Applied Neuroscience

Columbia Science Honors Program Fall 2016

> Artificial Intelligence in Robotics



Soft Robotics: Baymax



Baymax of *Big Hero 6* is based on the robotics research of Professor Chris Atkeson at Carnegie Mellon University

Baymax is capable of maintaining a good **spoken dialogue system** and **learning** new knowledge for better **understanding** and **interacting** with people.

Spoken Dialogue Systems: intelligent agents that are able to help users finish tasks more efficiently through speech interactions







Apple Siri Microsoft Cortana Microsoft XBOX Kinect Google Now

Test Your Understanding: Review

- 1. What is the name given to the regenerating electrical wave that propagates from its initiation site at the cell body down the axon?
- 2. Which structure of a neuron receives electrical information?
- 3. What structure of a neuron transmits electrical information?
- 4. What state exists in a neuron when the concentration gradient is equal and opposite to the electrical gradient across the membrane?
- 5. What equation is used to determine the equilibrium potential for an ion?
- 6. What ion primarily establishes the resting membrane potential in vertebrates?
- 7. What ion moves rapidly into the neuron during an action potential?
- 8. What method did Hodgkin and Huxley use to experimentally control the membrane potential of a neuron in order to study ionic permeability of membranes?
- 9. What technique was developed that allowed for the study of a single ion channel?
- 10. What specialized structures link electrical synapses?
- 11. What ion is important in neurotransmission at the synapse?

Artificial Intelligence in Robotics

Objective: Applications of Neural Networks in Robotics

Agenda:

1. IBM

IBM Watson

2. Google *Google Brain Google DeepMind*3. Tour of Yuste Lab





Approaches to Build Machines

1. Classical Artificial Intelligence

- 1960: All used to solve problems that human brains perform easily
 - Understanding text
 - Recognizing objects in images
- **Result:** disappointing and slow
 - Researchers believed that this suggested that computers need vast amounts of knowledge to be "intelligent"
- Solution: "Expert System"
 - Researchers combined computer programs with rules to solve problems by asking a series of questions



IBM Watson is a modern version of Classic AI *How so?* IBM Watson uses sophisticated knowledge created for a particular issue.

What can we conclude about Classic AI? Classic AI is highly tuned to answer a specific problem with encoded rules.

Test Your Understanding

Which of the following systems mimics human thinking?

- A. Artificial Intelligence
- B. Intelligent Agent
- C. Bot
- D. Database Management System

Which AI System finds and identifies patterns, for instance, in the words you use?

- A. Expert System
- **B. Intelligent System**
- C. Neural Network
- D. Fuzzy Logic

Test Your Understanding

Which of the following systems mimics human thinking?

- A. Artificial Intelligence
- B. Intelligent Agent
- C. Bot
- D. Database Management System

Which AI System finds and identifies patterns, for instance, in the words you use?

- A. Expert System
- B. Intelligent System
- C. Neural Network
- D. Fuzzy Logic

Approaches to Build Machines

2. Simple Neural Network Approach

 1943: Warren McCulloch and Walter Pitts created the first computational model of neural networks

Artificial Neural Networks (ANN): large collection of neural units with links that can be enforcing or inhibitory in their effect on activation state of the neuronal ensemble

• Did we understand the biology of real neurons when ANN was proposed?

No: This implies that ANNs do not have much in common with biological neurons.

What did researchers aim for in that case with ANN?
 A system that can learn from data without human supervision rather than biological realism.

Artificial Neural Networks



Neural Network Architecture

Each layer consists of **nodes**, represented by small circles.

1. Input Layer: nodes are passive and do not modify the data

2. Hidden Layer:

values from the input layer are sent to hidden nodes, where they are multiplied by **weights** (pre-determined). Before leaving the hidden node, this input is passed through a *sigmoid* function.

In this way, the output can only be between 0 and 1.

3. Output Layer:

active nodes of output layer combine and modify the data to produce the final two output values of this network

Sigmoid Function in Neural Networks



$$\frac{ds(x)}{dx} = \frac{1}{1 + e^{-x}} = s(x)(1 - s(x))$$

$$= (\frac{1}{1+e^{-x}})^2 \frac{d}{dx}(1+e^{-x})$$

$$= (\frac{1}{1+e^{-x}})(\frac{1}{1+e^{-x}})(-e^{-x})$$



The sigmoid function is critical in finding the proper weights to use.

Can neural networks be made without a sigmoid function?

Yes: the three layers would *collapse* into only two layers (AND, OR, NOT). *What does this imply?*

The summations and weights of the hidden layer combine into a single layer, creating a two layer network.

