

CMSC 373 Artificial Intelligence

Fall 2023

21-WrapUp

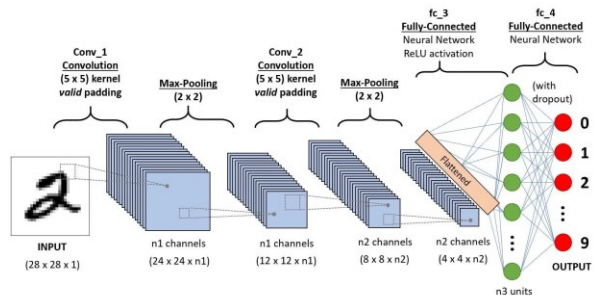
Deepak Kumar
Bryn Mawr College

1

Review for Exam 3

• Convolution Networks

- AlexNet
- Data Augmentation
- Data/Batch Normalization
- Cognitron, NeoCognitron
- Convolutions
- Deep Learning
- Dropout
- Feature Maps (Patches/Kernels/Filters)
- Image Classification
- Local patterns
- LeNet
- Maxpooling
- Object detection
- Overfitting
- VGGNet
- Visual Cortex (Hubel & Wiesel)




2

2

Review for Exam 3

• DeepQLearning

- Action
- AlphaGo Zero
- AlphaGo
- Atari
- Atari Learning Environment
- Breakout
- DeepMind
- DQN – Deep Q Network
- Episode
- Monte Carlo Tree Search
- Q Learning
- Q Table
- Reward
- State
- Unsupervised Learning
- Reinforcement Learning




3

3

Review for Exam 3

• Neural Networks for Natural Language

- Encoding
- De-Biasing
- Decoding
- Geometric Space
- GloVe
- Language Modeling
- LSTM
- NETTalk
- One-Hot Encoding
- OOV Words
- Pragmatics
- RNNs
- Syntax
- Semantic Space
- Semantics
- Sentiment Analysis
- Standardizing
- Summarization
- Text Classification
- Thought Vectors
- Tokenizing
- Vectorizing
- Word Embedding
- Word2Vec



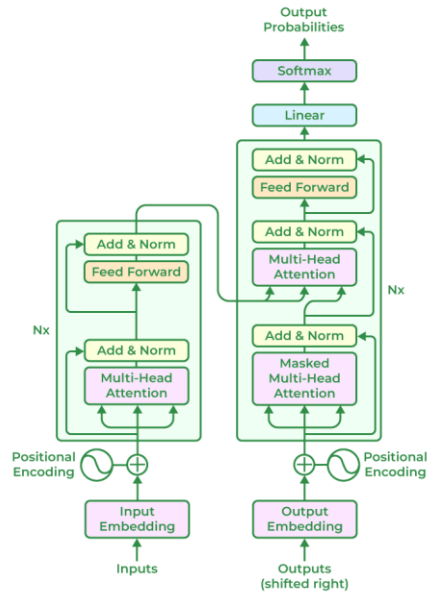
4

4

Review for Exam 3

- **Transformers**

- Assistant Model
- Attention
- Base Model
- Encoder
- Decoder
- Geometric Space
- GloVe
- Head
- Large Language Models (LLMs)
- Language Modeling
- Multi-Head Attention
- N-Grams
- Positional Encoding
- RLHF
- RNNs
- Self-Attention
- Semantic Space
- Sequence to Sequence Models
- Transformers
- Word Embedding
- Word2Vec




5

5

Review for Exam 3

- **Codes Bias**

- **The Future of AI**



6

6

Course Wrap Up

7

7

AI – History, Foundations, Applications, Implications

• History/Landmarks

- Turing Test, 1950
- Dartmouth Summer School (1956)
- Perceptron, 1958
- The seasons of AI
- Lighthill Report (1970s)
- Backpropagation (1980s)
- Japan's Fifth Generation Project (1990s)
- Deep Blue (1997)
- Google Translate (2000s)
- IBM Deep Blue (2012)
- Google's Cat recognition (2012)
- ImageNet (2012)
- ChatGPT (2022)

8

8

AI – History, Foundations, Applications, Implications

• Foundations

- Strong vs Weak AI
- Symbolic and Subsymbolic Approaches
- Search Algorithms (Heuristics. Depth-first, Breadth-First, Best-First. A*, Minimax, Alpha-Beta Pruning, MCTS)
- Knowledge Representation & Reasoning (Logic, Rule-Based, Frames, Semantic Networks, CD, Knowledge Graphs)
- Planning Systems (STRIPS)
- Expert Systems
- Forward/Backward Chaining
- Commonsense Knowledge
- Embodied Intelligence (Intelligence Without Representation, Behavioral AI)
- Agent-Based AI (rational Agents)
- Neural networks (McCulloch-Pitts, Perceptron, MLP, NN Learning, Backpropagation, Gradient Descent, SGD)
- Deep Learning (RNNs, Convolutional networks)
- Ethical AI

9

9

AI – History, Foundations, Applications, Implications

• Systems/Applications (partial list)

- ChatBots (from ELIZA to ChatGPT)
- SHRDLU, SHAKEY, STRIPS, MYCIN, R1/XCON, ZOOKEEPER, PROLOG,
- Game Playing
- GPS Navigation
- CYC
- Amazon, Wikipedia, Google's use of Knowledge Graphs
- Robots (Genghis, Cog, Boston Dynamics Big Dog, Spot, Mars Rovers)
- Alexa, Google's Assistant, Cortana
- MNIST Digit Recognition
- Keras, Google Colab
- Object Recognition
- Face Recognition

10


10

AI – History, Foundations, Applications, Implications

• Implications

- Ethical AI
- Biased AI
- Trustworthy AI
- Beneficial AI
- Future of AI
- AI Legislation

Landmarks in AI History



Wrap Up

Topics (Ambitious/Tentative List)

What is AI	Knowledge Graphs
Turing Test	Commonsense Reasoning
Symbolic versus Subsymbolic AI	Perceptrons
Narrow versus AGI	Neural Networks
ELIZA	Backpropagation
Winograd Schemas	Deep Learning
The Winters of AI	Convolution Neural Networks
Perceptron	ImageNet
Planning/Problem Solving	Bias in AI Systems
SHRDLU	Robots
SHAKEY	Reinforcement Learning
STRIPS	Deep Q-Learning
Search Algorithms	AlphaGo
Heuristics	Natural Language
Knowledge Representation &	Understanding
Reasoning	Recurrent Neural Networks
Logical Reasoning	Word2Vec
Expert Systems	Watson
MYCIN	ChatGPT
R1/XCON	Explainable AI
DENDRAL	Neuro-Symbolic Systems
CYC	Are we there yet?

13

13

ALL AI is Narrow AI

14

14

Narrow tasks are not the first step towards AGI

15

15

Easy things are hard

16

16

Hard things are harder

17

17

Wishful mnemonics are dangerous

Thought Vectors
Hallucinations
Dream Machines

18

18

Nowhere to go without embodiment

19

19

Alignment with human values

20

20

Are we there yet?

No.

21

21



22

Thank you!!!

23