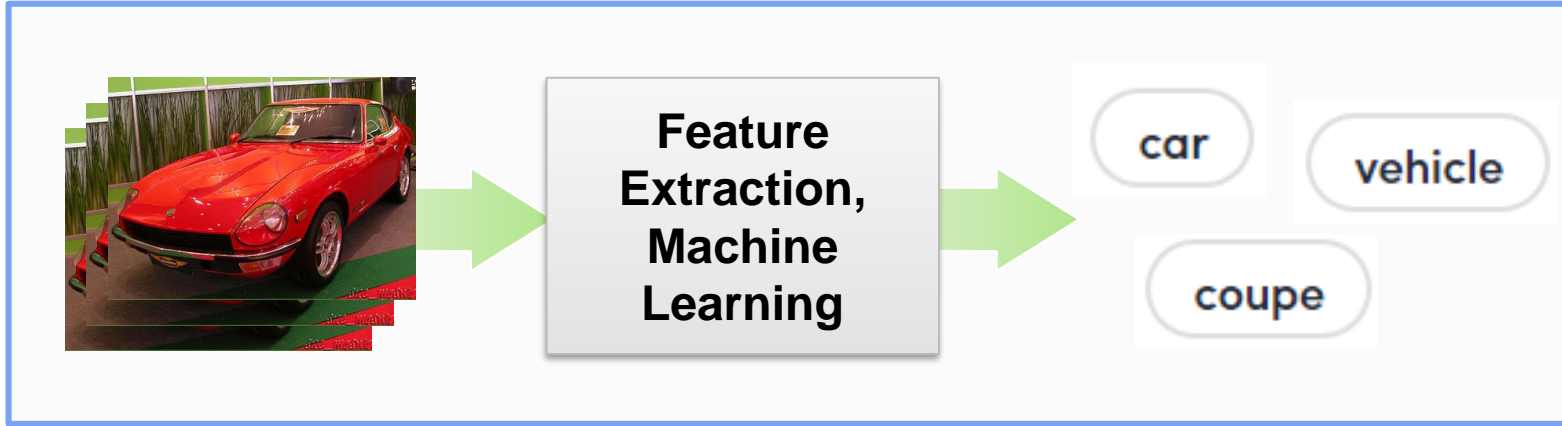


communication
沟通




WHY DEEP LEARNING

Algorithms that Learn from Examples



Traditional Approach

- Requires domain experts
- Time consuming
- Error prone
- Not scalable to new problems



Deep Learning Approach

- ✓ Learn from data
- ✓ Easy to extend
- ✓ Efficient & scalable

WHY DEEP LEARNING

Scale Matters


Millions to Billions of parameters

Data Matters

Learn with more data

Productivity Matters

SW + HW tools speed experiments



DEEP NEURAL NET


Function approximator

$$y_j = f \left(\sum_i w_{ij} x_i \right)$$

One layer

$$f(x) = \begin{cases} 0, & x < 0 \\ x, & x \geq 0 \end{cases}$$

nonlinearity



Deep Neural Net

Stacked layers learn progressively more useful features

Can be practically trained on huge datasets

SUPERVISED LEARNING

Learning mappings from labeled data



→ YES



→ NO

Learning $X \rightarrow Y$ mappings is hugely useful

SUPERVISED LEARNING

Learning mappings from labeled data

Image classification

Speech recognition

Speech synthesis

Recommendation systems

Natural language understanding

(Game state, action) → reward



Most surprisingly: these mappings can generalize

EXAMPLES

And explanations

➔ Content Creation

Also, See Andrew Edelsten's talk

User Interfaces

Game AI



CLASSIFICATION

[He et al.] arXiv:1512.03385

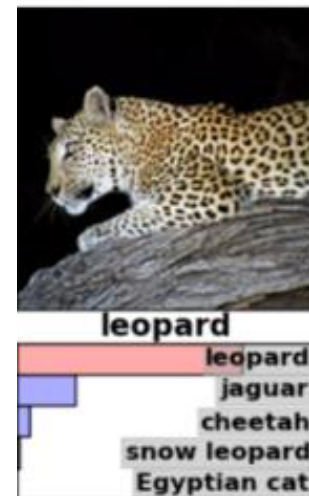
Where modern deep learning got its start:
Imagenet

Image classification useful for a bunch of tasks


Pretrained models widely available:

<https://github.com/KaimingHe/deep-residual-networks>

Transfer learning, perceptual losses super useful



34-layer residual



CONVOLUTIONAL NEURAL NETWORK

Convolution gives location invariance

Weight sharing a powerful technique


Terms you might hear:

Striding (skip outputs periodically)

Feature map (output of neural network layer)

Pooling (reduce size of feature map)

Dense layers (Fully connected)




COLORIZATION

[Zhang et al.] arXiv:1603.0851

Convolutional neural network to predict color from black and white images

Lots of cool old films and photos out there




Ansel Adams
photographs

Automatically
colorized

COLORIZATION



SUPERRESOLUTION




4x upsampling

Generative Adversarial Network for superresolution

These could have lots of interesting applications to games

Marco Foco, Dmitry Korobchenko will talk about this next!