Deep Learning Neural Networks

 H_3

 H_2

 H_1



Deep Learning Neural Network (DNN): an artificial neural network with multiple hidden layers of units between the input and output layers

DNN was enabled by access to fast computers and vast amounts of data for training.

Professor Geoffrey Hinton at the University of Toronto is a leader in this field and demonstrated the first use of **back propagation** in training multi-layer neural nets.

Deep Learning has successfully addressed:

- Image Classification
- Language Translation
- Identification of Spam in E-Mail

Test Your Understanding

True or False:

- The activation values of the hidden layer units in a neural network, with the sigmoid function applied at every layer, are always in the range (0,1).
- The outputs of a neural network represent probabilities and must sum to 1.
- Any logical function over binary-valued (0 or 1) inputs x_1 and x_2 can be *approximately* represented using some neural network.

Test Your Understanding

Which of the following statements are true?

- The activation values of the hidden layer units in a neural network, with the sigmoid function applied at every layer, are always in the range (0,1).
- The outputs of a neural network represent probabilities and must sum to 1.
- Any logical function over binary-valued (0 or 1) inputs x₁ and x₂ can be approximately represented using some neural network. Explanation: Since we can build the basic AND, OR, and NOT functions with the two layer network, we can approximately represent any logical function by composing these basic functions over multiple layers.

What Machine Learning Can Do

Input A	Response B	Application
Picture	Are there human faces? (0 or 1)	Photo tagging
Loan application	Will they repay the loan? (0 or 1)	Loan approvals
Ad plus user information	Will user click on ad? (0 or 1)	Targeted online ads
English sentence	French sentence	Language translation
Sensors from hard disk	Is it about to fail?	Preventive maintenance
Car camera and other sensors	Position of other cars	Self-driving cars





Source: Professor Andrew Ng at Stanford University

- Founder of Coursera
- Machine Learning Expert
- Co-Founder of Google Brain

Google Brain

hello

Google Brain is a Deep Learning Research Project led by Jeff Dean:

- Google Brain is implemented in:
 - Android OS Speech Recognition System
 - Photo Search for Google+

tim was for

Video Recommendation in YouTube

Google Brain: Robotic Arm Development

- Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection ISER, 2016
- Unsupervised Learning for Physical Interaction through Video Prediction NIPS, 2016



Deep Learning for Robots: Learning from Large-Scale Interaction

What makes robotic behavior distinct from human behavior? Robots follow a **sense-plan-act** paradigm:

1. Sense:

Robot observes the world around it and formulates an internal model

2. Plan:

Robot constructs a plan of action.

3. Act:

Robot executes its plan.

With this paradigm,

- 1. Perception is imprecise.
- 2. Models are incorrect.
- 3. No plan survives first contact with reality.

In contrast, humans move quickly, reflexively, and often with remarkably little advance planning. How?

Humans rely on highly developed and intelligence feedback

mechanisms that use sensory cues to correct mistakes and compensate for perturbations.

Learning Hand-Eye Coordination for Robotic Grasping



Google Brain: Robotic Arm Development

Google uses 14 robots to learn grasping in parallel *How does each robot develop hand-eye coordination?*

- 1. A deep **convolutional neural network (CNN)** is used to predict the outcome of a grasp, given a camera image and a potential motor command.
- 2. This CNN is then deployed on the robots the next day, in the inner loop of a servoing mechanism that continually adjusts the robot's motion is constantly predicting, by observing the motion of its own hand.

This is an example of **feedback** in a neural network.

- 3. With continuous feedback after over 800,000 grasp attempts, the robot exhibits intelligent reactive grasp behaviors.
 - Robot observes its own gripper and corrects its motions in real time.
 - Robot exhibits pre-grasp behaviors

 (*i.e.* isolating a single object from a group)

 These behaviors emerged naturally from learning, rather than

being programmed into the system.

Learning Hand-Eye Coordination for Robotic Grasping



Google Brain: Cat Recognition

Building High-level Features Using Large Scale Unsupervised Learning



TECHNOLOGY



How Many Computers to Identify a Cat? 16,000

Deep Learning Net Architecture

Deep learning aims to take machine learning to the next level, letting the program choose its own inputs.

Professor Hinton creates artificial programs that simulate the way the human brain would operate using a **4 plane model**:

- 1. Concept Plane: Inputs are fed into program.
- 2. Pattern Plane: Patterns are detected in the data.
- 3. Prime Plane: Decides how the patterns should be treated.
- 4. Action Plane: System is altered and performance analyzed.