GENERATIVE ADVERSARIAL NETWORK



Ming-Yu Liu

Exciting technique for unsupervised learning

Discriminator teaches generator how to create convincing output

FLUID SIMULATION

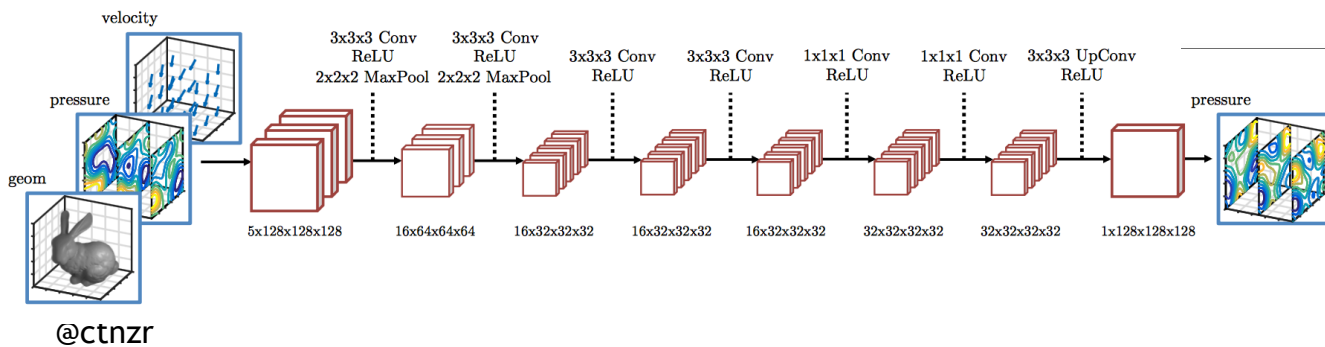
[[Tompson et al](#)] arXiv:1607.03597

Approximate solution to Euler equations using CNN

Use semi-supervised training with traditional solver to create training data

Accelerating Eulerian Fluid Simulation With Convolutional Networks

Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, Ken Perlin



EXAMPLES

And explanations

Content Creation

→ User Interfaces

Game AI



SPEECH RECOGNITION

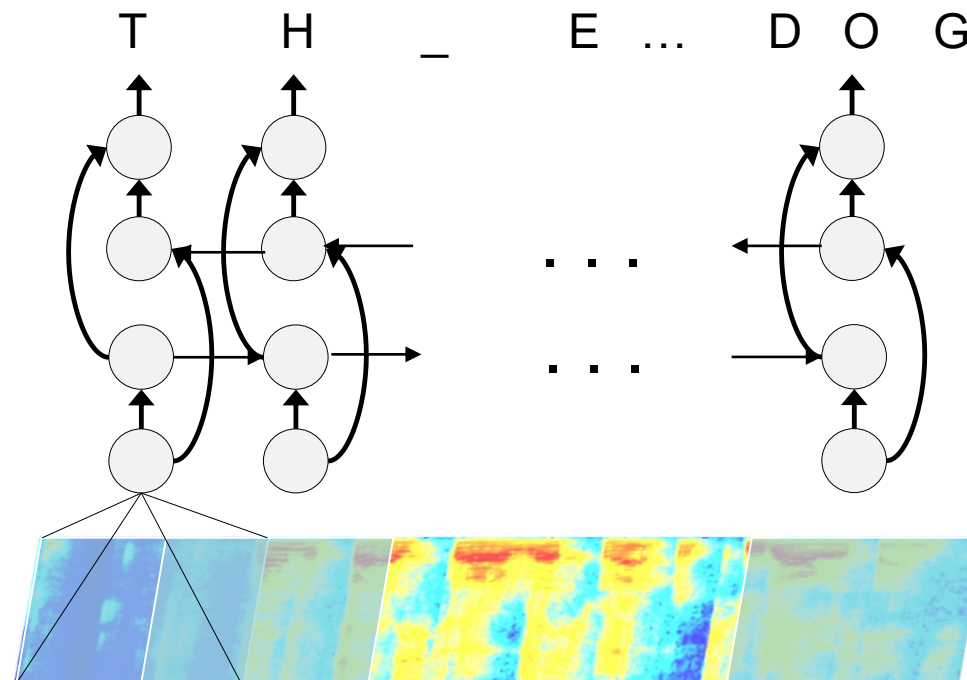
[\[Amodei et al.\] arXiv:1512.02595](#)

Beats human accuracy for some
speech recognition tasks

Trained on 12000 hours of data (1.4 Y)

Recurrent Neural Network

Long-Short-Term-Memory (LSTM)