By this method, new inputs are selected and the process repeats. This removes the need for human beings to set inputs.

"It was as if a person could suddenly cram in, say, the equivalent of five hours of skiing practice in ten minutes." **Geoffrey Hinton**

Approaches to Build Machines

3. Biological Neural Network Approach

With classic AI and simple neural networks, we cannot handle data that is constantly evolving. We are limited to finding patterns in large, **static** datasets.

What is an intelligent system?

Human Brain:

We hope that by studying how the brain works, we can learn what intelligence is and what properties of the brain are essential for any intelligent system.

What do we understand so far?

- 1. Brain represents information using SDRs
- 2. Memory is primarily a sequence of patterns
- 3. Behavior is an essential part of learning
- 4. Learning *must* be continuous

Cortical Learning Algorithm

Retina

Touch

Cochlea

High Velocity Data Stream

The **neo-cortex** builds online models from streaming data based on:

- Prediction
- Anomaly Detection
- Behavior

		•····•
Retina	Machine Learning Algorithm	
Cochlea		
Touch		$\bigwedge \land \land \land \land$

Computer

uses:

- Hierarchy
- Sequence
 Memory

Hierarchical temporary memory (HTM):

unsupervised learning method model developed by Jeff Hawkins that is based on the structural and algorithmic properties of the neo-cortex *This bio-mimetic model is based on the memory-prediction theory.*

On Intelligence



- The brain is distributed, not hierarchical.
- The goal of the brain is survival through action, not sensory prediction.
- The action-focused parts of the brain use different strategies than the perceptual system.

Google DeepMind: All Systems Go



The game of **Go** originated in China more than 25,000 years ago.

Rules: Players take turns to place black or white stones on a board, trying to capture the opponent's stones or surround empty space to make points of territory.

Go is played primarily through intuition and feel.

Games are used as a testing ground to invent smart, flexible algorithms that can tackle problems, sometimes in ways similar to humans.

Games mastered by AI include tic-tac-toe, checkers, and chess.

Google DeepMind: All Systems Go

Despite its simple rules, **Go** is more complex than chess and possesses more possibilities than the total number of atoms in the visible universe. *How so?*

Go has both a larger board with more scope for play and longer games. On average, there are many more alternatives to consider per move as well.



Google DeepMind: AlphaGo

Why is **Go** so hard for computers to play?

While traditional AI methods construct a **search tree** over all possible positions, a brute force search is intractable in Go:

1. Search space is huge

2. "Impossible" for computers to evaluate who is winning

AlphaGo is a system that combines an advanced tree search with deep neural networks.



Google DeepMind: AlphaGo



Exhaustive Search in AlphaGo



How do we reduce this search in AlphaGo? Neural networks take a description of the Go board as an input and process it through 12 neural networks.

- The **policy network** selects the next move to play.
- The value network predicts the winner of the game.

Neural Networks of AlphaGo



- The value network predicts the winner of the game.
- The policy network selects the next move to play.

Neural Network Training Pipeline in AlphaGo



Supervised learning of policy networks Policy Network: 12 layer convolutional neural network Training Data: 30M positions from human expert games Training Algorithm: maximize likelihood by stochastic gradient descent

$$\Delta \sigma \propto \frac{\partial \log p_{\sigma}(a \mid s)}{\partial \sigma}$$

Training Time: 4 weeks on 50 GPUs using Google Cloud Computing Results: 57% accuracy on held out test data (state-of-the art was 44%)

How DeepMind AlphaGo Beat Lee Sedol

AlphaGo

HE FIRST COMPUTER PROGRAM TO EVER BEAT A PROFESSIONAL PLAYER AT THE GAME OF GO.

How DeepMind AlphaGo Beat Lee Sedol



Lee Sedol is a professional Go player in South Korea. *He received \$200,000 for the match against Google DeepMind AlphaGo.*

 $\Delta \sigma \propto \frac{\partial \log p_{\sigma}(a \mid s)}{\partial \sigma} z$

Reinforcement learning of policy networks **Policy Network:** 12 layer convolutional neural network **Training Data:** games of self-play between policy network **Training Algorithm:** maximize wins *z* by policy gradient reinforcement learning **Training Time:** 1 week on 50 GPUs using Google Cloud

Results: 80% v. supervised learning.

Future of Artificial Intelligence



What do you foresee as the future of artificial intelligence?

- Climate Change Prediction
- Building Baymax
- Improving care
- Rethink how we value people



Conclusion

Tour of Laboratory of Professor Rafael Yuste