NEURAL MACHINE TRANSLATION

[\[Wu et al.\]](#) arXiv:1609.08144

Significant improvement in machine translation

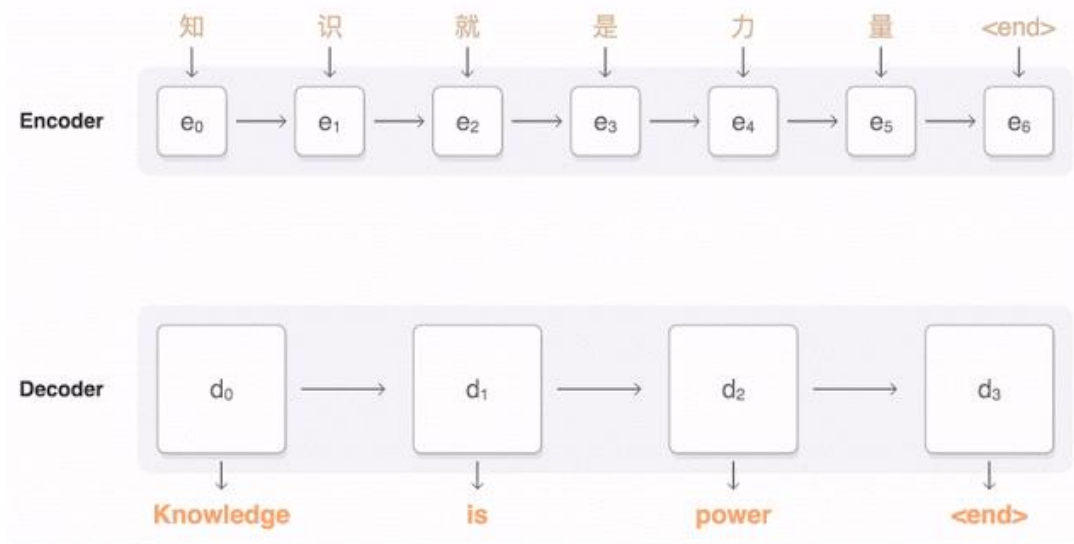
Google has deployed NMT for English to & from {French, German, Spanish, Portuguese, Chinese, Japanese, Korean, Turkish}

<i>Input sentence:</i>	<i>Translation (PBMT):</i>	<i>Translation (GNMT):</i>	<i>Translation (human):</i>
李克強此行將啟動中加總理年度對話機制，與加拿大總理杜魯多舉行兩國總理首次年度對話。	Li Keqiang premier added this line to start the annual dialogue mechanism with the Canadian Prime Minister Trudeau two prime ministers held its first annual session.	Li Keqiang will start the annual dialogue mechanism with Prime Minister Trudeau of Canada and hold the first annual dialogue between the two premiers.	Li Keqiang will initiate the annual dialogue mechanism between premiers of China and Canada during this visit, and hold the first annual dialogue with Premier Trudeau of Canada.

NEURAL MACHINE TRANSLATION

[\[Wu et al.\] arXiv:1609.08144](#)

Attentional sequence to sequence model (LSTM)




SPEECH SYNTHESIS: WAVENET

[van den Oord et al.] arXiv: 1609.03499

Audio generation using convolutional neural networks

Predict each sample directly

Cut scenes? NPCs that really talk?



Concatenative TTS



Wavenet

GESTURE RECOGNITION

[Molchanov et al., CVPR 2016]

Recurrent 3D CNN

RGB camera, depth camera,
stereo IR

What new games can we make with
better controls?

Online detection and classification
Real-time demo

EXAMPLES

And explanations

Content Creation

User Interfaces

→ Game AI



REINFORCEMENT LEARNING

Problem: Given

Current state

Possible actions

(Potentially delayed) Rewards

Learn policy for agent to maximize reward




Mnih et al. 2015

REINFORCEMENT LEARNING FOR DOOM

[Lample, Chaplot] arXiv:1609.05521

Deep Recurrent Q Network outperforms humans at single-player and deathmatch

Evaluation Metric	Single Player		Multiplayer	
	Human	Agent	Human	Agent
Number of objects	5.2	9.2	6.1	10.5
Number of kills	12.6	27.6	5.5	8.0
Number of deaths	8.3	5.0	11.2	6.0
Number of suicides	3.6	2.0	3.2	0.5
K/D Ratio	1.52	5.12	0.49	1.33



SUPER SMASH BROTHERS MELEE

[\[Firoiu, Whitney\]](#) arXiv:1702.06230

Reinforcement learning does better
than expert human players

Slox in this video is ranked #51

They beat 10 ranked players

Trained for Captain Falcon

Transfer learning to a few others



SUPER SMASH BROTHERS MELEE

How did they do it?

Trained on game state in an emulator (No pixel input)

No flowcharts/scripts

Although they think results might be improved with scripts

Ran ~50 emulators to generate
{state, action, reward} tuples during training




ENVIRONMENTS FOR RL

► OpenAI Universe



► DeepMind Lab



CONCLUSION

Deep learning is making new things possible

Lots of applications for games

- Content creation

- User interfaces

- Game AI

Can't wait to see what you all come up with!

Questions:
@ctnzs

